

Why Boltzmann Bubbles are Bad

Sinan Dogramaci

Draft of August 3rd, 7050 words

Abstract: Boltzmann gave a brilliant argument that entropy is almost sure to be higher in the future, but, given the time-symmetry of the laws of nature, the argument notoriously also has the skeptical implication that entropy is almost sure to have been higher in the past too. This skeptical result can be avoided if we can justify belief in what David Albert has named the *Past Hypothesis*. This paper argues that we cannot give a principled justification for our belief in the Past Hypothesis, not without generating skeptical implications about the future. The paper ends by observing that, although these skeptical results might be avoided if the laws are time-*asymmetric*, it is odd to draw such a concrete scientific conclusion only from philosophical anti-skepticism.

Key words: Past Hypothesis, Boltzmann, Skepticism, Simplicity

1. The Reversibility Problem

Why should we think entropy will increase in the future? Perhaps at some point in history it was possible for a scientist to reasonably infer, given their observations, that it is a fundamental dynamical law of nature that entropy increases until reaching equilibrium. But the modern scientific consensus is that the fundamental dynamics will not guarantee that entropy increases. What scientists

instead accept is the following line of argument, due to Boltzmann, for why, nevertheless, the future probably has higher entropy.¹

The basic argument is this.² Our observations only tell us what macrostate the world is in. The exact underlying microstate cannot be directly observed, and uncountably many possible microstates are consistent with the observed macrostate, so we can only make probabilistic claims about which microstate we're in. If it is rationally permissible to start with a uniform probability measure over microstates, as Boltzmann's argument presumes it is, and we conditionalize it on the observed macrostate as well as the dynamics, then, the argument concludes, we'll rationally arrive at a posterior which says there is a high probability that the future has higher entropy than the present.

How to picture the idea of the argument: if you add pepper to the salt shaker and you shake it, it's *possible* the precise movements of the particles will leave the salt and the pepper neatly separated, but as we know intuitively and as the argument rigorously shows, the *vast* majority of the possible precise ways this could go will leave them all mixed up, and those mixed states here represent the higher entropy states. There are various technical definitions given to "entropy", but it will not matter at all to the arguments of this paper how it is defined.

So, if entropy is below equilibrium, then it will probably increase. But there's a problem. Almost every proposed theory of the fundamental dynamical laws, at least since Newton, has it that the dynamical laws are *time-symmetric*: for any sequence of events the laws permit, they also permit a sequence of corresponding events in reverse. This raises the famous problem of reversibility.

¹ Boltzmann first gave the basic argument, but others, especially Gibbs, refined it and helped develop the new branch of science it is a part of, what Gibbs named "statistical mechanics". Myrvold [2021] gives an expert presentation of the history as well as the science.

² Albert [2000], Myrvold [2021], and Chen [2021] give accessible overviews of background material I will quickly summarize.

Why should we think entropy was *lower* in the *past*? What the uniform measure over microstates will imply, given a time-symmetric dynamics, is that entropy is likely to have been *higher* in the past.³ And that's a skeptical result. It implies that our beliefs about the past are radically mistaken. There has never been an ice age, the fossils you see in museums did not come from living dinosaurs, South America and Africa did not once interlock, and the cosmic microwave background radiation that we see now is not the result of the big bang.

The skeptical result might be blocked if we could deny that the dynamics are time-symmetric in the way that makes it likely, given our present evidence, that entropy was higher in the past. But this is not the standard approach to solving the reversibility problem. The standard response to the skeptical threat posed by reversibility, widely accepted among scientists and philosophers, is rather to endorse what Albert [2000] has named “the Past Hypothesis”, or the PH.⁴ The PH says the universe was born in a kind of low entropy state, specifically a state that is likely to evolve into a universe with the level of entropy, as well as the organized correlations among its macrostates, that we now observe. If we can justify believing the PH, this will block the skeptical result. If the uniform probability measure is conditionalized on the PH, as well as on our observations, then the dynamics imply that entropy is likely to increase from low to high, until reaching equilibrium, over the entire lifetime of the universe, as described by the second law of thermodynamics.⁵

³ For the purposes of the present discussion, we can assume a specification of a unique uniform Liouville measure. Different choices of background structure and parameters yield different uniform measures, but any natural choice raises the same skeptical reversibility issues that we will be discussing.

⁴ Loewer et al. [2023] demonstrates the wide-ranging, interdisciplinary endorsement of the PH. Boltzmann initially endorsed this kind of view, but later he came to accept a skeptical view.

⁵ Albert [2000] argues at length that, assuming a time-symmetric Newtonian framework, the right response to the skeptical threat posed by reversibility is to endorse the PH, and then, at the end of the book (ch.7, esp. pp.160-1), he discusses a time asymmetry in GRW-style collapse theories of

But what *justifies* belief in, or conditionalization on, the PH? Chen [2021] has made a compelling case for the skeptical view that there is no justification. I will likewise argue that there is no justification, at least on the assumption of a standard time-symmetric dynamics, an assumption I will keep in place throughout this paper (until the very end when I raise the question whether the paper's arguments constitute a reductio of the assumption). I will make it clear below how my argument, if correct, strengthens the case beyond what Chen aimed to show.

On the resulting skeptical view, while our beliefs about local matters might be justified (I have hands, I've been in this room for more than five minutes), our ordinary beliefs about the more remote parts of space and time are mistaken (there were never dinosaurs, there was no big bang, and nearly everywhere the universe is in the dead-looking state of thermodynamic equilibrium). This is what Chen describes as the “Boltzmann Bubble” skeptical hypothesis.

The skeptical conclusion is unbelievable, but I hope that by making the argument for it as clear, simple, and strong as possible, it will be easier for others to correctly diagnose where, if anywhere, it can be blocked. And if the reasoning toward the skeptical conclusion cannot be blocked anywhere, then philosophy may offer a lesson for practicing physicists. A time *asymmetric* dynamics might be the only option.

2. *An Anti-Skeptical Bayesian Framework*

I will assume a Bayesian model of belief and rational scientific reasoning: I assume that we have credences, that rational credences are probabilistically coherent, and that if rational prior credences are conditionalized on your total

quantum mechanics, but even then Albert emphasizes that we still must assume the PH. I will return to collapse theories at the very end of the paper.

evidence, then the resulting posterior credences are also rational (or justified—same thing).

I make the anti-skeptical assumption that at least some possible priors are rational. Some Bayesians are extreme permissivists. They say any coherent credences are rational priors. I assume this extreme form of permissivism is false, since if it were true, belief in the PH would be justifiable but so would belief in the flat earth theory. I don't want to save my own beliefs from skepticism at the cost of counting the flat earther's beliefs as equally rational. So, I don't agree that every coherent prior is rational.⁶

I also don't assume that there must be some rational prior that leads to exactly my actual posteriors when conditionalized on my actual total evidence. That would imply my actual posterior credences are rationally justifiable in a way that I don't believe they are. Like anyone, I have some irrational credences.

Our anti-skeptical assumptions must be more principled. We want some principles, or some kind of elegant theory, that will allow us to have rational credences but also allow us to have some irrational credences as a result of our performance errors, our deviations from what the principles or the theory demands.

The part of the theory we are mainly going to be looking into here will be the part that constrains the rational priors, since we're assuming conditionalization is the (entirety of the) part of the theory that concerns how priors rationally evolve into posteriors. Though there are some controversies over and some unresolved details about conditionalization, these can mostly be ignored here.⁷ But one question about conditionalization I'll assume a particular answer to, though, is the question of what its inputs are. I'll make the "dogmatist" anti-skeptical assumption

⁶ As Myrvold [2021, p.25] colorfully points out, not only is extreme permissivism itself hard to believe, it is hard to believe that anyone sincerely believes extreme permissivism!

⁷ See, for example, Cassell [2025, sec. 5] and references therein.

that the empirical evidence we can conditionalize on includes lots of information about our external local environment. I have hands, there's a cat over there, someone is telling me about their trip to Australia, and so on. I make this assumption not only because I think it's plausible, but because my aim in this paper is to argue that there is no plausible and principled set of anti-skeptical assumptions that will solve the reversibility problem, so I want to be as generous as possible with my anti-skeptical assumptions. So I assume the evidence we get and conditionalize on is strong enough to rule out skeptical scenarios about the local environment. I'm sure I'm not a brain in a vat or an isolated "Boltzmann Brain", since I conditionalize on lots of facts to the contrary. I can also retain my rational credences over time, so I can be sure I did not just pop into existence seconds ago, again since I conditionalized on lots of facts to the contrary.

So, if we view conditionalization and its inputs in this "dogmatist" way, it can get us a lot of anti-skeptical mileage.⁸ But we can't rely just on the nature of conditionalization to get us all the way to rational high posterior credence in the PH given our evidence. We obviously don't observe and then conditionalize on the PH. And if we start with the uniform prior, then conditionalization on our total evidence will leave the PH with a low posterior. Therefore, on our Bayesian model, any argument that the PH can enjoy a rational high posterior will have to involve some argument that the uniform prior is not the rational prior, or at least not the only rational prior. And this will have to be a non-trivially strong argument, and it will have to be an apriori argument.

⁸ Emery [2017] uses similar dogmatist looking principles to argue against certain quantum mechanical theories, ones that deviate too much from the contents of our ordinary observations. Also see Emery [2024, esp. sec. 3] for more general discussion, and more comparison between the general epistemological views of Pryor [2000] and Huemer [2001] and her narrower applications of a similar idea to physics as well as metaphysics. For general discussion of dogmatism and Bayesianism, see Pryor [2013] and Lasonen-Aarnio & Huemer [2024].

It must be non-trivially strong because the uniform prior is based on a measure that is, as Albert says, a “very natural-looking measure...(as a matter of fact) the simplest imaginable measure”.⁹ It thus requires some substantial argument to overturn the uniform prior’s strong initial claim to being the rational prior.

And the argument must be, at least partly, *apriori*. No argument in favor of one prior over another can possibly be given *just* by citing empirical evidence. This is a simple point, but it gives crucial guidance to our task here. If we start out thinking one prior is the best candidate for being the rational prior, and we decide we want to try to persuade ourselves and others that actually it’s *not* rational, or it’s not uniquely rational, and some alternative is a rational prior, then we must cite some *apriori* considerations that recommend the alternative prior. It would make no sense to say that our empirical evidence alone suffices to show us that the rival prior is actually the rational one. In the Bayesian framework, the evaluation of empirical evidence completely depends on a prior. If you decided the uniform prior isn’t as attractive as it initially seemed because you don’t think it gives a plausible

⁹ See p.351 of Loewer et al [2023]; the same claim is also on p.4 of Albert [2015].

See also Loewer [2012, p.124]: “It is reasonable to posit the uniform distribution since it is the simplest until evidence is adduced against it.” But I should clarify that Loewer here is talking about objective chances, not rational credences. Loewer is saying it’s reasonable or rational to believe or to have high credence, at least initially, that the objective chance facts obey the uniform distribution. (Albert also says he generally interprets probabilities objectively—see Albert [2000, p.81].) In this paper, I am *assuming* there is such a thing as rational credence, and I am asking: What plausible theory of rational credences escapes the skeptical challenge posed by reversibility? I am happy to suppose statistical mechanics concerns objective chances. But then my question is: What plausible theory of rational credences permits us to believe the objective chance facts are such as to make the PH true? Extreme permissivism would permit us to believe the PH, but it is very implausible. A uniform prior is very plausibly rational, but won’t permit us to believe the PH.

As Meacham [2023] explains, the reversibility challenge can be formulated whether we interpret the probabilities of statistical mechanics as objective chances or as rational credences. Meacham thinks we can escape the reversibility problem if we endorse the objective chance interpretation. But Meacham does not answer my question: Again, what plausible theory of rational credences permits belief in the PH, or in objective chance facts that make true the PH? He suggests something like the extreme permissivist view (see pp.161-2), but I rule that out.

evaluation of your empirical evidence, well then you must have already independently decided that some other prior is rational and already adopted it. If any independent considerations really can favor that other prior, they must be apriori, at least in part.^{10, 11}

This considerably narrows down our search for what could justify believing the PH. What we are going to have to do in the end is find some principled apriori considerations for thinking the uniform prior is not rational, or not uniquely rational. What could it be?

3. *The Appeal of Simplicity*

Here is what I think it must be. What we have to do is argue that the uniform prior does not correctly respect the simplicity of hypotheses, in the following way.

The uniform prior gives the PH a low posterior on our total empirical evidence (which I'll call E). From that, it follows that there must be some other hypothesis that assigns the same likelihood to E that the PH does, but has a higher prior than the PH. Suppose, unrealistically but for the sake of a clear preliminary illustration, that the PH assigns a likelihood of 1 to E, ie. $\Pr(E|PH) = 1$. Then the hypothesis $\sim PH \& E$ also assigns E a likelihood of 1 (trivially), and it will follow that it must have a higher prior than the PH.¹² But obviously $\sim PH \& E$ is a strange

¹⁰ Meacham [2023, p.157] makes a similar point.

¹¹ Philosophers can give the misleading impression that the PH is justified merely by virtue of its relation to our evidence. (Albert [2000, pp.83-5, 93-6] initially suggests that kind of idea, but he ultimately endorses a more complex argument (that we'll soon discuss).) Wallace [2023, p.104] says our justification for believing the PH is that it "explains a great deal of empirical data", and inferences to it are justified "by virtue of their explanatory power". In a recent podcast, Maudlin says "The only way you can explain what you see now and all the stars and what they're doing is that entropy was lower." (at 2:06: <https://www.youtube.com/watch?v=3riyyEmWwoY>) What Wallace and Maudlin are saying would be obviously false if "explain" just meant "assign high likelihood to our evidence", but what I think they intend, and what is anyway far more plausible, is that a theory adequately *explains* evidence only if it assigns the evidence high likelihood *and* the theory itself is suitably simple. And this is the kind of view I'll be engaging with below.

¹² The proof is trivial but a Venn diagram can make the claim most easy to see.

kind of thing to treat seriously as a scientific hypothesis. It's patently inelegant; it's ad hoc; it's grossly complex. Looking at it, then, we could convincingly argue, apriori, that the rational prior probability distribution, or some rational prior probability distribution, will give $\sim$ PH&E a lower prior than the PH. In other words, the simplicity of the PH could constitute an apriori justification for assigning it a higher prior, and if the uniform prior fails to do so, then so much the worse for the uniform prior and its claim to being (uniquely) rational.

So, if we've got to argue apriori that the uniform measure does not provide the uniquely rational prior probability distribution, and some other prior distribution that gives the PH a high posterior probability on our evidence is rational, then what I think we will have to do is argue in favor of that other prior distribution on the basis of how the two distributions treat the intrinsic features of the hypotheses at issue here (as opposed to the likelihood they assign E). We are interested in whatever feature(s) of a hypothesis is relevant to its prior, and am going to use the ordinary label "simplicity" for this, which is plausibly the apt label to use here but I will anyway stipulate that I'm using it as a catch-all for whatever feature of a hypothesis might make it plausible, apriori, that the hypothesis must, or may, enjoy a high prior. I don't have an analysis of simplicity on hand or generally know how to measure it, but we'll see that those details ultimately won't matter at all to the arguments of this paper. If it could be argued that the PH is simpler than the relevant rivals, and therefore deserves a higher prior than what the uniform measure would give it, then we would have the resources for responding to the Boltzmannian skeptical threat.

The example of the hypothesis $\sim$ PH&E was just a first illustration. Will other rivals that predict E as well as PH does be as easily dismissed as unsimple and therefore assigned a low prior? This would require some argument. If we lift the unrealistic assumption that $\Pr(E|PH) = 1$, it will still be true, given that

$\Pr(\text{PH}|\text{E}) < \frac{1}{2}$, that there must exist another hypothesis—call it skeptical hypothesis S—such that $\Pr(\text{E}|\text{S}) \geq \Pr(\text{E}|\text{PH})$, and $\Pr(\text{S}) > \Pr(\text{PH})$, ie. S predicts our evidence at least as well as PH does, and S has a higher prior.¹³ Now it is not, like it was with $\sim\text{PH}\&\text{E}$, obvious just from its surface syntax that this S hypothesis is grossly complex and ad hoc, but we can at least see now how there is room to argue apriori for departing from the uniform measure and instead using an alternative that gives PH a higher prior than its rival S on the basis of PH’s greater simplicity.

What Chen [2021] does is construct a specific skeptical hypothesis that plays this S role, and he then proposes that his hypothesis S is intuitively no less simple than PH. If he’s right, that would give us a version of the skeptical threat that we cannot respond to by arguing, apriori, that PH is simpler than its rival. In the next section I’ll review Chen’s skeptical hypothesis, and then after that I will give my own argument that we *cannot* answer the skeptical challenge by arguing that the PH is simpler than Chen’s hypothesis, or indeed any other S-style skeptical hypothesis about a high entropy past.

4. *Chen’s Hypothesis*

Chen [2021] compares the PH with the hypothesis that the universe was initially in thermodynamic equilibrium and it then underwent a fluctuation—a kind of random abrupt decrease in entropy that the laws of nature allow to happen—that made the kind of evidence we actually observe very likely. Let’s now use S for Chen’s fluctuation hypothesis.¹⁴ S does not predict our actual evidence, E, with certainty, but neither does the PH. What the PH predicts with high probability is that entropy is at about the level we observe it to be at, and also that our

¹³ The proof is straightforward.

¹⁴ Chen uses the labels $T_F + \text{MEH}'$ and also T_F^* for this; see his sec. 3.1. Though I am focusing on Chen’s sharpened formulation, the basic idea of a fluctuation hypothesis about our past has been familiar since Boltzmann.

observations will exhibit correlations of a certain kind that they actually do exhibit: shoe-boxes with a left shoe in them tend to contain a right shoe as well; apartments tend to have a spatula in the kitchen and not in the tub; portrait photos tend to be caused by faces that resemble the photos; to give a few of the illustrative examples Albert [2000] and Chen [2021] used. Chen formulates S so that it also predicts such correlations just as well as the PH does.¹⁵

In order to do this, he defines S in a way that makes reference to the PH: he defines S as the minimal fluctuation that produces, in our observed environment, the level of entropy and the correlations among macrophenomena that the PH produces. Does this automatically make S more complex than the PH? Chen argues it does not, and his argument seems right. The fact that we refer to theory T1 when we define theory T2 does not imply that T2 is any less simple than T1. As Chen argues (p.554), it's only the intrinsic properties of a theory that should matter to the evaluation of its simplicity. If I happen to see an event, *e*, and I have to use the definite description 'the cause of *e*' in order to introduce a name for another event, *c*, this does not imply or in any way suggest that *e* is simpler than *c*.

The vivid image of what Chen's hypothesis S says is that we are living in, as he nicely puts it, a "Boltzmann Bubble". For most of the time, the universe is in thermodynamic equilibrium, the kind of dead state that scientists predict our actual world will eventually evolve into. According to S, most of space is presently in the dead state, but around us there is a small region in which the kind of lower entropy macrostate we're familiar with exists. The present level of entropy in this region is at a local minimum: entropy even in our region was at equilibrium shortly before our observations began, and our region will very soon evaporate and return to equilibrium after the present moment.

¹⁵ In this paper, I am bracketing concerns, important as they are, about whether we have a clear formulation of the PH, a clear formulation of the aspects of our ordinary observations that the PH is meant to predict, and whether the former does successfully predict the latter. Such concerns are pressed by, eg., Winsberg (2004) and Earman (2006).

Since Chen's S and the PH, by design, assign the same likelihood to our evidence E, what we need is some argument, again an apriori argument, that the prior probability of S is lower than the PH. The uniform measure assigns the PH a lower prior than S. So, again, we need to argue against the (exclusive) rationality of the uniform measure's priors by looking at the relevant features of the PH and S and arguing that the PH exhibits the kind of simplicity that makes it deserve a higher prior than the more complex hypothesis S.

On this question, Chen just makes the intuitively true observation that S is not, or at least not obviously, less simple than the PH, or at least not much less simple.¹⁶ But anti-skeptics might be unpersuaded by this. Since skepticism is on the line, anti-skeptics might feel comfortable responding to Chen at this juncture by just insisting that PH must be simpler than S after all, even if that wasn't initially obvious—evaluations of simplicity, or apriori plausibility, are after all extremely hard to settle. Unfortunately though, I am going to show that the case for skepticism is stronger than what Chen showed. Below, in the section after the next one, I will give a new argument that we cannot escape the skeptical threat just by claiming that the PH is simpler than S. But first, in the next section, I'll address an influential argument from Albert [2000] for why we must be justified in believing the PH.

5. Albert's Argument

Boltzmann's argument shows that the future will surely have higher entropy, and the reversibility problem is that his argument also shows that the past will surely have higher entropy. The skeptical argument here relies on the uniform prior, and it assumes a time-symmetric dynamics. We are mostly focusing on the role of the uniform prior in this argument and its skeptical implications. Albert

¹⁶ See pp.553-4.

[2000, pp.114-119, 122], however, suggests that this argument should be criticized for its assumption about the dynamics. To give his proposed answer to the skeptical challenge here, Albert begins with the starting premise (see p.116) that, if we doubt the PH, then we must mistrust the alleged data that we rely on to justify our belief in whatever dynamics we accept. (The idea of the premise is that, if we doubt the PH, then we must think the scientific books and papers we think we see are either pure hallucinations or, even if we dogmatically trust our immediate perceptual experiences, these pages we see are only the random products of a fluctuation, not reliable historical records of the data we need to infer the dynamics we actually accept.) Given that premise, Albert can then say to the skeptic: if your purported skeptical argument (i.e. the reversibility problem) for doubting the PH depends on confidence in the dynamics, but confidence in the dynamics depends on confidence in the PH (i.e. the premise), well then this purported skeptical argument cannot after all be any convincing argument for doubting the PH!¹⁷

The reply to this, however, is that the premise is not strictly correct. It is possible to have empirically supported rational confidence in the dynamics while doubting the PH. Doubting the PH does not force you to believe that you are a Boltzmann Brain whose experiences and memories are all hallucinations, or that the books you've read are the product of a fluctuation. It's possible to doubt the PH while believing you are in a Boltzmann Bubble. As Chen (sec. 3.5) proposes, we could suppose the bubble's history extends far enough into the past, even a few centuries for example, to ensure that all the empirical evidence we need in order to infer the dynamics is entirely veridical. So, Albert's proposed strategy for saving our belief in the PH does not go through as stated.

¹⁷ Albert's idea here has been adapted and applied by Carroll [2021] into an argument that we are not Boltzmann Brains. (And the title of the present paper is an adaptation of the title of Carroll's very influential paper.) Elga [2025] gives compelling criticisms of this way of responding to the skeptical threat posed by Boltzmann Brains.

Here is a variation of Albert's argument that's also worth considering. (Albert did not explicitly argue in this way, but his general discussion suggests it.) Using a uniform measure over microstates as your rational prior credence function can seem very tempting at first, but it turns out that having too much uniformity in your priors can seriously inhibit inductive reasoning. In many cases, if you start with priors that are "uniform" in one or another relevant sense, then evidence you collect will not at all confirm any hypotheses about what the more general patterns in the universe are.¹⁸ The standard uniform measure over microstates will stop us from rationally extrapolating correlations in the intuitively rational way. For example, finding a left shoe in the bargain bin won't confirm that there is also a right shoe somewhere in there, and this is so even though you frequently found left shoes accompanied by right shoes in the past. This is the kind of correlation that someone using the uniform prior doesn't expect to see in the world.¹⁹ But extrapolating correlations is what a rational inductive reasoner should do.²⁰ So, even if a uniform measure over microstates initially looks like the most natural option, the assumption of anti-skepticism about induction would, on its own, be enough to entail that it must be rational to start with a prior that looks very different. Just as we are happy to treat Cartesian skeptical scenarios as apriori implausible, we ought to likewise treat Humean skeptical scenarios as apriori implausible, and thus we have here an apriori argument against the uniform prior.

¹⁸ A familiar but still striking illustration of this kind of thing is the elementary fact that, if you start with a uniform prior over all the possible sequences of outcomes from tossing a coin, then no matter what you see, your observations will not confirm any hypothesis about the bias of the coin.

¹⁹ The reason why is more boring than the coin example: the uniform prior puts you inside the smallest possible Boltzmann Bubble it can, so you won't expect to find any solid object at all when you reach back into the bargain bin, let alone a right shoe.

²⁰ Here I am making the argument in terms of our inductive discoveries of correlations in macrostates rather than our inductive discovery of the dynamical laws themselves because the standard uniform measure over microstates turns out to have certainty of the dynamical laws built into it. A deeper, lengthier, and far more technical examination of the issues here would look into the important question of what kind of rational prior would allow us to inductively discover the dynamical laws, but I will settle for just making the point I make in the text by using the example of macrostate correlations.

Furthermore, if we use as our prior the uniform measure conditionalized on the PH, in other words if we start out in a sense assuming the PH, then we *can* get a prior that extrapolates ordinary correlations in the intuitive way.

However, this unfortunately is not, just on its own, a compelling reason for believing or assuming the PH. It's true that anti-skepticism about induction does force us to reject starting with the uniform prior. But it does not force us to start with a prior that assumes the PH. There are skeptical hypotheses we could conditionalize the uniform measure on so as to produce a prior which, when also conditionalized on our evidence, will confirm our actual scientific theory of what the world is presently like and will be like in the future, but won't confirm the PH and our ordinary beliefs about the past. Chen's S was already meant to predict correlations like those between left shoes and right shoes. More generally, we can always take the hypothesis, F, that a dramatic thermodynamic fluctuation from equilibrium all the way down to our current entropy took place at some recent time, and we could conjoin F with $(E \supset T)$, where E is our evidence and T is a statement of whatever theory we want E to confirm.

Now, obviously the resulting hypothesis, $F \& (E \supset T)$, looks very suspiciously ad hoc and so doesn't plausibly enjoy as high a prior as PH. My aim here, though, is just to show that there really is no alternative but to argue that the hypotheses that function as skeptical rivals to the PH are apriori unsimple and so deserve a very low prior, something low enough that a genuinely rational prior probability distribution will leave it low even as a posterior conditionalized on our evidence.

6. The Skeptical Implications of Simplicity

What would be so wrong then with saying the PH is simple and therefore more probable than the sum probability of all the alternatives that are also consistent with our evidence E (and consistent with the dynamics)? Simplicity is

not as well understood as we might wish, and again I have no particular theory of simplicity that I endorse, but isn't it clearly worth appealing to simplicity if it saves us from skepticism? What is the problem with that?

The problem is that it does not save us from skepticism. It just changes what the relevant skeptical hypothesis we're threatened by is. The problem is that we now must allow by parity that the skeptical hypothesis that the *future* will have *low* entropy is likewise simple and so likelier than the alternative on which the future has high entropy. Here's why.

We have some total evidence, E , which we will conditionalize a rational prior on. If we start with the uniform prior, and conditionalize it on E , that would make us confident that the graph of entropy probably goes upwards not only as we look toward the future but also as we look back toward the past. The uniform prior gives us a skeptical view of the past. But if we start with a high enough prior for the PH, our posterior will imply the anti-skeptical view that, about x years ago, the universe was in a state that had very low entropy, and that was very likely to gradually increase in entropy and evolve into the kind of state we observe ourselves to live in today, a state with the level of entropy we observe today, and with the sorts of correlations we observe (left shoes accompanied by right shoes and so on). We are now considering an anti-skeptical strategy that recommends assigning the desired sufficiently high prior to the PH on the basis of the simplicity of the PH. But the problem is this. The PH just says that, x years in the past, the universe was in a state of low entropy y , a state that, over x years, will very probably evolve into a state with the entropy and the correlations we now observe. But if this is a good reason to assign the PH a high prior, then there is equally strong reason to assign a high prior to a mirror image hypothesis about the future. Let the Future Hypothesis, the FH, be the temporal mirror of whatever the PH says about our past. What the PH says the universe was like x years ago, the FH says the

same about the universe x years into the future. However the PH makes our presently observed correlations likely, the FH likewise makes those same correlations likely—since the dynamics is time-symmetric, the FH can easily do this. If the PH deserves a high prior due to simplicity, then the FH deserves the same high prior. Both attribute low entropies at similar lengths of time from now. Given the time symmetric dynamics, they will assign the same likelihoods to the present level of entropy and the observed correlations found in our evidence E . So, they will have the same posteriors. And that's bad.

It's bad because the FH is a skeptical hypothesis. It implies that our ordinary beliefs about the future are radically mistaken. For, just as the PH and our total evidence, E , confirm the plausible hypotheses we ordinarily believe about the past, the FH and E will confirm parallel but absurd hypotheses about the future. The future that the FH has us predict will resemble something like a movie being run backwards. Adults will shrink into children, omelets will unfry and assemble into eggs, and eventually the universe will compress itself back into the kind of state it was born in (this last prediction going against current cosmology, if not widely-held ordinary belief). This is all absurd. But if the PH has properties that make it simple and thus have a high prior probability, then the FH has the same properties and so is just as simple and has just as high a prior probability. We were tempted (or desperate!) to say that hypotheses that attribute states with low entropy are simpler and thus more apriori probable (ie. have a higher prior) so that we could make the PH more probable than S and rescue our ordinary beliefs about our past, but now we see this comes at the cost of making the FH similarly probable and thus sacrificing our ordinary beliefs about our future.

Price [1996] does not directly discuss the justification for believing the PH, but he does sympathetically entertain this idea that the universe symmetrically grows and declines in entropy, a so-called “Gold” universe, a universe where the

graph of its entropy over time looks like a rainbow with the PH and the FH describing each endpoint. I am not endorsing Price's own motivations for giving serious credence to that view, which he does not treat as a radical skeptical hypothesis. For example, unlike Price I am happy to proceed under the assumption that it is a brute fact that time, or our experience of time, flows from past to future. Nevertheless, if my arguments in this paper are good, then they do help vindicate Price's main idea that scientists and philosophers have consistently begged the question when they try to argue for, or just insist on, one or another asymmetry between the past and the future.

To be clear, it's not part of the standard idea of a Gold universe that entropy is at its peak *today* in particular; rather, the peak is usually taken to occur in the remote future. But what I am arguing here is that, if our present evidence E together with the high prior of simple theories justifies high confidence in a view about the entropy of the recent or the remote past, then E and high prior of simple theories also justifies believing a parallel view about the entropy of the near or the distant future as well. This is why any reason we have to believe the PH is also a reason to believe the corresponding FH, and the resulting view is one where entropy is at its peak today. Our ordinary beliefs imply that entropy was slightly lower in the recent past and greatly lower in the far past, and this anti-skeptical picture of the downward sloping curve of (the graph of) entropy as we look back into the past is implied by, and thus happily saved by, the PH. But if PH's simplicity earns it a high prior, and its relation to E gives it a high posterior, then we must give an equally high prior and posterior to the corresponding FH, and then that will, disastrously, imply a corresponding downward sloping curve of entropy as we look into the future from our present position in time. If the PH rescues our belief that fossilized bones and other archeological debris originally came from

living animals, the FH delivers an absurd belief in a future where dead bones and decayed matter will amazingly reassemble themselves and recreate life.

In a more recent paper, Chen [2025] in fact says that the PH is privileged by its simplicity. He says (p.974): “While many boundary conditions exhibit detailed correlations, only a select few qualify as fundamental laws—namely, those that are both simple and explanatory. One might ask why we prefer the Past Hypothesis, a macroscopic description, over a precise microscopic specification of the universe’s initial condition. The reason is that the Past Hypothesis is far simpler while still providing a powerful explanation for a wide range of temporally asymmetric regularities.” Chen is not here proposing what justifies believing the PH specifically over the skeptical alternative S that he engaged with in his earlier paper, but it may still be useful to make it clear why, even if things like predicted correlations, simplicity, and explanatory power make the PH a law, this does not block the skeptical threat posed by the FH.

To begin with, recall that skeptical hypothesis S was designed to assign E at least as high a likelihood as PH does, so the only way for PH to have a higher posterior than S is for PH to have a higher prior than S. So if PH gives a better explanation of E than S does, then that cannot be a matter of the likelihood it assigns E being higher. Perhaps we can insist that PH still gives a better explanation in some sense not clearly reflected in the Bayesian’s likelihoods. Perhaps PH explains better partly because it is simpler, or partly because it posits certain correlations in the universe. Perhaps being a boundary condition is part of what makes the PH simple or explanatorily powerful. And perhaps somehow all that really does make the PH a law. But, unfortunately, none of this undermines the claim of parity with the FH. Does the PH in any way offer explanatory power, or exhibit simplicity, or posit correlations, or play some role as a boundary condition, or qualify as a law, in a way that the FH will not equally well make a claim to? No,

the FH does everything the PH does in these respects. The FH posits low entropy, assigns likelihood to correlations, and if we want we can say it even imposes a boundary condition just as much as the PH does: let the universe *end* according to the FH, just as the PH claims it began. And since (as we assume throughout) the laws are time-symmetric, the FH will predict/retrodict E, and even nomically entail E, just as well as PH does. So, the two have perfectly equal claim to enjoying a high prior by virtue of all of these sorts of features.

Some readers are going to object that the PH and FH don't explain our total evidence E equally well, as follows. (This is also suggested by Chen's mention of "temporally asymmetric regularities" above.) Suppose we know by our dogmatic trust in observation and memory that entropy *increased* over some observed time span, from *a* to *b*. (We know that ice cubes are melting in warm water, and so on.) Could we argue that the PH explains such a fact, an increase in entropy from *a* to *b*, while the FH does not? No, we can't, and it is important to see why this objection (that I have found many people tempted by) doesn't get off the ground. What the PH helps explain is why entropy is where it is at time *a*, the start of the observed timespan, but the problem is that, given the level of entropy at time *a*, the PH then plays no role at all in explaining how entropy proceeded to increase from *a* to *b*. It's Boltzmann's argument that explains the increase from *a* to *b*, and does so equally well completely regardless of whether the PH is true or false.

We also observe correlations occurring over a span of time, but this again doesn't favor the PH. After pulling a left shoe out of the bargain bin, we usually dig up a right shoe soon after. But the PH doesn't explain or predict (specifically, it doesn't assign a high likelihood to) this kind of thing any more than the FH does. The PH says, about time *a*, that things like left shoes and right shoes are to be found in bins together, and the FH says the same thing about time *b*. There is no temporal asymmetry here that either hypothesis explains better than the other does.

7. *Conclusions*

Recently philosophers have been gripped by the skeptical threat posed by the empirical evidence that there will be so many Boltzmann Brains that you probably are one yourself.²¹ What I've tried to show in this paper is that the skeptical threat posed by the empirical evidence that we live in a Boltzmann Bubble is far more serious. To make the Boltzmann Brain threat interesting, we must make internalist and undogmatist epistemological assumptions. Externalists and dogmatists can easily claim justification for believing they are not Boltzmann Brains. What makes the threat posed by Boltzmann Bubbles so severe is that it's not at all easy to see what strong yet principled assumptions could allow us to answer it.

Should we give in to an extreme permissivist view that tolerates any coherent view no matter how wild? As I said earlier, this would let us say what we want to say, but it also lets others, flat earthers and such, say things that it is absurd to treat as equally reasonable.

Should we say, perhaps, that the datum that entropy has increased over some observed timespan is enough to justify an extrapolation all the way to the PH, an extrapolation from the humanly observable present to a time 14 billion years into the past? I have no very sophisticated argument against this, but I take it that this will strike others in the way it strikes me: a flimsy inference to rest an entire cosmological theory on. And if we think the trendline of entropy can be so easily extrapolated from our observed measurements of it, then we take away from the brilliance and significance usually credited to Boltzmann's own argument for why entropy will increase into the future. I'm not aware of anyone having ever endorsed anything like this position, but if someone now decides this is the best resolution of

²¹ Especially due to Carroll [2021].

the skeptical problem here, it would be useful to see it explicitly endorsed, defended, and developed. Take that as the conclusion of this paper if you want.

But what seems to me the best, or rather the least bad, escape route from this skeptical threat is a different response. It seems to me we need to give up a key assumption we've stuck to throughout this paper. I mean the assumption that the dynamics is time symmetric in the way nearly every plausible theory of the dynamics has been since Newton. The escape then is to endorse some kind of time asymmetric theory, perhaps a kind of GRW-style collapse theory, some theory on which the FH arguably cannot explain things as well as the PH can, and so the FH doesn't earn the high prior the PH earns itself.

But it is uncomfortable to imagine a philosophical examination like this one implying any particular answer to an empirical question like that. What if it just turns out that the empirical evidence strongly supports a time symmetric dynamics? Should we take the threat of skepticism to imply that this just can't happen? These are questions for further research.

Bibliography

- Albert, David, 2000, *Time and Chance*, Harvard
- Albert, David, 2015, *After Physics*, Harvard
- Carroll, Sean, 2021, "Why Boltzmann Brains Are Bad", in Dasgupta et al, *Current Controversies in Philosophy of Science*, Routledge
- Cassell, Lisa, 2025, "Conditionalization", in Sylvan et al, *A Companion to Epistemology*, 3rd ed., Wiley-Blackwell
- Chen, Eddy Keming, 2021, "Time's Arrow and Self-Locating Probability", *PPR*
- Chen, Eddy Keming, 2025, "The Simplicity of Physical Laws", *Nous*

- Earman, John, 2006, “The Past Hypothesis: Not Even False”, *Studies in History and Philosophy of Modern Physics*
- Elga, Adam, 2025, “Boltzmann Brains and Cognitive Instability”, *PPR*
- Emery, Nina, 2017, “Against Radical Quantum Ontologies”, *PPR*
- Emery, Nina, 2024, “Mooreanism in Metaphysics from Mooreanism in Physics”, *Inquiry*
- Huemer, Michael, 2001, *Skepticism and the Veil of Perception*, Rowman & Littlefield
- Lasonen-Aarnio, Maria, & Michael Huemer, 2024, “Does Fundamental Evidence Consist of Seemings?”, in Roeber et al, *Contemporary Debates in Epistemology*, 3rd ed., Routledge
- Loewer, Barry, 2012, “Two Accounts of Laws and Time”, *Philosophical Studies*
- Loewer, Barry, Brad Weslake, and Eric Winsberg, editors, 2023, *The Probability Map of the Universe*, Harvard
- Meacham, Chris, 2023, “The Meta-Reversibility Objection”, in Loewer et al [2023].
- Myrvold, Wayne, 2021, *Beyond Chance and Credence*, Oxford
- Price, Huw, 1996, *Time’s Arrow and Archimedes’ Point*, Oxford
- Pryor, James, 2000, “The Skeptic and the Dogmatist”, *Nous*
- Pryor, James, 2013, “Problems for Credulism”, in Tucker, *Seemings and Justification*, Oxford
- Wallace, David, 2023, “The Logic of the Past Hypothesis”, in Loewer et al [2023]
- Winsberg, Eric, 2004, “Can Conditioning on the ‘Past Hypothesis’ Militate Against the Reversibility Objections?”, *Philosophy of Science*