

1% Skepticism

Eric Schwitzgebel
Department of Philosophy
University of California at Riverside
Riverside, CA 92521-0201

eschwitz at domain: ucr.edu

July 1, 2015

1% Skepticism

Abstract:

A 1% skeptic is someone who has about a 99% credence in non-skeptical realism and about a 1% credence that some radically skeptical scenario obtains. The first half of this essay defends the epistemic rationality of 1% skepticism, endorsing modest versions of dream skepticism, simulation skepticism, cosmological skepticism, and wildcard skepticism. The second half of the essay explores the practical behavioral consequences of 1% skepticism.

1% Skepticism

Certainly there is no practical problem regarding skepticism about the external world. For example, no one is paralyzed from action by reading about skeptical considerations or evaluating skeptical arguments. Even if one cannot figure out where a particular skeptical argument goes wrong, life goes on just the same. Similarly, there is no “existential” problem here. Reading skeptical arguments does not throw one into a state of existential dread. One is not typically disturbed or disconcerted for any length of time. One does not feel any less at home in the world, or go about worrying that one’s life might be no more than a dream (Greco 2008, p. 109).

[W]hen they suspended judgement, tranquility followed as it were fortuitously, as a shadow follows a body.... [T]he aim of Sceptics is tranquility in matters of opinion and moderation of feeling in matters forced upon us (Sextus Empiricus, c. 200 CE/1994, I.xii, p. 30).

I have about a 1% credence in the disjunction of all radically skeptical scenarios combined. In other words, I am about 99% confident that I am awake, not massively deluded, and have existed for decades in roughly the form I think I have existed, doing roughly the sorts of things I think I have been doing; and I find it about 1% subjectively likely that, instead, some radically skeptical possibility obtains – for example, that I am a radically deceived brain in a vat, or that I am currently dreaming, or that I first came into existence only a few moments ago.

I aim to persuade you to embrace a similar credence distribution, on the assumption that you currently have a much lower credence in skepticism, neglecting skeptical possibilities entirely in your daily decision-making. Since 1% is non-trivial in affairs of this magnitude, the practical stakes are substantial. The 1% skeptic ought to live life a bit differently from the 0% skeptic.

I won't be insisting on precisely 1%. "Somewhere around 0.1% to 1%, plus or minus an order of magnitude" will do, or "highly unlikely but not as unlikely as a plane crash".

1. Grounds for Doubt: Dreams.

Sitting here alone in my office, I am almost certain that I am now awake.

I can justify this near certainty, I think, on phenomenological grounds. It's highly unlikely that I would be having experiences *like this* if I were asleep. My confidence might be defensible for other reasons too, as I'll soon discuss.

My current experience has both general and specific features that I think warrant the conclusion that I am awake. The general features are two.

First general feature: I currently have detailed sensory experience in multiple modalities. That is, I currently have a visual sensory experience of my computer screen, of the clutter on my desk, and of my window overlooking the 60/215 freeway; simultaneously, I am having an auditory experience of the hum of the freeway and the tapping of my fingers on the keyboard. Perhaps I also currently have tactile experience of my fingers on the keys and my back against the chair, taste experience of coffee lingering in my mouth, a proprioceptive general sense of my

bodily posture, etc.¹ But according to the theories of dream experience that I favor, the phenomenology of dreaming is *imagistic* rather than *sensory* and also rather sketchy or sparse in detail, including for example in the specification of color – more like the images that occur in a daydream than like the normal sensory experiences of waking life.² On such views, the experience of dreaming that I am on the field at Waterloo is more like the experience of imagining that I am on the field at Waterloo than it is like the experience of standing on the field of Waterloo taking in the sights and sounds; and dreaming that I'm in my office is more like the rather sketchy imagery experience I have at home, lying in bed at night, thinking about going to the office tomorrow, than it is like my current experience of seeing my computer screen, hearing my fingers on the keyboard, etc.

Second general feature: Everything appears to be mundane, stable, continuous, and coherently integrated with my memories of the past. I seem to remember how I got here (by walking across campus after lunch), and this memory seems to fit nicely with my memories of what I was doing this morning and with my memories of what I was doing yesterday and last week and last year. Nothing seems weirdly gappy or inexplicably strange. I share Descartes' view that if this were a dream, it would probably involve discontinuities in perception and memory that would be evident once I thought to look for them.³

¹ For discussion of how sparse or abundant our sensory experience is, see Dennett 1991; Searle 1992; Block 2007; Schwitzgebel 2011b. I am assuming that experience is not sparse for me at this moment.

² See McGinn 2004 and Ichikawa 2009 for defense of this imagery theory of dream experience.

³ Descartes 1641/1984, p. [89-90] 61-62. I agree with Domhoff 2003 (p. 45, 152-154) that the bizarreness of dreams might tend to be overrated, especially if mundane dreams are less likely to be remembered than bizarre ones. But even if, as Domhoff suggests, dreams are not much more bizarre or discontinuous than ordinary waking relaxed thought, waking relaxed thought, it seems plausible to say, involves frequent scene discontinuities and speculative

Some *specific features* of my current experience also bolster my confidence that I am awake. I think of written text as generally unstable in dream experiences. Words won't stay put; they change and scatter away. But here I am, visually experiencing a stable page. And pain is not as vividly felt as in waking life; but I have just pinched myself and experienced the vivid sting. Light switches don't change the ambient lighting; but I have just flicked the switches, or seemed to, changing the apparent ambient light. If you, reader, are like me, you might have your own favorite tests.⁴

But here's the question. Are these general and specific features of my experience sufficient to justify not merely *very high* confidence that I am awake – say 99.9% confidence – but full-on 100% confidence? I think not. I'm not all *that* sure that I couldn't dream of a vivid pinch or a stable page. One worry: I seem to recall “false awakenings” in which I seemed to judge myself awake in a mundane, stable world. Likewise, I'm not all that sure that my favorite theory of dreams is correct. Other dream theorists have held that dreams are sometimes highly realistic and even phenomenologically indistinguishable from waking life – for example, Revonsuo (1995), Hobson, Pace-Schott, and Stickgold (2000), and Windt (2010) – eminent disagreement! I would not risk a thousand dollars for the sake of one dollar on the denial of the possibility that I often have experiences much like this, in sleep. I'm not sure I'd even risk \$1000 for \$250.

But even if I can't point to a feature of my current experience that seems to warrant 100% confidence in my wakefulness, might there still be a philosophical argument that would rationally deliver 100% confidence? Maybe externalist reliabilism about justification is true, for

associations, including hypotheticals that are not consistent with earlier hypotheticals, in a manner that is rather different from mundane sensory experience in a stable environment.

⁴ See LaBerge and Rheingold 1990 for an extended discussion of “dreamsigns” and methods of “state testing”.

example.⁵ And if so, maybe my opinion that I am awake is counterfactually connected in the right way, with extremely high reliability, to the fact that I am indeed awake, even though I can't fully articulate my grounds; and then maybe if I felt absolutely certain, I would be justified in that absolute certainty.⁶ Or maybe any successful referent of "I" necessarily picks out a waking person, so that if I succeed in referring to myself I must indeed be awake.⁷ But for me at least, such philosophical arguments don't deliver 100% confidence. I'm insufficiently confident in the philosophical truth of the required theories.⁸

Is it, maybe, just *constitutive of being rational* that I assume with 100% confidence that I am awake? For example, maybe rationality requires me to treat my wakefulness as an unchallengeable framework proposition or "hinge" assumption in Wittgenstein's (1951/1969) sense. I feel some of the allure of Wittgenstein. But the thought that such a theory of rationality might be correct, though somewhat comforting, does not vault my confidence to the ceiling.

All the reasons I can think of for being confident that I'm awake seem to admit of some doubt; and even stacking those reasons together, doubt remains. To put a number on my doubt suggests more precision than I really intend, but the non-numerical terms of ordinary English

⁵ See Goldman 2008 for a review.

⁶ See Sosa 2007 for a somewhat more subtle version of this argument. I am sympathetic with the critiques of Sosa in Ichikawa 2008 and Brown 2009. But to be clear: Even if I think Sosa's argument *probably* works, my central argument should still succeed, as long as it is rational for me to retain some skeptical distance from philosophical arguments of this sort and withhold 100% credence on that basis.

⁷ For a similar argument, see Valberg 2007 against "immanent" dream skepticism. As Valberg emphasizes, such an argument is consistent with the "transcendent" possibility that "THIS" is all a dream. The present section could be accordingly recast.

⁸ One way of thinking about this last claim is that I have a non-100% higher-order credence that it would be rational to assign a 100% credence to the proposition that I am awake. I might then apply some version of a weighted rational reflection principle, as in Elga 2013. (See also Christensen 2010; Lasonen-Aarnio 2014; Sliwa and Horowitz forthcoming.) The reasoning at the end of this paragraph fails if some such anti-skeptical consideration ought to give me 100% confidence in my wakefulness, in a way that is immune to undermining by higher-order doubts.

have the complementary vice of insufficient clarity. So, with caveats: Given my reflections above, a 90% credence that I am awake seems unreasonably low. I am much more confident that I am awake than that a coin flipped four times will come up heads at least once. On the other hand, a 99.999% credence seems unreasonably high, now that I've paused to think things through. Neither my apparent phenomenological grounds nor my dubious philosophical theorizing about the nature of justification seems to license so extreme a credence. So I'll split the difference: a 99.9% credence seems about right, give or take an order of magnitude.

To think of it another way: Multiplying a 20% credence that I'm wrong about the general features of dreams by a 20% credence, conditional upon my being wrong about dreams in general, that I commonly have mundane office-like experiences like this in dreams, would yield a 4% credence that I commonly have mundane experiences like this in my dreams – that is, a 96% credence that I don't commonly have experiences like this in my dreams. This seems a rather high credence, really, for that particular theoretical proposition, given how poorly we understand dream phenomenology and the range of expert opinions about it. Once I admit a 4% credence that this broad type of phenomenology is common in dreams, it's hard for me to see a good epistemic path down to a 0.001% or lower credence that I am now dreaming.

2. Grounds for Doubt: Simulation Skepticism.

John Searle has argued that digital computers could never have conscious experience (Searle 1980, 1984). There's some chance he's right about that. But he might be wrong. And if he is wrong, it might be possible to create conscious beings who live entirely within simulated environments inside of computers – like Moriarty in Star Trek's "Ship in a Bottle" (Echevarria 1993) or like the "citizens" who live long, complex lives within powerful, sheltered computers in

Greg Egan's *Diaspora* (1997). And *if* such beings exist, some of them might be ignorant of their nature as computationally constructed beings in artificial environments. They might think that their world is not the computational creation of some other set of beings. They might even think that they live on "Earth" in the early "21st century" – and if some further technological and philosophical assumptions are granted, they might even have experiences that are in every way subjectively indistinguishable from the experiences of 21st-century Earthlings in non-simulated environments.

Normally I assume that I'm not a being instantiated in someone else's computational device. Might that assumption be wrong?

Nick Bostrom (2003) argues that we should assign about a one-third credence to being simulated beings, "sims", of this sort. He invites us to imagine the possibility of a technologically advanced civilization able to cheaply run vast numbers of "ancestor simulations" that contain conscious, self-reflective, and philosophical beings with experiences and attitudes similar to our own. Bostrom suggests that we give approximately equal epistemic weight to three possibilities: (a.) that such technologically advanced civilizations do not arise, (b.) that such civilizations do arise but choose not to run vastly many ancestor simulations, or (c.) that such civilizations do arise and the world as a whole contains vastly many more sims than non-sims. Since in the third scenario the vast majority of beings who are in a subjective and epistemic situation relevantly similar to our own are sims, Bostrom argues that our credence that we ourselves are sims should approximately equal our credence in the third scenario, about 1/3.

One might object to the starting assumptions, as Searle would. Or one might, with Weatherson (2003), press technical philosophical challenges against Bostrom's use of an indifference principle in making that final argumentative move. More generally, one might

challenge Bostrom's assertion that sims and non-sims are in relevantly similar epistemic situations. Or one might object that to the extent we assign credence to the possibility that we are sims, that credence assignment undercuts our grounds for conjecturing about the future of computation (Birch 2013; Crawford 2013). Or one might object, on technological grounds, that Bostrom overestimates the probability that civilizations ever progress to a point where huge numbers of conscious ancestor simulations are run. Legitimate concerns, all. And yet none of these concerns seems to imply that we should assign *zero* credence to our being sims in Bostrom's sense. What if we assigned about a 0.1 to 1% credence? Is that plainly too high? Given the amazing (seeming-)trajectory of computation through the 20th and 21st centuries, and given the philosophical and empirical tenuousness of objections to Bostrom's argument, it seems reasonable to reserve a non-trivial credence for the possibility that the universe contains many simulated beings. And if so, it seems reasonable to reserve some non-trivial sub-portion of that credence for the possibility that we are among those beings.

As David Chalmers (2003/2010, 2012) has emphasized, the simulation possibility needn't be seen as a *skeptical* possibility. Chalmers analogizes to Berkeleyian idealism. Berkeley (1710-1713/1965) holds that the world is fundamentally composed of minds and their ideas, coordinated by God – and yet, Berkeley insists, this is no form of skepticism. Cups, houses, rivers, brains, and fingers all exist just the same, according to Berkeley, and can be depended upon; only they are metaphysically constituted rather differently than people tend to suppose. Likewise, if the simulation hypothesis is correct, we and our surroundings might be fundamentally constituted by computational processes in high-tech computers but still have reality enough that most of our commonsense beliefs qualify as true.

To count as a radically skeptical scenario in the same sense that the dream scenario is radically skeptical or the “brain in a vat” scenario is radically skeptical, we must be in a *small* sim or an *unstable* sim – maybe a short-term sim booting up only a few minutes ago in our subjective time (with all our seeming-memories in place, etc.) and doomed for deletion soon, or maybe a spatially small sim containing only this room or this city, or maybe a sim with inconstant laws that are due soon for a world-destroying change.

Conditionally upon assuming that I am in a sim, how much of my credence should I distribute to the possibility that I’m in a large, stable sim, and how much should I distribute to the possibility that I am instead in a small or unstable sim? Philosophical advocates of the simulation hypothesis have tended to emphasize the more optimistic, less skeptical possibilities.⁹ But it’s unclear what would justify a high degree of optimism. *Maybe* the best way to develop conscious beings within simulations is to evolve them up slowly in giant, stable sim-planets that endure for thousands or millions or billions of sim-years. But maybe it’s just as easy, or easier, to cut and copy and splice and spawn variants off a template, creating person after person within small sims, or lots of copies of pre-fab worlds that endure briefly as scientific experiments or toys, like the millions of sample cities pre-packaged with the 21st-century computer game SimCity.

Our age is awestruck by digital computers, but we should also bear in mind that simulation might take another form entirely: A simulation could conceivably be created from ordinarily-structured analog physical materials, at a scale that is small relative to the size and power of the designers – a world in miniature in a sandbox, as it were. Or there might be future technologies as different from electronic computers as electronic computers are from clockwork

⁹ Bostrom 2003 (reaffirmed in dialogue with me: Schwitzgebel 2011a); Chalmers 2003/2010; and especially Steinhart 2014.

and which can give rise to conscious beings in a manner we can't now even begin to understand. Simulation skepticism thus needn't entirely depend on hypotheses about computing technology. As long as we might well be artificially created playthings, and as long as in our role as playthings we might well be radically mistaken in our ordinary day-to-day confidence about the recent past or the near future or the existence of Milwaukee, then we might well be sims in the sense relevant to sim-skepticism.

Bostrom's argument for a one-third credence that we are sims should be salted with caveats both philosophical and technological – and yet it seems that we have positive empirical and philosophical reason to assign *some* non-trivial credence to the possibility that the universe contains many sims, and conditionally upon that to the possibility that we are among the sims, and conditionally upon that to the possibility that we (or you, or I) are in a sim-environment small enough or unstable enough to count as a skeptical scenario. Multiplying these credences together, it seems reasonable to be quite confident that we are not in such a small or unstable sim-environment, but it's hard to see grounds for being *hugely* confident about that. Again, a reasonable credence might be about 99.9% plus or minus an order of magnitude.

3. Grounds for Doubt: Cosmological Skepticism.

According to mainstream physical theory, there is an extremely small but finite chance that a molecule-for-molecule duplicate of you (within some arbitrarily small error tolerance) could spontaneously congeal, by random chance, from disorganized matter. This is sometimes

called the *Boltzmann brain* scenario, after physicist Ludwig Boltzmann, who conjectured that our whole galaxy might have originated from random fluctuation, given infinite time to do so.¹⁰

It seems reasonable to assign some sizeable credence to the following philosophical hypothetical: *If* such a Boltzmann brain or “freak observer” emerged, it could have some philosophical and cosmological thoughts, including possibly the thought that it is living on a full-sized, long-enduring planet that contains philosophers and physicists entertaining the Boltzmann brain hypothesis. (If, given some kind of “content externalism”, e.g., Putnam 1981, it would require weeks or years of history to have a thought with that content, imagine a freak system with that spatiotemporal size or more.) The standard view is that it is vastly more likely for a relatively small freak system to emerge than for a relatively large one to emerge. If so, almost all freak observers who think they are living on long-enduring planets and who think they will survive to see tomorrow are in fact mistaken: They are a blip in the chaos that will soon consume them.

Might I be a doomed freak observer?

Well, how many freak observers are there in the world, compared to human-like observers who have arisen by what I think of as the more normal process of millions of years of evolution in a large, stable system? If the world is one way – for example, if it is infinite, and prone to fluctuations, and Big Bangs are extremely rare – then the ratio of freaks to normals might be large as the size of the patch under consideration goes to infinity. If the world is another way – for example, if it is limited to a few Hubbles radius and duration after which it settles into an unfluctuating state, then there might not be a single freak anywhere. If the world is still another way, the ratio might be about 50-50. Do I have any compelling reason to think

¹⁰ See Boltzmann 1897; Bostrom 2002; Carroll 2010; De Simone, Guth, Linde, Noorbala, Salem, and Vilenkin 2010; Crawford 2013; Boddy, Carroll, and Pollack 2014.

I'm almost certainly in a world with a highly favorable ratio of normals to freaks? Given the general dubiety of (what I think of as) current cosmological theorizing, it seems rash to be hugely confident about which sort of world this is.

But maybe my rational credence here needn't depend on such cosmological ratios. Suppose that there are a million doomed freak duplicates of me who think they are "Eric Schwitzgebel" writing about the philosophy of Boltzmann brains, etc., for every one evolved-up "Eric Schwitzgebel" on a large, stable rock. Maybe there's a good philosophical argument that the evolved-up Erics should rationally assign 99.999% or more credence to being non-freaks, even after they have embraced cosmological dubiety about the ratio of freaks to normals. Maybe it's just constitutive of my rationality that I'm entirely sure I'm not a freak; or maybe it's an unchallengeable framework assumption; or maybe some there's some externally-secured forward-looking reliability, which a freak cannot have, that warrants stable-Eric's supreme confidence – but again, as in the dream case, it seems unreasonable to be highly certain that such philosophical arguments succeed. In the face of both cosmological and philosophical doubt, I see no epistemically responsible path to supreme certainty.

Here's another cosmological possibility: Some divine entity made the world. One reason to take the possibility seriously is that atheistic cosmology had not yet produced a stable scientific consensus about the ultimate origin of the universe.

But although it seems possible that some entity intentionally designed and launched the universe, I'm inclined to think that a sober look at the evidence does not compel the conclusion that if such a deity exists, it is benevolent. The creator or creators of the universe might be perfectly happy to have created lots of doomed freaks. They might not mind deceiving us, might even enjoy doing so or regard it as a moral obligation. Maybe God is not a very skilled

constructor of universes and I am one of a million trial runs, alone in my room, like an artist's quick practice sketch. Maybe God is a sadistic adolescent who has made a temporary Earth so that he can watch us fight like ants in a jar. Maybe God is a giant computer running every possible set of instructions, most of which, after the N^{th} instruction, produce only chaotic results at the $N+1^{\text{th}}$. (Some of these scenarios overlap the sim scenarios.) Theology is an uncertain business.

It seems unreasonable to maintain extremely high confidence – 99.999% or more – that our position in the universe is approximately what we think it is. The Boltzmann brain hypothesis, the sadistic adolescent deity hypothesis, the trial-run hypothesis – those are only some of the possibilities. I appear to have a very limited perspective on the world. I am like a flea on the back of a dog, watching a hair grow, and saying “ah, so that's the nature of the cosmos!” It is rational, I think, to reserve a small credence – again I'd suggest about 0.1% – for epistemically catastrophic cosmological possibilities, in which I am radically wrong about huge portions of what I normally regard to be utterly obvious.

4. Grounds for Doubt: Wildcard Skepticism.

These three skeptical scenarios – dream skepticism, simulation skepticism, and cosmological skepticism – are the only specific skeptical worries that currently draw a non-trivial portion of my credence. I think I have *grounds for doubt* in each case. People dream frequently and on some leading theories dreams and waking life are often indistinguishable. Starting from a fairly commonplace set of 21st-century mainstream Anglophone cultural attitudes, one can find positive reasons to assign a small but non-trivial credence to simulation doubts and cosmological doubts.

In contrast, I currently see no good reason to assign even a 0.0001% credence to the hypothesis that aliens envatted my brain last night and are now feeding it fake input. Even if I think such a thing is *possible*, nothing in my existing network of beliefs points toward any but an extremely tiny chance of its being true. I also find it difficult to assign much credence to the possibility that I'm a deluded madman living in a ditch or asylum, hallucinating an office and believing I am the not-very-famous philosopher Eric Schwitzgebel. No positive argument gives these doubts much grip. (Maybe if I thought I was as immensely famous and interesting as Wittgenstein, I would have reasonable grounds for a non-trivial sliver of madman doubt.) But I also think that if I knew more, my epistemic situation might change. Maybe I'm underestimating the evidence for envatting aliens; maybe I'm overestimating the difference between my epistemic situation and the situation of a madman. Maybe there are other types of scenarios I'm not even considering but which I should consider, and which, if I did properly consider them, would toss me into further doubt. So I want to reserve a small portion of my credence for the possibility that there is some skeptical scenario that I am overlooking or not taking seriously enough. Call this *wildcard skepticism*.¹¹

5. Grounds for Doubt: Conclusion.

A *radically skeptical scenario*, let's say, is any scenario in which a large chunk of the ordinary beliefs that we tend to take for granted is radically false. I offer the examples above as a gloss: If I'm dreaming or in a small sim, or if I'm doomed freak, or if I'm the brief creation of an indifferent deity, then I am in a skeptical scenario. So also if I'm a recently envatted brain, or

¹¹ Compare the "catch-all" hypothesis in formal theories of evidence: Shimony 1970; Earman 1992; Wenmackers and Romejin 2014.

if radical solipsism is true, or if induction utterly fails and the future is radically unlike the past. Call the view that no such radically skeptical scenario holds *non-skeptical realism*.

I suggest that a final credence in the disjunction of all radically skeptical scenarios combined should be around 0.1% to 1%, plus or minus an order of magnitude. This is the view I am calling *1% skepticism*.

I am unconcerned about exact numbers. Nor am I concerned if you reject numerical approaches to confidence. I intend the numbers for convenience only, to gesture toward a degree of confidence much higher than an indifferent shrug but also considerably lower than nosebleed heights of nearly absolute certainty. If numbers suggest too much precision, then I think I can make do with “highly unlikely but not extremely unlikely”, “low probability, but not as low as winning the state lottery with a single ticket”.

I am unconcerned about defending against complaints that my view is insufficiently skeptical. There are few genuine radical skeptics. I assume that sincere opposition will come almost entirely from philosophers much more confident than I in non-skeptical realism. In one sense, this isn't a skeptical position at all: I'm allowing a 99%-99.9% credence in non-skeptical realism, which anti-skeptical philosophers sometimes tell me is entirely consistent with their anti-skeptical views. That's fine; but other anti-skeptical philosophers tell me that they think 0.1%-1% is far too high a credence in the disjunction of all radically skeptical scenarios. Those are my intended opponents.

You might object that skeptical possibilities are impossible to assess in any rigorous way, that the proposed credence numbers are unscientific, indefensible by statistics based on past measurements, insufficiently determinable – that they even to some extent undermine their own

apparent basis (see the next section). Right! Now what? Virtual certainty doesn't seem the proper reaction.

Although my thesis concerns skepticism, the reader might notice that I have never once used the word "know" or its cognates. I'm not concerned, in this essay, about application of the term "knowledge". Maybe I know some things that I believe with less than 99.9% credence.¹² Maybe I know that I am awake despite not having perfect credence in my wakefulness. My concern here is rational credence, not knowledge.

6. Is 1% Skepticism Self-Refuting?

If I am a Boltzmann brain, then my cosmological beliefs, including my evidence for the belief that I might be a Boltzmann brain, are not caused in the proper way by scientific evidence for those beliefs. If I am in a small simulation, my evidence about the fundamental nature of the world, including the nature of simulations, is dubiously grounded. If I am dreaming, my seeming memories about the phenomenology of dreaming and wakefulness might themselves be mere dream errors. The scenarios I've described can, it seems, at least partly undermine the evidence that seems to support them.¹³

Such apparently self-undermining evidence does not, however, defeat skepticism. Montaigne compares skeptical affirmation to rhubarb, which flushes everything else out of the body and then itself last.¹⁴ Suppose your only evidence about the outside world were through video feeds. If you were to discover video-feed evidence that the video feeds were unreliable,

¹² See Vogel 1999; Hill and Schechter 2007. Following Radford 1966, I argue in Myers-Schulz and Schwitzgebel 2013 that knowledge might not even require belief.

¹³ For versions of this self-undermining argument, see Moore 1959; Maitzen 2010; Crawford 2013.

¹⁴ Montaigne 1580/1595/2003, p. 590. See also Diogenes Laertius 3rd c. CE/1972 on Pyrrho.

the proper response would not be to discard that evidence as self-defeating and retain high confidence in one's feed. It would be surprising if it were generally the case that evidence that one is in a risky epistemic position somehow bites its own tail, disappearing through self-defeat to leave certainty in its place. If something in your experience suggests that you might be dreaming – maybe you seem to be levitating off your chair – it would presumably not be sensible as a general policy to dismiss that evidence on grounds that if you are dreaming such evidence is unreliable as evidence that you are dreaming. If something hints that you might be in a simulation or if scientific cosmological evidence converges toward ever more bizarre and dubious options, the proper response would be more doubt, not less.

What is plausibly self-defeating or epistemically unstable is high credence in *one specific skeptical scenario*. Assigning more than 50% credence to being a freak observer based on seemingly excellent cosmological evidence in support of that fact probably is cognitively unstable.¹⁵ Evidence that this is a dream might be equally compatible with sim-skepticism or cosmological skepticism or madman skepticism. It might be hard to justify high confidence in just one of these possibilities. But 1% skepticism is a different matter, since it tolerates a diversity of scenarios and trades only in low credences, imprecisely specified.

7. Should I Try to Fly, on the Off-Chance That This Is a Dream-Body?

Philosophers sometimes say that skepticism is unlivable and can have no permanent practical effect on those who attempt to endorse it. See the Greco quote in the epigraph, for example. Possibly this is Hume's view too, in Part One of his *Treatise* (1740/1978). I disagree.

¹⁵ For a version of this argument, see Carroll 2010. See Rinard 2015 for a related argument regarding external world skepticism in general, which I believe I can avoid for a related reason: my relatively high credence in non-skeptical realism.

Like Sextus (the other epigraph quote), I think radical skepticism, including 1% skepticism, need not be behaviorally inert. To explore this idea, I will use examples from my own life, which I have chosen because they are real, as best I can recall them – and thus hopefully also realistic – but you may interpret them as merely hypothetical if you prefer.

I begin with a whimsical case. I was revising the dreaming section of this essay. Classes had just been released for winter break, and I was walking to the science library to borrow more books on dreaming. I had just been reading Thompson (2015) on false awakenings. There, on the wide-open, empty path through east campus, I spread my arms, looked at the sky, and added a leap to one of my steps, in an attempt to fly.

My thinking was this: I was almost certainly awake – but only *almost* certainly. At that moment, my credence that I was dreaming was higher than usual for me, maybe around 0.3% (though I didn't conceptualize it numerically at the time). I figured that *if* I was dreaming, it would be nice to fly around instead of trudging. On the other hand, if I was not dreaming, it seemed no big deal to leap, and in fact kind of fun – perhaps a bit embarrassing if someone saw me, but no one seemed to be around.

I'll model this thinking with a decision matrix. I don't intend the numbers to be precise, nor do I mean to imply that I was at the time in an especially mathematical frame of mind.

Call dream flying a gain of 100, waking leap-and-fail a loss of 0.1, continuing to walk in the dream a loss of 1 (since why bother with the trip if it's just a dream), and dreaming leap-and-fail a loss of 1.05 (the loss of the walk plus a little more for the leap-and-fail disappointment) – all relative to a default of zero for walking, awake, to the library. For simplicity, I will assume that if I'm dreaming, things are not overall much better or worse than if I'm awake (for example,

I can get the books tomorrow). Let's assign a 50% likelihood of successfully flying, conditional upon its being a dream, since I don't always succeed when I try to fly in my dreams.

Here's the payoff matrix:

	dreaming (0.3%)	not dreaming (99.7%)
leap	50% fly: +100 50% not fly: -1.05	-0.1
don't leap	-1	0

The expected value formula yields approximately +0.05 as the expected gain for leaping and -0.003 as an expected loss for not leaping.

It was an unusual occasion. Normally more is lost for trying to fly (e.g., embarrassment or distraction from more important tasks). Normally my credence in the dream possibility, when the issue even occurs to me, is lower than 0.3%. And I'm assuming a high relative payoff for flying. Also the decision is not repeatable. If I even once try to fly and fail, that will presumably influence my sense of the various probabilities and payoffs. For example, my first leap-and-fail should presumably reduce my credence that I am dreaming and/or my credence that if this is a dream I can fly.

You might say that if the effect of skeptical reflections is to make one attempt to fly across campus, that is not a good result! However, this is a poor objection if it assumes hindsight certainty that I was not dreaming. A similar argument could be mounted against having bought car insurance. A better version of this objection considers the value or disvalue of the type of psychological state induced by 1% skepticism, conditional upon the acknowledged 99%-99.9% subjective probability of non-skeptical realism. If skeptical doubt sufficiently impairs my

approach to what I am 99%-99.9% confident is my real life, then there would be pragmatic reason to reject it.

I'm not sure that in this particular case it played out so badly. The leap-and-fail was silly, but also whimsically fun. I briefly punctured my usual professional mien of self-serious confidence. I would have not leapt in that way, for that reason, adding that particular weird color to my day, and now to my recollections of that winter break, if I hadn't been thinking so much about dream skepticism.¹⁶

8. Weeding or Borges?

It was Sunday and my wife and children were at temple. I was sitting at my backyard table, in the shade, drinking tea – a beautiful spring day. Borges' *Labyrinths*, a favorite book, lay open before me; but then I noticed Bermuda grass sprouting amid the daisies. I hadn't done any weeding in several weeks. Should I finally stop procrastinating that chore? Or should I seize current pleasure and continue to defer the weeds? As I recall it, I was right on the cusp between these two options – and rationally so, let's assume. I remember this not as a case of weakness of will but rather as a case of rational equipoise.

Suddenly, skeptical possibilities came to mind. This might be a dream, or a short-term sim, or some other sort of brief world, or I might in some other way be radically mistaken about my position in the universe. Weeding might not have the long-term benefits I hoped. I might wake and find that weeds still needed doing. I might soon be unmade, I and the weeds dissolving into chaos. None of this I regarded as *likely*. But I figured that if my decision

¹⁶ Maybe if I incorporate such an expectation into the original decision, the value in the not-dreaming/leap cell would become positive. It's not clear, however, that standard decision theory can non-paradoxically incorporate outcomes that reflect one's happiness with the decision process itself.

situation, before considering skepticism, was one of almost exact indifference between short-term benefit and long-term benefit, then this new skeptical consideration should ever-so-slightly tilt me toward the short-term benefit. I lifted the Borges and enjoyed my tea.

Before considering the skeptical scenarios, my decision situation might have been something like this:

Borges: short-term expected value 2, long-term expected value 1, total expected value 3.

Weeding: short-term expected value -1, long-term expected value 4, total expected value 3.

With 1% skepticism, the decision situation shifts to something like this:

	some skeptical scenario obtains (1%)	the world is approximately how I think it is (99%)
Borges	short-term gain: 2; long-term gain: 50% likely to be 1 (e.g., if the sim endures a while), 50% likely to be 0	short-term gain: 2 long-term gain: 1
weeding	short-term gain: -1 long-term gain: 50% likely to be 4, 50% likely to be 0	short-term gain: -1 long-term gain: 4

This yields an expectation of +2.995 for Borges vs. +2.980 for weeding.

Of course I'm oversimplifying. If I'm a Boltzmann brain, I probably won't get through even the first sentence of Borges. If I'm a sim, there's some remote chance that the Player will bestow massive pleasure on me as soon as I start weeding. Etc. But the underlying thought is, I believe, a plausible one: To the extent radically skeptical scenarios render the future highly uncertain, they tilt the decision theoretical scales slightly toward short-term benefits over long-term benefits that require short-term sacrifices.

It has been gently suggested to me that my backyard reasoning was just post-hoc rationalization – that I was already set upon choosing the Borges and just fishing for some justifying excuse. Maybe so! This tosses us back into the pragmatic question about the psychological effects of 1% skepticism: Would it be a good thing, psychologically, for me to have some convenient gentle rationalization to justify just a smidgen more *carpe diem*? In my own case, probably so. Maybe this is true of you too, if you're the type to read this deep into a philosophy article?

9. How to Disregard Extremely Remote Possibilities.

If radical skepticism is on the table, by what rights do I simplify away from extremely remote possibilities? Maybe it's reasonable to allow a $1/10^{50}$ credence in the existence of a Player who gives me at least 10^{50} lifetimes' worth of pleasure if I choose to weed in the Borges scenario. Might my decision whether to weed then be entirely driven by that remote possibility?

I think not, for three reasons.

First, *symmetry*: My credences about such extremely remote possibilities appear to be approximately symmetrical and canceling. In general, I'm not inclined to think that my prospects will be particularly better or worse due to their influence on extremely unlikely deities,

considered as a group, if I pull weeds than if I do not. More specifically, I can imagine a variety of unlikely deities who punish and reward actions in complementary ways – one punishing what the other rewards and vice versa. (Similarly for other remote possibilities of huge benefit or suffering, e.g., rising to a 10^{100} -year Elysium if I step right rather than left.) This indifference among the specifics is partly guided by my general sense that extremely remote possibilities of this sort don't greatly diminish or enhance the expected value of such actions. I see no reason not to be guided by that general sense – no argumentative pressure to take such asymmetries seriously.

Second, *diminishing returns*: Bernard Williams (1973) famously thought that extreme longevity would be a tedious thing. I tend to agree instead with John Fischer (1994; Fischer and Mitchell-Yellin 2014) that extreme longevity needn't be so bad – but it's by no means clear that 10^{20} years of bliss is 10^{20} times more valuable than a single year of bliss.¹⁷ (Similarly for bad outcomes and for extreme but instantaneous outcomes.) Value might be very far from linear with temporal bliss-extension for such magnitudes. And as long as one's credence in remote outcomes declines sharply enough to offset any increasing value in the outcomes, extremely remote possibilities will be decision-theoretically negligible.

Third, *loss aversion*: I'm loss averse (Kahneman and Tversky 1979): I'll take a bit of a risk to avoid a sure or almost-sure loss. And my life as I think it is, given non-skeptical realism, is the reference point from which I determine what counts as a loss. If I somehow arrived at a $1/10^{50}$ credence in a deity who would give me 10^{50} lifetimes of pleasure if I avoided chocolate

¹⁷ One dilemma: If I achieve that bliss by repeating similar experiences over and over again, forgetting that I have done so, then this is something like what I've elsewhere called "goldfish-pool immortality" (Schwitzgebel 2014a); and it seems reasonable not to regard goldfish-pool cases as additively choiceworthy. Alternatively, if I remember all 10^{20} years, then I seem to have become something radically different in cognitive function than what I presently am, so I might be choosing my extinction.

for the rest of my life (or alternatively a deity who would give me 10^{50} units of pain if I didn't avoid chocolate for the rest of my life), and if there were no countervailing considerations or symmetrical chocolate-rewarding deities, then on a risk-neutral utility function it might be rational for me to forego chocolate evermore. But foregoing chocolate would be a loss relative to my reference point; and since I'm loss averse rather than risk neutral, I might be willing to forego the possible gain (or risk the further loss) so as to avoid that almost-certain loss of life-long chocolate pleasure. I might reasonably decline a gamble with a 99.99999% chance of death and a 0.00001% chance of 10^{100} lifetimes' worth of pleasure, even bracketing diminishing returns. I might even reasonably decide that at some level of improbability – $1/10^{30}$? – no finite positive or negative outcome could lead me to take a substantial almost-certain loss. And if the time and cognitive effort of sweating over decisions of this sort itself counts as a sufficient loss, then I can simply disregard any possibility where my credence is below that threshold.

These considerations synergize: the more symmetry and the more diminishing returns, the easier it is for loss aversion to inspire disregard. Decisions at credence $1/10^{50}$ are one thing, decisions at credence $1/10^3$ quite another.¹⁸

10. My Airplane Descends Through the Fog.

I'm a passenger in a jumbo jet descending through turbulent night fog into New York City. I'm not usually nervous about flying, but the turbulence has me on edge. I estimate the odds of dying in a jet crash with a major U.S. airline to be small – but descent in difficult weather is one of the riskiest situations in flight. So maybe my odds of death in the next few

¹⁸ Similar issues arise in the extensive literature on Pascal's wager (Pascal 1670/2005), especially finite-valued versions of it (Hájek 2003; Bostrom 2009).

minutes are about one in a million? I can't say those are odds I'm *entirely* happy about. One in a million is on my decision-theoretic radar in a way that one in 10^{50} is not.

But then I think: Maybe I'm a short-term sim! Based on the reasoning in Section 2, my credence that I'm a short-term sim should be about one in a thousand. In many of those short-term scenarios, my life will end soon. So my credence that I die today because I'm a short-term sim should be at least an order of magnitude higher than my credence that I die in a mundane plane crash.

Should I be alarmed or comforted by these reflections?

Maybe I should be alarmed. Once I began to consider simulation scenarios, my estimated odds of near-term death rose substantially. On the other hand, though, I'm accustomed to facing the simulation possibility with equanimity. The odds are lowish, and I don't see much to do about it. The descent of the airplane adds only a small bit to the odds of my near-term death, on this way of thinking – within rounding error – and in an airplane, as in the sim, no action seems likely to improve my chances. I should just sit tight. The airplane's turbulent descent has made no material change in my prospects or decision options.

At the time, as I now seem to recall it, I was neither alarmed nor much comforted by these reflections. The likelihood of imminent death given the simulation possibility seemed less emotionally real than what I would have said were the much smaller odds of mundane imminent death.

It's a common pattern in our lives – reaching a conclusion based on theoretical reasoning but then failing to be moved by it at a “gut” level. The Stoic sincerely judges that death is not bad, but quakes in fear on the battlefield; the implicit racist sincerely judges that her black

students are every bit as capable as her white students, but her habitual reactions fail to change.¹⁹ Relatedly, the academic philosopher might defend her position passionately at a conference, staking her career on it in all sincerity, but feel reluctant to wager on it in a non-academic context.²⁰ I am trying to carry my 1% skepticism out of its academic context. That works nicely enough when the act is safe, optional, low-stakes, playful; but this is something different.

Some philosophical traditions recommend “spiritual exercises” aimed at bringing one’s spontaneous emotional reactions into line with one’s philosophical opinions – for example, in the Stoic tradition, vividly imagining death with equanimity.²¹ Maybe I should do something similar, if I really wish to live as a 1% skeptic? One concern: Any realistic program of spiritual exercises presumably would require considerable time and world stability to succeed, so it’s likely to fail unless non-skeptical realism is true; and there’s something semi-paradoxical about launching a program that will probably succeed in lowering my credence in non-skeptical realism only if non-skeptical realism is true.

11. Agnosticism.

I used to say that my confidence in the non-existence of God was about the same as my confidence that my car was still parked where I left it rather than stolen or towed – a credence of maybe 99.9%, given the safety of my usual parking spots, a credence sufficient, I thought, to warrant the label atheist. But reflection on skeptical possibilities has converted me to agnosticism.

¹⁹ Philosophical discussions include Gendler 2008; Frankish 2010; Schwitzgebel 2010; Brownstein and Saul (eds.) forthcoming.

²⁰ Lycan 2013; DeRose forthcoming.

²¹ On the philosophical relevance of spiritual exercises, see Hadot 1995; Stalnaker 2006. Special thanks to my former student Ted Preston, whose 2005 PhD. dissertation is a cross-cultural exploration of this topic.

Cosmological skepticism seems to leave plenty of room for the possibility of a god or gods – 50%? 10%? – conditional upon accepting it. That puts the 1% cosmological skeptic already in the ballpark of a 0.1% credence in a divinity on those grounds alone. Furthermore, if I am in a simulation, the power that the managers of the simulation likely have over me – the power to stop the sim, to erase or radically alter me, to resurrect me, to work miracles – are sufficient that I should probably regard them as gods.²² And of course plenty of non-skeptical scenarios involve a god or gods (some non-skeptical sim scenarios, some non-skeptical Big Bang cosmologies). It would be odd to assign a 1% credence to the disjunction of all radically skeptical scenarios, and some substantial sub-portion of that credence to skeptical scenarios involving gods, while assigning only negligible credence to *non*-skeptical scenarios involving gods. Once I embrace 1% skepticism, I can no longer sustain my formerly high levels of atheistic self-assurance.

Reflection on radically skeptical possibilities might increase open-minded philosophical tolerance in general. In Schwitzgebel (1996) I argued that this was Zhuangzi's (3rd c. BCE/2009) aim in defending skepticism; it also seems plausibly to have been among Montaigne's (1580/1588/1958) aims. I intend this claim as a testable empirical conjecture about the psychology of philosophy. Ordinary intelligent readers who engage with good, accessible skeptical texts like those of Zhuangzi, Sextus, Montaigne, Bayle, and Hume, will, if I am right, tend to experience a reduction of philosophical self-confidence, even if they don't accept the skeptic's conclusions. Given the diversity of philosophical opinion, most philosophers who confidently hold specific philosophical views about specific matters of contentious fact must be

²² Bostrom 2003; Schwitzgebel 2014; Steinhart 2014.

mistaken. So a skeptically-induced reduction in philosophical self-confidence might tend to be good epistemic medicine.

12. “I Think There’s About a 99.8% Chance That You Exist” Said the Skeptic.

Alone in my backyard or walking across an empty campus, it can seem quite reasonable to me to reserve a 0.1% to 1% credence in the disjunction of all skeptical scenarios, and conditionally upon that to have about a 10% to 50% credence in the nonexistence of the people I normally regard as existing – my family, my readers, the audience at a talk.

But now I imagine myself before an actual live audience. Can I say to them, sincerely, that I doubt their existence – that I think there’s a small chance that I’m dreaming right now, that I think there’s a small chance they might merely be mock-up sprites, mere visual input in a small me-centered sim, lacking real conscious experience? This seems, somehow, even stranger than the run-of-the-mill strangeness of dream skepticism in solitary moments.

I’ve tried it on my teenage son. He has heard my arguments for 1% skepticism. One day, driving him to school, a propos of nothing, I said, “I’m *almost* certain that you exist.” A joke, of course. How could he have heard it, or how could I have meant it, in any other way?

One possible source of strangeness is this: My audience would be fully aware that they aren’t mere mock-up sprites, just as I also would invest much less than a 0.1% credence in my being a mindless mock-up sprite. So it’s tempting to say that in some sense the audience would see that my doubts are misplaced.

But in non-skeptical cases, at least, we can view people as reasonable in having substantial credences in possibilities we confidently dismiss, if we recognize an informational asymmetry. The blackjack dealer who sees she has a 20 doesn’t think the player a fool for

standing on a 19. Even if the dealer sincerely tells the player she has a 20, she might think the player reasonable to confess substantial doubt about the dealer's truthfulness. So why do radically skeptical cases seem different?

One possible clue is this: It doesn't feel wrong in quite the same way to say "I think that *we might all* be part of a short-term sim". Being *together* in skeptical doubt seems fine – in the right context, it might even be kind of friendly, kind of fun. Maybe, then, the issue is a matter of respect – a matter of treating one's interlocutor as an equal partner, metaphysically, epistemically, morally? There's something offensive, perhaps, or inegalitarian, or silencing, about saying "I'm certain that *I* exist, but I have some doubts about whether you do".

I feel the problem most keenly in the presence of the people I love. I can't doubt that we are in this world together. It seems wrong – merely a pose, possibly an offensive pose – to say to my dying father, in seeming sincerity at the end of a philosophical discussion about death and God and doubt, "I think there's a 99.8% chance that you exist". It throws a wall up between us. At least I felt that it did, despite my father's intellectual sophistication, the one time I tried it. (He forgave me, I think, more than I have been able to forgive myself.)

Can friend-doubting be done a different way? Maybe I could say, "For these reasons, *you* should doubt *me*. And I will doubt you, too, just a tiny bit, so that we are doubting together. Very likely, the world exists just as we think it does; or even if it doesn't, even if nothing exists beyond this room, still I am more sure of you than of almost anything else."

Even with this gentler, more egalitarian approach, the skeptical decision calculus might rationally justify a small tilt toward selfish goods over sacrifice for others who might not exist. To illustrate with a simplified model: If, before skeptical reflection I am indifferent between X pleasure for myself versus Y pleasure for you, reducing my credence in your existence to .998

might shift me to indifference between only $.998 * X$ pleasure for myself versus Y pleasure for you. I find that slightly selfish potential consequence unwelcome. Is that where 1% skepticism should lead?

On the other hand, selfish choices are often long-term choices, e.g., for money and career. Kind, charitable engagement is often more pleasant short-term, and doubt about the reality of someone you are interacting with might be more difficult to justify than doubt about the future. If pre-skepticism I am indifferent between Option A with gains for me of X_1 short-term plus X_2 long-term versus Option B with gains for you of Y_1 short-term plus Y_2 long-term, then embracing a $.998$ credence in your existence and $.995$ credence in the future, I might model the choice as Option A = $X_1 + .995 * X_2$ versus Option B = $.998 * Y_1 + (\text{approx.}) .993 * Y_2$, rationalizing a discount of my long-term gains for your short-term benefit.

In what direction does skeptical doubt tend to move one's character? I'm aware of no direct empirical evidence. But I think of the great humane skeptics in the history of philosophy: Zhuangzi, Montaigne, Hume. The moral character that shines through their philosophical work seems unharmed by their doubts.

If we're evaluating 1% skepticism in part for its psychological effects conditional on the assumption that non-skeptical realism is correct, then there is a risk here, a risk that I will doubt others selfishly or disrespectfully, alienating myself from them. But I believe that this risk can be managed. Maybe, even, it can be reversed: In confessing my skepticism to you, I make myself vulnerable. I show you my weird, nerdy doubts, which you might laugh at, or dismiss, or

join me in. If you join me, or even just engage me seriously, we will have connected in a way that I treasure.²³

²³ For helpful discussion and comments, thanks to Stephanie Allen, Taylor Cyr, Kirkland Gable, Max McCoy, Carlos Narziss, Christian Pillsbury, Micaela Quintana, Sam Richards, Melanie Rosen, Davy Schwitzgebel, Sylvia Wenmackers, Nathan Westbrook, Eric Winsberg, and Ben Yelle; readers of the relevant posts at New APPS, The Splintered Mind, and in social media; and the audience at University of Miami.

I dedicate this essay to the memory of my father.

References

- Berkeley, George (1710-1713/1965). *Principles, dialogues, and philosophical correspondence*, ed. C.M. Turbayne. New York: Macmillan.
- Birch, Jonathan (2013). On the “simulation argument” and selective skepticism. *Erkenntnis*, 78, 95-107.
- Block, Ned (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481-548.
- Boltzmann, Ludwig (1887). Zu Hr. Zermelo’s Abhandlung “Ueber die mechanische Erklärung irreversibler Vorgänge”. *Annalen der Physik*, 296 (2), 392-398.
- Boddy, Kimberly K., Sean Carroll, and Jason Pollack (2014). De Sitter space without quantum fluctuations. Manuscript. URL <http://arxiv.org/abs/1405.0298> (accessed Nov. 25, 2014).
- Bostrom, Nick (2003). Are we living in a computer simulation? *Philosophical Quarterly*, 53, 243-255.
- Bostrom, Nick (2009). Pascal’s mugging. *Analysis*, 69, 443-445.
- Brown, Jessica (2009). Sosa on skepticism. *Philosophical Studies*, 143, 397-405.
- Brownstein, Michael, and Jennifer Saul, eds. (forthcoming). *Implicit bias and philosophy*. Oxford: Oxford.
- Carroll, Sean (2010). *From eternity to here*. New York: Penguin.
- Chalmers, David (2003/2010). The Matrix as metaphysics. In *The Character of Consciousness*. Oxford: Oxford.
- Chalmers, David (2012). *Constructing the world*. Oxford: Oxford.
- Christensen, David (2010). Rational reflection. *Philosophical Perspectives*, 24, 121-140.
- Crawford, Lyle (2013). Freak observers and the simulation argument. *Ratio*, 26, 250-264.

- De Simone, Andrea, Alan H. Guth, Andrei Linde, Mahdiyar Noorbala, Michael P. Salem, and Alexander Vilenkin (2010). Boltzmann brains and the scale-factor cutoff measure of the multiverse. *Physical Review D*, 82, 063520.
- Dennett, Daniel C. (1991). *Consciousness explained*. Boston, MA: Little, Brown and Company.
- DeRose, Keith (forthcoming). *Knowledge, skepticism, and context*, vol. 2.
- Descartes, René (1641/1984). *Meditations on first philosophy*. In *The Philosophical Writings of Descartes*, trans. J. Cottingham, R. Stoothoff, and D. Murdoch. Cambridge: Cambridge.
- Diogenes Laertius (3rd c. CE/1974). *Lives, teachings, and sayings of famous philosophers*, trans. R.D. Hicks. Cambridge, MA: Harvard.
- Domhoff, G. William (2003). *The scientific study of dreams*. Washington, DC: American Psychological Association.
- Earman, John (1992). *Bayes or bust?* Cambridge, MA: MIT.
- Echevarria, René (1993). *Star Trek: The next generation: Ship in a Bottle*, dir. A. Singer. Episode 6x12.
- Egan, Greg (1997). *Diaspora*. Millennium.
- Elga, Adam (2013). The puzzle of the unmarked clock and the new rational reflection principle. *Philosophical Studies*, 164, 127-139.
- Fischer, John M. (1994). Why immortality is not so bad. *International Journal of Philosophical Studies*, 2, 257-270.
- Fischer, John M., and Benjamin Mitchell-Yellin (2014). Immortality and boredom. *Journal of Ethics*, 18, 353-372.
- Frankish, Keith (2010). Dual-process and dual-systems theories of reasoning. *Philosophy Compass*, 5, 914-926.

- Gendler, Tamar S. (2008). Alief and belief. *Journal of Philosophy*, 105, 634-663.
- Goldman, Alvin (2008). Reliabilism. *Stanford Encyclopedia of Philosophy*, Spring 2011 edition. URL: <http://plato.stanford.edu/archives/spr2011/entries/reliabilism>
- Greco, John (2008). Skepticism about the external world. In *The Oxford handbook of skepticism*, ed. J. Greco. Oxford: Oxford. p. 108-128.
- Hadot, Pierre (1995). *Philosophy as a way of life*, trans. A.I. Davidson. Malden, MA: Blackwell.
- Hájek, Alan (2003). Waging war on Pascal's wager. *Philosophical Review*, 112, 27-56.
- Hill, Christopher S., and Joshua Schechter (2007). Hawthorne's lottery puzzle and the nature of belief. *Philosophical Issues*, 17, 102-122.
- Hobson, J. Allan, Edward F. Pace-Schott, and Robert Stickgold (2000). Dreaming and the brain. *Behavioral and Brain Sciences*, 793-842.
- Hume, David (1740/1978). *A treatise of human nature*, ed. L.A. Selby-Bigge and P.H. Nidditch. Oxford: Oxford.
- Ichikawa, Jonathan (2008). Skepticism and the imagination model of dreaming. *Philosophical Quarterly*, 58, 519-527.
- Ichikawa, Jonathan (2009). Dreaming and imagination. *Mind & Language*, 24, 103-121.
- Kahneman, Daniel, and Amos Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-292.
- Lasonen-Aarnio, Maria (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88, 314-345.
- LaBerge, Stephen, and Howard Rheingold (1990). *Exploring the world of lucid dreaming*. New York: Ballantine.

- Lycan, William G. (2013). On two main themes in Gutting's *What Philosophers Know*.
Southern Journal of Philosophy, 51, 112-120.
- Maitzen, Stephen (2010). A dilemma for skeptics. *Teorema*, 29, 23-34.
- McGinn, Colin (2004). *Mindsight*. Cambridge, MA: Harvard.
- Montaigne, Michel de (1580/1595/2003). Apology for Raymond Sebond. In *The Complete Essays*, ed. M.A. Screech. London: Penguin.
- Moore, G.E. (1959). Certainty. In *Philosophical Papers*. London: Allen and Unwin.
- Myers-Schulz, Blake, and Eric Schwitzgebel (2013). Knowing that P without believing that P.
Noûs, 47, 371-384.
- Pascal, Blaise (1670/2005). *Pensées*, trans. R. Ariew. Indianapolis: Hackett.
- Putnam, Hilary (1981). *Reason, truth and history*. Cambridge: Cambridge.
- Radford, Colin (1966). Knowledge – by examples. *Analysis*, 27, 1-11.
- Revonsuo, Antti (1995). Consciousness, dreams and virtual realities. *Philosophical Psychology*, 8, 35-58.
- Rinard, Susanna (2015). *Reasoning one's way out of skepticism*. Manuscript.
- Schwitzgebel, Eric (1996). Zhuangzi's attitude toward language and his skepticism. In P. Kjellberg and P.J. Ivanhoe, eds., *Essays on skepticism, relativism, and ethics in the Zhuangzi*. Albany, NY: SUNY.
- Schwitzgebel, Eric (2010). Acting contrary to our professed beliefs, or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, 91, 531-553.
- Schwitzgebel, Eric (2011a). Bostrom's response to my discussion of the simulation argument. Blog post at The Splintered Mind, Sep. 2. URL:

<http://schwitzsplinters.blogspot.com/2011/09/bostroms-response-to-my-discussion-of.html>

Schwitzgebel, Eric (2011b). *Perplexities of consciousness*. Cambridge, MA: MIT.

Schwitzgebel, Eric (2014a). Goldfish-pool immortality. Blog post at The Splintered Mind, May 30. URL: <http://schwitzsplinters.blogspot.com/2014/05/goldfish-pool-immortality.html>

Schwitzgebel, Eric (2014b). Our possible imminent divinity. Blog post at The Splintered Mind, Jan. 2. URL: <http://schwitzsplinters.blogspot.com/2014/01/our-possible-imminent-divinity.html>

Searle, John (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417-457.

Searle, John (1984). *Minds, brains, and science*. Cambridge, MA: Harvard.

Searle, John (1992). *The rediscovery of the mind*. Cambridge, MA: MIT.

Sextus Empiricus (c. 200 CE/1994). *Outlines of skepticism*, trans. J. Annas and J. Barnes. Cambridge: Cambridge.

Shimony, Abner. Scientific inference. In R.G. Colodny, ed., *The nature and function of scientific theories*. Pittsburgh: University of Pittsburgh.

Sliwa, Paulina, and Sophie Horowitz (forthcoming). Respecting *all* the evidence. *Philosophical Studies*.

Sosa, Ernest (2007). *A virtue epistemology*. Oxford: Oxford.

Stalnaker, Aaron (2006). *Overcoming our evil*. Washington, DC: Georgetown University.

Steinhart, Eric C. (2014). *Your digital afterlives*. New York: Palgrave.

Thompson, Evan (2015). *Dreaming, waking, being*. New York: Columbia.

Valberg, J.J. (2007). *Dream, death, and self*. Princeton: Princeton.

Vogel, Jonathan (1999). The new Relevant Alternatives Theory. *Philosophical Perspectives*, 13, 155-180.

Weatherson, Brian (2003). Are you a sim? *Philosophical Quarterly*, 53, 425-431.

Wenmackers, Sylvia, and Jan-Willem Romeijn (2014). A new theory about old evidence:
Humbly catching them all. Manuscript.

Williams, Bernard (1973). *Problems of the self*. Cambridge: Cambridge.

Windt, Jennifer M. (2010). The immersive spatiotemporal hallucination model of dreaming.
Phenomenology & the Cognitive Sciences, 9, 295-316.

Wittgenstein, Ludwig (1951/1969). *On certainty*, ed. G.E.M. Anscombe and G.H. von Wright.
New York: Harper.