

ARTICLE

Does the Best System Need the Past Hypothesis?

Chris Dorst

University of Florida, Gainesville, FL, USA
Email: cdorst@ufl.edu

(Received 18 August 2022; revised 16 June 2023; accepted 05 October 2023; first published online 23 October 2023)

Abstract

Many philosophers sympathetic with Humeanism about laws have thought that the fundamental laws will include not only the traditional dynamical equations, but also two additional principles: the Past Hypothesis (PH) and the Statistical Postulate (SP). PH says that the universe began in a particular very-low-entropy macrostate $M(0)$, and SP posits a uniform probability distribution over the microstates compatible with $M(0)$. This view is arguably vindicated by the orthodox Humean Best System Account (BSA). However, I argue that recent developments of the BSA render the Past Hypothesis otiose. In particular, *Pragmatic* Humeanism does not support the idea that PH is a law.

1. Introduction

Orthodox Boltzmannian statistical mechanics seems to imply a radical skepticism about the past. If we assume that the fundamental dynamical laws are those of classical mechanics, then their time-reversal invariance suggests that the past, just like the future, was of higher entropy than the present; the world actually fluctuated into its current state, which is a local entropy minimum. In that case, it is almost certain that all of our beliefs and memories about the past are false, and the apparent records we have of it are misleading.

Of course, as Albert (2015, 5) puts it, “we are as sure as we are of anything that that’s not right.” This surety is partly what has led Albert and other philosophers to suggest that we include a principle that he calls the “Past Hypothesis” in our fundamental physical theory and accord it the status of a law of nature.¹ Doing so, they argue, secures the reliability of our memories, records, and beliefs about the past, while maintaining the reliability of orthodox statistical mechanics concerning predictions about the future.

¹ Boltzmann (1896/1964) originally suggested the Past Hypothesis. Feynman (1965/2017) is the first I am aware of who posited it as a law.

More specifically, Albert (2000, 2015) and Loewer (2001, 2004, 2012) have advocated the following package of principles as the fundamental physical theory of the world:

- the fundamental dynamical laws (which I will here assume to be those of classical mechanics);
- the **Past Hypothesis** (PH), namely the claim that the initial macrostate of the universe was one of extremely low entropy;
- the **Statistical Postulate** (SP), namely that there is a uniform probability distribution (according to the natural measure) over the microstates compatible with the macrostate described by the Past Hypothesis.

They call this package the *Mentaculus*, and suggest that its components all be regarded as fundamental laws of nature.²

It is uncontroversial that dynamical principles like the Hamiltonian equations of motion are candidates for fundamental physical laws. But the nomological status of principles like PH and SP is less clear. This is for at least three reasons. First, they concern the initial conditions of the universe, and it is not clear that laws can do that. Second, SP posits non-dynamical probabilities, and some accounts of laws may struggle to make sense of this. Third, PH refers to the property of being *low entropy*, which is vague and therefore may be unsuited to figuring in fundamental laws.

However, there is at least one account of laws on which none of these worries appear to gain traction, namely David Lewis's Best System Account (BSA).³ On that account, the laws comprise a system that provides an optimally efficient summary of all the particular matters of fact that obtain in the entire history of the universe, i.e. the "Humean mosaic." Given this background metaphysics, Loewer (2001, 2004) and Albert (2015) argue that the BSA can make sense of non-dynamical laws that posit probabilities over initial conditions even if the dynamical laws are fundamentally deterministic. In the same spirit, Chen (2022b) argues that the vagueness of "low entropy" does not prohibit PH from being a fundamental law according to the BSA.⁴

If these authors are right, the BSA does not *preclude* PH and SP from counting as fundamental laws. But what positive considerations suggest that they are BSA laws? The basic idea is straightforward: the *Mentaculus* is far more informative about the character of the mosaic than are the dynamical laws by themselves, and yet it is not significantly more complicated. As such, it is an efficient summary of the mosaic and

² In one respect this cannot be right, since the fundamental dynamical laws of this world are not those of classical mechanics. Chen (2022a, 2023) discusses how to implement PH in a quantum setting. While I will follow Albert and Loewer in focusing on the classical case, I think that my arguments here will carry over to the case of quantum mechanics and thus apply to Chen's Quantum *Mentaculus* and *Wentaculus*. It's also worth noting that, contra classical mechanics, our current best physics is not strictly time-reversal invariant. However, this lack of invariance is subtle, and it does not appear to account for the manifest temporal asymmetries with which we are familiar—hence the role for the Past Hypothesis. See Roberts (2022) for a lucid and helpful discussion.

³ See Lewis (1973, 1986, 1994).

⁴ See Hicks and Schaffer (2017) for a related argument that BSA laws may refer to non-fundamental properties. This helps to preempt a worry raised by Cohen and Callender (2009, 10) that PH cannot be a member of the best system because it is massively complex when translated into the language of microphysics.

therefore it is reasonable to think that it qualifies as the best system and its members qualify as laws. However, I shall argue here that this is incorrect; the best system does not include PH (though it does include a variant of SP).

My argument will not be that PH is omitted from the best system because it fails to do one or another of the many things that Albert and Loewer have argued that it does. For example, I will assume here that PH does help to secure the reliability of retrodictions (a point about which Leeds (2003) is skeptical); that it helps to explain the entropy increase of subsystems rather than just the universe as a whole (contra Winsberg (2004)); and that the entropy of the very early universe is indeed well-defined and thus that we can coherently formulate PH in the first place (contra Earman (2006)).⁵

Rather, my argument will be that the traditional conception of the BSA is indefensible, for reasons that have been articulated at length by some recent commentators. In particular, the BSA faces what I call the *Pragmatic Objection*, which is that it is unclear why creatures like us should care about discovering the best system so conceived. Some recent developments of the BSA have been expressly designed to avoid this objection, especially the accounts articulated in Hicks (2018), Dorst (2019a), and Jaag and Loew (2020). These accounts seek to devise a conception of the best system that is best *for creatures like us*, and in doing so they converge on roughly the same idea: the best system is one that is maximally effective at *amplifying* our information about the mosaic. I will argue that such a view makes the Past Hypothesis otiose, and therefore that it will not be included in the best system.

Here is how this paper proceeds. In section 2 I explain the Past Hypothesis in more detail, as well as the argument for its nomological status on the orthodox BSA. In section 3 I elaborate the Pragmatic Objection against the orthodox BSA and show how this motivates a shift to a more explicitly pragmatic view called the Best Predictive System Account (BPSA). In section 4, then, I argue that the Past Hypothesis is not a law on the BPSA. In section 5 I address some concerns about the resulting picture. In section 6 I conclude by drawing some morals about the explanatory limitations of physical laws from a Humean perspective.

2. The Past Hypothesis and the orthodox BSA

Understanding the Past Hypothesis requires a brief discussion about some of the foundations of Boltzmannian statistical mechanics (SM). At any given time t , the precise microstate of an n -particle system is given by a point X_t in a $6n$ -dimensional phase space Γ , which represents all possible instantaneous states of the system. The deterministic Hamiltonian equations of motion define a vector field on the phase space that fixes the dynamics of the phase point⁶ so that its temporal evolution corresponds to a curve in phase space.

Creatures like us are unable to discern the precise microstates of macroscopic systems. Instead, we characterize such systems in terms of macrovariables like temperature, volume, density, and pressure. Specifying the values of these

⁵ See Gryb (2021) for a comprehensive and insightful discussion of these difficulties, as well as several others.

⁶ Where it does not lead to confusion, I will be sloppy about the distinction between the representations (e.g. the phase point) and the items being represented (e.g. the microstate).

macrovariables confines the system's microstate to a subregion Γ_R of Γ ; the volume of any such region is given by the Liouville measure μ_L . Intuitively, knowing the macrostate of a system constrains its microstate to be one of those that realizes the macrostate in question, but typically there are many different microstates that do so. The basic assumption of Boltzmannian SM is that macrostates supervene on microstates: there can be no change in the system's macrostate without a corresponding change in its microstate.

For many systems there is a dominant macrostate, called the equilibrium macrostate Γ_{eq} , which takes up the majority of the volume of Γ according to the measure μ_L , i.e. $\mu_L(\Gamma_{\text{eq}})/\mu_L(\Gamma) \approx 1$. A system in thermal equilibrium has a microstate within Γ_{eq} .

The Boltzmann entropy S_B of a system with phase point X is given by

$$S_B(X) = k_B \log(\mu_L(\Gamma_X)), \quad (1)$$

where k_B is the Boltzmann constant and Γ_X is the macrostate containing X . A system's entropy is thus proportional to the phase space volume of the macrostate that it occupies. It follows that the equilibrium macrostate (if it exists) is the highest-entropy macrostate.

Given this setup, part of what Boltzmann was able to make plausible is that if a system is currently in a non-maximum-entropy macrostate, it is overwhelmingly likely that it will evolve through macrostates of increasing entropy and eventually reach the equilibrium macrostate. A hand-wavy argument for this is as follows. Suppose that at time t a system is in a non-maximum-entropy macrostate $M(t)$, corresponding to phase space region $\Gamma_{M(t)}$. Impose a probability distribution that is uniform on μ_L over the microstates in $\Gamma_{M(t)}$. Then the probability that the system will evolve into Γ_{eq} is overwhelmingly high. That is because most of the microstates in $\Gamma_{M(t)}$ are ones that deterministically increase in entropy toward the future until they reach Γ_{eq} . So given a uniform probability measure over the microstates in $\Gamma_{M(t)}$, it is overwhelmingly probable that a system in $M(t)$ lies on an entropy-increasing trajectory. Intuitively, unless the dynamics conspire to confine it to a small subregion of Γ , its phase point will aimlessly "wander" around Γ and find itself in macrostates whose phase space volume (read: entropy) is larger and larger until it reaches the largest-volume (equilibrium) macrostate, where it is overwhelmingly likely to stay.⁷

Thus we appear to have an account of the pervasive regularity that systems evolve to higher entropy in the future. The hope is that this sort of reasoning can be applied not just to gases in boxes (a typical starting point in these sorts of discussions), but to arbitrary macroscopic systems and even the universe as a whole. Processes as diverse as book pages becoming yellowed, hair turning grey, and ice melting in glasses of water would then all be accounted for by this explanation of entropy increase.

The problem, however, is that the Hamiltonian equations of motion are time-reversal invariant, and therefore all of this reasoning works equally well toward the past. That is, given a system in a non-maximum-entropy macrostate at time t_1 , we can use these arguments to predict that it will very likely be in a higher-entropy macrostate at t_2 . But we can also use these arguments to retrodict that it was very likely in a higher-entropy macrostate at t_0 . The same reasoning that leads us to infer

⁷ See, e.g., Albert (2000) or Frigg and Werndl (2022) for more detailed discussions of this reasoning.

that the half-melted ice cube in a glass of water will almost certainly be more melted in a few minutes also leads us to infer that it was almost certainly more melted a few minutes ago, and fluctuated into its current lower-entropy, less-melted state.

Applying this reasoning to the universe as a whole produces even more problematic results. Given that the universe is currently in a non-maximum-entropy macrostate, we can infer that it is overwhelmingly likely to evolve to higher-entropy macrostates in the future. But we can also infer that it is overwhelmingly likely to have evolved from higher-entropy states in the past. If that were right, it would mean that all of the records we appear to have of the past—memories, photographs, written accounts, etc.—are not accurate. They were produced not as a result of the events they appear to record, but by chaotic molecular fluctuations.

This problem is often called the *Reversibility Objection*:

Reversibility Objection: If Boltzmannian statistical mechanics predicts that systems in non-equilibrium states will increase in entropy toward the future, then it also predicts that they will increase in entropy toward the past.

This is problematic, of course, because we know that those systems were *not* in higher-entropy states toward the past. Given the abundance of non-equilibrium systems in our environment, then, Boltzmannian statistical mechanics licenses an enormous variety of egregiously incorrect inferences about the past.

A standard way of fixing this is to add the Past Hypothesis to the fundamental principles of Boltzmannian SM. That is, we posit that the universe started in a very-low-entropy macrostate $M(0)$ corresponding to a very-small-volume region Γ_0 of the universe's phase space. We then apply the aforementioned uniform probability distribution, on the Liouville measure μ_L , over the microstates in Γ_0 (this is the Statistical Postulate). In doing so, we arrive at Albert and Loewer's Mentaculus.⁸

With these additions, the thought is that Boltzmannian SM still makes all the correct predictions that it made without PH and SP, but it no longer makes the problematic retrodictions about the past being higher entropy. In particular, when we go to retrodict past states based on the present macrostate $M(t)$, we conditionalize not only on $M(t)$, but also on the fact that the universe began in $M(0)$. Then, since entropy-increasing trajectories are so much more common than entropy-decreasing trajectories, the probability that we reached $M(t)$ on an entropy-decreasing trajectory is minuscule; instead, it is far more likely that we got here on a trajectory that was more or less uniformly increasing in entropy since it began in $M(0)$. This blocks the conclusion that our present records of the past probably coalesced out of molecular chaos, and instead suggests that they were produced in roughly the manner that we think they were.

Many authors have been skeptical about one or another aspect of this project, but here I'm going to assume that the Mentaculus does indeed allow us to derive the universal increase in entropy over time as well as the various temporal asymmetries of our experience. That is, I'm going to suppose that it does everything that Albert and

⁸ There are other ways one might try to characterize the atypicality of the initial state rather than in terms of entropy, e.g. Penrose's (1989) Weyl curvature hypothesis.

Loewer hope that it does. If so, how should we think about the nomological status of the Past Hypothesis and Statistical Postulate?

It is difficult to address this question in a vacuum, namely, without a metaphysical account of laws of nature on hand. Many people who have regarded PH and SP as laws have done so on a Humean Best System framework.⁹ The argument for the lawhood of PH and SP according to the BSA is fairly straightforward. The best system is the one that achieves the best balance between simplicity, strength (i.e. informativeness), and fit (i.e. how probable the best system says the mosaic is) with the actual mosaic. The intuition here is that the best system should provide a concise and informative summary of what happens in the mosaic. Here is Lange making the point by imagining a conversation with God:

You: Describe the universe please, Lord.

God: Right now, there's a particle in state Ψ_1 and another particle in state Ψ_2 and I'll get to the other particles in a moment, but in exactly 150 million years and 3 seconds, there will be a particle in state Ψ_3 and . . .

You (checking watch): Lord, I have an appointment in a few minutes.

God: Alright, I'll describe the universe in the manner that is as brief and informative as it is possible simultaneously to be—by giving you the members of the “Best System.”

You: Do tell . . . (Lange 2009, 101–2).¹⁰

Given this picture, the argument that PH and SP qualify as BSA laws is as follows. The dynamical laws by themselves are time-reversal invariant, so they capture nothing about the pervasive temporal asymmetries (collectively, the “arrow of time”) in the mosaic. But add PH and SP and suddenly all those asymmetries come into focus. So the Mentaculus is far more informative about the character of the mosaic than are the dynamical laws by themselves, and yet it is not significantly more complicated. Plausibly, then, it qualifies as the best system, and its members as laws (cf. Loewer 2007).

One might admit that the Mentaculus is the best system and still demur from attributing nomological status to PH and SP. Lewis himself suggested that particular facts such as initial conditions could conceivably be included in the best system, but that only the regularities of the best system qualify as laws (1983, 367). However, many commentators have pointed out that PH and SP look and act like laws in a number of important respects—they support counterfactuals, explain other laws, underwrite inductive inferences, etc.¹¹ In short, in addition to figuring in the best system, they also play some standard law roles, so they should be counted as laws.

Of course, if PH and SP don't make it into the best system in the first place, then regardless of how much they act like laws in other respects, they will not count as laws in the best systems framework. Ultimately, I think that PH will not make it in, though a modified version of SP will. To get clearer on why that is, we need to look at the motivations behind Lewis's conception of the BSA.

⁹ See, e.g., Albert (2000, 2015), Loewer (2001, 2004, 2007, 2012), North (2011), and Chen (2022a).

¹⁰ See Albert (2015, 23) and Beebe (2000, 574) for similar expressions of this idea.

¹¹ See, e.g., Callender (2004), Chen (2023), Loewer (2012), and North (2011).

3. From the BSA to the BPSA

Hall (n.d.) has persuasively argued that there is a unique challenge for Humean accounts of lawhood that doesn't arise for non-Humean accounts. If Humeanism is right and the laws are mere patterns in the particular matters of fact, then it is unclear why we should be so interested in discovering them. In Hall's words, the question is why laws are "distinctively appropriate targets of scientific inquiry," or DATSIs. While non-Humeans can appeal to the laws' exalted metaphysical status here, Humeanism has no such luxury. Patterns in the mosaic are a dime a dozen, and nothing in the Humean viewpoint singles out some of those patterns as having any special metaphysical status. What Humeanism needs, then, is an alternative explanation of why the laws are DATSIs. And as Hall suggests, the most natural place to look is to their instrumental value: what makes the laws a worthy target of our investigations is their practical utility to creatures like us.

A number of recent authors have argued that the BSA fails to address this concern. This is because an efficient summary of the mosaic lacks significant practical utility.¹² Such a summary might report facts such as the universe's total mass or energy, the average lifespans of stars and sizes of galaxies, the total number of particles, etc. Indeed, one would expect the best system to contain a good deal of statistical facts reported in the form of averages, standard deviations, and so forth; such measures are designed to condense large amounts of information into an easily digestible form.

While these sorts of facts about the mosaic might be academically interesting, it is hard to see how they would be much use to an agent embedded within the mosaic attempting to navigate it. How would knowing the universe's total energy or the average galaxy size help you find your way around and pursue your goals? Coupled with Hall's argument that the Humean should appeal to the practical utility of the laws to explain their status at DATSIs, this gives us the *Pragmatic Objection* to the BSA:

Pragmatic Objection: The BSA fails to explain why the laws are distinctively appropriate targets of scientific inquiry.

This is a serious difficulty. Discovering the laws of nature is one of the central aims of science, and the BSA makes that aim appear misguided. The issue is that the BSA's standards of simplicity, strength, and fit are designed to achieve a goal—producing an efficient summary—that is just not that useful. What's needed, then, is a modification of these systematizing standards.

The requisite modifications have, I think, been proposed in the accounts developed by Hicks (2018), Dorst (2019a), and Jaag and Loew (2020). Here I focus primarily on my Best Predictive System Account (BPSA). The idea behind the BPSA is that we can address the Pragmatic Objection by tuning the systematizing standards to generate, not a maximally informative and concise summary of the mosaic, but rather an

¹² The ensuing discussion summarizes arguments from Hicks (2018), Dorst (2019a), and Jaag and Loew (2020). For critical evaluations of these arguments, see, e.g., Friend (2022), Sánchez (2022), and the contributions to Loew, Jaag, and Hicks (2023).

optimal predictive system for creatures in our epistemic situation. The laws are those patterns picked out by the best predictive system (BPS).¹³

In Dorst (2019a), I argued that such a system would be responsive to a number of “predictive desiderata” (see also Callender (2017, chapters 7 and 8), Jaag and Loew (2020), and Loewer (2020) for related discussions). Rather than recapitulating each of these desiderata and their motivations in detail, here I will simply describe the general type of system they are designed to produce. If nature is kind, the best predictive system will be one that has an input/output form that enables it to amplify a given chunk of information about the mosaic into a great deal more information. Any particular system might be better at amplifying some kinds of chunks rather than others; the best predictive system will be one that is optimized to amplifying the kinds of chunks that we typically have access to. That is, we will be able to plug in a chunk of information that we already know (or have the capacity to figure out), and the system will output a great deal more information in return.

This motivates two general kinds of predictive desiderata: “output maximizers” and “input constraints.” Output maximizers are aimed, straightforwardly, at ensuring that the BPS can output a lot of information that is useful to us. By contrast, input constraints aim at restricting the input information that is required to generate a given output. More specifically, they try to ensure that the BPS doesn’t require information that it would be prohibitively difficult for us to ascertain. Thus they will tend to reflect our epistemic limitations. For example, given that it is practically impossible for us to discern the precise microstate of typical macroscopic systems, Jaag and Loew (2020) suggest that the laws should be “error tolerant”: they should return approximately accurate predictions given approximately accurate inputs.

To some extent, output maximizers and input constraints conflict, much like strength and simplicity in the orthodox BSA.¹⁴ The system that gives us the laws is the one that achieves the best balance of these desiderata, and the system with the “best balance” is the one with the highest predictive utility.

The benefits of shifting from the BSA to the BPSA are manifold,¹⁵ but for our purposes it will suffice to note that the BPSA straightforwardly addresses the Pragmatic Objection to the orthodox BSA. That objection arose because an efficient summary of the Humean mosaic is not sufficiently useful to creatures like us, and therefore doesn’t explain why the laws are DATSIs. By contrast, a system of principles that are maximally predictively useful to creatures in our epistemic situation would clearly be *quite* useful to us, and easily explains why the laws are DATSIs.

If the BPSA is the right account of laws, then the actual laws will tend to conform to its systematizing standards. This is indeed what we find if we look at putative actual laws from physical practice—including, most notably for present purposes, the dynamical laws of classical mechanics; these laws are *highly* effective at amplifying our information about the mosaic.

¹³ Note that in this context, the word “predictive” is used in a generic sense that is not concerned with temporal direction.

¹⁴ I discuss this point further, and try to put it to work in defending the BPSA, in Dorst (2023).

¹⁵ For example, Dorst (2019b) argues that it helps the Humean avoid worries about explanatory circularity; both Loew and Jaag (2020) and Dorst (2022) develop accounts of why laws so conceived would support counterfactuals, as we expect laws to do; and Dorst (2021) argues that it dissolves the quantum measurement problem.

In sum, the BPSA improves on the BSA in both explaining why the laws are DATSIs and in selecting principles that better align with putative actual laws of nature found in scientific practice. It does so by bringing the epistemic situations of the users of that system into clearer view, and giving them a more prominent role in shaping the character of the best system. Given this change in perspective, however, we have to consider anew the question of whether the BPSA would deem PH and SP to be laws.

4. Why the Past Hypothesis is not a BPSA law

4.1 Predictions and probabilities

The BPSA clearly supports the idea that the classical dynamical laws would make it into the best system, given an appropriate mosaic. So let us assume that they do. What we need to consider is whether adding PH and SP to the classical dynamical laws would constitute an improvement on the predictive utility of the total system.

Jaag and Loew (2020, 2545) tentatively suggest that it might. To do so, they draw on an argument advanced by Albert (2015, chapter 1). Suppose we are told that a system is in a certain macrostate, but we are given no probability distribution over the possible microstates compatible with that macrostate. Then the dynamical laws by themselves allow many behaviors that would strike us as unexpected and bizarre, such as a rock spontaneously disassembling into statuettes of the British royal family or reciting the Gettysburg Address (ibid., 1). Albert's point is that if we were given a probability distribution over the microstates compatible with the rock's initial macrostate, we could discount these absurd possibilities and thus make better predictions about its behavior.¹⁶

Strictly speaking, this argument does not show that either PH or SP are required for predictive purposes. What it suggests is that, for such purposes, we need a probability distribution over the microstates of the rock compatible with its macrostate *at the time we are trying to predict its behavior*. As we have currently formulated it, SP applies to the microstates compatible with the initial macrostate $M(0)$ posited by PH, but nothing in Albert's argument requires any claim whatsoever about the initial conditions of the universe.

To be sure, PH and SP together might do the job that this argument suggests needs to be done. Given the SP-licensed probability distribution over the microstates compatible with $M(0)$, it is plausible that a typical macroscopic system that "branches off" from the rest of the universe at some later time t will be overwhelmingly likely to increase in entropy, in the ways we would roughly expect, toward the future of t .¹⁷ But it's not clear that PH and SP would do the job better than other principles. Consider, for example, a modified version of SP that can be applied at any time:

SP*: Given a system in macrostate $M(t)$, the probability at t that the system's microstate lies in $\Gamma_R \subseteq \Gamma_{M(t)} = \mu_L(\Gamma_R)/\mu_L(\Gamma_{M(t)})$, where Γ_R is a measurable subset of $\Gamma_{M(t)}$.¹⁸

¹⁶ Similar points are made by Ismael (2009, 91–2) and Loewer (2012, 129).

¹⁷ See Loewer (2012, 124–5) for a "seat of the pants" argument to this effect.

¹⁸ See, e.g., Frigg (2008, 672) for a similar construal of the statistical postulate.

SP* just imposes a uniform probability distribution, on the Liouville measure, over the microstates compatible with the rock's macrostate at t . This makes it incredibly unlikely that the rock will exhibit any of the bizarre behaviors Albert imagines to the future of t . So SP* by itself can do the job required by Albert's argument.

Let's abbreviate the system consisting of just the classical dynamical laws and SP* as \mathcal{S} , and the system consisting of the classical dynamical laws, SP, and PH (i.e. the Mentaculus) as \mathcal{S}^+ . Note that it is not entirely clear how to understand \mathcal{S} as a summary (as it would have to be on the traditional BSA), since it does not give consistent results about the probabilities for a given system's trajectory over a certain interval. Rather, it gives different probabilities at different times, and thus fails to provide a coherent picture of the contents of the mosaic. But this problem evaporates if \mathcal{S} is viewed as a system meant to amplify our information about the mosaic rather than summarize it. In that case, we choose the time at which to apply it, and supply the requisite information about a physical system's macrostate at that time. \mathcal{S} then provides us with probabilities about that physical system's macrostate at other times.

Of course, \mathcal{S} will tell us that any given physical system is very likely to have decreased in entropy to get to its macrostate at whatever t we choose. So the natural thing to say at this point is that while \mathcal{S} might be fine as far as predicting future behaviors of macroscopic systems like the rock, we still need \mathcal{S}^+ to secure all the retrodictions that we would otherwise get wrong if we relied on \mathcal{S} . Compared to \mathcal{S} , \mathcal{S}^+ licenses equally good inferences about the future and much more accurate inferences about the past, so it is surely a better predictive system, all things considered.

I think this is wrong; \mathcal{S} has a stronger claim to being the best predictive system than does \mathcal{S}^+ . To see why, it will help to look at the BPSA from a slightly different perspective.

4.2 The BPSA from the original position

It is tempting to explicate the BPSA by recasting Lange's one-on-one conversation with God:

You: Lord, I have no idea what I'm doing here. I just spent four hours at Target. Could you help me out?

God: Sure. Right now, there's a particle in state Ψ_1 , and on the other side of the universe there's a particle in state Ψ_2 , and when the second particle moves in a trajectory described by this fancy equation, the first particle will move in a trajectory described by this other fancy equation . . .

You (frustrated): Lord, I can't see the other side of the universe.

God: Good point. In that case, I'll give you a set of principles that would be maximally predictively useful to you.

You: Do tell . . .

Ultimately, however, I think this can't be quite right.

Granted, if that is how the conversation really went, then it is plausible that the set of principles God provides you would satisfy the sorts of predictive desiderata that are employed in the BPSA and that are characteristic of putative actual laws. But God's

principles would also be likely to possess features that are *not* characteristic of putative actual laws.

In particular, the principles God provides you could be responsive to your idiosyncratic epistemic profile, which contains all sorts of information that we would not expect to be referenced by laws of nature. Maybe it includes, for example, the brand and model of your current toothbrush, or the paths your family walked when you were a child, or melodies tied to particular people from your past. Your familiarity with these sorts of things would make it fair game for God's principles to appeal to them. But I would be pretty surprised if the laws of nature made reference to the Sonicare 4100 (available at Target), or to that two-mile loop around the lake that I used to walk, or to Joplin's "The Entertainer." In short, while it may be reasonable to expect the laws to be responsive to your general epistemic situation, it is not reasonable to expect them to be personalized.¹⁹

I think this suggests that the proper analogy for the motivations of the view is not a one-on-one conversation with God, but something more like a Rawlsian original position. Imagine that you are behind a veil of ignorance that prevents you from knowing which person you are in the mosaic, while at the same time you know the entirety of the mosaic. Using that knowledge your task is to design a set of principles that—when the veil of ignorance is lifted and you figure out who you are (but also lose your knowledge of the entirety of the mosaic)—would be most predictively useful to you.²⁰

Why is this a good way to think about the BPS? My entirely unoriginal suggestion is that science is fundamentally a communal, human enterprise, and human beings occupy all sorts of different positions and work toward all sorts of different goals. Accordingly, science aims to find predictive principles that are useful for everyone, not just for some particular person. The veil of ignorance is thus a way of securing the laws' *general* utility, *contra* what might result from a one-on-one conversation with God.

As with any original position construction, there will be questions about precisely who gets included and, correspondingly, the precise group of people that you might end up being once the veil of ignorance is lifted—in short, who are the "creatures like us" that inform the structure of the BPS? In this context there is no obvious role for a social contract, so the first of these questions is easy to answer: you are the only person in the original position. The second question is harder. You know that you are a human and not, say, a North Atlantic lobster (let's suppose). But on the other hand, can we also presuppose that you know that you're not an early Cro-Magnon, or someone with severe Alzheimers? As these examples illustrate, the exact extent of your ignorance may influence the BPS you end up designing.

So we will have to make some choices about who the veil of ignorance covers. The process of doing so can be governed by reflective equilibrium, weighing who we think

¹⁹ Thanks to Marc Lange for helping me appreciate this point.

²⁰ It is important to understand your task in the original position counterfactually. If you know the entire mosaic, then you know everything that ever happens, including what theories scientists come up with and how those theories are used to affect people's lives. So your task is not to produce a system that you *will be given* once you find yourself in the mosaic. Rather, your task is to produce a system that you *would most want to have access to* no matter who you end up being. Whether you ever actually have access to this system is irrelevant.

ought to be included against whether the resultant principles come out looking like laws of nature. I shall leave the details of these considerations aside here, and merely note in passing that they raise some interesting normative questions about the origins and ethical status of our concept of lawhood.

4.3 The status of PH and SP

Given this setup, we can more clearly state the assumptions that the BPS makes about our epistemic situation: it assumes precisely those features that are common to the epistemic profile of everyone that you might end up being once the veil of ignorance is lifted. We might call this a “generalized epistemic profile” (GEP). The GEP includes not only items of occurrent knowledge, but also facts about the kinds of information we are in a position to ascertain. Given each of our epistemic idiosyncrasies, none of our epistemic profiles *exactly* agrees with the GEP, though all of them *roughly* do, at least in their general character. Thus we can say that the BPS is designed to be optimally predictively useful to someone whose epistemic state is characterized by the GEP.

This implies that the process of evaluating a candidate system for its predictive utility is more complex than has previously been recognized. Consider an arbitrary fact f that would be predictively useful to us (f could be, say, a dynamical principle or a particular matter of fact). It is not necessarily true that a candidate system including f is automatically a better predictive system than one excluding it, for if f is already contained in the GEP, then nothing other than redundancy is gained by including it in the candidate system as well. Doing so would inflate the system without any corresponding benefit in the amount of information we can extract about the mosaic, making it a less efficient predictive system overall. Note that when I talk about the information *we* can extract from a system, the “*we*” is important; it refers to agents whose epistemic position is (roughly) characterized by the GEP. If a candidate system doesn’t include f but the GEP does, then *we* can still extract the relevant information about the mosaic. Essentially the GEP can “cover” for the candidate system, allowing the system to be leaner than it could be otherwise.

Thus, for any fact (or set of facts) that would be predictively useful, before inferring that it will be included in the BPS, we need to ask whether it is already included in the GEP. Of course, this requires us to get clear about some of the contents of the GEP. For example, the GEP will be characterized by the epistemic limitations motivating the input constraint *predictive desiderata*, such as the fact that we cannot readily discern the precise microstate of typical macroscopic systems, or the fact that we tend to have access to spatiotemporally local information.

But, more importantly for present purposes, the GEP will also include knowledge of the arrow of time. Each of us, after all, is incredibly well acquainted with the temporal asymmetries. As Albert suggests, some crude familiarity with them “will very plausibly have been hard-wired into the cognitive apparatus of any well-adapted biological species by means of a combination of natural selection and everyday experience and explicit study and God knows what else” (2015, 39). Clearly, any summary of our epistemic situation that *left out* knowledge of the arrow of time would be radically incomplete.

But if the relevant facts about the arrow of time are already included in the GEP, the BPS itself doesn't need to secure them. One way to see this is to imagine yourself trying to design the BPS from the original position. You look around and notice that everyone you might end up being, once the veil is lifted, already knows the sorts of facts about the arrow of time that S^+ conveys and S does not. Indeed, they are so confident in their beliefs about these facts that essentially nothing could lead them to revise those beliefs; these are facts, remember, about which they are *as sure as they are of anything*. They will be fixed points in any inferences drawn from the BPS. So why bother repeating them?

Moreover, if the GEP contains information about the arrow of time, and if proponents of PH are correct in thinking that the arrow must ultimately be traceable to a universal entropy gradient, then system S in combination with the GEP will already rule out microhistories that involve significant entropy increases toward the past: we know that the past was different from the future in such-and-such ways, and the only way for those differences to obtain given S is for the past to have been lower entropy. If so, the combination of S and the GEP will already imply that the initial state was one of remarkably low entropy. There is, again, no need to add that claim to the system itself.²¹

What becomes of the Reversibility Objection, on this view? When we encountered it in section 2, the way we put it was that orthodox Boltzmannian SM (i.e. system S) makes an enormous variety of incorrect predictions about the past. But systems—sets of propositions—don't make predictions by themselves; intentional creatures—*agents*—make predictions using these systems. The inferences that a given agent will draw from a given set of propositions depend not just on the propositions, but also on the agent: their innate psychologies, background knowledge, etc. And no agent whose epistemic profile is characterized by the GEP would use S to seriously infer that the universe coalesced out of molecular chaos into its present state. We know this, of course, because *we* are such agents, and *we have not drawn this inference*.

Consider a few analogies. Why doesn't the best predictive system include, say, the Peano axioms, or an explication of *modus ponens*²², or an arrow schema telling the users in which direction to *read* its propositions?²³ Without them, wouldn't a system like S be pretty useless? Wouldn't their inclusion therefore massively increase S 's amplifying power, allowing the derivation of many more truths about the mosaic? The answer is that they *would* make it into a system designed for creatures whose epistemic profile is sufficiently impoverished; for such creatures, it would be quite useful to be told the kinds of foundational logicomathematical principles and linguistic conventions that we take for granted. But including such principles in a

²¹ The position I am recommending here is somewhat reminiscent of the Special Science Localism suggested by Callender (2011), according to which there is no *primordial* statistical mechanical probability distribution. Rather, those probabilities are applied by agents in particular contexts when they want to predict the behavior of (particular sorts of) chunks of the mosaic. Where the accounts primarily differ is that the present one maintains the hope that the statistical mechanical probabilities could be global rather than merely local (notwithstanding the fact that their application toward the past is constrained by our knowledge of the arrow of time). If so, they could be applied to any chunk of the mosaic, including to the initial total state.

²² Cf. Carroll (1895).

²³ Cf. *Philosophical Investigations*, §86 (Wittgenstein, 2009).

predictive system designed for us would be pointless. No one ever worried that without the appropriate arrow schemas, \mathcal{S} 's principles might be read incorrectly, and thus used to infer all sorts of nonsense about the mosaic. Likewise, there is no serious worry that without PH, \mathcal{S} might be used to infer all sorts of falsities about the past. Neither of these are genuine possibilities for creatures like us.

As this discussion makes clear, altering the GEP can affect the character of the BPS. For example, we could imagine what the BPS would be like if the GEP wasn't subject to some of the epistemic limitations that motivate the input constraint predictive desiderata. Likewise, we could imagine removing knowledge of the arrow of time from the GEP. Doing so would produce a role for PH in the best predictive system: if we did not already know about the world's temporal asymmetries, then \mathcal{S}^+ would be far superior to \mathcal{S} in terms of its predictive utility, for it would allow us to learn all sorts of facts about the past that we would be ignorant of using \mathcal{S} .

But this is trivial. Given sufficient creativity in the design of a GEP, we can imagine a BPS with *all sorts* of surprising features. For example, consider what the BPS would be like for creatures who know more about the future than about the past, or who only know spatiotemporally diffuse information, or who know nothing at all. While it is difficult to imagine what the lives of such creatures would be like, there is comparably little difficulty in conceiving of a predictive system tuned to their epistemic conditions. It would be very different from a predictive system tuned to our own—quite dissimilar to either \mathcal{S} or \mathcal{S}^+ —and we would have very little reason to care about it aside from academic curiosity. The same, I suggest, is true of a predictive system designed for creatures who have no knowledge of time's arrow. The epistemic lives of such creatures would be very different from ours, and there is little reason for us to care about the best predictive system for them.

In other words, by changing the GEP so that it no longer characterizes our epistemic situation, we re-encounter the Pragmatic Objection that was the downfall of the orthodox BSA. Doing so makes it unclear why the laws, so conceived, are distinctively appropriate targets of scientific inquiry. The fact that we can imagine different GEPs, some of which would make \mathcal{S}^+ the best predictive system, is therefore of little consequence.

5. Loose ends

The argument I just advanced raises a number of significant questions. This section addresses several of them.

Does this argument imply that any principles that agree with our intuitive judgments or expectations cannot be members of the best predictive system?

The reason one might think this is that if everyone you might end up being (once the veil is lifted) correctly expects a given type of physical system to behave in a certain way, then information to that effect needn't be included in the best predictive system because it would already be part of the GEP. But the above argument does not imply this. For example, everyone correctly expects a cannonball fired at a positive angle to the horizontal to follow an arcing trajectory and fall back to the ground. So why should Newton's equations of motion, which predict such behavior, make it into the BPS? The answer is obvious: Newton's equations of motion provide far more information than our naive expectations. Given facts like the mass of the cannonball

and the force generated by the cannon explosion, they allow us to make a far more precise prediction about the trajectory of the cannonball than we could get by relying on our (widely shared) expectations. In general, a principle's agreement with our intuitive expectations isn't grounds for excluding it from the BPS if that principle also goes beyond those expectations in certain ways, e.g. by precisifying them or otherwise extending them beyond the point where they give out.

The system \mathcal{S} is empirically incoherent, since it implies that our present evidence for it is unreliable. By contrast, \mathcal{S}^+ is empirically coherent given its inclusion of PH and SP. Isn't this reason to prefer \mathcal{S}^+ to \mathcal{S} as the best predictive system?

In short, no. This worry makes the same mistake we have already seen: it tries to evaluate \mathcal{S} in a vacuum, whereas \mathcal{S} is really designed to function in the context of the GEP. And the total body of knowledge encompassed by both the GEP and \mathcal{S} is not empirically incoherent.

Perhaps the worry is rather that \mathcal{S} is *inconsistent* with the GEP, because they reach inconsistent verdicts about the past. But of course, they are not *strictly* inconsistent; it's not as if \mathcal{S} says that it is impossible for the past to have been lower entropy, only that it is monumentally unlikely. In more mundane contexts, we routinely accept many pairs of claims, each of which makes the other improbable when considered in isolation. Indeed, the traditional Mentaclus itself is like this, since PH is massively unlikely on the basis of the classical dynamical laws. So it's hard for me to see why the near inconsistency of \mathcal{S} and the GEP would be a problem.

Price (1996, 2004) has argued that we ought to try to explain why the universe's initial state had remarkably low entropy. If PH is part of the best system, then we have such an explanation: it was required by law.²⁴ But if PH isn't in the best system and therefore isn't a law, this explanation is foreclosed. Isn't this an explanatory deficiency of \mathcal{S} as compared to \mathcal{S}^+ ?²⁵

The appropriate response here is to keep firmly in mind that explanatory considerations are not part of the criteria for membership in the best system, at least not according to the BPSA. If one candidate system is better than another, that is because it is a better *predictive* system for creatures like us, not because it is more explanatory. Of course, insofar as the explanatory superiority of one system can be parlayed into predictive superiority, that system will be better. (In Dorst (2019b) I suggested that this will often be the case, providing an indirect account of the value of seeking explanations in the first place: they lead us to predictively superior theories.) But as we've seen, \mathcal{S}^+ is not predictively superior to \mathcal{S} , so its ability to offer an explanation of the initial low entropy does not help it in the competition.

This accords nicely with the position of Callender (2004), who questions why we ought to *demand* an explanation of initial low entropy. Certainly we should not eschew one if there are independent reasons to find it plausible. But that is a far cry from deciding between two theories solely on the basis of their (in)abilities to explain such a state. As Callender puts it, "to know *beforehand*, as it were . . . that the [low-entropy initial state] needs explanation seems too strong" (207, Callender's emphasis).

If these arguments succeed, they show that \mathcal{S} is a better predictive system than \mathcal{S}^+ , since the latter contains redundant information. But could the same be said of \mathcal{S} ? It contains both the classical dynamical laws and SP^ ; what about a system, which we might call \mathcal{S}^- , which*

²⁴ Cf. Callender (2004, section 6).

²⁵ Thanks to an anonymous referee for raising this concern.

consists of the dynamical laws alone? Is there a case to be made that it is superior to \mathcal{S} for reasons similar to the ones that render \mathcal{S} superior to \mathcal{S}^+ ?

In order for \mathcal{S}^- to be superior to \mathcal{S} , it would have to reveal at least as much information about the mosaic (when combined with the GEP). If that were the case, then it would be a better predictive system because it is more compact and yet returns the same outputs, making it better at amplifying our information. The problem is that \mathcal{S}^- does not reveal as much information about the mosaic as does \mathcal{S} , for the former does not provide us with any probabilities. It deems propositions possible or impossible based on whether they are consistent with the dynamical laws, but it does not assign probabilities to these possibilities. This wouldn't be problematic if we could derive the probabilities from the GEP, for then the combination of \mathcal{S}^- and the GEP together would be just as informative as the combination of \mathcal{S} and the GEP (the SP^* -licensed probabilities from \mathcal{S} would be redundant). But I do not see any way of getting the probabilities out of the GEP. The relevant probabilities are ones that are uniform, on the Liouville measure μ_L , over the microstates compatible with a given macrostate. What in the GEP directs us to apply a uniform probability distribution as opposed to a non-uniform one? Even confining ourselves to uniform probability distributions, altering the measure would still alter the resulting probabilities. What in the GEP fixes the correct measure? As far as I can see, nothing in our shared epistemic state does anything to underwrite either a uniform probability distribution or the Liouville measure as the correct grounds for making predictions on the basis of the dynamical laws. These have to be supplied by the predictive system itself, rendering \mathcal{S} superior to \mathcal{S}^- .

What should we make of the probabilities posited by SP^* ? Are they objective, epistemic, or some mixture thereof?

This question has been raised with respect to the probabilities employed by the Mentaculus.²⁶ The basic puzzle is that the Hamiltonian equations of motion are deterministic, so it is hard to see how to fit probabilities into the picture unless they are epistemic. But if they are epistemic, then it is hard to see how they can figure into a fundamental theory that is supposed to license explanations of the phenomena within its scope. Of course, the same sorts of questions can be raised about the probabilities employed by \mathcal{S} , since it uses the same deterministic equations of motion as \mathcal{S}^+ .

This is a large question, and I doubt that anything I can say here will be decisive, but here is my answer: the probabilities employed by \mathcal{S} are *objective enough*. By that I mean that they can play the explanatory roles required of them.

My reason for thinking this is that these probabilities have the same ontological status as any other structures that are posited by the best predictive system, and such structures often have significant explanatory roles. To summarize a large literature very briefly, a number of philosophers (e.g. Callender (2017), Cohen and Callender (2009), Loewer (2007, 2020), Hall (n.d.), and Miller (2014)) have suggested that the best system may introduce what we might call "manufactured structures" along with the laws to help achieve a better overall balance of the systematizing standards. For example, Hall (n.d.) proposes doing this with the magnitudes of mass and charge.

²⁶ For discussion, see, e.g., Frigg (2008), Frigg and Hoefer (2015), Ismael (2009), Loewer (2001), and Winsberg (2008).

Consider a mosaic in which the only fundamental non-modal facts concern the locations of particles at times. We might then allow candidate systems to posit additional magnitudes possessed by these particles, where the motivation for making such posits is for the candidate system to achieve a better overall balance of the relevant standards. In that event, as Hall puts it, “what would make it the case that there *are* masses and charges is just that there is a candidate system that says so, and that, partly by saying so, manages to achieve an optimal combination of simplicity and informativeness” (27).

Although this maneuver was originally proposed in the context of the orthodox BSA, it works equally well in the context of the BPSA. The BPS would then be allowed to posit structures that are not fundamental elements of the mosaic, and claims about such structures would be made true by the fact that the candidate system that posits them is the best predictive system.

I do not mean to take a stand here on exactly which structures might be legitimately introduced by the BPS in this manner.²⁷ All that I would point out is that such structures often have significant explanatory roles to play within the relevant theories. Mass and charge, for example, are supposed to help explain particle motions within classical mechanics. If they can be made to do this even though they are really manufactured structures (on this view), then surely the probabilities employed by SP*—which gain entry into the BPS in the same manner, i.e. by producing a better predictive system overall—can play those explanatory roles as well. Undoubtedly this will require a rather deflationary notion of explanation in order to work, but that’s something Humeans should already be on board with.

Still, you might push back: “Are these probabilities *really* objective? Aren’t they at least somewhat epistemic?” After all, if one uses them while ignoring parts of the GEP, one can derive all sorts of falsities about the past. But of course, they *aren’t designed* to be used while ignoring the GEP. They are manufactured structures tuned to the epistemic conditions of creatures like us, just like mass and charge would be on Hall’s proposal, or like the wavefunction on Miller’s, or indeed like the laws themselves on the BPSA. Does that make them epistemic as opposed to objective? The best I can do here is to echo Dennett (1991, 51): I think the view itself is clearer than either of the labels, so I shall leave that question to anyone who still finds illumination in them.

6. Conclusion: Why physics can’t explain everything

Frisch (2014) (from which this section borrows its title) argues that the pragmatic motivations behind Albert and Loewer’s conception of the BSA make it hard to see why the Mentaculus would qualify as the best system. Without questioning whether the Mentaculus would be *part* of the best system, he suggests that various special science laws would likely be included as well. This, in turn, leads him to some broader conclusions regarding limitations on the explanatory ambitions of physics.

In a sense, I have argued for the converse claim. Whereas Frisch thinks that the best system will consist of the Mentaculus *and then some*, I have argued that it will consist of less than the Mentaculus.²⁸ In particular, I have suggested that in the

²⁷ Miller (2014) and Bhogal and Perry (2017) compellingly use this maneuver with respect to the quantum wavefunction.

²⁸ I’ve reached no verdict here about Frisch’s claim that special science laws would also make it in.

context of the BPSA, the Past Hypothesis is otiose: it adds nothing that is not already captured by the combination of \mathcal{S} and the GEP.

Despite drawing the converse conclusion about the Mentaculus, I agree with Frisch in thinking that there are lessons here about the explanatory ambitions of physics (or, more accurately, of physical laws) from a Humean perspective. According to the BPSA, the laws are “designed” for creatures whose epistemic lives occur *in medias res*: some kinds of facts we have already established; many others we have not. If this is the right way to think of the laws, then they may take for granted the same sorts of facts that we do. But then the expectation that we should be able to explain everything by appeal to the laws alone is unrealistic; the laws cannot be expected to explain the sorts of facts that they take for granted. In particular, I don’t think we should expect the laws to license an explanation of the arrow of time. Facts about the arrow are already assumed in the epistemic milieu that provides the practical need—and hence, on the Humean view, the metaphysical foundations—for the laws in the first place.

Acknowledgments. Thanks to Eddy Keming Chen, Josh Hunt, Marc Lange, and Barry Loewer for helpful discussions. Also thanks to two anonymous referees for this journal, whose feedback substantially improved the paper.

References

- Albert, David. 2000. *Time and Chance*. Cambridge, MA: Harvard University Press.
- Albert, David. 2015. *After Physics*. Cambridge, MA: Harvard University Press.
- Beebe, Helen. 2000. “The Non-Governing Conception of Laws of Nature.” *Philosophy and Phenomenological Research* 61 (3):571–94. <https://doi.org/10.2307/2653613>.
- Bhagal, Harjit and Zee Perry. 2017. “What the Humean Should Say About Entanglement.” *Noûs* 51 (1): 74–94. <https://doi.org/10.1111/nous.12095>.
- Boltzmann, Ludwig. (1896) 1964. *Lectures on Gas Theory*. Berkeley, CA: University of California Press.
- Callender, Craig. 2004. “Measures, Explanations and the Past: Should ‘Special’ Initial Conditions Be Explained?” *The British Journal for the Philosophy of Science* 55 (2):195–217. <https://doi.org/10.1093/bjps/55.2.195>.
- Callender, Craig. 2011. “The Past Histories of Molecules.” In *Probabilities in Physics*, edited by Claus Beisbart and Stephen Hartmann, 83–114. New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199577439.003.0004>.
- Callender, Craig. 2017. *What Makes Time Special?* Oxford: Oxford University Press.
- Carroll, Lewis. 1895. “What the Tortoise Said to Achilles.” *Mind* IV (14):278–80.
- Chen, Eddy Keming. 2022a. “From Time Asymmetry to Quantum Entanglement: A Humean Unification.” *Noûs* 56 (1):227–55. <https://doi.org/10.1111/nous.12355>.
- Chen, Eddy Keming. 2022b. “Fundamental Nomic Vagueness.” *Philosophical Review* 131 (1):1–49.
- Chen, Eddy Keming. 2023. “The Past Hypothesis and the Nature of Physical Laws.” In *Time’s Arrow and the Probability Structure of the World*, edited by Barry Lewer, Eric Winsberg, and Brad Weslake, 204–48. Cambridge, MA: Harvard University Press.
- Cohen, Jonathan and Craig Callender. 2009. “A Better Best System Account of Lawhood.” *Philosophical Studies* 145:1–34. <https://doi.org/10.1007/s11098-009-9389-3>.
- Dennett, Daniel. 1991. “Real Patterns.” *The Journal of Philosophy* 88 (1):27–51. <https://doi.org/10.2307/2027085>.
- Dorst, Chris. 2019a. “Towards a Best Predictive System Account of Laws of Nature.” *British Journal for the Philosophy of Science* 70 (3):877–900. <https://doi.org/10.1093/bjps/axy016>.
- Dorst, Chris. 2019b. “Humean Laws, Explanatory Circularity, and the Aim of Scientific Explanation.” *Philosophical Studies* 176:2657–79. <https://doi.org/10.1007/s11098-018-1145-0>.
- Dorst, Chris. 2021. “There Is No Measurement Problem for Humeans.” *Noûs* 57 (2):263–89. <https://doi.org/10.1111/nous.12399>.

- Dorst, Chris. 2022. "Why Do the Laws Support Counterfactuals?" *Erkenntnis* 87:545–66. <https://doi.org/10.1007/s10670-019-00207-1>.
- Dorst, Chris. 2023. "Predictive Infelicities and the Instability of Predictive Optimality." In *Humean Laws for Human Agents*, edited by Michael Hicks, Christian Loew, and Siegfried Jaag, 193–214. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780192893819.003.0010>.
- Earman, John. 2006. "The 'Past Hypothesis': Not Even False." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 37 (3):399–430. <https://doi.org/10.1016/j.shpsb.2006.03.002>.
- Feynman, Richard. (1965) 2017. *The Character of Physical Law*. Cambridge, MA: MIT Press.
- Friend, Toby. 2022. "The Humean Pragmatic Turn and the Case for Revisionary Best Systems Accounts." *European Journal for Philosophy of Science* 12:11. <https://doi.org/10.1007/s13194-021-00436-8>.
- Frigg, Roman. 2008. "Chance in Boltzmannian Statistical Mechanics." *Philosophy of Science* 75 (5):670–81. <https://doi.org/10.1086/594513>.
- Frigg, Roman and Carl Hoefer. 2015. "The Best Humean System for Statistical Mechanics." *Erkenntnis* 80:551–74. <https://www.jstor.org/stable/24735114>.
- Frigg, Roman and Charlotte Werndl. 2022. "Equilibrium in Boltzmannian Statistical Mechanics." In *The Routledge Companion to Philosophy of Physics*, edited by Eleanor Knox and Alastair Wilson, 403–14. New York: Routledge.
- Frisch, Mathias. 2014. "Why Physics Can't Explain Everything." In *Chance and Temporal Asymmetry*, edited by Alastair Wilson, 221–40. New York: Oxford University Press.
- Gryb, Sean. 2021. "New Difficulties for the Past Hypothesis." *Philosophy of Science* 88 (3):511–32. <https://doi.org/10.1086/712879>.
- Hall, Ned. n.d. "Humean Reductionism about Laws of Nature." <https://philpapers.org/archive/HALHRA.pdf>.
- Hicks, Michael. 2018. "Dynamic Humeanism." *British Journal for the Philosophy of Science* 69 (4):987–1007. <https://doi.org/10.1093/bjps/axx006>.
- Hicks, Michael and Jonathan Schaffer. 2017. "Derivative Properties in Fundamental Laws." *British Journal for the Philosophy of Science* 68 (2):411–50. <https://doi.org/10.1093/bjps/axv039>.
- Ismael, Jenann. 2009. "Probability in Deterministic Physics." *The Journal of Philosophy* 106 (2):89–108. <https://www.jstor.org/stable/20620154>.
- Jaag, Siegfried and Christian Loew. 2020. "Making Best Systems Best for Us." *Synthese* 197:2525–50. <https://doi.org/10.1007/s11229-018-1829-1>.
- Lange, Marc. 2009. *Laws and Lawmakers*. New York: Oxford University Press.
- Leeds, Stephen. 2003. "Foundations of Statistical Mechanics—Two Approaches." *Philosophy of Science* 70 (1):126–44. <https://doi.org/10.1086/367873>.
- Lewis, David. 1973. *Counterfactuals*. Malden, MA: Blackwell.
- Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61 (4):343–77. <https://doi.org/10.1080/00048408312341131>.
- Lewis, David. 1986. *Philosophical Papers*, Vol. II. New York: Oxford University Press.
- Lewis, David. 1994. "Humean Supervenience Debugged." *Mind* 103 (412):474–90. <https://www.jstor.org/stable/2254396>.
- Loew, Christian and Siegfried Jaag. 2020. "Humean Laws and (Nested) Counterfactuals." *Philosophical Quarterly* 70 (278):93–113. <https://doi.org/10.1093/pq/pqz037>.
- Loew, Christian, Siegfried Jaag, and Michael Townsen Hicks (eds.) 2023. *Humean Laws for Human Agents*. Oxford: Oxford University Press.
- Loewer, Barry. 2001. "Determinism and Chance." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 32 (4):609–20. [https://doi.org/10.1016/S1355-2198\(01\)00028-4](https://doi.org/10.1016/S1355-2198(01)00028-4).
- Loewer, Barry. 2004. "David Lewis's Humean Theory of Objective Chance." *Philosophy of Science* 71 (5):1115–25. <https://doi.org/10.1086/428015>.
- Loewer, Barry. 2007. "Counterfactuals and the Second Law." In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by Huw Price and Richard Corry, 293–326. Oxford: Clarendon Press.
- Loewer, Barry. 2012. "Two Accounts of Laws and Time." *Philosophical Studies* 160:115–37. <https://doi.org/10.1007/s11098-012-9911-x>.

- Loewer, Barry. 2020. "The Package Deal Account of Laws and Properties." *Synthese* 199:1065–89. <https://doi.org/10.1007/s11229-020-02765-2>.
- Miller, Elizabeth. 2014. "Quantum Entanglement, Bohmian Mechanics, and Humean Supervenience." *Australasian Journal of Philosophy* 92 (3):567–83. <https://doi.org/10.1080/00048402.2013.832786>.
- North, Jill. 2011. "Time in Thermodynamics." In *The Oxford Handbook of Philosophy of Time*, edited by Craig Callender, 312–50. New York: Oxford University Press.
- Penrose, Roger. 1989. *The Emperor's New Mind*. New York: Oxford University Press.
- Price, Huw. 1996. *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time*. New York: Oxford University Press.
- Price, Huw. 2004. "On the Origins of the Arrow of Time: Why There is Still a Puzzle about the Low-Entropy Past." In *Contemporary Debates in Philosophy of Science*, edited by Christopher Hitchcock, 219–39. Oxford: Blackwell.
- Roberts, Bryan. 2022. *Reversing the Arrow of Time*. New York: Oxford University Press.
- Sánchez, Verónica Gómez. 2022. "From Nomic Humeanism to Normative Relativism." *Philosophical Perspectives* 36 (1):118–39. <https://doi.org/10.1111/phpe.12163>.
- Winsberg, Eric. 2004. "Can Conditioning on the 'Past Hypothesis' Militate Against the Reversibility Objections?" *Philosophy of Science* 71 (4):489–504. <https://doi.org/10.1086/423749>.
- Winsberg, Eric. 2008. "Laws and Chances in Statistical Mechanics." *Studies in History and Philosophy of Modern Physics* 39 (4):872–88. <https://doi.org/10.1016/j.shpsb.2008.05.005>.
- Wittgenstein, Ludwig. 2009. *Philosophical Investigations*, 4th ed. Translated by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Malden, MA: Wiley-Blackwell.

Cite this article: Dorst, Chris. 2024. "Does the Best System Need the Past Hypothesis?" *Philosophy of Science* 91 (2):410–429. <https://doi.org/10.1017/psa.2023.152>