

What Makes a Good Skeptical Thought Experiment?

Peter Kung, Pomona College

March 22, 2009

Under review at *Synthese*

Abstract

How do skeptical thought experiments work? The skeptic invites us to “entertain” a vividly described possibility, then argues that skeptical conclusions follow from our inability to rule that possibility out. I contend that the argument’s soundness hinges on whether the skeptical possibilities generate *legitimate doubts*: we need to separate the issue of whether we *can* rule a possibility out from whether we *have* to rule it out. I canvass a number of familiar interpretations of “possible,” including logical, *a priori*, and epistemic possibility, and argue that it is much harder than is generally acknowledged for a skeptical possibility to raise legitimate doubt.

The first time you read Descartes’ *Meditations*, or watch *The Matrix*, skeptical worries naturally flourish. You worry that you may not know the ordinary things you took yourself to know. How do skeptical thought experiments work? At first blush the procedure seems simple enough. The skeptic asks us to entertain a vividly described scenario and invites us to “feel the force” of the skeptical hypothesis. Can’t we just *see* that the possibility raises doubt about our claim to know that things are as they seem, that we need to “rule out” the scenario in order to know?

sdf

It *would* be great if we could rule out all skeptical scenarios. But even if we cannot, the fact that we cannot by itself does not show that we *have* to rule them out. It would be great if I could dunk a basketball; the fact that I cannot does nothing to establish that I have to do it. Dunking a basketball is something I don’t have to do, that I can’t do, and that I nonetheless would like to be able to do. Maybe something similar is true for ruling out skeptical possibilities. It might be something that we’d like to be able to do, that we can’t do, but unless the possibilities raise *legitimate* doubt, it isn’t something that we have to do.

In this paper I explore what it takes for a skeptical possibility to raise legitimate doubt. I'll argue that some skeptical possibilities do not raise legitimate doubt and hence do not have to be ruled out, even if we still would like to be able to rule them out. I am certainly not the first to explore this idea. Wittgenstein, Austin, Rorty, Williams, and many others argue that "skeptical doubts" are "illegitimate" and fail to raise genuine challenges to our knowledge and knowledge claims.¹ My tack will be different from theirs. Many of those authors take the notion of a skeptical *possibility* for granted; I want to focus on the sense in which skeptical scenarios are possible. Do skeptical scenarios need to be epistemically possible to raise legitimate doubt? *A priori* possible? Does a scenario's logical possibility suffice to raise legitimate doubt? Carefully attending to different ways a skeptical scenario may be said to be possible will reveal that we do not, in fact, "feel the force" of a skeptical scenario if possibility is interpreted in many of the common ways (including all the ways just mentioned). The payoff is greater precision in our understanding of what's at stake when skeptics introduce radical skeptical possibilities. While this paper does not argue for an anti-skeptical conclusion, we will see that it is harder than is generally acknowledged for a skeptical scenario to raise legitimate doubt.

Here's the plan. In section one I explain how skeptical scenarios figure in skeptical reasoning, and motivate our search for what I'll call a *possibility condition* on skeptical scenarios. In section two I explain what burden of proof the skeptic bears. With the preliminaries out of the way and the aim clear, in sections three through seven I consider and reject five candidate possibility conditions: epistemic possibility, logical possibility, *a priori* possibility, non-certainty, and subjective indistinguishability. I conclude in section eight that only a skeptical scenario's metaphysical possibility has any hope of raising legitimate doubts.

¹ See Wittgenstein (1969), Austin (1946), Rorty (1979), Williams (1995). My strategy also bears a passing resemblance to that of relevant alternative theorists like Dretske (1970) and Goldman (1976). But this is only a passing resemblance. Relevant alternative theorists propose that we are entitled to ignore irrelevant alternative possibilities simply because they are...irrelevant. By contrast I consider all alternative possibilities and show, on independently plausible grounds, that many alternative possibilities do not raise legitimate doubts.

1. Skeptical Scenarios and Skeptical Arguments

We begin by reminding ourselves what skeptical scenarios are and how they figure in skeptical arguments. Let *O* be a proposition that subject *S* putatively knows.² A skeptical scenario describes a situation in which *S* forms the belief that *O* despite the fact that *O* is false.³ Not just any story will do. The scenario must meet an *explanatory constraint*: it must “satisfactorily explain” how *S* comes to have her evidence, and hence her belief that *O*, even though *O* is false.⁴ It’s not easy to articulate what counts as a “satisfactory explanation,” but for our purposes that won’t matter. I’ll assume this constraint is met.

A “satisfactorily explanatory” (and by now familiar) example comes from *The Matrix*. In *The Matrix* scenario, machines have enslaved human beings and keep them in blissfully ignorant submission by connecting them to an enormous computer network, the Matrix, which runs an almost perfect simulation of late twentieth-century earth. Humans attached to the Matrix have experiences as of, and hence think they are in, late twentieth-century society, but in reality they are floating in pods with the Matrix feeding signals directly into their brains. Now take my claim to know that I have shaken hands with Justin Timberlake. In the Matrix scenario I falsely believe that I have shaken hands with Justin Timberlake: I’ve been floating in my pod my whole life, and have never touched another human being.

Skeptics use skeptical scenarios to support key premises in their arguments. Take the closure version of the skeptical argument.

- I. You do not know that the skeptical scenario does not obtain.
- II. If you do not know that the skeptical scenario does not obtain, then you do not know that *O*.
- III. You do not know that *O*.

² In this paper I focus knowledge and skeptical challenges to knowledge claims, but all considerations should apply to justification as well.

³ One can also offer a skeptical challenge using a scenario in which *O* is true but the way *S* forms the belief that *O* is inconsistent with *S* knowing that *O*. Consider my belief that I have hands: in the Matrix scenario, human beings still have hands, so my belief is true. Yet many concede that in the scenario I don’t know that I have hands. This second type of skeptical scenario won’t make a difference to our discussion of possibility, so for simplicity we’ll consider only skeptical scenarios in which *O* is false.

⁴ For a discussion of why the explanatory constraint is critical, see, e.g., Cross (forthcoming). Thanks to Jim Kreines for pointing me to Cross’s paper.

A well-chosen skeptical scenario makes premise I) plausible, and it does so by raising a *possibility* — the possibility described by the scenario — that you are allegedly unable to rule out. Let's put this thought in terms of a *possibility condition*.

(PC) a skeptical scenario must be possible.

The upshot is that a skeptical scenario's possibility raises doubt that you need to assuage if you are to know that O. This way of posing the skeptical challenge, in terms of ruling out possibilities, occurs throughout the skepticism literature.⁵

Given the way I intend us to understand “skeptical challenge,” it is obvious that successful skeptical challenges can be raised against our knowledge claims. Defense lawyers do it all the time. Suppose we're in a café and I claim to know that Mary-Kate has been sitting in three tables over for the last twenty-five minutes. You know that my evidence is visual. You point out that her twin sister Ashley lives nearby. Perhaps it's Ashley not Mary-Kate sitting three tables over? Your scenario is a possibility that raises doubt about my knowledge claim. In this case you have reason to believe the possibility is actually true. But I don't assume that evidence of actuality is a general requirement. Our target question is this: how should we understand the possibility condition? What sort of possibility is needed to raise legitimate doubt and drive a successful skeptical challenge? Is it epistemic possibility? Logical possibility? Metaphysical possibility? We are interested in finding the weakest condition that is, *ceteris paribus*, sufficient to raise legitimate doubts. That will help us identify the strongest skeptical challenge.

Lest the reader think that nothing hangs on the issue, a number of prominent epistemological views seem to *assume* a metaphysical possibility condition on skeptical scenarios. I have in mind sensitivity, safety, relevant alternatives, and some contextualist views, views analyzed in a possible worlds framework.⁶ Relevant alternatives are those

⁵ To select several examples more or less at random, “ruling out possibilities” talk occurs in introductory epistemology essays like DeRose (1999) and BonJour (2002) and recent books for specialists like Conee & Feldman (2004), Greco (2000), and Pritchard (2005). Contextualists (DeRose 1995, Lewis 1996), sensitivity (Nozick 1981), safety (Sosa 2000), and relevant alternative theorists (Dretske 1970, Goldman 1976) also routinely talk this way. These views have many contemporary proponents.

⁶ See the references in the previous note. I myself do not endorse any of these views; I am merely pointing out that they presume a certain answer to the question of what sort of possibility is needed to raise legitimate doubt.

that are true in “nearby” possible worlds; skeptical hypotheses are irrelevant because they are true only in “remote” possible worlds. It’s most plausible that the worlds in question are metaphysically possible worlds. Sensitivity and safety are analyzed with subjunctive conditionals, and those conditionals are usually glossed with possible world semantics. To determine whether I know that I have hair, for example, the sensitivity view tells us to evaluate the conditional, “If I were to lack hair, I would not believe that I had hair.” The conditional’s truth depends on whether, in the nearby possible worlds where I lack hair, I believe that I lack hair. Again, these are metaphysically possible worlds. To drive this point home, consider how these views explain the barn country case (Goldman 1976). These views hold that Henry’s true belief that there is a barn in front of him does not count as knowledge because, in a nearby possible world, Henry is looking at a barn façade; there, his belief is false. For this analysis of the barn case to make sense, nearness had better not be a function of Henry’s subjective state. It’s a metaphysical fact that holds regardless of Henry’s beliefs or evidence. Nearness and remoteness are metaphysical notions.

These considerations show that at least some philosophers assume that skeptical scenarios must be metaphysically possible. If we suspect that something less than metaphysical possibility will serve the skeptic’s purpose equally well, and hence that there’s a problem with views that rest on the assumption, then that should be made explicit. This paper is an explicit discussion of the possibility condition.⁷

Let me be clear about what I am *not* doing. I am not offering an anti-skeptical argument, nor am I offering a theory of knowledge or justification that explains why, e.g., perception confers knowledge or justification. My concern is with the conditions a skeptical scenario must meet to generate legitimate doubt. For all that I say here, there may very well be skeptical scenarios that meet the possibility condition that we are unable to rule out. So this paper does not have an anti-skeptical conclusion. Nonetheless, by making clear how strong the skeptical challenge is, I hope to make clear what burdens anti-skeptics need and need not meet.

⁷ A recent discussion of the possibility condition is Beebe (forthcoming). I discuss Beebe’s arguments at several points below.

2. Burden of Proof

We will need some ground rules to evaluate various candidates for the possibility condition, ground rules about who has the burden of proof. A background assumption is that our investigation takes place in the context of what Pryor (2000) calls the “modest anti-skeptical project,” the project of assuring ourselves using premises that we find plausible that we have knowledge or justification. If I can make it plausible, for example, that a skeptical scenario’s mere logical possibility raises no legitimate doubt, then I take myself to have shown that the possibility condition demands more than logical possibility. It is always open to the skeptic to insist that, nonetheless, we must rule out any logically possible scenario. But doing so would not be required for the modest anti-skeptical project to succeed.

Let me raise another important burden-of-proof issue by examining an argument by Beebe (forthcoming). Beebe questions whether in a skeptical scenario it must be metaphysically possible for O to be false. He argues no: skeptical scenarios can be metaphysically impossible.

The falsity of [the metaphysical possibility condition] 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 [the metaphysical possibility condition], 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. (Beebe forthcoming, pp. 5–6)

I agree that it’s absurd to think that a skeptical challenge can be raised only against the necessarily false belief. But I don’t believe Beebe has successfully shown that skeptical scenarios needn’t be metaphysically possible.

If a skeptic brings to our attention a legitimate doubt about our claim to know that O, then even the modest anti-skeptical project requires that we silence that doubt. But what should we make of “legitimate doubts” that no one is aware of? It seems that for

us theorists to be justified in believing that a skeptical scenario raises legitimate doubts, we have to at least be aware that the scenario meets the possibility condition, whatever that condition turns out to be. To see this, let's reconstruct the theist's skeptical challenge to the atheist as follows:

1. It is metaphysically possible that God exists.
2. If it is metaphysically possible that God exists, then there are legitimate doubts about atheism.
3. There are legitimate doubts about atheism.⁸

Let's grant premise 2) for the moment. Even if we grant premise 2), to be justified in believing conclusion 3), you have to be justified in believing premise 1). If 1) is true but nobody is justified in believing it, then it is hard to see why 1) should raise doubt in any actual believer's or any theorist's mind. What matters is whether the atheist and theist skeptics are justified in believing that it is metaphysically possible that God does or does not exist. If neither modal belief is justified then neither has successfully raised a skeptical challenge against the other.

Can the skeptic retreat to a weaker claim about premise 1), something like, "it cannot be ruled out that 1) is true" and still raise legitimate doubt? I think not. First, that move would arguably transform the metaphysical possibility requirement into something weaker, perhaps an epistemic or non-certainty possibility requirement, which we will discuss below. Second, it is implausible. Suppose Yuval the Skeptic peers thoughtfully at the argument above and then declares, "Ah, but written down on a parchment in this box is a scenario that may or may not show that it is metaphysically possible that God exists." The atheist can't rule out that Yuval's parchment establishes 1). Yet I take it so far Yuval has raised no legitimate doubt about whether atheism is true; it would be very strange for the atheist to withdraw belief in response. Absent some further argument, there is no reason to take failing to rule out the existence of legitimate doubts to itself raise a legitimate doubt. The skeptic's retreat to a weaker version of premise 1) fails.

⁸ Compare to the Academic Skeptical argument in Graham (2007, p. 29).

The lesson of the theist/atheist standoff is that we need some kind of luminosity constraint on the possibility condition.⁹ We'll state the general version of the luminosity constraint in terms of "appreciation": whatever the possibility condition on skeptical scenarios amount to, to raise legitimate doubt we must be able to *appreciate* that a scenario meets that condition. What appreciation amounts to will vary depending on the possibility condition. As we just saw, for a metaphysical possibility condition, appreciation amounts to justification: you need to be justified in believing that the scenario is metaphysically possible. If the possibility condition is itself epistemic, for instance an *a priori* justification possibility condition, then mere awareness might be enough. Perhaps you don't also have to be justified in believing that it is *a priori* possible.¹⁰

We've arrived at two burden-of-proof ground rules. First, we're holding ourselves only to the modest anti-skeptical project. Second, whatever the possibility condition on skeptical scenarios amount to, to raise legitimate doubt we must appreciate that a scenario meets that condition.

3. Epistemic Possibility

Let's start our tour through the candidates for the possibility condition with *epistemic possibility*. This isn't really a single candidate but a family of them. What makes them a family is that they all invoke an epistemic notion on the right hand side: P is epistemically possible iff S does not know/is not justified in believing/has no reasons to believe/... that P is false.¹¹

⁹ Although I'm borrowing his term, this isn't Williamson's (2000) definition of luminosity. His luminosity requires knowledge; I'm allowing luminosity to be something weaker than knowledge.

¹⁰ I aim to avoid complicating the paper with internalism vs. externalism issues however the luminosity constraint is one place it does arise. If you have a pure externalist epistemic possibility condition, then the luminosity constraint will be more demanding. That's because a pure externalist epistemic property is one that you can possess even if you are in no way aware that you possess it. The way I'm understanding "raising legitimate doubt," doubts are the sort of thing you must be aware of.

¹¹ For different ways to construe epistemic possibility, see the literature on concessive knowledge attributions (sentences like "I know that Christopher Dodd won't be the 2008 Democratic nominee, but it is possible that he will be"): DeRose (1991, 1999), Hawthorne (2004), Rysiew (2001), and Stanley (2004). The label "concessive knowledge attributions" is due to Rysiew (2001).

One might think that epistemic notions are uniquely suited to raise skeptical challenges. But just the opposite is true. An epistemic possibility condition renders skeptical scenarios impotent.

Let's work with a concrete case, our Matrix scenario. The proposition you claim to know, O, is that you have shaken hands with Justin Timberlake. The skeptical scenario is that you're in the Matrix; you've been floating in your pod your whole life and have never touched another human being. And let's use a knowledge version of the epistemic possibility condition: the Matrix scenario is true for all you know.

Recall our aim: we are looking for the weakest possibility condition that is, *ceteris paribus*, sufficient to raise legitimate doubts. The luminosity constraint requires that we must be able to appreciate that the scenario meets the possibility condition. This means, in the present case, that the Matrix scenario is epistemically possible, that we can appreciate that it is epistemically possible, and that it being epistemically possible raises legitimate doubt about your claim to know that O.

The first issue, then, is whether the Matrix scenario is epistemically possible. Here we immediately encounter a problem. No anti-skeptic should grant that skeptical scenarios are possible in this sense.

To claim that the Matrix scenario is epistemically possible is to claim that a subject lacks [insert epistemic property here] with respect to the scenario; in our example, the epistemic property is knowledge. You shouldn't concede that it is true for all you know that you are in the Matrix; since the Matrix scenario is one in which O is false, that is tantamount to conceding that you do not know that O. That's precisely what is at issue. If you were ignorant of whether you are in the Matrix, the issue whether there are legitimate doubts about O has already been settled: yes, there are, that's why you are ignorant! Someone asserting that you are ignorant whether you are in the Matrix does nothing *by itself* to raise legitimate doubts. Remember the skeptical thought experiment was supposed to *convince* us that there are legitimate doubts about O, not start with it as a

premise.¹² Hence the anti-skeptic should simply deny that skeptical scenarios are epistemically possible.

It won't help for the skeptic to reply that what's crucial is whether the scenario is *in fact* epistemically possible and not whether the anti-skeptic is willing to concede that it is epistemically possible. The reply would, in effect, deny the luminosity constraint, so it would inherit the implausibility of denying luminosity: If the Matrix scenario is in fact epistemically possible though we theorists can't appreciate that it is, it is hard to see why it should raise any doubt. And there's a further implausible consequence of denying luminosity: if the scenario is epistemically *impossible* though we theorists can't appreciate that either, then as Beebe (forthcoming) points out, "no skeptical challenge could ever be raised against a belief that [in fact] counts as knowledge." If you in fact know that you have hands, then by the condition, the Matrix scenario cannot raise any doubts. But that seems wrong. "My knowledge of [a] fact may aid me in responding to... challenges, but it cannot prevent... challenges from being raised" (p. 18). We should be able to raise legitimate doubts about things that you in fact know, particularly if don't have the very highest confidence in those things.¹³

Epistemic possibility fails to be an adequate interpretation of the possibility condition. Let's turn to some others.

4. Logical Possibility

Next we'll consider *logical possibility*.¹⁴ Our evidence leaves open the question whether O in the sense that it is logically possible for O to be false, or, if this turns out to be a different claim, logically possible that we don't know that O . I'll make two points in this section. First, I'll explain in general terms why the logical possibility is insufficient to raise legitimate doubt. Second, we'll look at another argument of Beebe's that logical possibility isn't even necessary, and see that it doesn't quite succeed.

¹² Pryor (2000, pp. 522–23) makes a similar point.

¹³ See also Lewis (1996, p. 552–53).

¹⁴ Schiffer (2004) endorses this constraint.

4.1. Wouldn't It Be Great?

Let's take a very general look at what the logical possibility condition demands. Our diagnosis of epistemic possibility should make us suspicious from the outset. If a skeptical scenario being true for all you know fails to raise doubt, it is hard to see why a skeptical scenario being true for all you know via logic should raise doubt.¹⁵ Let's see why it doesn't.

Consider a commonsense "open question" gloss. Does the fact that a matter is not settled by logic alone show that the question is left open in any interesting sense? Does it raise any doubt about whether the question can be answered? It should only if either there is no other way to settle the matter or if there is some special reason to think that it ought to be settled by logic alone. The skeptic supplies no special reason for the latter and is not entitled to assume the former.

To understand why, we need to consider what logical possibility is and what it is not. Van Inwagen (1998) points out that so-called logical possibility isn't really a distinct "kind" of possibility.

But there is no such thing as logical possibility — not, at least, if it is really supposed to be a species of possibility. Belief in the reality of "logical possibility" may be based, at least in part, on a faulty inference from the reality of logical impossibility, which is real enough. Logical impossibility is an epistemological category: the logically impossible is that which can be seen to be impossible on the basis of logical considerations alone — or, to be liberal, logical and semantical considerations alone... What I dispute is the contention that if a concept or state of affairs is not logically impossible, then it is "logically possible." It hardly follows that, because a certain thing cannot be proved to be impossible by a certain method, it is therefore possible in any sense of 'possible' whatever. (p. 71)

Suppose we think of logic as something we use to establish the truth or falsity of claims. Some things we can establish are true in this way; other things we can establish are false (and hence impossible) in this way; van Inwagen's example of the latter is round squares. It might be true that we cannot establish by logic and semantic considerations alone that, given our experiential evidence, skeptical scenarios are false. To this the anti-skeptic

¹⁵ Thanks to Masahiro Yamada for putting the point this way.

should simply respond: Perhaps true, but why should this be a concern? Why should this raise any doubts in my mind? Just as it would be great if I could dunk a basketball, though alas I cannot, it would of course be nice if we could rule out skeptical scenarios using logic alone; our perceptual knowledge would then be on a reassuringly firm logical foundation. But the fact that we cannot shows only that one way of establishing facts is unable to determine whether skeptical scenarios are true or false. Unless we are given some reason to think that we *should* be able to establish the falsity of a skeptical scenario using logic alone, that there is some reason to privilege logic here, the mere logical possibility of a skeptical scenario by itself has no skeptical force whatsoever.

This diagnosis — “it would be great if we could rule out skeptical scenarios [by way W], but the fact that we can’t raises no doubts about O” — will recur. Let’s call it the it-would-be-great problem. We’ll see the it-would-be-great problem in our next candidate, *a priori* possibility. But before we turn to *a priori* possibility, let’s consider whether logical possibility is necessary to raise legitimate doubt.

4.2. Is Logical Possibility Necessary?

Logical possibility isn’t sufficient to raise legitimate doubt. But is it necessary? Beebe (forthcoming, §IV) argues for a strong conclusion: effective skeptical scenarios can be logically impossible. Beebe’s springboard is Wittgenstein’s contention that other ways of inferring or calculating are, in some sense, possible, even if we can’t fully understand what those ways are. He takes Wittgenstein’s contention to ground a skeptical scenario in which we possess deviant *a priori* belief-forming mechanisms.

“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 consists in”(p. 14).

Let’s call this the miscalculation scenario, and the possibility it presents the miscalculation possibility. Beebe contends that “since there is an incompatibility between my having genuine knowledge” of *a priori* matters and the miscalculation possibility, we “must be in a position to rule out this possibility” if we are to have *a priori* knowledge. But we can’t rule out the miscalculation possibility. If we were in the miscalculation

scenario, we would falsely believe that we weren't; our intuitive evidence would be exactly similar to what it is now. So we have no *a priori* knowledge.

Perhaps the miscalculation scenario is a good one. But even if it is, it does not show that successful skeptical scenarios can be logically impossible. We need to remember the lessons of section 2 about burden of proof and pay attention to what it takes to rule out a possibility.

Take the belief that $2+3=5$. The miscalculation scenario is supposed to be one in which our most compelling intuitions about calculating are mistaken; it's supposed to be true that $2+3\neq 5$. But it continues to seem very clear to us that in the scenario $2+3=5$. Hence I take myself to reasonably believe that the miscalculation scenario is impossible. It looks like, regardless whether the scenario raises legitimate doubt, we need not worry about it.

Beebe anticipates this move.¹⁶ He notes that the belief that the miscalculation scenario is impossible because $2+3=5$ is an *a priori* belief, and that's the exactly the belief that the skeptic is challenging. You can't appeal to *a priori* beliefs to block an *a priori* skeptical challenge that questions those very beliefs.

There's something enticing about Beebe's response. We'd really like to be able to rule out the miscalculation scenario without resting on the very beliefs that the scenario is intended to call into question. However what Beebe is doing is in effect insisting on the ambitious anti-skeptical project. It would be great if we could vindicate our *a priori* knowledge without availing ourselves of those *a priori* things we find overwhelmingly plausible. Alas we can't. But the modest anti-skeptic will insist that if we find ourselves satisfied, upon careful reflection, that a scenario is false (e.g., because we find it implausible that $2+3$ could be something other than 5), then that's a perfectly respectable modest anti-skeptical response.

Notice that this type of response doesn't generalize; we couldn't use it to respond to external world skeptical scenarios. Although commonsense tells us, for example, that we have hair, we have no problem entertaining a Matrix scenario in which we are bald.

¹⁶ Beebe (forthcoming, pp. 15–16).

When we consider the Matrix scenario, we don't "import" our commonsense belief that we have hair and judge that, even in the Matrix, we have hair.¹⁷ But that's exactly what we do with our simple arithmetical beliefs. When we consider the miscalculation scenario, it continues to seem to us that *in the scenario* $2+3=5$. Although it's supposed to be true in the scenario that $2+3\neq 5$, when we ask ourselves what $2+3$ really is in the scenario, we still come up with the answer 5.¹⁸

Notice also that the response leaves Beebe in an unstable position. He alleges that the miscalculation scenario is an example of a good though logically impossible skeptical scenario. But in the enticing response he seems to suggest that, given the dialectical context, we aren't entitled to believe that the miscalculation scenario is logically impossible. If that's true, then Beebe can't claim to have offered a counterexample to the logical possibility condition. Is it enough that we are not entitled to believe that it *is* possible? No, for then we fail the luminosity constraint: we're not in a position to appreciate that the scenario meets the possibility condition.¹⁹

I've argued that the logical possibility condition is incorrect; logical possibility is insufficient to raise legitimate doubt. That argument leaves it open that logical possibility is still necessary. Beebe argues that it isn't necessary, that logically impossible scenarios can raise legitimate doubt. We've just seen that Beebe's argument does not establish this. Logical possibility isn't the weakest condition sufficient to raise legitimate doubt, but it may still be necessary.

5. *A Priori* Possibility

There's a theme emerging here: while it may be true that a skeptical scenario cannot be ruled out in a particular way, that on its own does nothing to raise legitimate doubt,

¹⁷ The import/export language is from Gendler (2000).

¹⁸ It's this fact that seems to lead to "imaginative resistance" when we try to imagine counter-arithmetical facts. See Nichols (2006a).

¹⁹ There's still the question whether the miscalculation scenario raises legitimate doubt. I don't see why it should (in the context of the modest anti-skeptical project). The problem is related to the one that distinguishes these *a priori* skeptical scenarios from external world ones. In the counter-arithmetical scenario, we don't understand what it would be for the scenario to *really* be true. We can say the words, but regardless 2 and 3 continue to seem to sum to 5, even *in the scenario*. That's because, as Beebe himself concedes, possibilities like the one presented by the miscalculation scenario "may not be real possibilities for creatures like us and *they may not be fully intelligible to us*" (p. 13, italics mine).

unless the skeptic provides an argument that we *have* to rule it out that way. This it-would-be-great diagnosis applies equally well to our next candidate, *a priori possibility*.

And once again our diagnosis of epistemic possibility should make us suspicious from the outset. If a skeptical scenario being true for all you know fails to raise doubt, it is hard to see why a skeptical scenario being true for all you know *a priori* should raise doubt. Let's take a closer look.

P is *a priori* possible iff it is not a priori that not-P. The *a priori* possibility condition would then be something like: a (satisfactorily explanatory) scenario raises legitimate doubt iff it is not *a priori* that the scenario is false. We have to be careful to separate the issue of what it takes for a scenario to *raise* legitimate doubt from what we have to do to *rule out* a legitimate doubt. It's important to take the first issue first: does the fact it is not *a priori* that a scenario is false *raise* legitimate doubt?²⁰ To this question the anti-skeptic should again give the it-would-be-great response: It would be great if we could show *a priori* that skeptical scenarios were false. But why does the fact that we cannot raise legitimate doubt? Absent a reason to think that it *ought* to be *a priori* that skeptical scenarios are false, the fact raises no legitimate doubt.

Now I can imagine some readers who went along with the it-would-be-great reasoning for logical possibility putting on the breaks here.²¹ "Whoa. Look, this condition says that you can't rule out the scenario *a priori*. But you can't rule it out *a posteriori* either because intuitively, doing so would be circular. The circularity problem is what's different about this *a priori* case vs. the earlier non-certainty and logical possibility cases. Hence you have no way of ruling out the scenario, period."

This objection concerns the *second* issue, about what we can or cannot do to rule out a scenario. It says nothing about the *first* issue, whether we need to rule the scenario out in the first place. We need to worry about ruling the scenario out only if it raises legitimate doubt; the it-would-be-great response says it does not. The ambitious anti-skeptical project might demand that we (non-circularly) rule out every scenario. But it

²⁰ It's easy to see why we often confuse the two issues. *A priori* possibility is sometimes expressed as "P cannot be ruled out *a priori*." But the fact that a scenario cannot be ruled out *a priori* does not by itself show that it has to be ruled out at all.

²¹ Thanks to Yuval Avnur for pressing me to be clearer here.

seems within the bound of a modest anti-skeptical project to require only that we rule out scenarios that raise legitimate doubt.

Well, wouldn't it be weird to concede that, even if a skeptical scenario raised no legitimate doubt, it couldn't *somehow* be ruled out? Perhaps that would be weird, but there is a response available here, one that, admittedly, not everyone will find convincing. It's a Moorean response. If a scenario raises a legitimate doubt, then appealing to facts that the scenario calls into question to assuage those doubts would be blatantly circular. However once it is conceded that a scenario raises no legitimate doubt, then *that* circularity worry, at least, disappears. One might be able to appeal to *a posteriori* considerations plus deduction to conclude that the scenario does not obtain, since there has been no legitimate doubt cast on the *a posteriori* considerations. There remain, of course, a host of other issues to address concerning the Moorean response before we can deem it successful.²² But it is an option that remains on the table.

Regardless of the viability of the Moorean response, I conclude that *a priori* possibility fails to be sufficient to raise legitimate doubt.

There is a second way to understand the *a priori* possibility condition that might seem worrisome, one that raises questions about the connection between our putative experiential evidence E and P's falsity. P is *a priori* possible iff no connection between E and the falsity of P can be established *a priori*

This is a theorist's concern, and not a subject's concern. We theorists have to establish a connection between E and the propositions E is supposed to be evidence for. In the present philosophical context it's extremely unlikely that we theorists will be using empirical methods to do so. That leaves *a priori* methods. *If* we were incapable of showing *a priori* that E provided evidence for not-P, then that would raise legitimate doubts about E's evidential value. Hence if a skeptical scenario is *a priori* possible in this sense, it follows that we have no evidence against it.

Notice, however, that an anti-skeptic's metaepistemological project just *is* to explain why experience provides evidence (or justification, or confers knowledge; I'll

²² I attempt to address some of them in Kung & Yamada (2009a, 2009b) and Kung (2009).

suppress the parenthetical from now on). That is an *a priori* metaepistemological project. Because the anti-skeptic purports to have *a priori* arguments that perceptual experience confers justification, she will charge that a skeptic's assertion that there is no such argument simply descends from questions about legitimate doubt to the question of whether experience provides justification at all.

That last question is a perfectly legitimate one, and an anti-skeptic must have an answer to it, but it is not the question we are addressing here. Recall what I said I am not doing: I am not offering an anti-skeptical argument, nor am I explaining why perceptual experience or intuitions provide evidence. I conceive the dialectic this way: the question about experience providing evidence is the more fundamental one, and in developing their theories, anti-skeptics attempt to give *a priori* arguments that experiences *does* provide evidence. Of course skeptics should and do object at that point. But even if we grant for the sake of argument the anti-skeptics have done their job, skeptical worries remain. It is still possible to raise doubt about the things that (we are granting) we justifiably believe, and those doubts might make us realize, despite the starting position we granted to the anti-skeptic, that we cannot be justified after all. This is simply to say that the second way to understand the *a priori* possibility condition raises a good question, but it is not a question about skeptical scenarios' ability to raise legitimate doubt.

6. Non-Certainty

Let us now turn to *non-certainty*. Philosophers rarely endorse a “non-certainty” condition because it is clearly too weak.²³ A non-certainty condition is most plausibly understood as a catch-all for one of the other conditions (or some combination of conditions) we have already examined. But let's take a look at this condition in its own right.

The non-certainty condition would be something like the following: Even if you have an answer to a question, there may be some sense in which the question is left open if there is any conceivable way in which you might be mistaken. Unless you are absolutely certain of the answer, the question is, in some sense, still open, and, the thought goes,

²³ An exception is Unger (1975). Remember a *weak* possibility constraint makes it very *easy* to raise legitimate doubts, and hence generates a very *strong* skeptical argument.

that raises some legitimate doubt about the answer. For example, it is possible in the present sense that I am a robot. I've got lots of evidence that I'm not a robot and I don't believe that I am a robot. I might even know that I am not a robot. But I'm not absolutely certain of it. Because I can't rule out that possibility for certain, there is — allegedly — some nagging doubt about whether I am a human being. (Let's assume that being a human being entails not being a robot, although that is also something I'm not completely certain of.)

Fallibilists will of course reject the non-certainty condition. But we don't want to assume fallibilism. We want to know whether non-certainty that about a scenario raises legitimate doubt about O.

As our open question gloss shows, all the skeptic points out is that there is a “conceivable way” for O to be false. Insofar as we can make sense of this, it is hard to see why this fact alone should shake our confidence. There is very little that is absolutely certain. I take myself to know that π is irrational. We have a proof of this fact; I've gone through the proof. But I am not absolutely certain of it. There are things I am more certain of; that π is irrational is not as evident to me as the fact that $5+7=12$. There “could” be a subtle flaw in the proof that π is irrational, “could” in the sense that I can at least entertain the thought that I am mistaken, even though I am confident that I am not. I can tell an intelligible, not blatantly contradictory story in which I am mistaken.²⁴ Of course I can't specify any of the relevant details in my story: my story won't say what the flaw *is*. Without that detail it's hard to see why the scenario constitutes a reason to doubt: I still have no clue about a possible flaw in my reasoning. What the scenario might make apparent that I am not in the very best epistemic position with respect to my conclusions. But failing to be in the very best epistemic position does not, by itself, constitute a reason to doubt one's conclusions. It would be great if we were absolutely certain that skeptical scenarios are false, but unless we are given reason to think we *should* be absolutely certain, the fact that we are not does nothing to raise skeptical worries.

²⁴ By contrast I am unable to tell an intelligible, not blatantly contradictory story in which $5+7\neq 12$. Gendler's (2000) Tower of Goldbach fable is an intelligible but blatantly contradictory story in which $5+7\neq 12$ that she claims we can nonetheless imagine. Many of her readers disagree.

7. Subjective Indistinguishability

The final candidate we will examine is *subjective indistinguishability*; we will be able to deal with this candidate fairly quickly because its plausibility depends on previous arguments that we have rejected above.

Subjective indistinguishability is the intuitive and familiar idea that a skeptical scenario is, subjectively speaking, identical, or at least very similar, to our actual situation. Things *seem* the same to us in the skeptical scenario, though crucial facts are different. We don't have our usual connection to the facts. O is false. Let's grant straightaway that something like subjective indistinguishability is a necessary condition on a successful skeptical scenario. Arguably it should be thought of as part of the explanatory constraint. To "satisfactorily explain" how subject S comes to have her evidence will involve a story about why S's experiences in the scenario match the ones she is actually having.

The more pressing question for our inquiry is whether subjective indistinguishability alone suffices to raise legitimate doubt. The answer is pretty clearly no.

Lewis defines what it is for a possibility to be subjectively indistinguishable as follows: a possibility W is uneliminated for S iff S's experiences and memories in W match S's experiences and memories in actuality.²⁵ First, notice that Lewis's definition explicitly notes that indistinguishable possibilities are *uneliminated*; it says nothing about whether they *have* to be eliminated. Mere indistinguishability doesn't speak to the latter issue. Second, the definition as stated says nothing about what the relevant type of possibility is. We still need to know the sense in which W is possible. Does W range over epistemically possible worlds? Logically possible worlds? Metaphysically possible worlds?

²⁶ Answering that question will take us back to the issues we discussed in previous sections.

Third, if W is understood to range over logically and metaphysically impossible worlds as well as possible worlds, as Beebe (forthcoming, §VI) intends, then we will

²⁵ Lewis (1996, p. 424–25). See also Beebe (forthcoming, §VI).

²⁶ Thanks to Masahiro Yamada for pointing out this avenue of response. Lewis himself clearly intends metaphysical possibility. But it isn't something that he argues for, and some skeptics will want a wider range of worlds.

encounter the problem we faced in section 4.2 with logically impossible skeptical scenarios. You can consider a scenario in which you have been created with faulty arithmetical intuitions, and though it seems to you that $2+3=5$, in fact $2+3\neq 5$, however it is hard to see why that scenario should raise any doubt in your mind about whether $2+3=5$, and in any case it seems clear that you can rule out the scenario.

Can the skeptic adopt this slightly weaker position: “Subject indistinguishability suffices to raise legitimate doubt unless the scenario is obviously (or substitute some other adverb) impossible. The $2+3\neq 5$ scenario is obviously impossible, but other scenarios aren’t, and those other scenarios, provided they are subjectively indistinguishable, do raise legitimate doubt.” This is a fairly common reaction to the arguments in this paper.²⁷ But it is not compelling because the skeptic amounts to an assertion that the scenario cannot be ruled out. The subjectively indistinguishable scenario is one that cannot be ruled out via logic, or via arithmetic, or with certainty, or via whatever other methods the skeptic is willing to allow. Hence it does not address the target issue in this paper; it does not *identify* a property that skeptical scenarios have that *enables* them to raise legitimate doubt. It points out (or rather asserts) that we cannot rule out the scenario without explaining why we *have* to rule it out.

I conclude that while subjective indistinguishability is a necessary condition on skeptical scenarios, it is not sufficient to raise legitimate doubt.

8. Concluding Thoughts on Metaphysical Possibility

It is time to sum up. We set out to investigate the sense in which skeptical scenarios have to be possible. Our aim was to find the weakest possibility condition that is sufficient, *ceteris paribus*, to raise legitimate doubts, and we established that raising legitimate doubt requires meeting the luminosity constraint. We have looked at a number of candidates — epistemic possibility, non-certainty, logical possibility, a priori possibility, and experiential possibility — and found them all deficient.

We’re now in a position to identify a common problem with our candidates. They each try to define the possibility condition in *epistemological* or *psychological* terms, in

²⁷ Thanks to audiences at the CSPA and the Claremont Colleges WIP, in particular Jim Kreines, for pressing this line of response for the skeptic.

terms of *methods* (logical proof, *a priori* reasoning), *epistemic properties* (justification, reasons, knowledge), or *psychological states* (certainty, subjective indistinguishability). Those epistemological notions seem ill-suited for explaining when a scenario should raise doubts because explanations involving them will either raise the very issue we were debating or leave us wondering why the inability to rule out the scenario in the proposed way is supposed to raise legitimate doubt. They seem to confuse what it takes to rule a legitimate doubt out with what it takes to raise a legitimate doubt in the first place. Attempts to use hybrids or combinations of these epistemological or psychological notions will similarly founder.

This suggests that we should look to a *metaphysical* notion rather than an epistemological one. And in fact the most promising proposal is metaphysical possibility: if a skeptical scenario is metaphysically possible — if it really *could be true*, period, and not just could be true for all we know *a priori*, or for all we know for certain — that seems to give us some reason to worry (provided we meet the luminosity constraint). It's this thought that seems to motivate a number of authors to take external world skeptical scenarios (as opposed to the *a priori* variety) seriously. Not only is “[t]he hypothesis that all of our present experiences are the deceptions of an evil demon... not absurd,” writes Pryor,

[i]t seems to be a genuine metaphysical possibility. So we can't reject that hypothesis out of hand. If we do know that we're not being deceived by an evil demon, it's plausible that that knowledge would have to rest on things we know about our environment on the basis of perception. (p. 524)

BonJour explicitly requires that skeptical scenarios be “genuinely possible”:

The versions of skepticism in question [that reply on skeptical hypotheses] are committed to the positive claims (a) that the hypotheses in question are genuinely possible, and (b) that all of the various relevant sorts of evidence could have existed in the same way even if the skeptical hypotheses were true, with both of these claims presumably being alleged to be established on an *a priori* basis. (p. 262)

If a skeptical scenario is another way the world *could* be, then there are two genuine alternatives, two metaphysically possible hypotheses about how the world is, the external

world hypothesis and the skeptical hypothesis. If you appreciate that your evidence is *metaphysically consistent* with both, and because everything would seem the same (or relevantly similar) if the skeptical hypothesis were true, then your evidence appears powerless to choose between the two.

Nothing I have said so far suggests that metaphysically possible skeptical scenarios won't pose a significant skeptical challenge. They do. It seems metaphysically possible that Barack Obama has an identical twin brother. According to the suggestion we're now entertaining, that raises legitimate doubt about whether it was Barack Obama you saw at the campaign rally the other day. Can you rule out that it wasn't his twin? If not, then the skeptic has successfully raised a skeptical challenge. However posing the skeptical challenge with a metaphysical possibility condition opens up a new way to think about and perhaps answer the skeptical challenge, via *modal* epistemology. We could diffuse a skeptical scenario by showing that we were not justified in believing it to be metaphysically possible. I think that for many types of skeptical scenarios it is possible to make such an argument. But I leave that for other work.²⁸

²⁸ See Kung (2009). I have enjoyed fruitful and extensive discussion with Yuval Avnir, Peter J. Graham, and Masahiro Yamada. Thanks as well to my colleagues Steve Davis, Mike Green, Paul Hurley, Amy Kind, Jim Kreines, Laura Perini, Alex Rajczi, Peter Thielke, and Rivka Weinberg. as well. I received excellent feedback on earlier drafts at the 2007 Southern California Epistemology Workshop and the 2008 Central States Philosophy Association Meeting, particularly my commentator Evan Fales. Thanks to audiences of both. Special thanks to James Beebe, who pointed me to his "Constraints on skeptical hypotheses" at the CSPA.

9. References

- Austin, J. (1946). Other minds. *Proceedings of the Aristotelian Society, Supplement*, 20, 148–87.
- Beebe, J. (forthcoming). Constraints on skeptical hypotheses. *Philosophical Quarterly*.
- BonJour, L. (2002). *Epistemology: Classic problems and contemporary responses*. Lanham, MD: Rowman & Littlefield.
- Brueckner, A. (2005). Fallibilism, underdetermination, and skepticism. *Philosophy and Phenomenological Research*, 71(2), 384–91.
- Chalmers, D. (2002). Does conceivability entail possibility? In T. Gendler and J. Hawthorne (Eds.), (2002), 145–200.
- Cross, T. (forthcoming). “Skeptical success.” Oxford volume edited by T. Gendler.
- DeRose, K. (1991). Epistemic possibilities. *Philosophical Review*, 100(4), 581–605.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104(1), 1–52.
- DeRose, K. (1999). Introduction. In K. DeRose and T. Warfield (Eds.), *Skepticism: a Contemporary Reader* (pp. 1–24). New York: Oxford University Press.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111(2), 167–203.
- Dretske, F.I. (1970). Epistemic Operators. *Journal of Philosophy*, 67 (24), 1007–23.
- Gendler, T.S. (2000). The puzzle of imaginative resistance. *Journal of Philosophy*, 97(2), 55–81.
- Goldman, A.I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy* 73(20), 771–91.
- Graham, P. (2007). The theoretical diagnosis of skepticism. *Synthese*, 158(1), 19–39.
- Greco, J. (2000). *Putting skeptics in their place*. Cambridge University Press.
- Klein, P. (2008). Skepticism. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, Fall 2008 Edition, forthcoming URL = <http://plato.stanford.edu/archives/fall2008/entries/skepticism/>.
- Kung, P. (2009). On the possibility of skeptical scenarios. Manuscript submitted for publication.
- Kung, P. & Yamada, M. (2009a). There is no easy bootstrapping problem. Manuscript submitted for publication.

- Kung, P. & Yamada, M. (2009b). Justification without vindication. Manuscript, Claremont, CA.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567. Reprinted in Lewis, D. (1999). *Papers in metaphysics and epistemology* (pp. 418–45). Cambridge University Press.
- Nichols, S. (2006a). Imaginative blocks and impossibility: an essay in modal psychology. In S. Nichols (Ed.) *The architecture of the imagination*. Oxford: Oxford University Press.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–49.
- Rorty, R. (1979). *Philosophy and the mirror of nature*. Princeton University Press.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 126–31.
- Schiffer, S. (2004). Skepticism and the vagaries of justified belief. *Philosophical Studies*, 119, 161–84.
- Sosa, E. (2000). Skepticism and contextualism. *Philosophical Issues*, 10, 1–18.
- Stanley, J. (2005). Fallibilism and concessive knowledge attributions. *Analysis*, 65(2), 126–31.
- Unger, P. (1975). *Ignorance: The case for scepticism*. Oxford University Press.
- van Inwagen, P. (1998). Modal epistemology. *Philosophical Studies*, 92, 67–84.
- Williams, M. (1995). *Unnatural doubts*. Princeton University Press.
- Williamson, T. (2000). *Knowledge and its limits*. New York: Oxford University Press.
- Wittgenstein, L. (1969). *On certainty*. G.E.M. Anscombe & G.H. von Wright (eds.), D. Paul & G.E.M. Anscombe (trans.). New York: Harper & Row.