Advancing the Boltzmann Brain Skeptical Challenge

John Pittard, forthcoming in *The Philosophical Review* [Please cite published version]

Abstract: Some leading cosmological models predict that in the future history of the universe, there will be vast numbers of Boltzmann Brains ("brains" formed by chance in otherwise empty space), including many that are physical duplicates of your own brain and that (arguably) have conscious experiences indistinguishable from your own experiences. In this article, I advance the skeptical challenge posed by this prediction in two ways. First, I challenge the common assumption that scientific evidence for Boltzmann Brain superabundance presents a serious skeptical worry only if that evidence gives one reason to assign significant credence to the hypothesis that one is a Boltzmann Brain. On this assumption, Boltzmann Brain-related skeptical worries are easily swept aside. I develop a new and more formidable Boltzmann Brain skeptical argument that is not committed to this assumption. Second, I take up recent objections to Center Indifference, a principle underlying Boltzmann Brain skeptical arguments. Dogramaci and Schoenfield (2025) have argued that Center Indifference imposes absurd constraints on one's credences for certain scientific theories. I show that their arguments turn on some subtle confusions and that Center Indifference does not have the absurd implications that they attribute to it.

1. Introduction

A "Boltzmann Brain" is an isolated "brain" that comes into existence by chance in an otherwise disordered region of space through some random process of particle accretion or through some quantum fluctuation. Such a brain could, for some period, have intrinsic physical states that are

exactly like those that might be had by some brain of a human being having normal human experiences. Some leading cosmological models predict (at least on certain plausible assumptions) that over the course of the universe's history, the number of randomly formed Boltzmann Brains will vastly outnumber the number of "normal" brains (i.e., ones belonging to creatures who are the product of natural selection, who genuinely perceive their environment, and so on) (Carroll 2021). The reason for this is as follows. Long after the ordered physical universe we observe around us has succumbed to heat death, chance events will occasionally produce localized low-entropy regions with highly ordered physical matter. The chance formation of an isolated Boltzmann Brain, while predicted to be extremely rare, is vastly more probable than the chance formation of a significantly larger region of order of the sort that is required for the existence of intelligent creatures who have normal brains and who reliably form accurate beliefs about their environment. While normal brains may comprise all (or nearly all) of the brains that have hitherto existed in the low-entropy universe that we observe today, if the universe persists for long enough in a future high-entropy state, it may be expected that Boltzmann Brains will eventually be predominant among all the brains that have at some point existed. On some mainstream cosmological models (including, arguably, "the current best-fit model" known as ΛCDM (Carroll 2021, 7)), the expected number of Boltzmann Brains is so large as to all but ensure that Boltzmann Brains will comprise a large majority of those brains that at some point pass through a sequence of physical states that perfectly mirror the states of your own brain over the last several seconds.

Scientific support for cosmological models predicting such an abundance of Boltzmann Brains will pose a skeptical problem if one thinks, as many do, that *phenomenal internalism* is true (Saad 2024, 4). For our purposes, phenomenal internalism is the thesis which says that if two brains in our universe have the exact same internal physical properties, then there can be no qualitative differences in the phenomenal states of the two brains. Given phenomenal internalism, cosmological

models that predict a superabundance of Boltzmann Brains would thereby give us a reason to think that most "observers" (i.e., conscious entities) with experiences that are (internally) like yours are deceived Boltzmann Brains. And it would seem that it is not rational for you to agree with this prediction while continuing to trust your faculties of perception and memory.

Let's give a label to the troubling hypothesis that will be our focus here:

> SUPERABUNDANCE: Boltzmann Brains will comprise a large majority of all brains that at some point pass through a sequence of physical states that perfectly mirrors the recent sequence of physical states of your own brain. During these intervals of overlap, Boltzmann Brains have experiences that are phenomenally equivalent to (and indiscernible from) your recent experiences.

When I speak of "scientific support" for SUPERABUNDANCE, I do not mean to imply that there are scientists who believe that SUPERABUNDANCE is true. The sense in which science supports SUPERABUNDANCE is that some of the leading cosmological models predict an eventual superabundance of Boltzmann Brains and that, as physicist Brian Greene has said, it is "surprisingly difficult" (Sample 2020) to formulate plausible models that avoid this implication. While leading models arguably support SUPERABUNDANCE, it is worth noting that cosmology is an unsettled and rapidly evolving area of inquiry. Some cosmologists contend that the ΛCDM model, despite its significant successes, needs to be fundamentally rethought (Binney et al. 2025). For purposes of this discussion, however, I will consider the epistemic implications of learning (or seeming to learn) that there is robust scientific consensus that the best supported models predict SUPERABUNDANCE (at least on the assumption that phenomenal internalism is true).

While it may seem intuitive that such scientific support of SUPERABUNDANCE generates a skeptical challenge, accurately characterizing the nature of this challenge is not a straightforward task. Indeed, I will argue that most discussions of Boltzmann Brain skepticism have failed to correctly characterize the skeptical threat, making it too easy to sweep aside any skeptical concerns. These discussions seem to assume that scientific evidence for SUPERABUNDANCE poses a skeptical worry for you only if, in light of that evidence, you should think it likely that you are a Boltzmann Brain. If this assumption were correct, you would be right to dismiss skeptical worries posed by scientific support for SUPERABUNDANCE, as it is easy to see that you would not be rational in thinking that you are probably a Boltzmann Brain. But as I will explain, the skeptical possibility that poses a genuine concern is not the possibility that you are a Boltzmann Brain, but rather the possibility that your appearances concerning the scientific evidence are deeply misleading. When it appears to you that scientific evidence strongly supports a hypothesis like SUPERABUNDANCE that is in a certain sense self-undermining (since it defeats one's beliefs about the scientific evidence), you thereby have reason to think that you have somehow been misled about the nature of the scientific evidence. Most attempts to dismiss Boltzmann Brain-related skeptical worries simply do not address this more formidable (if more subtle) skeptical challenge.

After arguing for the revised and improved Boltzmann Brain skeptical argument (Sections 2–4), I turn in Section 5 to recent objections to a crucial assumption of the argument. A key commitment of Boltzmann Brain skeptical arguments, both in their traditional formulations and in the improved version I develop here, is a principle of "Center Indifference." Center Indifference says, roughly, that on the supposition that the world contains two observers O1 and O2 whose phenomenal states are internally indistinguishable, your credence that you are O1 should be equal to your credence that you are O2. An implication of Center Indifference is that it is not rational for you to confidently maintain that you are not a Boltzmann Brain while also affirming that Boltzmann Brains will

comprise the majority of all those observers who will have had a phenomenal state that is indistinguishable from your present phenomenal state. Dogramaci and Shoenfield (2025) have recently argued that Center Indifference should be rejected since it entails some absurd constraints on how strongly one's evidence could support certain scientific hypotheses. Most worryingly, they argue that if Center Indifference is correct, then one could not rationally assign a high credence to SUPERABUNDANCE even on the condition that one is an Ordinary Observer (i.e., an observer that is *not* a Boltzmann Brain) with perfectly reliable cognitive faculties. As I will show, their case against Center Indifference rests on some subtle confusions. Center Indifference does not have the worrying implications that they attribute to it.

My aim in this paper is not to provide a full assessment of the skeptical challenge posed by scientific support for SUPERABUNDANCE, but to set the stage for future discussions that proceed from a clear understanding of the nature of that challenge and of the central epistemological questions it raises. One such question concerns the conditions under which a rational presumption against skepticism may be defeated. In the wake of (apparent) scientific evidence for SUPERABUNDANCE, we face a situation where a rational presumption against skepticism stands in tension with our trust in standard methods of scientific inference. The question of how one should navigate such tension is largely unaddressed in the epistemological literature that engages more traditional skeptical problems.

2. The failure of the traditional argument for Boltzmann Brain skepticism

In this section, I consider how the Boltzmann Brain skeptical challenge has standardly been articulated and why this standard way of formulating the challenge fails. Here is one way to formulate the challenge as it has typically been understood:

Traditional BB Skeptical Argument

1. In light of the apparent scientific support for SUPERABUNDANCE, you should have a very high credence for SUPERABUNDANCE.
2. Your credence that you are a Boltzmann Brain, conditional on SUPERABUNDANCE, should be extremely high.
3. If (1) and (2), then your credence that you are a Boltzmann Brain should be high.
4. Thus, your credence that you are a Boltzmann Brain should be high. [From 1–3.]
5. If your credence that you are a Boltzmann Brain should be high, then you are not justified in holding beliefs based on sense perception or memory.
6. Therefore, you are not justified in holding beliefs based on sense perception or memory. [From 4 and 5.][1]

The only premises in this argument that could be contested in a plausible and substantive way are premises (1) and (2).[2] I will briefly describe the arguments for each of these premises.

---

[1] This sort of argument for Boltzmann Brain-related skepticism may be found in Crawford (2013, 253–254), Dogramaci (2020, 3718–19), Carroll (2021, 13–16), Kotzen (2021, 29–30), Avni (2023), Dogramaci and Schoenfield (2025, 3), and Page (2024, 61–62).

[2] (3) is secure so long as a "very high" credence multiplied by an "extremely high" credence is a "high" credence. Quibbling with this assumption on account of the vagueness of these terms does not threaten the central thrust of the argument. (5) appears secure so long as you know that the apparent sense impressions and apparent memories of a Boltzmann Brain are essentially random and thus extremely unlikely to be accurate.

The argument for (1), in brief, is that you have a strong *pro tanto* reason to trust the predictions of well-supported scientific models. Absent some special reason to think that in this case well-supported scientific models should not be trusted, and absent any reason to think that scientists have not in fact said what it appears to you that they have said, you have strong reason to think that SUPERABUNDANCE is probably true. Now, one might resist (1) by insisting that in this case we *do* have some reason to doubt some claim which, according to scientists, is predicted by well-supported cosmological models. This sort of response will be discussed later. For now, I merely aim to gesture at the considerations offered in support of (1).

Premise (2) may be motivated by appealing to Center Indifference, the principle briefly described in the introduction. Center Indifference pertains to one's credences for self-locating propositions. Self-locating propositions have content that is not limited to purely "third-personal" content pertaining to the character of the actual world but that includes (or entails) first-personal content pertaining to one's identity or "location" within that world. The proposition expressed by "There is a country called China" is an example of a non-self-locating proposition, whereas the proposition expressed when I say "I have never been to China" is self-locating since, in uttering this sentence, I locate myself among that subset of individuals who have never been to China. Center indifference principles have been formulated in different ways and defended by various philosophers (Bostrom 2003, 249–250; Elga 2004). I will use Builes's (2024, 780) formulation, which makes use of the notion of a "centered world" (Lewis 1979), which is a designated individual and time within a particular possible world. Builes labels two centered worlds *similar* just in case they are associated with the same possible world. Having introduced this terminology, here is how Builes (2024, 780) formulates the principle:

> Center Indifference (**CI**)**:** For any two similar centered worlds c_1 and c_2, if both c_1 and c_2 are compatible with your evidence, then it is rationally required to set $Cr(c_1 \mid c_1 \text{ or } c_2)=1/2$.

In support of CI, note that the competing possibilities *c1* and *c2* pick out the very same possible world and thus agree on all of the uncentered facts. For this reason, the theoretical virtues that normally guide us in choosing between empirically adequate theories (e.g., the theoretical virtue of simplicity) would seem to have no bearing here (Builes 2024, 787). Given that the relevant locations are both occupied by *someone* and that your evidence does not discriminate between these locations, it would seem to be rationally arbitrary to assign a higher credence to one of these centered possibilities than the other.

Premise (2) would seem to straightforwardly follow from CI if we assume that your evidence is compatible with all and only those centered worlds involving an "internal duplicate" of yours, that is, an individual whose experience (at the designated time) is phenomenally equivalent to your current experience. To see this, consider first a scenario where at some point in the world's history, there will be 1,000 internal duplicates of yours and 999 of these will be Boltzmann Brains. If your evidence is compatible with all and only those centered worlds that involve an internal duplicate of yours, then CI requires that, conditional on this scenario, your credence that you are one of the Boltzmann Brains is 999/1,000. SUPERABUNDANCE is not so precise as this scenario, but SUPERABUNDANCE does entail that Boltzmann Brains will comprise the "vast majority" of your internal duplicates. So, by the same sort of reasoning, CI requires that, conditional on SUPERABUNDANCE, you should assign a very high credence to the hypothesis that you are a Boltzmann Brain.

This reasoning from CI to (2) would likely be resisted by proponents of certain externalist accounts of evidence (e.g., Williamson 2000, chap. 9) which say that your evidence does not consist

only of internally accessible facts about your phenomenology, but may also include a rich set of external facts (e.g., the fact that you are sitting in a chair). On such an account, your evidence would include facts that are not compatible with the hypothesis that you are a Boltzmann Brain. For present purposes, I will set such externalist accounts aside and simply assume that the correct account of evidence is a phenomenal account according to which one's empirical evidence consists of propositions describing the phenomenal character of one's present experience.[3] This assumption is appropriate in the present context since the responses to Boltzmann Brain skepticism that I will be considering do not challenge the phenomenal account or presuppose that an externalist theory of evidence is correct.[4]

Taken on their own, premises (1) and (2) are each plausible. But there is very good reason to think that the Traditional BB Skeptical Argument that relies on these premises is unsound, leading us to the conclusion that (1) and (2) cannot both be true. The reason for thinking that the argument is unsound is that there is a clear and compelling argument against (4), the intermediate conclusion

---

[3] Dogramaci and Schoenfield, after characterizing a phenomenal theory of empirical evidence (2025, 7–8), develop their arguments against CI on the assumption, for sake of argument, that the phenomenal theory is correct (18–19). This assumption is meant to be charitable to those who appeal to CI to motivate Boltzmann Brain-related skepticism.

[4] A second assumption needed to move from CI to (2) is that a single possible world can contain observers who exist at disparate times. Builes (2024, 783–84, 789–792) notes that one might challenge this assumption by appealing to presentism (the view that only present things have existence). Here, I will simply assume that individuals at disparate times may occupy the same possible world (in whatever sense of "possible world" is relevant to CI).

which says that your credence that you are a Boltzmann Brain should be high.[5] If you are rationally required to believe that you are probably a Boltzmann Brain, this would only be because you are rationally required to think it probable that human scientists have said that leading cosmological models predict a large number of Boltzmann Brains.[6] (I take it for granted that, without this belief about scientific testimony, there are no other empirical considerations available to you that might require you to believe that you are probably a Boltzmann Brain.) However, if you are rationally required to believe that you are probably a Boltzmann Brain, then rationality requires that you do *not* think it probable that human scientists have said that cosmological models predict a great many Boltzmann Brains.[7] For on the supposition that you are a Boltzmann Brain, you know nothing about

---

[5] The following argument against (4) is adapted from Dogramaci (2020). A similar argument against (4) is developed by Carroll (2021, 16) in his claim that believing that one is a Boltzmann Brain is "cognitively unstable."

[6] Avni (2023) disagrees with this claim and argues that someone with apparent memories of having received strong scientific evidence for SUPERABUNDANCE is rationally required to think it probable that they are a Boltzmann Brain. Avni argues for this even while acknowledging that, by giving a high credence to the Boltzmann Brain hypothesis, one's scientific evidence is defeated. Avni's argument relies on the premise that you should have a high credence for SUPERABUNDANCE conditional on your being an Ordinary Observer. This superficially plausible premise should be rejected, for reasons that will become clear shortly. In footnote 14, I return to Avni's argument.

[7] On the supposition that you are a Boltzmann Brain, perhaps it is probable that some non-human scientific community somewhere in the universe has reported that the best cosmological model predicts SUPERABUNDANCE. But the only rational basis you might have for thinking that the best

what human scientists have said since your apparent memories are the product of an isolated "brain" assembled by chance. Thus, if you are rationally required to believe that you are probably a Boltzmann Brain, then it is the case both that you must think it probable that human scientists have said that leading cosmological models predict a great many Boltzmann Brains (since this is the only available consideration that might require you to believe that you are a Boltzmann Brain) *and* that you should *not* think it probable that human scientists have said this (since all beliefs about human scientific testimony are defeated by a high credence for the claim that you are a Boltzmann Brain). Since it is incoherent to suppose that you both should and should not hold some belief about human scientific testimony, we can conclude that, contrary to (4), it is *not* the case that you should have a high credence that you are a Boltzmann Brain.[8] Since (4) follows from premises (1)–(3), and since (3) is clearly unproblematic, we should conclude that either (1) is mistaken or (2) is mistaken (or both are).

---

cosmological model predicts SUPERABUNDANCE is your having learned through testimony that scientists have said that there is scientific support for SUPERABUNDANCE. I have focused claims about what *human* scientists have said to focus on the sort of claim that you could learn by testimony rather than a more generic claim that might be inferred from an antecedent belief that you are a Boltzmann Brain.

[8] I take it for granted here that you are not presented with a rational dilemma where you both ought and ought not be in some doxastic state. While some might question this assumption, it's difficult to see how the claim that you *are* in a dilemma could serve to advance a skeptical argument. For it's not clear why one should think that the best response to such a dilemma would involve adopting a radically skeptical stance.

In seeing reason to reject the Traditional BB Skeptical Argument, and (4) in particular, should we thereby conclude that there are no genuine skeptical concerns posed by apparent scientific support for SUPERABUNDANCE? It is tempting to answer affirmatively on the basis of the following assumption:

> **The BB Worry Assumption**: Skepticism is rationally required in response to learning about scientific support for SUPERABUNDANCE only if rationally responding to this evidence requires that one give significant credence to the hypothesis that one is a Boltzmann Brain.

I call this the "BB Worry Assumption" since it expresses the idea that the skeptical threat posed by scientific support for SUPERABUNDANCE arises entirely from the "worry" that one is a BB (i.e., a Boltzmann Brain). If this assumption was correct, then a successful argument against (4) would suffice to defuse skeptical worries posed by scientific evidence for Boltzmann Brain cosmologies.

Let's use the label "non-BB skeptical hypothesis" for any skeptical hypothesis that does *not* involve one's being a Boltzmann Brain. Those who make the BB Worry Assumption dismiss the possibility that, upon learning of apparent scientific support for SUPERABUNDANCE, one should adopt a high credence for some non-BB skeptical hypothesis (even while one's credence that one is a Boltzmann Brain should remain low). For example, Dogramaci and Schoenfield (2025, 5) say that "what you're ultimately really worried about, when confronting the skeptical challenge posed by evidence that we live in a BBU [i.e., a universe where Boltzmann Brains are superabundant], is not really the probability of any cosmological model, but rather the probability that *you* are a BB [i.e., a Boltzmann Brain]." They go on to suggest that when "facing down this new skeptical challenge posed by evidence for BBs," we may ignore non-BB skeptical hypotheses since we may assume that these other skeptical hypotheses "have already been assigned negligible credence by some traditional

anti-skeptical philosophical considerations" (5). But even if every non-BB skeptical hypothesis is rightly assigned a negligible prior credence, a non-BB skeptical hypothesis could still be relevant in the present context if that hypothesis was significantly confirmed by the evidence of apparent scientific support for SUPERABUNDANCE. In ignoring this possibility and supposing that the only relevant worry posed by this evidence is that one is a Boltzmann Brain, Dogramaci and Schoenfield implicitly make the BB Worry Assumption. A number of others writing on Boltzmann Brain-related skeptical concerns also seem to implicitly accept this assumption, since the only skeptical hypothesis given attention is the hypothesis that one is a Boltzmann Brain.[9]

In the next section, I argue that the BB Worry Assumption is mistaken. If the apparent scientific support for SUPERABUNDANCE is sufficiently strong, this supplies the basis for a formidable skeptical argument that is in no way committed to the idea that you should think it probable that you are a

---

[9] This applies to the discussions cited in footnote 1. Two authors who do *not* make the BB Worry Assumption include Elga (2025), who I discuss shortly, and Wallace (2023), discussed in footnote 17. While those who make the BB Worry Assumption overlook what I take to be the more serious skeptical challenge, this is not to say that their discussions have no important bearing on this challenge. Dogramaci and Schoenfield, for example, begin their discussion of Boltzmann Brain skeptical concerns by arguing, on Bayesian grounds, that your evidence does not confirm the hypothesis that you are a Boltzmann Brain. They take this result to be sufficient to resolve any relevant skeptical concern. I do not agree that this result is sufficient to block any relevant skeptical concerns, since I will argue that apparent scientific support for SUPERABUNDANCE raises a skeptical worry that is not tied to the possibility that one is a Boltzmann Brain. However, Dogramaci and Schoenfield then go on to develop objections to CI that, if cogent, would block any skeptical threat posed by apparent scientific support for SUPERABUNDANCE.

Boltzmann Brain. In a recent paper, Elga (2025, 134–135) explains how, even if apparent scientific evidence for SUPERABUNDANCE does not make it rational to assign a high credence to the specific hypothesis that one is a Boltzmann Brain, such evidence could in principle make it rational to assign a high credence to a disjunction of skeptical scenarios. This could happen if the evidence one receives is extremely unlikely to be possessed by a reliable human observer in a universe where SUPERABUNDANCE is *false*, and where receiving this evidence is not so improbable conditional on one's being an unreliable observer subject to some sort of systematic deception. In such a case, suggests Elga, one might reasonably conclude that some skeptical scenario or other probably obtains even though one does not give significant credence to any specific skeptical hypothesis. The argument I develop in the next two sections essentially claims that this formal possibility described by Elga may realistically obtain if evidence continues to mount for cosmological models that predict SUPERABUNDANCE.

3. The *real* skeptical challenge posed by evidence of Boltzmann Brain superabundance

I will now argue that apparent scientific evidence for SUPERABUNDANCE may raise worries about the reliability of your cognitive faculties that are not tied to the Boltzmann Brain hypothesis or to any other specific skeptical scenario. The improved Boltzmann Brain skeptical argument I develop here references the following proposition:

> **WORRYING CONSENSUS**: There is a consensus among human cosmologists that well-supported cosmological models predict that SUPERABUNDANCE is true (at least on the standard assumption of phenomenal internalism).

Let *E* be your present total evidence, which I will assume includes a number of (apparent) memories that give the clear and strong appearance that WORRYING CONSENSUS is true. *E* might include, for example, apparent memories of having read various articles where well-known and highly reputable scientists assert the truth of WORRYING CONSENSUS and describe multiple lines of evidence in support of cosmological models predicting SUPERABUNDANCE.

The following possibilities are mutually exclusive and collectively exhaustive, so I will assume that your credences for these possibilities sum to 1:

i. SUPERABUNDANCE is *true*.

ii. **MISLEADING SCIENCE**: SUPERABUNDANCE is *false* but WORRYING CONSENSUS is *true*.

iii. **MISLEADING APPEARANCES**: SUPERABUNDANCE is *false* and WORRYING CONSENSUS is *false*, but at some point you receive evidence that gives the strong appearance that WORRYING CONSENSUS is true.

iv. **NO CONCERN**: SUPERABUNDANCE is *false*, WORRYING CONSENSUS is *false*, and at no point do you receive evidence that gives the strong appearance that WORRYING CONSENSUS is true.

Presumably, before receiving evidence which represents WORRYING CONSENSUS as being true, it would be reasonable for you to assign a very high credence to NO CONCERN. Prior to receiving the (apparent) testimony of cosmologists about the predicted Boltzmann Brain superabundance, you likely have no reason to think that SUPERABUNDANCE is true or that it will be a prediction of a well-supported cosmological models. But upon receiving evidence that strongly represents WORRYING CONSENSUS as being true, you learn that NO CONCERN is false. At this point, the only remaining possibilities are options (i)–(iii).

Having introduced this terminology, I can now explain the argument. First, affirming SUPERABUNDANCE is not rationally compatible with continued trust that your overall outlook (constituted by your many beliefs) is a reliable guide to reality. Given CI, your credence that you are a Boltzmann Brain must be high conditional on SUPERABUNDANCE. And if you are a Boltzmann Brain, then your memories and sense impressions are completely misleading and, as a result, your overall outlook is likely to be extremely inaccurate.

Second, MISLEADING APPEARANCES is also an unacceptable option for the non-skeptic. On the supposition that there is *not* a consensus among cosmologists that well-supported models predict SUPERABUNDANCE, despite your having many apparent memories of scientists reporting that there is such a consensus, then you must acknowledge that you are subject to some serious sort of cognitive error or deception, one that should shake your confidence in other aspects of your overall outlook. Perhaps, for example, there is some vast, coordinated conspiracy to misrepresent scientific views on Boltzmann Brains. And conditional on there being some such conspiracy, you would not be reasonable in maintaining trust in other matters that have been reported from the same or similar sources. Or perhaps scientists have never said anything about Boltzmann Brains, but you are suffering from some sort of hallucination or are in the midst of some ultrarealistic dream.[10] To the extent that you give credence to such possibilities, you should significantly reduce your confidence in your overall outlook. While some of the ways that might account for MISLEADING APPEARANCES might be less epistemically catastrophic than others, it seems clear that believing MISLEADING APPEARANCES is not an option that comports with a normal level of epistemic self-trust. Someone

[10] Elga (2025, 135) also identifies these skeptical scenarios as hypotheses that might receive evidential confirmation in the event that your evidence is sufficiently improbable conditional on your being a reliable observer in a universe where SUPERABUNDANCE is false.

who commends that we simply disbelieve all of the supposed news stories where scientists talk about Boltzmann Brains would hardly be viewed as commending a commonsense, non-skeptical option.

To rationally resist skepticism, then, your only option is to assign a high credence to MISLEADING SCIENCE.[11] You must hold that even though there is scientific consensus that well-supported cosmological models predict SUPERABUNDANCE (at least on the standard assumption of phenomenal internalism), SUPERABUNDANCE is false. In assigning a high credence to MISLEADING SCIENCE, you would, presumably, assign a high credence to the following disjunction: the prediction of a superabundance of Boltzmann Brains by the leading cosmological models is a product of one or more errors in these models, or scientists are somehow mistaken in thinking that these models predict a superabundance of Boltzmann Brains, or phenomenal internalism is false (so that a Boltzmann Brain that has the same internal physical properties as your brain would not have a phenomenal state that is just like your phenomenal state).

While MISLEADING SCIENCE is the only position which offers a chance of escaping skepticism, it would arguably be unreasonable for you to assign it a high credence (at least if the apparent scientific support for SUPERABUNDANCE is extremely strong). Doing so, it may be argued, would either involve groundless opposition to the weight of scientific authority or an implausible rejection of

---

[11] Some accounts of how MISLEADING SCIENCE could be true would involve significant concessions to skepticism. For example, if you held that scientific support of SUPERABUNDANCE should be rejected as misleading because the scientific enterprise as a whole is wildly misguided and riddled with errors, this would be to take a significant step in a skeptical direction. To avoid significant concessions to skepticism, you must affirm MISLEADING SCIENCE without taking this result to be symptomatic of some radical failure infecting our scientific theorizing.

phenomenal internalism. To reject phenomenal internalism in favor of an externalist theory of consciousness is to posit a strange sort of non-local causation, one where phenomenal states of a given brain are directly sensitive to features of the past or the spatiotemporal environment beyond the brain. And despite the fallibility of science, in general you should think it highly unlikely that a firm prediction of an extremely well-supported scientific model is false.

The improved BB skeptical argument just sketched may be formulated as follows:

Improved BB Skeptical Argument

7. You are rational in maintaining robust confidence in your overall outlook only if it is rational for you to assign a high credence to MISLEADING SCIENCE.
8. It is not rational for you to assign a high credence to MISLEADING SCIENCE.
9. Therefore, it is *not* the case that you are rational in maintaining robust confidence in your overall outlook.

As already explained, the case for (7) rests on CI, since it is CI (or some principle like it) which entails that you cannot rationally affirm SUPERABUNDANCE while remaining confident that you are not a Boltzmann Brain. And the case for (8) rests in large part on a requirement that one give due weight to the findings and predictions of science (absent any good evidence that the scientific evidence is in this case misleading).

In the following sections, I will develop the case for (8) and address objections to CI that, if cogent, would undermine the case for (7). But first, it is worth spelling out why this argument represents an advance over the Traditional Argument. Crucially, the Improved BB Skeptical Argument is *not* committed to the claim that you should assign a high credence to the proposition

that you are a Boltzmann Brain. The argument entails that a rational position will give significant credence to the disjunction of SUPERABUNDANCE and MISLEADING APPEARANCES, but so long as you may give significant credence to the latter proposition, there will be no rational requirement that you think it probable that you are a Boltzmann Brain.

It might be objected that, conditional on either SUPERABUNDANCE or MISLEADING APPEARANCES being true, you should give greater credence to SUPERABUNDANCE. After all, the only specific skeptical scenario rendered salient by the apparent scientific consensus is the scenario where SUPERABUNDANCE is true and you are a Boltzmann Brain. If this was right, then while the (so-called) Improved BB Skeptical Argument does not explicitly state that you should be confident that you are a Boltzmann Brain, this would be an implication of the argument. In this case, the BB Worry Assumption would be correct.

But on reflection, the claim that you should give more credence to SUPERABUNDANCE than to MISLEADING APPEARANCES is implausible.[12] MISLEADING APPEARANCES is essentially the disjunction of every scenario *other* than SUPERABUNDANCE that could explain how you have come to receive evidence which strongly represents WORRYING CONSENSUS as being true even though WORRYING CONSENSUS is in fact false. Among the skeptical possibilities encompassed within MISLEADING APPEARANCES are scenarios involving "Truman Show" type conspiracies, ultrarealistic dreams, existence within a realistic but deeply deceptive computer simulation, and so on. Now, consider first which of the following hypotheses has greater probability: (a) SUPERABUNDANCE is true and you are a randomly assembled Boltzmann Brain which just by chance has a rich and coherent body of

---

[12] The explanation of why this is implausible may also be found in Elga's (2025, 130–2) discussion of the "Labelscramble" case (especially Version 3), which Elga uses to shed light on Boltzmann Brain concerns.

evidence that, by happenstance, represents SUPERABUNDANCE as being supported by scientific evidence; or (b) SUPERABUNDANCE is false and you are involved in some other sort of scenario which has resulted in your having a rich and coherent body of evidence that incorrectly represents SUPERABUNDANCE as being supported by scientific evidence. Option (b), which encompasses a much, much wider range of possibilities, should be assigned a significantly higher credence.[13] Thus, we should reject the suggestion that the Improved BB Skeptical Argument is implicitly committed to the BB Worry Assumption. The Improved BB Skeptical Argument says that the majority of your credence should be assigned to the disjunction SUPERABUNDANCE or MISLEADING APPEARANCES, and since the latter should receive greater credence than the former, this result does not require you to think it probable that you are a Boltzmann Brain.[14]

---

[13] It should be noted that option (a) is only one of the two possibilities encompassed within SUPERABUNDANCE. Your credence for SUPERABUNDANCE should be your credence for option (a) *plus* your credence for the possibility that SUPERABUNDANCE is true and you are an Ordinary Observer. However, CI entails that this additional possibility must be assigned a much lower credence than (a), which means that your credence for SUPERABUNDANCE will be only slightly higher than your credence for (a). Since your credence for MISLEADING APPEARANCES should be assigned a *much* higher credence than your credence for (a), we can conclude that your credence for MISLEADING APPEARANCES should also be higher than your credence for SUPERABUNDANCE.

[14] We are now in a position to appreciate where Avni's (2023) argument (that you should believe you are a Boltzmann Brain) goes wrong. A crucial premise of Avni's argument is that, conditional on your being an Ordinary Observer and on your having apparent memories of strong scientific evidence for SUPERABUNDANCE, you should be confident in SUPERABUNDANCE. (This is the premise

Since the Improved BB Skeptical Argument does not imply that you should think that you are probably a Boltzmann Brain, it avoids the specific self-undermining objection was pressed against premise (4) of the Traditional BB Skeptical Argument. But might the Improved BB Skeptical Argument be subject to another sort of self-undermining problem? One might argue that, conditional on MISLEADING APPEARANCES, it is probable that you are suffering from serious cognitive problems that compromise not only your perceptual evidence, but also the rational faculties involved in your philosophical reasoning. In this case, you arguably would not be rational in giving significant credence to MISLEADING APPEARANCES, since arriving at the conclusion that MISLEADING APPEARANCES is probably true would undermine trust in the very reasoning capacities that led you to that conclusion.

While this self-undermining objection raises important issues that cannot be thoroughly explored here, one way of responding is to call into question the suggestion that MISLEADING

---

that Avni labels SCIENCE.) Here's why this premise should be rejected. Let BB stand for the proposition that you are a Boltzmann Brain, OO stand for the proposition that you are an Ordinary Observer, and E stand for your total evidence, including your apparent memories of scientific support for SUPERABUNDANCE. For the reasons just explained, it is at least *permissible* for you to assign greater a priori probability to (MISLEADING APPEARANCES & OO & E) than to (SUPERABUNDANCE & BB & E). But given CI (which Avni accepts), we know that (SUPERABUNDANCE & BB & E) has greater a priori probability than (SUPERABUNDANCE & OO & E). Thus, by transitivity, we may conclude that it is at least permissible to assign greater a priori probability to (MISLEADING APPEARANCES & OO & E) than to (SUPERABUNDANCE & OO & E). But this in turn entails that, contrary to Avni's "SCIENCE" premise, it is *not* the case that, conditional on OO & E, you should be highly confident in SUPERABUNDANCE.

APPEARANCES should be attended by serious doubts concerning your capacities for philosophical reasoning. Many of the skeptical scenarios encompassed within MISLEADING APPEARANCES are ones that target your faculties of perception or memory but leave your reasoning abilities and capacities for introspection unscathed. For this reason, it is not obviously incoherent to assign a high credence to MISLEADING APPEARANCES while maintaining confidence in the reasoning that serves as the basis for that credence.

I have argued that the skeptical hypothesis that should be of concern is *not* the hypothesis that you are a Boltzmann Brain, but rather MISLEADING APPEARANCES. When the skeptical challenge is reframed in this way, many of the responses that have been given to Boltzmann Brain skeptical worries are rendered irrelevant. Generally, responses to the Boltzmann Brain skeptical challenge have labored to show that we are not rationally pressured to believe that one is a Boltzmann Brain, either because believing this would be self-undermining in some way (Carroll 2021; Dogramaci 2020) or because the coherence and ordered character of our experience constitutes strong empirical evidence against this possibility (Kotzen 2021; Dogramaci and Schoenfield 2025). A proponent of the Improved BB Skeptical Argument may simply grant these claims. The fact that one has decisive reason to reject the hypothesis that one is a Boltzmann Brain does nothing to ease the worry that MISLEADING APPEARANCES is the most plausible explanation of one's situation.

If the most common responses to Boltzmann Brain skeptical worries do not pertain to the Improved BB Skeptical Argument, what avenues of response *are* available? There are two principal options. First, one could contest (7) by arguing for the rejection of CI (and any principle relevantly like it). If CI may be reasonably be rejected, then you might have the option of simply accepting SUPERABUNDANCE, all the while maintaining confidence that you are an Ordinary Observer who can learn things by way of scientific testimony. This stance is defended by Dogramaci and Schoenfield, and I will take up their objections to CI in Section 5.

Second, one could contest (8) by arguing that you may rationally assign a high credence to MISLEADING SCIENCE (either by maintaining that the scientific evidence is misleading or by rejecting phenomenal internalism). I will not in this paper attempt anything like a full assessment of (8). My main aims here are to present a more formidable Boltzmann Brain skeptical argument that is not committed to the BB Worry Assumption and to answer what I take to be misguided arguments against CI. But I do want to show that a strong case may be made for (8). After all, if (8) was obviously false—that is, if it was clear that you should assign MISLEADING SCIENCE a much higher credence than MISLEADING APPEARANCES—then one might doubt whether the skeptical argument developed in this section is really an improvement on the Traditional BB Skeptical Argument. In the next section, I will attempt to show that the case for (8) at least has significant prima facie plausibility.

4. On the respective merits of MISLEADING SCIENCE and MISLEADING APPEARANCES

Whether your appearances concerning the scientific support for SUPERABUNDANCE are better explained by MISLEADING SCIENCE or by MISLEADING APPEARANCES will depend on how strong the scientific support for SUPERABUNDANCE appears to be. The greater the degree of support that some hypothetical body of scientific evidence gives to SUPERABUNDANCE, the more improbable it is that this body of evidence would be obtained by an otherwise reliable human scientific community in a universe where SUPERABUNDANCE is *false*. In a case where it appears to you that there is *incredibly strong* scientific support for SUPERABUNDANCE, it could very well be that the probability that human scientists have actually obtained this evidence in a universe where SUPERABUNDANCE is false is lower than the probability that scientists have *not* obtained such evidence and your appearances to the contrary are misleading. Consider, for example, a case where SUPERABUNDANCE appears to have a

degree of scientific support that is on a par with the degree of scientific support that is (in fact) enjoyed by the claim that there are multiple galaxies. In such a case, it would arguably be more reasonable to conclude that your appearances concerning the scientific evidence are seriously misleading than to think that the entire community of human scientists has been so seriously misled by the cosmological data.

I readily concede that, at present, the apparent scientific support for SUPERABUNDANCE is too weak to generate much pressure to assign significant credence to MISLEADING APPEARANCES. As noted in the introduction, present cosmology is far from settled and some experts have expressed doubts concerning the adequacy of the ΛCDM model. Given this, it would not be especially surprising to learn that the leading cosmological models are mistaken in their prediction of SUPERABUNDANCE. But might there be realistic scenarios where the (apparent) scientific evidence continues to strengthen, even to the point where it ceases to be reasonable to assign more credence to MISLEADING SCIENCE than to MISLEADING APPEARANCES?

Some might argue that such a scenario is not a realistic possibility since there is an a priori rational presumption in favor of epistemic self-trust that makes it reasonable for us to conclude that the scientific evidence is misleading whenever this is the only position that avoids skepticism.[15] However, even if we affirm that you have a priori justification to assign a very low prior probability to the disjunction of all skeptical hypotheses, you may still be required to assign a high probability to certain skeptical hypotheses *conditional* on various worrying sorts of evidence that you might receive. And if you then receive such worrying evidence, the initial presumption against skepticism may be

[15] For arguments that we have a priori justification to affirm our cognitive reliability, see, for example, Cohen (2010) and DeRose (2017, chap. 7).

defeated. Consider the following example. Imagine that throughout the day I gather coins from various places—some I receive as change at the dry cleaners, others I find under my couch cushions, others I remove from public fountains, and so on. Nothing seems unusual about any of the coins I've collected in this haphazard way. In the afternoon, after I've collected one hundred coins, I walk into some nearby woods far away from any observers and I proceed to flip each of the collected coins exactly one time. To my amazement, I observe that every single coin lands tails. In light of this experience, it seems that I cannot completely discount the following possibility:

> **SKEPTICAL SCENARIO:** My experience of seeming to see all 100 flipped coins land tails is explained by my being subject to some sort of skeptical scenario (an ultrarealistic dream, for example, or the manipulation of my sense impressions and/or memories by some deceptive agent).

Plausibly, SKEPTICAL SCENARIO should be given more credence than the hypothesis that the 100 coins have landed tails purely by chance.[16] Granted, there are non-skeptical explanations of the result that do not appeal to mere chance. Perhaps some agent (God? a spirit? an extremely sophisticated human trickster?) has manipulated the outcome of the tosses in some manner that does not involve

---

[16] Assuming that the coins are fair, the probability that all 100 tosses would by chance have the same result—either all heads or all tails—is vanishingly small: $1/2^{99}$. By contrast, the probability of witnessing a uniform result when *dreaming* about tossing coins, or when being given deceptive coin toss experiences by a manipulating agent, does not seem nearly so low. Even if skeptical scenarios should be given tiny priors, it is arguable that this is a scenario where a skeptical hypothesis should receive more credence than the non-skeptical, "mere chance" hypothesis.

tampering with my perceptual faculties or memory (or any other actions that would result in some skeptical scenario). But these non-skeptical explanations seem incredibly far-fetched (though just *how* far-fetched will, admittedly, be a contested matter). Arguably, they are *more* far-fetched than the claim that I am caught up in some sort of ultrarealistic dream or hallucination. Quite plausibly, the most reasonable response in this situation would be for me to assign a high credence to SKEPTICAL SCENARIO, despite the availability of non-skeptical alternatives.

Matters are arguably similar in a case where it appears to you that multiple strands of scientific evidence provide extremely strong support to SUPERABUNDANCE (support that is on a par, say, with the evidence for multiple galaxies). Because you cannot rationally believe SUPERABUNDANCE, the bulk of your credence must be divided between a position which posits an enormously improbable confluence of misleading scientific evidence (MISLEADING SCIENCE) and one that succumbs to skepticism (MISLEADING APPEARANCES). A presumption against skepticism is thus set in opposition to the presumptive reliability of science. It is not implausible that in such a case, a defeasible presumption against skepticism is at least partially defeated.[17]

---

[17] In a recent piece addressing Boltzmann Brain cosmology, David Wallace (2023) appeals to an a priori presumption against skepticism to argue that one should never give significant credence to a proposition like MISLEADING APPEARANCES. He draws a distinction between the *primary* scientific evidence (e.g., experimental results) and the *proximal* evidence by which we come to form beliefs about the primary evidence (e.g., journal articles, memories of scientific testimony, etc.). In Wallace's discussion, E stands for the proximal evidence for a scientific theory (essentially, what I am calling the *appearances*), and H stands for the hypothesis that E obtains and that the primary evidence is "approximately what the proximal evidence says that it is" (297). Wallace notes that E&~H (which

One might argue that the apparent scientific evidence for SUPERABUNDANCE could never be so strong as to push one to endorse MISLEADING APPEARANCES since the case for SUPERABUNDANCE relies on philosophical presuppositions that may be questioned and that are beyond the purview of science. For example, SUPERABUNDANCE would be false if phenomenal internalism is false or if there is a God who, wanting to avoid rampant deception, will annihilate the physical universe before there is sufficient time for Boltzmann Brains to arise.[18] If the disjunction of such philosophical and theological theories (which rule out deceived Boltzmann Brains) has greater probability than MISLEADING APPEARANCES, then there would not be significant pressure to move towards skepticism even in the face of mounting scientific evidence for models that predict superabundant Boltzmann Brains. But consider those thinkers who (rightly or wrongly) think that externalist theories of consciousness and appeals to divine providence (and other such "non-scientific" escape routes) are collectively less plausible than the various skeptical ways of explaining the apparent

---

is essentially equivalent to MISLEADING APPEARANCES) is a skeptical hypothesis and then says the following: "Assuming that we give low priors to skeptical scenarios, Pr(~H | E)<<1" (297) (where "<<" means much, much less than). But it is simply incorrect that assigning a low prior to E&~H entails that Pr(~H | E) is low. So long as my prior for E is very low, Pr(E&~H) will be low even if Pr(~H | E) is high. An a priori presumption against skepticism, then, will not deliver the result that Pr(~H | E) is low. To get that result, one would need to posit that the initial presumption against skepticism is *empirically indefeasible*, a much stronger claim that is not especially plausible.

[18] See Saad (2024) for a development of a radically externalist theory of consciousness that blocks skeptical worries posed by Boltzmann Brains or other sorts of Boltzmann observers.

scientific evidence (i.e. an ultrarealistic dream, a deceptive simulation, and so on). At least for these thinkers, MISLEADING APPEARANCES may be the only credible option when it appears that there is incredibly strong scientific support for cosmological models that predict superabundant Boltzmann Brains.

5. Defending Center Indifference against recent objections

The Improved BB Skeptical Argument, like the Traditional BB Skeptical Argument, crucially depends on CI. If we did not affirm any constraint like CI, then one could simply "follow the science" and believe SUPERABUNDANCE. CI gives us reason to think that this option is self-undermining and thus irrational.

Dogramaci and Schoenfield have recently argued that CI has absurd consequences and should be rejected. According to them, it is not *only* the case that CI rules out confidence in SUPERABUNDANCE, a hypothesis they label "BBU" (for "Boltzmann Brain Universe"). A further implication of CI, they allege, is that there are severe constraints on the extent to which your evidence may confirm BBU *even on the supposition that you are an Ordinary Observer* (2025, 20). So, letting "OO" stand for the hypothesis that you are an Ordinary Observer and Pr be some credence function that is rational for you, CI implies (according to Dogramaci and Schoenfield) that, for *any* body of evidence E, there are strong upper limits on Pr(BBU | E & OO). And this seems to them to be an absurd result. Supposing OO is true, you are reasonable (according to Dogramaci and Schoenfield) in trusting your cognitive faculties and following the scientific evidence wherever it leads. So, while one might plausibly argue that your *unconditional* confidence in BBU must be limited (since confidence in BBU requires high credence for BB which in turn defeats any evidential grounds for confidence in BBU), it seems that there should be no special limits on how confident

you may be in BBU *conditional* on your being an Ordinary Observer. If CI entailed that there was such a limit, this would arguably give us a strong reason to question CI.

In light of the argument in the earlier sections of this paper, one might think that I should be happy to accept that there are strong upper limits on Pr(BBU | E & OO). After all, I have argued that when it appears to you that there is strong scientific support for BBU, this raises skeptical worries that are not tied to the possibility that you are a Boltzmann Brain. Receiving such evidence gives you reason to think that, even if you are not a Boltzmann Brain, your appearances concerning the scientific evidence are not accurate. If that is right, then it is not surprising that Pr(BBU | E & OO) should not be high.

This response to Dogramaci and Schoenfield's worry, however, would not be an adequate defense of CI. The reason is that their formal argument allegedly showing that CI constrains Pr(BBU | E & OO) would still go through even if we took it as given that Ordinary Observers are perfectly reliable. The counterintuitive result that they derive cannot be explained away by noting that we must take seriously the possibility that an Ordinary Observer is somehow deceived, since the premises they put forward and the inferences they draw would appear no less legitimate on the assumption that you know with certainty that MISLEADING APPEARANCES is false and that the appearances of Ordinary Observers are accurate indicators of the scientific evidence. And if CI implied that there are strong constraints on your credence for BBU *even conditional on your being a reliable Ordinary Observer with accurate appearances about the scientific evidence*, this *would* be absurd.

As I will now show, Dogramaci and Schoenfield are mistaken in thinking that CI supports an artificial limit on your credence for BBU given OO. In the discussion of their argument to follow, I will assume that you have somehow ruled out with certainty any skeptical scenarios that do *not* involve your being a Boltzmann Brain. So OO can be taken to entail not only that you are an

Ordinary Observer, but also that you are an Ordinary Observer whose appearances concerning the scientific evidence are perfectly accurate.

To construct their argument, Dogramaci and Schoenfield use H_E as a label for "the hypothesis that the vast majority of brains with total evidence E are BBs" (i.e., are Boltzmann Brains) (22).[19] For sake of concreteness, they subsequently suggest that "vast majority" could be interpreted as meaning at least 99%, and I will adopt this interpretation. $\sim H_E$, they then assert, "is the hypothesis that some correspondingly smaller proportion (anything less than a vast majority) of brains with E are BBs."[20]

---

[19] Dogramaci and Schoenfield's argument that I will be discussing here is on pp. 22–25 of their article. I will avoid repetitious citations in my exposition of their argument. There is one worry about the way that Dogramaci and Schoenfield characterize H_E. On the phenomenal account of evidence which they are here taking granted (at least for sake of argument), my evidence consists in centered propositions about *my own* present phenomenal experience (7). On this view, a distinct observer whose phenomenal state is internally indistinguishable from my phenomenal state would arguably not count as having the same evidence as me. For that observer would have as evidence information about *their* phenomenal experience rather than the information that I have about *my* phenomenal experience. A simple fix here would be to characterize H_E as a thesis about the proportion of Boltzmann Brains among those observers whose phenomenal state is indistinguishable from the phenomenal state you would have when E is your total evidence. To ease discussion, however, I will follow Dogramaci and Schoenfield in speaking as though a single body of evidence may be had by multiple observers.

[20] Strictly speaking, this is not a precise characterization of $\sim H_E$ since one way for H_E to be false is for there to be no brains with evidence E. This inaccuracy in the characterization of $\sim H_E$ has little

They use BB to abbreviate the hypothesis "I'm a Boltzmann Brain" and OO to abbreviate the hypothesis "I'm an Ordinary Observer." Since I have been framing the relevant Boltzmann Brain arguments in second-person rather than first-person terms, I will use the label BB to stand for the proposition that *you* are a Boltzmann Brain, and OO for the proposition that *you* are an Ordinary Observer (whose appearances concerning the scientific evidence are perfectly reliable).

Dogramaci and Schoenfield note, correctly, that an immediate implication of CI is the following:

Lemma: $Pr(BB \mid E \,\&\, \sim H_E) < Pr(BB \mid E \,\&\, H_E)$

In other words, given CI, it follows that your confidence that you are a Boltzmann Brain must be higher given H_E (which is a thesis about the proportion of Boltzmann Brains among observers with evidence E, which by supposition is *your* evidence) than given $\sim H_E$. To draw out absurd implications from Lemma (and thus from CI), they rely on the following premise, which they take to be uncontroversial:

Defeat: $Pr(H_E \mid BB) = Pr(H_E \mid BB \,\&\, E) < \frac{1}{2}$

Their rationale for accepting Defeat is as follows. First, on the supposition that you are a BB, your evidence is completely random and therefore (it is claimed) cannot be taken to have any bearing on

---

bearing on the overall argument, which is mostly concerned with probabilities conditional on E. Since E is phenomenal evidence, the condition where E is true is also a condition where there is some brain—namely yours—that has E; this being the case, $\sim H_E$ *is* equivalent to the claim that the proportion of brains with E that are Boltzmann Brains is smaller than 99%.

H_E or any other cosmological hypothesis. Essentially, your being a Boltzmann Brain would render your evidence useless as a guide to the character of the external world. This, they claim, explains why $Pr(H_E \mid BB) = Pr(H_E \mid BB \,\&\, E)$. Second, Dogramaci and Schoenfield argue that $Pr(H_E \mid BB)$ should be significantly less than ½ on account of the following considerations: (i) H_E is "an antecedently wildly improbable hypothesis" (22); and (ii) the single data point of your being a Boltzmann Brain would at best only slightly boost the probability of H_E. In support of the latter claim, note that your being a Boltzmann Brain would not do much to confirm H_E over other similar hypotheses about the proportion of Boltzmann Brains among those who share your evidence (e.g., the hypothesis that the Boltzmann Brain proportion is between 98% and 99%, or between 97% and 98%, etc.).

Dogramaci and Schoenfield show that Lemma and Defeat collectively entail certain results that they think are absurd. Since they think Defeat should be accepted as uncontroversial, their *reductio* argument is meant to give us reason to reject Lemma and thus CI (which entails Lemma). Here (using their labels) are the first of the putatively problematic results that follow from the conjunction of Lemma and Defeat:

Result 1: $Pr(H_E \mid E) < ½$

Result 2: $Pr(H_E \mid E \,\&\, OO) < ½$

I will not rehearse the formal proofs offered by Dogramaci and Schoenfield, as I believe that they have correctly shown that Result 1 and Result 2 follow from Lemma and Defeat. Rather, I will argue that we should reject Defeat. But first, let's consider what Dogramaci and Schoenfield say about Results 1 and 2.

Result 1, which holds for *any* body of evidence E, says that your credence for H_E could never rationally exceed ½. Dogramaci and Schoenfield write that "Result 1 looks like a weird kind of scientific revisionism: from philosophical premises, we've proved that our scientific evidence cannot strongly confirm a cosmological model according to which most brains like our own are BBs" (23). While Dogramaci and Schoenfield think that Result 1 is an implausible constraint, they concede that defenders of CI may not share their inclination to view Result 1 as objectionable. Indeed, in earlier sections of this paper, I defended something like Result 1, though by a different argumentative path. Since (assuming CI is correct) confidence in SUPERABUNDANCE would defeat any evidential grounds one might have for such confidence, one cannot justifiably be confident in the truth of SUPERABUNDANCE. By the very same reasoning, it also follows that one cannot have evidential grounds for confidence in H_E, a hypothesis that affirms Boltzmann Brain superabundance at least among those brains with evidence just like yours. Result 1, then, does not appear problematic from the perspective of someone who endorses CI.

Result 2, on the other hand, *is* absurd, at least given our present assumption that Ordinary Observers are perfectly reliable. When reasoning on the supposition that you are an OO who has perfectly reliable faculties of perception and memory, there is no reason to think that there should be some strong upper limit on your credence for H_E.

I will argue that there is no pressure on the proponent of CI to accept Result 2. Before diagnosing where this argument against CI errs, it is helpful to note that the argument Dogramaci and Schoenfield have offered generalizes in some very implausible ways, impugning uncontroversial principles of statistical reasoning that have nothing to do with centered propositions. For example, suppose that you are about to randomly sample a single "Sim" from a massive virtual world run on a supercomputer. You know that there are two kinds of Sims: Ordinary Observer Sims and Boltzmann Brain Sims. Ordinary Observer Sims are completely reliable in the beliefs they form

about their simulated world, while Boltzmann Brain Sims have beliefs that are formed in a completely random way. Before randomly sampling a Sim, you do not know anything about the proportion of Boltzmann Brain Sims within the simulation. Upon sampling a Sim, you are initially able to learn its belief state and nothing more. Let B stand for the specific belief state had by the randomly selected Sim, and let E* stand for the information you learn when you learn that a randomly sampled Sim has belief state B. Similarly, let H_{E^*} stand for the proposition that at least 99% of the Sims with B are Boltzmann Brain Sims; let BB* stand for the proposition that the randomly sampled Sim is a Boltzmann Brain Sim; and let OO* stand for the proposition that the randomly sampled Sim is an Ordinary Observer Sim. Finally, consider the following:

Lemma*: $Pr(BB^* \mid E^* \,\&\, \sim H_{E^*}) < Pr(BB^* \mid E^* \,\&\, H_{E^*})$

Defeat*: $Pr(H_{E^*} \mid BB^*) = Pr(H_{E^*} \mid BB^* \,\&\, E^*) < \frac{1}{2}$

Whatever one thinks about CI and the norms that apply when reasoning about self-locating propositions, Lemma* should clearly be accepted. Lemma* is the uncontroversial claim that your confidence that the randomly sampled Sim is a Boltzmann Brain Sim should be higher on the supposition that more than 99% of Sims like this one are Boltzmann Brain Sims than on the supposition that some smaller percentage of the Sims like this one are Boltzmann Brain Sims.

Furthermore, the case for Defeat* would appear to be at least as strong as the case for Defeat. For we may suppose that H_{E^*} has a very low antecedent probability. And while BB* would be a single evidential datapoint in favor of a large quantity of Boltzmann Brain Sims, there is no reason to think that it would boost H_{E^*} over ½. Finally, if it's really the case that the evidence of a Boltzmann Brain is completely useless as a guide to physical reality (so that it has no bearing on H_E), then it

would also seem that the belief set of a Boltzmann Brain Sim, being perfectly random, is useless as a guide to its virtual world (so that $Pr(H_{E^*} \mid BB^*) = Pr(H_{E^*} \mid BB^* \,\&\, E^*)$).

Lemma* and Defeat* collectively entail $Pr(H_{E^*} \mid E^* \text{ and } OO^*) < ½$. This result is absurd. On the supposition that the randomly sampled Sim is a perfectly reliable Ordinary Observer Sim, it should be possible for us to learn any number of things from the belief state of the Ordinary Observer Sim, including facts that would be strong evidence for H_{E^*}. Since Lemma* and Defeat* collectively entail an absurd result, and since Lemma* is obviously correct (as even someone who rejects CI should concede), we should conclude that Defeat* is false. But this in turn should lead us to suspect that Defeat is also false, since any reasoning that might seem to support Defeat would equally apply to Defeat*.

While I think that this *reductio* argument should convince us that Dogramaci and Schoenfield's argument against CI is flawed in some way, I will now show where their argument goes wrong and why Defeat should in fact be rejected. Each part of Dogramaci and Schoenfield's argument for Defeat is mistaken, resting on subtle confusions. One crucial misstep comes in their argument that $Pr(H_E \mid BB \,\&\, E) < ½$. Recall that in arguing for this, they suggest that H_E is an "antecedently wildly improbable hypothesis." Why think this? Presumably, one reason for thinking that H_E is very improbable is that, even conditional on there being one or more Boltzmann Brains, there is no special reason to think that it is especially likely that these Boltzmann Brains are so numerous as to make it probable that they constitute over 99% of the observers with your evidence. They might constitute 2% of such observers, or 49%, or 97%, etc. The highly specific hypothesis that Boltzmann Brains constitute over 99% of such observers is quite improbable, given the large number of alternatives that appear to be no less probable.

This sort of reasoning might be convincing if we were considering a hypothesis which concerned the overall proportion of Boltzmann Brains among all observers. But that is not what H_E

is. H_E is a hypothesis about the proportion of Boltzmann Brains among observers with evidence E. And for certain specifications of E, $Pr(H_E \mid BB \& E)$ will be extremely high *even if the overall proportion of Boltzmann Brains among all observers is extremely low*. Indeed, as I will explain, for values of E such that E is evidence that strongly supports H_E (conditional on OO), $Pr(H_E \mid BB \& E)$ will be high, making the Defeat premise false.

To explain, note that there are at least two very different sorts of reasons that could make it the case that $Pr(H_E \mid BB \& E)$ is very high:

(a) Boltzmann Brains are so abundant that, for *any* set of evidence, it is probable that Boltzmann Brains with that evidence will vastly outnumber Ordinary Observers with that evidence.

(b) E includes observations that are *extremely* unlikely to be made by a reliable Ordinary Observer in any universe where Boltzmann Brains are *not* superabundant; for this reason, even in a universe where Boltzmann Brains are not very abundant, if there is an observer with evidence E, it is most likely that E is possessed only by one or more Boltzmann Brains and not by any Ordinary Observers.

It is obvious how (a) could make it the case that $Pr(H_E \mid BB \& E)$ is high (for any possible E). But it might be less obvious how a high value for $Pr(H_E \mid BB \& E)$ could result from E being specified in a way that makes (b) true. To illustrate this possibility, suppose my evidence E includes an apparent memory that yesterday I flipped exactly 100 fair coins and that the number of coins which landed tails was 100. The corresponding hypothesis H_E will likely be true on the condition that Boltzmann Brains are so numerous that they comprise the vast majority of brains with orderly and coherent experience (experience of the sort that an Ordinary Observer might have). But now suppose that

Boltzmann Brains comprise only 10% of all brains with orderly and coherent experience. Even on this condition, H_E might be highly probable. It is extremely improbable that a perfectly reliable Ordinary Observer would have E, since it is overwhelmingly unlikely that 100 coin tosses would yield anywhere close to 100 tails results. However, given that a *Boltzmann Brain* has an apparent memory about the number of tails resulting from a hundred tosses, there is no reason to think that 100 is an especially unlikely number to feature in that apparent memory. Since a Boltzmann Brain's apparent memory is randomly formed, the "remembered" number of tails should not be expected to reflect the objective probabilities that would pertain to an actual coin toss experiment. Thus, even if Boltzmann Brains comprise only 10% of all brains with coherent and orderly experience, we can expect that Boltzmann Brains will comprise a much higher percentage of those brains with evidence E in particular. H_E, then, could be quite probable even in a universe where Ordinary Observers greatly outnumber Boltzmann Brains with coherent and orderly experience. At least for *this* specification of E, there is reason to question Dogramaci and Schoenfield's claim that H_E is an "antecedently wildly improbable hypothesis," a claim that undergirds their case for Defeat.

The above example casts doubt on Defeat, but a more detailed discussion, with greater mathematical explicitness, is needed to show that Defeat should be rejected in the specific case of evidence that would strongly support H_E (conditional on OO). One example of such evidence, according to Dogramaci and Schoenfield, would be when your total evidence E includes your apparently having "just observed some very large sample of brains with evidence E, and they are all BBs" (24).[21] The case I will now discuss takes inspiration from this suggestion. While I affirm that

---

[21] It should be noted that Dogramaci and Schoenfield contend that their objections to CI go through even on a phenomenal conception of evidence (18). So they should be happy to characterize this

such evidence would support a high credence for H_E conditional on OO, this affirmation poses no problem for the proponent of CI. For on this way of specifying E, the Defeat premise is false.

Let's imagine that you have, somehow, randomly sampled one thousand brains from the set of all brains throughout time and space whose total (phenomenal) evidence is coherent and of the sort that might be had by a normally functioning and reliable Ordinary Observer. (I will call this sort of evidence "normal evidence"; normal evidence excludes chaotic experiences and other phenomenal states that are outside of the range of experiences accessible to perfectly reliable Ordinary Observers.[22]) Using advanced scientific equipment, you are running a test which will tell you exactly how many of the sampled brains are Boltzmann Brains. This number will be displayed on a machine's electronic display. Let E0 be the total evidence you will have if and when you see the machine display the number 0, let E1 be the total evidence you will have if and when you see the machine display the number 1, and so on. Assuming you are an Ordinary Observer, your total evidence in a moment's time will be one of the 1,001 members of the set {E0, E1,…, E1000}.

Consider, now, E1000 and the corresponding hypothesis H_{E1000}. It's tempting to think that $Pr(H_{E1000} \mid BB \& E)$ should be high if and only if the probability of overall Boltzmann Brain superabundance is high. But that is not the case. $Pr(H_{E1000} \mid BB \& E)$ will be high even on an

---

evidence in terms of an *apparent* sampling of brains, with the assumption that for an OO, these appearances would reliably represent the external facts.

[22] By focusing on brains with normal evidence, we assure that the sample has a clear and straightforward bearing on H_E. The mere fact that Boltzmann Brains vastly outnumber Ordinary Observers would not by itself provide strong evidence for H_E, since it could still be the case that Ordinary Observers vastly outnumber Boltzmann Brains when we limit our attention to those brains that have orderly experiences of the sort that might be had by a reliable Ordinary Observer.

assumption that Boltzmann Brains are quite rare in comparison to Ordinary Observers. To explain, let BBU_x designate the hypothesis that, among all brains with normal evidence, the expected ratio of Boltzmann Brains to Ordinary Observers is x. So BBU_{50}, for example, designates the hypothesis that among brains with normal evidence, the ratio of Boltzmann Brains to Ordinary Observers is 50 to 1. Finally, for sake of simplicity, let's assume that among all of the brains whose total evidence is member of the set {E0, E1,…, E1000}, the expected proportion of Boltzmann Brains is no different than the expected proportion of Boltzmann Brains among all brains with normal evidence. Let's now consider the probability of H_{E1000} on the supposition of $BBU_{0.001}$, that is, on the supposition that among brains with normal evidence, there are 1,000 times as many Ordinary Observers as there are Boltzmann Brains. Even given this assumption, we should still expect that Boltzmann Brains will comprise the vast majority of brains that have evidence E1000. To see why, note first that among the reliable Ordinary Observers whose total evidence is from the set {E0, E1,…, E1000}, the proportion who get evidence E1000 will be extremely small. Given $BBU_{0.001}$, the probability that a brain randomly sampled from the brains with normal evidence is a Boltzmann Brain is only 1/1,001. The probability that *all one thousand* of the brains in the sample are Boltzmann Brains is thus $1/1{,}001^{1{,}000}$, or less than one out of $10^{3{,}000}$.

Next, among those Boltzmann Brains whose total evidence is from the set {E0, E1,…, E1000}, how many can be expected to receive evidence E1000? Well, since a Boltzmann Brain's visual experience is the product of a randomly constructed brain state and is not influenced by the actual proportion of Boltzmann Brains, we should expect that a Boltzmann Brain is just as likely to have evidence E1000 as any other member of the set. So, for a Boltzmann Brain whose total evidence is a member of {E0, E1,…, E1000}, there is a 1/1,001 chance that its evidence is E1000.

A Boltzmann Brain whose evidence will be from {E0, E1,…, E1000} is, then, approximately $10^{2{,}997}$ more likely to receive evidence E1000 than a reliable Ordinary Observer whose evidence will

be from {E0, E1,…, E1000}. So, even when we take into account that the expected number of Ordinary Observers with evidence from {E0, E1,…, E1000} is one thousand times more than the expected number of Boltzmann Brains with evidence from that set, the expected number of Boltzmann Brains with evidence E1000 in particular is still approximately $10^{2,994}$ times larger than the expected number of Ordinary Observers with E1000! Now, this does *not* mean that H_{E1000} is expected to be true. If brains are sufficiently sparse, the most likely scenario might be one where no brain ever has evidence E1000. If the probability that there is some Ordinary Observer with E1000 is vanishingly small, then it could be highly probable that no Boltzmann Brain will have E1000 even though a Boltzmann Brain having this evidence is $10^{2,994}$ more probable than an Ordinary Observer having that evidence. And if no brain has E1000, then H_{E1000} is false. What we *can* conclude, though, is that conditional on some brain having E1000, H_{E1000} is almost certainly true. (Note that H_{E1000} would be true in a scenario where there is exactly one brain with E1000 and that brain is a Boltzmann Brain, making the relevant Boltzmann Brain percentage 100%.) In other words, we may conclude that (H_{E1000} | $BBU_{0.001}$ & E1000) is extremely high. And, conditional on $BBU_{0.001}$ and E1000, the additional fact that you are a Boltzmann Brain would only add to the probability of H_{E1000}. (This is because learning that some brain with E1000 is a Boltzmann Brain would be some evidence favoring a higher proportion of Boltzmann Brains among brains with E1000.) Thus, we can conclude that Pr(H_{E1000} | $BBU_{0.001}$ & BB & E1000) is extremely high.

Thus far, I've argued that Pr(H_{E1000} | $BBU_{0.001}$ & BB & E1000) is extremely high. But to contest Defeat, I will now show that the lesson learned from the $BBU_{0.001}$ scenario is sufficiently general, so that we cannot rule out the reasonability of Pr(H_{E1000} | BB & E1000) > ½. Confining the mathematical details to the footnotes, here are the key results. For any BBU_x hypothesis, we may calculate what the following ratio:

The key ratio:

$$\frac{\text{The expected value, given } BBU_x\text{, of the number of BBs with E1000}}{\text{The expected value, given } BBU_x\text{, of the number of all brains with E1000}}$$

It is reasonable to assume that if the key ratio is greater than 0.99 for some BBU_x, then $(H_{E1000} \mid E1000 \,\&\, BBU_x) > ½$. As it turns out, the key ratio *is* greater than 0.99 for any BBU_x hypothesis where $0 < x < 155$ or where $x > 98{,}094$.[23] For BBU_x hypotheses where $155 < x < 98{,}094$, the key

[23] We may calculate the key ratio as follows. Let N be the expected value (conditional on BBU_x) of the number of OOs with evidence from the set {E0, E1,…, E1000}. Nx will then be the expected value (conditional on BBU_x) of the number of BBs with evidence from the set {E0, E1,…, E1000}. The number of BBs expected to have E1000 in particular will then be Nx/1,001 (as there are 1,001 possibilities, each assumed to be equally likely). The number of OOs expected to have E1000 will be $N \cdot (x/(x+1))^{1,000}$ (i.e., N multiplied by the probability that each of the thousand brains in a random sample are Boltzmann Brains). Thus, the key ratio, r, will be:

$$r = \frac{\frac{Nx}{1{,}001}}{\frac{Nx}{1001} + N\left(\frac{x}{x+1}\right)^{1,000}}$$

This in turn simplifies to:

$$r = \frac{x}{x + 1{,}001\left(\frac{x}{x+1}\right)^{1,000}}$$

While the above form of the equation works well for calculating values of the key ratio in a spreadsheet, to find the local maxima and minima, it is perhaps easier to begin by rewriting the key ratio as follows:

$$r = \frac{(x+1)^{1,000}}{(x+1)^{1,000} + 1{,}001x^{999}}$$

ratio is less than 0.99. The key ratio is at its lowest for BBU_{999}, where the value of the ratio is only 0.73.[24] But even in this case, it could be reasonable to think that H_{E1000} is probably true (conditional on E1000). Note that a key ratio of 0.73 does *not* entail that it is likely that Boltzmann Brains comprise approximately 73% of brains with E1000. This would only follow if we expected there to be sufficiently many brains with E1000. A key ratio of 0.73 might instead indicate something like the following: in the extremely unlikely scenario that the world contains some brain with total evidence E1000, there is about a 73% chance that there is only one brain with E1000 and that this is a Boltzmann Brain (so that Boltzmann Brains comprise 100% of the brains with E1000) and about a 27% chance that there is only one brain with E1000 and that this is an Ordinary Observer (so that Ordinary Observers comprise 100% of the brains with E1000). Something like this latter possibility would hold if brains are expected to be sparse in comparison to the numerous sets of total evidence that are possible for a brain to have. If brains are expected to be sparse in this way, then even on the supposition of BBU_{999}, where the key ratio reaches its low of 0.73, H_{E1000} is likely to be true in the event that some brain has evidence E1000.

At the local maxima or minima, the derivative of this equation, $\frac{dr}{dx}$, will be equal to zero. Using the quotient rule and algebraic simplification, we arrive at the following:

$$\frac{dr}{dx} = \frac{x^{998}(x+1)^{999}(1{,}001x - 999{,}999)}{((x+1)^{1{,}000} + 1{,}001x^{999})^2}$$

This derivative is equal to 0 at $x = 0$, $x = -1$, and $x = 999$. Only the last value is relevant here, as we are concerned only with positive values for x. When we set x to 999 in the equation for the key ratio r, we find that r reaches a minimum of approximately 0.73. This minimum may also be confirmed (without the somewhat involved calculus) using a spreadsheet to calculate r for different values of x.

[24] See previous note for mathematical explanation.

Thus, so long as you could reasonably think that brains are likely to be sparse in the relevant way, you could reasonably hold that, for *every* positive value of x, $Pr(H_{E1000} \mid E1000 \,\&\, BBU_x) > ½$. From this, we may further conclude that you may reasonably hold that $Pr(H_{E1000} \mid BB \,\&\, E1000) > ½$. For on the suppositions of E1000 and BBU_x (for any positive x), the additional fact that you are a Boltzmann Brain would only add to the probability of H_{E1000}.[25] Thus, there is no reason to accept the Defeat premise in the case of evidence E1000. That is, there is no reason to think that $Pr(H_{E1000} \mid BB) = Pr(H_{E1000} \mid BB \,\&\, E1000) < ½$. And without the Defeat premise, there is no reason to think that CI places constraints on $Pr(H_{E1000} \mid E1000 \,\&\, OO)$.

Let me sum up the discussion of this section so far. Dogramaci and Schoenfield claim that CI amounts to a "weird scientific revisionism" since CI implausibly entails that, *even on the supposition that you are a reliable Ordinary Observer*, your evidence cannot support a high credence for BBU, the hypothesis of Boltzmann Brain superabundance. In arguing for this, Dogramaci and Schoenfield use the proposition H_E as a proxy for BBU. This is a mistake. H_E is not a proposition about the proportion of brains that are Boltzmann Brains, or even a proposition about the proportion of brains with normal evidence that are Boltzmann Brains. Rather, H_E is a hypothesis about the proportion of Boltzmann Brains among those brains with evidence E. And in cases where E is evidence that is highly unlikely to be had by a reliable Ordinary Observer, it may be highly probable that the vast majority of brains with E are Boltzmann Brains *even if Boltzmann Brains make up a very small portion of the brains with normal evidence*. In advancing their argument, Dogramaci and Schoenfield assume that $Pr(H_E \mid BB \,\&\, E)$ is low. This assumption is revealed to be unjustified once we see how the probability of H_E may diverge dramatically from the probability of BBU, and why these values

---

[25] Note that we can set aside the possibility of BBU_0, as this is ruled out by the supposition that BB is true.

are especially apt to diverge when E is evidence that would (conditional on OO) strongly support H_E.

Thus far, I have addressed Dogramaci and Schoenfield's arguments against CI that rely on their Defeat premise. But they also point to two allegedly absurd results that may be derived from Lemma alone or from Lemma and a weaker premise. Here is the first result, which is entailed by Lemma and uncontroversial principles of Bayesian confirmation:

Result 2.1: $Pr(H_E \mid E \,\&\, OO) < Pr(H_E \mid E \,\&\, BB)$

The next result follows from Result 2.1 together with only one component of Defeat, a component which may seem uncontroversial. In my argument against Defeat above, I contested the claim that $Pr(H_E \mid BB \,\&\, E) < ½$. Dogramaci and Schoenfield say that even if one rejects this inequality, one should accept the other component which says that $Pr(H_E \mid BB) = Pr(H_E \mid BB \,\&\, E)$. As they say, one "cannot deny that the condition that I'm a BB neutralizes the force of my empirical evidence" (24). This equality component of Defeat, together with Result 2.1, entails the following:

Result 2.2: $Pr(H_E \mid E \,\&\, OO) < Pr(H_E \mid BB)$

Dogramaci and Schoenfield find both of these results to be problematic, but the weirdness of Result 2.2 is perhaps more evident. The mere fact that you are a Boltzmann Brain is at best very weak evidence concerning the portion of Boltzmann Brains among brains with evidence E. So why should it be the case that H_E is more strongly supported by BB than by the conjunction E & OO?

Here again, the objection given to CI rests on a subtle confusion. Perhaps surprisingly, your being a Boltzmann Brain does *not* completely "neutralize" the force of your empirical evidence. That

is, *contra* Dogramaci and Schoenfield, it is not the case that $Pr(H_E \mid BB) = Pr(H_E \mid BB \& E)$. Rather, $Pr(H_E \mid BB \& E) > Pr(H_E \mid BB)$. To see this, it will be helpful to distinguish between observer-*dense* worlds and observer-*sparse* worlds. In observer-dense worlds, the ratio of brains to total possible evidence sets is very high and a typical evidence set is had by many observers (whether BBs or OOs). In observer-sparse worlds, the ratio of brains to total possible evidence sets is very low, most of the possible evidence sets are not had by *any* observer, and those evidence sets that *are* had by an observer are typically had by only one observer (whether a BB or an OO). Conditional on the world being observer-sparse, the conjunction of BB and E is extremely strong evidence for H_E, while BB alone is not strong evidence for H_E. For suppose you know only that you are a BB and that the world is sparse. In this case, you should think it likely that no brain has evidence E and thus that H_E is probably false. (If no brain has evidence E, then H_E is false since it is *not* the case that the vast majority of brains with E are Boltzmann Brains.) But if you know that you are a BB with evidence E and that the world is sparse, then you should think that H_E is probably true. For in a sparse world, it is probable that no *other* brain has E, in which case there will be exactly one brain that has E (namely you), and since this brain is a Boltzmann Brain, 100% of brains with E will be Boltzmann Brains (making H_E true). In other words, $Pr(H_E \mid BB \& E \& \text{the world is sparse}) >> Pr(H_E \mid BB \& \text{the world is sparse})$. And since you cannot rule out a priori that the world is sparse, it should also be the case that $Pr(H_E \mid BB \& E) > Pr(H_E \mid BB)$.

So, the derivation of Result 2.2 relies on a false premise. What about Result 2.1? Result 2.1 essentially says that E will always be more supportive of H_E conditional on BB than conditional on OO. This result is, I suggest, correct and one that a proponent of CI must accept. It is possible to give an explanation that renders Result 2.1 more intuitive, but I will not offer this explanation here since it is somewhat involved and since, in any case, Result 2.1 is not *obviously* problematic. I will simply remind the reader that, in the case discussed above where you randomly sample a Sim from a

simulated world, we can derive from Lemma* the Result 2.1*: $Pr(H_{E^*} \mid E^* \,\&\, OO^*) < Pr(H_{E^*} \mid E^* \,\&\, BB^*)$. Since even the opponent of CI should accept Lemma*, they should also accept Result 2.1*. Result 2.1* is perfectly analogous to Result 2.1 and is no less surprising. Since we should accept Result 2.1*, it would be unreasonable to reject CI for the reason that it entails Result 2.1.[26]

---

[26] Dogramaci and Schoenfield give one final argument against CI that I address here. Let E be some evidence you might receive in a moment's time, and let an "evidential duplicate" of yours be an observer who at some point occupies an evidential position phenomenally equivalent to the one you *presently* occupy (prior to learning whether you will receive evidence E). Let "RBB" stand for the proposition that "a randomly-selected evidential duplicate of yours (which may or may not be you) is a BB." Dogramaci and Schoenfield claim (27–28) that CI is committed to the following: Pr(E | BB) = Pr(E | RBB). This alleged commitment is unreasonable. For suppose that E involves coherent and orderly experience. Your having orderly experience would be incredibly improbable conditional on BB (the claim that *you* are a Boltzmann Brain), but is not so improbable given RBB (since RBB does not decisively rule out your being an OO). In arguing that CI is committed to Pr(E | BB) = Pr(E | RBB), Dogramaci and Schoenfield appeal to what they say is the more general intuitive thought lying behind CI, namely, that "the reasoning that I would apply to *myself* must be the same as the reasoning I would apply to *a randomly chosen evidential duplicate of myself"* (25-26). They claim that this symmetry principle commits one to holding that (i) Pr(BB) = Pr(RBB) and (ii) Pr(BB & E) = Pr(RBB & E), claims which collectively entail Pr(E | BB) = Pr(E | RBB). But it is mistaken to think that (ii) is supported by the relevant symmetry principle. E is not just a fact that *someone* has such and such orderly experiences, but rather is a fact about *your experiences in particular.* (If it was not a fact about your experiences in particular, then there would be no reason to think that E & BB is less

6. Conclusion

I have argued that discussions of Boltzmann Brain-related skeptical worries have generally failed to identify the most serious skeptical worry posed by apparent scientific evidence for SUPERABUNDANCE. The possibility that is genuinely concerning is not that you are a Boltzmann Brain, but rather that those appearances which suggest that there is scientific support for SUPERABUNDANCE are deeply misleading. In the face of apparent scientific support for SUPERABUNDANCE, your credence must be divided between three options: SUPERABUNDANCE, MISLEADING SCIENCE, and MISLEADING APPEARANCES. If CI is correct, then any stance which gives significant credence to SUPERABUNDANCE may be ruled out as irrationally self-undermining. This leaves MISLEADING SCIENCE and MISLEADING APPEARANCES as the only options that could be worthy of significant credence. Even if you should think that a skeptical hypothesis like MISLEADING APPEARANCES has an extremely low a priori probability, you should *also* think it very unlikely that a clear prediction of a cosmological model that is extremely well-supported by scientific evidence will turn out to be wrong. Thus, to the extent that there appears to be *extremely strong* scientific support for future Boltzmann Brain superabundance (and to the extent that phenomenal externalism, divine prevention of Boltzmann Brains, or other such escapes from Boltzmann Brain skeptical concerns are judged to be highly implausible), there will be pressure to give significant credence to MISLEADING APPEARANCES. At present, the cosmological evidence may not be weighty enough to

---

probable than E & RBB.) So, the claim that is relevantly parallel to E & BB is *not* E & RBB, but the following: a randomly-selected evidential duplicate of yours is both a Boltzmann Brain *and* goes on to have experiences that are indistinguishable from the ones you would have upon receiving E. And it seems perfectly reasonable to hold that this latter proposition has the same probability as E & BB.

generate significant skeptical pressure. But in principle, it would seem that the apparent support for SUPERABUNDANCE could be strong enough that it would be more reasonable for you to conclude that you are deceived about the nature of the scientific evidence than to maintain that you are not deceived and that the evidence which seems to strongly to point to SUPERABUNDANCE is simply misleading.

When seemingly faced with an uncomfortable choice between skeptical doubt and a departure from norms of scientific inference, the argument against CI from Dogramaci and Schoenfield offers the promise of a third option that is blessedly unproblematic. If CI should be rejected on account of its absurd implications, then you could simply accept the prediction of SUPERABUNDANCE while maintaining confidence that you are an Ordinary Observer with reliable faculties of memory and perception. Unfortunately, the promise offered by this argument against CI is a mirage. More careful inspection shows that CI does not have the absurd implications attributed to it. This is not to say that CI is obviously correct, as there remain genuine philosophical concerns about the principle. But the seemingly decisive objections to CI we have explored here rest on one or more premises that appear plausible but are false.

References

Avni, Ron. 2023. "The Boltzmann Brains Puzzle." *Noûs* 57 (4): 958–72.

Binney, James, Roya Mohayaee, John Peacock, and Subir Sarkar. 2025. "Manifesto: Challenging the Standard Cosmological Model." *Philosophical Transactions of the Royal Society A* 383:1–36.

Bostrom, Nick. 2003. "Are We Living in a Computer Simulation?" *The Philosophical Quarterly* 53 (211): 243–55.

Builes, David. 2024. "Center Indifference and Skepticism." *Noûs* 58:778–98.

Carroll, Sean. 2021. "Why Boltzmann Brains Are Bad." In *Current Controversies in Philosophy of Science*, edited by Shamik Dasgupta, Brad Westlake, and Ravit Dotan, 7–20. New York: Routledge.

Cohen, Stewart. 2010. "Bootstrapping, Defeasible Reasoning, and A Priori Justification." *Philosophical Perspectives* 24 (1): 141–59.

Crawford, Lyle. 2013. "Freak Observers and the Simulation Argument." *Ratio* 26 (3): 250–64.

DeRose, Keith. 2017. *The Appearance of Ignorance*. Oxford: Oxford University Press.

Dogramaci, Sinan. 2020. "Does My Total Evidence Support That I'm a Boltzmann Brain?" *Philosophical Studies* 177:3717–23.

Dogramaci, Sinan, and Miriam Schoenfield. 2025. "Why I Am Not a Boltzmann Brain." *The Philosophical Review* 134 (1): 1–33.

Elga, Adam. 2004. "Defeating Dr. Evil with Self-Locating Belief." *Philosophy and Phenomenological Research* 69 (2): 383–96.

Elga, Adam. 2025. "Boltzmann Brains and Cognitive Instability." *Philosophy and Phenomenological Research* 111 (1): 127–136.

Kotzen, Matthew. 2021. "What Follows from the Possibility of Boltzmann Brains?" In *Current Controversies in Philosophy of Science*, edited by Shamik Dasgupta, Brad Westlake, and Ravit Dotan, 21–34. New York: Routledge.

Lewis, David. 1979. "Attitudes De Dicto and De Se." *The Philosophical Review* 88 (4): 513–43.

Page, Don N. 2024. "Bayes Keeps Boltzmann Brains at Bay." *Foundations of Physics* 54 (5): 62.

Saad, Bradford. 2024. "Lessons From the Void: What Boltzmann Brains Teach." *Analytic Philosophy* advance access:1–28.

Sample, Ian. 2020. "Physicist Brian Greene: 'Factual Information Is Not the Right Yardstick for Religion.'" *The Guardian*, February 8, 2020.

https://www.theguardian.com/science/2020/feb/06/brian-greene-theoretical-physicist-interview-until-the-end-of-time.

Wallace, David. 2023. "A Bayesian Analysis of Self-Undermining Arguments in Physics." *Analysis* 83 (2): 295–98.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.