

The Hallucination That Cannot Be: A Three-Axis Refutation of the Boltzmann Brain Problem

Lou-Malou Longtin

An Independent Researcher (Québec, Canada)
E-mail: loumaloulongtin@gmail.com
<https://orcid.org/0009-0007-9858-1243>

Longtin, Lou-Malou (2026) The Hallucination That Cannot Be: A Three-Axis Refutation of the Boltzmann Brain Problem. *Philosophy of Cosmology*, Volume 36, 135-144. <https://doi.org/10.29202/phil-cosm/36/4>

This paper develops a comprehensive refutation of the Boltzmann brain problem along three converging axes — semantic, ontological, and probabilistic — while maintaining the methodological humility appropriate to cosmological reasoning. First, the concept of “hallucination” presupposes a reference reality; absent correspondence, the term loses its content and cannot bear skeptical weight. Second, a “brain” requires a substrate — spatiotemporal location, energy gradients, causal connectivity — without which mentality is unintelligible; positing a mind in literal nothingness is incoherent, and positing an alien substrate collapses the “brain” metaphor itself. Third, within any sufficiently large ensemble of fluctuation-generated states, the combinatorial space of incoherence dwarfs the space of law-like, self-consistent structures; as a result, coherent, persisting observer-worlds constitute a measure-zero subset and are thus statistically negligible. Together, these considerations dissolve the threat that Boltzmann-brain observers dominate anthropic reasoning and underwrite the continued trustworthiness of scientific inference from within a law-governed cosmos.

This argument was developed through first-principles reasoning, with iterative dialogue using large language models to stress-test coherence and explore objections; the argument and conclusions remain the author’s own.

Keywords: Boltzmann brain, anthropic reasoning, hallucination, ontology, probability, cosmology, measure problem, cosmological inference

Received: 27 October 2025 / Accepted: 9 January 2026 / Published: 1 February 2026

Introduction

The Boltzmann brain paradox, first arising implicitly from Ludwig Boltzmann’s reflections on entropy and fluctuation (Boltzmann 1897), poses a troubling question for cosmology and epistemology alike. If the universe is sufficiently vast in space or eternal in

time, random fluctuations from thermal equilibrium might, by sheer chance, generate isolated self-aware observers. These hypothetical observers — so-called “Boltzmann brains” — would briefly flicker into existence, perhaps with false memories and coherent but illusory perceptions, before dissolving back into chaos. Because random fluctuations should vastly outnumber the intricate, low-entropy processes that yield evolved observers like us, a naïve application of anthropic reasoning suggests that we ourselves are overwhelmingly likely to be such fluctuation-born observers. This conclusion, if accepted, threatens the epistemic basis of cosmology: if our perceptions are the unanchored products of fluctuation, why trust any inference about the external world or its history?

The paradox has attracted considerable attention over the past two decades (Bostrom 2002; Dyson, Kleban, and Susskind 2002; Page 2008; Carroll 2017), not least because it connects deep issues in cosmology with foundational questions in epistemology and philosophy of mind. Standard responses typically attempt to resolve the problem by appealing to specific cosmological dynamics — for instance, arguing that de Sitter vacua decay too quickly to produce Boltzmann brains in significant numbers (Page 2006), or that eternal inflation selects for ordinary observers via a cutoff procedure (Garriga and Vilenkin 2001). However, these responses depend heavily on particular assumptions about the global structure of spacetime and the measure problem — the difficulty of defining meaningful probabilities in infinite or eternally inflating universes.

The present paper adopts a different approach. Rather than attempt to resolve the Boltzmann brain problem through cosmological engineering, it dismantles the problem itself. This paper argues that the paradox arises only because three distinct categories of confusion are left unaddressed: (i) a *semantic confusion*, in which the term “hallucination” is extended beyond its coherent scope; (ii) an *ontological confusion*, in which the notion of a “brain” is divorced from the conditions that make brains possible; and (iii) a *probabilistic confusion*, in which possibility is conflated with typicality in a way that neglects the combinatorial structure of state space. Once these confusions are disambiguated, the Boltzmann brain problem dissolves without remainder.

The argument is constructed to remain robust across all plausible resolutions of the cosmological measure problem. Whether one adopts a global cutoff approach (e.g., scale-factor measure) or a local approach (e.g., causal-diamond measure), the space of fluctuation-born coherent observers remains of measure zero relative to the total state space. The skeptical sting attributed to Boltzmann brains is thus philosophically ill-posed and probabilistically misplaced. Moreover, by integrating the three critiques into a mutually reinforcing structure, this paper offers a foundational refutation of the Boltzmann brain problem that does not depend on speculative features of high-energy physics or cosmological dynamics.

The structure of the paper is as follows. Collapse of “Hallucination” analyzes the semantic collapse of “hallucination” when severed from reference. The Substrate Paradox examines the ontological incoherence of positing a “brain” without a substrate. Statistical Collapse explores the statistical insignificance of coherent fluctuation observers within the vast combinatorial space of incoherence. Cosmological Implications considers the implications of these results for cosmological inference and methodological naturalism. Conclusions concludes. Counterarguments and Replies address potential counterarguments and objections, showing how the three-axis refutation withstands them.

Collapse of “Hallucination”

The Boltzmann brain paradox gains much of its rhetorical force from the image of a lone observer hallucinating a world that does not exist. Yet this framing already smuggles in a semantic confusion. The concept of “hallucination” is *contrastive*: it derives its meaning by deviation relative to a reference. To call a percept a hallucination is to say that it fails to correspond to the external world — but that presupposes there *is* an external world to which it could correspond. Without a reference reality, the term loses its footing. It ceases to mark a failure of representation and instead merely re-describes whatever appears. “Hallucination” without correspondence is like “error” without a standard or “fiction” without a reality: the word collapses into vacuity.

This collapse is not merely linguistic but conceptual. All ordinary uses of “hallucination” — from psychiatry to philosophy of perception — rely on a contrast between appearance and reality. A subject undergoing psychosis, for example, may report vivid percepts of nonexistent entities. Yet those percepts are identifiable as hallucinations only because they fail to match public, law-governed reality. Dreams, too, often exhibit internal coherence and causal flow, but we do not mistake them for waking perception precisely because they lack correspondence. Even drug-induced hallucinations, though sometimes richly structured, are epistemically parasitic on the world they distort. The reference point is indispensable.

A further illustration comes from the history of perceptual theory. The classical empiricists — Locke, Berkeley, and Hume — all recognized that perception involves an inferential bridge between sensory input and the external world. Even when that bridge is faulty, the *concept* of fault depends on the existence of the world as a standard. More recent work in cognitive neuroscience supports this picture: hallucinations are increasingly understood not as random noise but as the misapplication of predictive models (Friston 2010). The very possibility of misapplication presupposes that there is something — an external causal order — to be misapplied. Remove that order, and “hallucination” loses its anchor.

Once we sever the connection to reference, the skeptical force of the Boltzmann brain hypothesis dissipates. If there is no external world for the Boltzmann brain’s percepts to misrepresent, then nothing is a hallucination. The concept no longer does skeptical work. It becomes a purely internal description of phenomenology, stripped of its contrastive edge. A fluctuation-generated phenomenology could be maximally coherent from within — containing stable laws, memory traces, even apparent histories — and yet remain non-correspondent. The mere possibility of such internal coherence does not entail epistemic defeat.

This reveals a subtle equivocation at the heart of the Boltzmann brain argument. It relies simultaneously on *two* incompatible intuitions: (1) that hallucination is epistemically devastating because it involves a deviation from reality, and (2) that there may be no external reality at all. But these cannot both hold. If (2) is true, then (1) is meaningless: without a reality to deviate from, hallucination ceases to be a coherent category. And if (1) is to retain its skeptical force, then (2) must be false, because hallucination presupposes the very reality the scenario denies. The argument collapses under the weight of this contradiction.

The same point can be framed in terms of truth conditions. Truth, in the correspondence sense relevant to epistemology, depends on a world-to-mind relation. To call a belief true is to say that it matches the way the world is. If no world exists beyond the belief, then “true” and “false” lose their conventional meanings. They collapse into mere internal consistency. A Boltzmann brain’s “belief” that it inhabits a universe billions of years old cannot be false if there is no universe at all — but it cannot be true either. It simply *is*. The skeptical sting

arises only if we assume that hallucination retains its ordinary meaning in a context where its preconditions are absent. Once that assumption is withdrawn, the sting vanishes.

It might be objected that internal coherence alone could ground epistemic justification, even in the absence of external correspondence. This line of thought echoes coherentist theories of truth (e.g., Bradley 1914), according to which beliefs are justified by their mutual support rather than by correspondence to a world. Yet this objection misfires in the Boltzmann brain context. Coherentism explains *justification* within a network of beliefs, but it does not rescue the concept of hallucination. The very idea of a hallucination presupposes the failure of correspondence — not just internal inconsistency. Coherent hallucinations still misrepresent, but only *relative* to something beyond themselves. If there is nothing beyond them, then they are not hallucinations at all but simply phenomenological states. The skeptical bite vanishes once more.

Philosophers of perception have long recognized the contrastive nature of hallucination. As Martin (2004) and Fish (2009) argue in their disjunctivist accounts, hallucinations and veridical perceptions are not internally indistinguishable states; they are categorically different phenomena precisely because only the latter stand in the right kind of relation to the world. If that relation is absent, the classificatory scheme collapses. The Boltzmann brain scenario exploits precisely this collapse while pretending to preserve the original conceptual structure. This is a sleight of hand, not a paradox.

Moreover, this semantic dissolution has a profound epistemological consequence: it shifts the burden of proof. If the skeptic wishes to invoke hallucination as a reason to doubt our knowledge of the world, they must *first* secure the existence of the very world that gives hallucination its meaning. But once that is granted, the Boltzmann brain scenario ceases to threaten epistemic justification, because our perceptions are then tethered to a reality that grounds the distinction between veridicality and illusion. The argument cannot have it both ways.

In short, the skeptical weight carried by the term “hallucination” depends on the very world whose existence the Boltzmann brain hypothesis calls into question. Remove that world, and the term evaporates. What remains is not a skeptical threat but a semantic dissolution. The concept ceases to apply, and with it, the paradox loses its force.

The Substrate Paradox

The second axis of the refutation targets the ontological presuppositions of the Boltzmann brain hypothesis. The notion of a “brain” — or, more generally, a physical system capable of supporting mentality — is not content-free. Brains are not Platonic forms. They are material entities with specific enabling conditions: spatiotemporal location, energy gradients, causal connectivity, and complex organization sustained over time. To imagine a brain existing in literal nothingness is not to imagine a rare physical possibility; it is to abandon the concept of a brain altogether.

The substrate paradox can be formulated as a dilemma. Either (a) the posited Boltzmann brain lacks any substrate, in which case the entity is unintelligible; or (b) it exists within an utterly alien substrate, in which case the “brain” metaphor ceases to track what exists.

Option (a) is straightforwardly incoherent. A physical system cannot exist without physical conditions. To speak of a “brain” arising from a fluctuation in the void is to speak nonsense. Even the most minimal notion of mentality presupposes some spatiotemporal structure within which states evolve and causally interact. Without space, time, or energy,

there is no physical medium to instantiate computation, no causal process to generate states, and no persistence to sustain consciousness. The idea collapses into contradiction.

Option (b) attempts to salvage the hypothesis by positing that the “brain” arises not in literal nothingness but in some exotic substrate radically different from familiar physical reality — for instance, a quantum vacuum fluctuation or a higher-dimensional brane. But this move undermines the hypothesis in a different way. The term “brain” is no longer doing descriptive work; it has become a loose metaphor. If the structure in question bears no relevant similarity to biological brains or their functional analogues, then calling it a “brain” is misleading. It risks anthropomorphic projection rather than scientific explanation.

This difficulty can be seen by analogy with other physical systems. To call a whirlpool a “galaxy” because both rotate would be an error of category; superficial resemblance is not sufficient for conceptual identity. Likewise, to label a momentary arrangement of particles in a quantum vacuum a “brain” requires that it share not just some superficial features but the deep causal and organizational structures that underwrite cognition. Without those features — neural dynamics, causal integration, energy consumption, temporal continuity — the label ceases to pick out what it does in ordinary contexts. The Boltzmann brain becomes a mere metaphor for an unknown configuration, not a meaningful description.

The point is not that mentality is impossible in non-biological substrates. Functionalist theories of mind (Putnam 1967; Fodor 1975) rightly allow that cognition could, in principle, be implemented in diverse physical systems. The point is rather that mentality cannot be substrate-free, and the “brain” metaphor cannot survive unscathed if the substrate is radically alien. Once we abandon the enabling conditions that make brains possible, we have ceased to talk about brains.

Moreover, the substrate paradox forces us to confront a hidden assumption in many skeptical treatments of the Boltzmann brain problem: that the very same term (“brain”) can be meaningfully applied across radically different ontological contexts without loss. This assumption is rarely defended, but it is crucial. If “brain” simply means “whatever physical pattern supports mentality,” then the skeptical argument reduces to a trivial tautology: mentality requires a physical pattern. But if “brain” retains its richer connotations — structure, dynamics, causal history — then the Boltzmann brain hypothesis either presupposes conditions that already imply an external world or becomes incoherent. Either way, it fails to deliver the skeptical conclusion it seeks.

The substrate paradox also exposes a deeper tension in the Boltzmann brain scenario: the conflation of physical possibility with conceptual coherence. It may be physically possible, in a sufficiently large spacetime, for a fluctuation to produce a complex configuration isomorphic to a functioning brain. But this possibility presupposes the very physical structure — spacetime, energy, causality — that the skeptical version of the scenario seeks to call into question. A “brain” arising from fluctuation presupposes a background within which fluctuation occurs. It cannot serve as a skeptical wedge against the existence of that background without sawing off the branch on which it sits.

In this sense, the Boltzmann brain scenario is parasitic on the very cosmology it seeks to undermine. It presupposes a law-governed spacetime capable of sustaining fluctuations while suggesting that such a spacetime may not exist. The scenario is not self-consistent. Once the physical preconditions of fluctuation are denied, the concept of a Boltzmann brain collapses into incoherence. Once those preconditions are affirmed, the skeptical argument loses its force, because the existence of a background spacetime capable of generating fluctuation already entails a world external to the brain.

The substrate paradox thus mirrors the semantic collapse of “hallucination.” Just as hallucination presupposes reference, so too “brain” presupposes substrate. Absent substrate, the concept dissolves; positing a radically alien substrate collapses the metaphor. In either case, the skeptical scenario ceases to be what it purports to be. It is not a paradox but a conceptual muddle.

Statistical Collapse

The third and final axis of the refutation is probabilistic. Even if one were to grant, for the sake of argument, that the semantic and ontological objections could be overcome, the Boltzmann brain hypothesis still fails on statistical grounds. The problem is not that Boltzmann brains are impossible, but that they are vanishingly improbable as coherent observers relative to the total space of fluctuation-generated states. Possibility does not entail typicality, and infinity does not confer representativeness.

Let an ensemble of fluctuation-generated states be given. Partition this ensemble into two classes:

N_{coh} : the number of coherent, law-governed observer-worlds

N_{tot} : the total number of fluctuation-generated states

The claim is that N_{coh} grows strictly sub-linearly relative to N_{tot} as state complexity increases. The reason is combinatorial. The space of incoherent, short-lived, or self-contradictory configurations expands explosively with the number of degrees of freedom. By contrast, the constraints required for sustained law-like structure — stable physical laws, consistent records, thermodynamic arrows, and reliable memory traces — sharply narrow the admissible states. The more structure a state exhibits, the smaller the region of state space it occupies.

This disparity in growth rates means that the relative frequency

$$P = \frac{N_{coh}}{N_{tot}}$$

tends to zero in the infinite limit. Coherent, law-governed states form a measure-zero subset of the total space of fluctuation-generated states. They are not merely rare; they are statistically negligible.

This conclusion is robust across competing solutions to the cosmological measure problem. Whether one adopts a *global cutoff* approach, such as the scale-factor measure (De Simone et al. 2008), or a *local* approach, such as the causal-diamond measure (Bousso 2006), the combinatorial dominance of incoherent states persists. No physically reasonable measure reverses this imbalance. The overwhelming majority of fluctuation-generated states are chaotic noise, ephemeral blips, or internally inconsistent fragments. The probability that a randomly sampled fluctuation will yield a coherent observer-world is effectively zero.

One might object that even a measure-zero subset can contain infinitely many elements. This is true, but irrelevant. The question is not whether Boltzmann brains exist but whether they dominate anthropic reasoning — whether a randomly selected observer is likely to be one. If coherent observers form a set of measure zero, then conditionalizing on coherence collapses the posterior weight on “we are fluctuation observers” toward zero. In other words, once we condition on the fact that our experience exhibits stable laws, consistent records, and an extended thermodynamic arrow, the Boltzmann brain hypothesis ceases to be the typicality-maximizing hypothesis.

Dyson, Kleban, and Susskind (2002) emphasize this point in a different register: the more ordered a fluctuation, the more suppressed its probability. Small fluctuations are common; large, structured fluctuations are exponentially rare. A brain-like fluctuation is already fantastically improbable; a coherent cosmos with billions of years of consistent history is exponentially more so. The anthropic argument that “we should expect to be Boltzmann brains” fails to account for this conditioning. We do not merely observe *consciousness*; we observe *structured, law-governed reality*. That fact shifts the posterior distribution dramatically.

This line of reasoning also dissolves the apparent force of the “infinite time” objection. Even if an eternal de Sitter spacetime produces infinitely many Boltzmann brains, it also produces infinitely many incoherent fluctuations. Infinity by itself proves nothing. What matters is the *relative measure* of coherent states, and that measure tends to zero. Infinity cannot compensate for vanishing measure. A dart thrown at a dartboard of infinite area will not land on a target of zero area, no matter how many darts are thrown.

The statistical collapse thus completes the three-axis refutation. Semantic collapse shows that the concept of hallucination loses its skeptical bite when divorced from reference. The substrate paradox shows that the concept of a brain collapses without physical enabling conditions. Statistical collapse shows that coherent fluctuation observers are of measure zero and therefore anthropically negligible. Taken together, these arguments dissolve the Boltzmann brain problem from three directions: conceptual, ontological, and probabilistic.

Cosmological Implications

If fluctuation observers are semantically unstable, ontologically incoherent, and statistically negligible, then the Boltzmann brain problem does not undermine methodological naturalism. The trustworthiness of scientific inference does not depend on the absence of fantastical possibilities; it depends on the structure of typicality under reasonable priors. Once coherence, persistence, and correspondence are included among the data to be explained, the posterior weight on fluctuation hypotheses collapses.

This result has significant implications for cosmology. It suggests that the epistemic foundations of cosmological inference are more secure than often supposed. We are entitled to treat our observations — stable regularities, publicly shareable measurements, theory-mediated predictions — as arising from a law-governed cosmos rather than a transient fluctuation. The skeptical challenge fails not because fluctuation observers are impossible, but because they cannot bear the explanatory burden placed upon them.

Moreover, this analysis reframes the cosmological measure problem. The problem is often portrayed as a looming epistemic threat: without a unique, physically motivated measure on the multiverse, we cannot assign probabilities to different kinds of observers, and anthropic reasoning collapses. But the three-axis refutation shows that much of the threat is illusory. Regardless of how the measure is defined, the relative weight of fluctuation observers remains vanishingly small once coherence is conditioned upon. The details of the measure affect quantitative estimates but not the qualitative conclusion.

This point bears emphasizing because it is frequently misunderstood. The Boltzmann brain problem is sometimes presented as a *reductio* of eternal inflation or de Sitter cosmology: if these models predict that Boltzmann brains vastly outnumber ordinary observers, then the models must be false. But the *reductio* depends on the assumption that we should expect to be typical members of the set of *all* observers. Once we refine the reference class to include only

coherent, law-governed observers, the conclusion changes. We should expect to be typical members of that class — and within that class, ordinary observers dominate.

The three-axis refutation also strengthens the case for methodological naturalism — the commitment to explaining phenomena by appeal to law-governed natural processes rather than miraculous coincidence. If the alternative to a law-governed cosmos is not a coherent hallucination but a combinatorial abyss of incoherence, then the epistemic superiority of naturalistic explanation is reinforced. Far from being undermined by cosmology, methodological naturalism is vindicated by it.

Finally, this result has implications beyond cosmology. It bears on debates in epistemology, philosophy of mind, and metaphysics. In epistemology, it shows that radical skeptical scenarios often fail not because they are impossible but because they are conceptually confused. In philosophy of mind, it underscores the dependence of mentality on physical substrate and the limits of functionalist abstraction. In metaphysics, it illustrates the importance of measure-theoretic reasoning in assessing modal claims. The Boltzmann brain problem thus serves as a case study in how careful conceptual analysis can dissolve apparent paradoxes.

Conclusions

The Boltzmann brain problem has long been regarded as a profound challenge to cosmology and epistemology. It suggests that, in a sufficiently vast or eternal universe, fluctuation-born observers will vastly outnumber evolved ones, rendering it overwhelmingly likely that we ourselves are such observers. If true, this conclusion would undermine the epistemic basis of science and our confidence in cosmological inference.

This paper has argued that the paradox dissolves when analyzed along three converging axes. First, the semantic collapse of “hallucination” shows that the concept loses its skeptical bite when divorced from reference. Second, the substrate paradox shows that the concept of a “brain” collapses without the physical enabling conditions that make mentality possible. Third, the statistical collapse shows that coherent fluctuation observers form a measure-zero subset of the space of possible fluctuations and are therefore anthropically negligible. Together, these arguments dismantle the Boltzmann brain problem at its root.

The resulting picture is one in which the epistemic foundations of cosmology remain intact. Ordinary scientific inference — from stable regularities, publicly shareable measurements, and theory-mediated predictions — remains licensed. The specter of Boltzmann brains does not undermine methodological naturalism or anthropic reasoning. Instead, it reveals the importance of conceptual clarity, ontological realism, and probabilistic rigor in cosmological thought.

What remains are ordinary burdens of theory choice in cosmology: to articulate measures, initial conditions, and dynamical laws that render observers and their records unsurprising in a law-governed cosmos. The Boltzmann brain problem is not one of these burdens. It is a pseudo-problem born of semantic slippage, ontological confusion, and probabilistic naiveté. Once these are corrected, the hallucination that cannot be vanishes.

Counterarguments and Replies

No philosophical argument is complete without engagement with potential objections. This section considers three of the most prominent counterarguments to the foregoing analysis and shows why they fail to undermine the three-axis refutation.

The Simulation Objection: Coherence Without Reference

One might object that the semantic collapse argument fails because a sufficiently detailed simulation could generate coherent phenomenology without external reference. If a Boltzmann brain's experiences are functionally identical to those of an evolved observer, why should the absence of a reference world matter?

The reply is twofold. First, the semantic point concerns the meaning of *hallucination*, not the possibility of phenomenology. A simulated world may be internally coherent, but its coherence does not make it veridical. If the simulation is causally disconnected from the external world, then its contents are not representations but self-contained structures. Calling them "hallucinations" misapplies the term. Second, even granting that phenomenology could be simulated, the substrate paradox reappears. The simulation itself would require a physical substrate — computational machinery embedded in spacetime — which reintroduces the very external world the skeptical scenario sought to deny.

The Measure Objection: Infinity Overwhelms Measure Zero

Another objection holds that even a measure-zero set can contain infinitely many elements. In an eternal de Sitter space, for example, there may be infinitely many Boltzmann brains, even if their relative measure is zero. Should we not then expect to be one?

This objection conflates cardinality with probability. The relevant question is not how many Boltzmann brains exist, but how likely it is that a randomly selected observer is one. If coherent observers form a measure-zero subset, then the probability of randomly sampling one from the ensemble approaches zero. Infinity does not change this. As Carroll (2017) emphasizes, typicality must be defined relative to a measure, not raw cardinality. Conditioning on coherence collapses the posterior weight on Boltzmann brains regardless of their absolute number.

The Indistinguishability Objection: Epistemic Equivalence

A final objection contends that even if Boltzmann brains are semantically unstable, ontologically dubious, and statistically negligible, we cannot *know* that we are not one. Our experiences would be indistinguishable from those of an evolved observer. The skeptical sting remains.

The reply is that indistinguishability is irrelevant to epistemic justification. As Wallace (2012) argues in a different context, the mere logical possibility of a skeptical scenario does not undermine knowledge unless that scenario is probable under reasonable priors. Once we condition on the coherence, persistence, and law-governed structure of our experience, the probability that we are Boltzmann brains becomes vanishingly small. Epistemic justification depends on probability, not bare possibility. We are no more undermined by the possibility that we are Boltzmann brains than by the possibility that we are brains in vats or victims of Cartesian demons.

Acknowledgements

AI-Assisted Drafting Disclosure

During drafting, the author used multiple commercially available LLM systems (including OpenAI's ChatGPT and others) as an intellectual sparring partner: to ironman counterarguments; test objections; reformulate passages to verify intended meaning; maintain a formal academic register; and role-play plausible reader responses. Portions of these exchanges were edited and integrated into the final text. The author did not treat the models

as epistemic authorities and remains solely responsible for the argument, structure, citations, and conclusions.

References

- Albrecht, A., & Sorbo, L. (2004). Can the universe afford inflation? *Physical Review D*, 70(6), 063528. <https://doi.org/10.1103/PhysRevD.70.063528>
- Boltzmann, L. (1897). On certain questions of the theory of gases. In *Populäre Schriften (Popular Writings)*. Leipzig: Barth.
- Bostrom, N. (2002). *Anthropic Bias: Observation Selection Effects*. New York: Routledge.
- Bousso, R. (2006). Holographic probabilities in eternal inflation. *Physical Review Letters*, 97(19), 191302. <https://doi.org/10.1103/PhysRevLett.97.191302>
- Carroll, S. M. (2017). Why Boltzmann Brains Are Bad. *arXiv preprint arXiv:1702.00850*.
- De Simone, A., Guth, A. H., Linde, A., Noorbala, M., Salem, M. P., & Vilenkin, A. (2008). Boltzmann brains and the scale-factor cutoff measure of the multiverse. *Physical Review D*, 82(6), 063520. <https://doi.org/10.1103/PhysRevD.82.063520>
- Dyson, L., Kleban, M., & Susskind, L. (2002). Disturbing implications of a cosmological constant. *Journal of High Energy Physics*, 2002(10), 011. <https://doi.org/10.1088/1126-6708/2002/10/011>
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Fodor, J. A. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Garriga, J., & Vilenkin, A. (2001). Many worlds in one. *Physical Review D*, 64(4), 043511. <https://doi.org/10.1103/PhysRevD.64.043511>
- Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, 120(1–3), 37–89. <https://doi.org/10.1023/B:PHIL.0000033751.66949.97>
- Page, D. N. (2008). Is our universe likely to decay within 20 billion years? *Physical Review D*, 78, 063535. <https://doi.org/10.1103/PhysRevD.78.063535>
- Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (Eds.), *Art, Mind, and Religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.