Objective Credences, Epistemic Chances, and Explanations of Time Asymmetry: Review of Myrvold's 'Beyond Chance and Credence'

Katie Robertson* and Alexander Franklin†

May 24, 2024

1 Introduction

The nature of probability in statistical mechanics is puzzling. On the one hand it's often been assumed that, if we were to know the exact positions and momenta of the particles or entities that underlie the statistical description, the probabilities would fall out of the picture and the system would evolve deterministically – this suggests that the probabilities are to be thought of as essentially epistemic. And yet the huge variety of different systems – from steam engines to stars – whose evolution is accurately described by statistical mechanics suggests an objectivity and agent-independence that seems to belie the epistemic descriptor. Wayne Myrvold's insightful and tremendously scholarly intervention sets out an intermediate position that might explain both features of the success of statistical physics. The core idea is that while at root these probabilities are credences, he demonstrates that the dynamics have the effect that "the physics swamps the prior" (Engel (1992), as cited in Myrvold (2021, p. 107)).[1]

Beyond Chance and Credence, however, has far more going for it than just providing a new account of the nature of probability in statistical physics. It is based on many years of serious historical engagement with the development of probability theory, and the book is filled with citations of fascinating philosophical analyses from various figures. These are woven together to address questions of the nature and significance of probability in physics and metaphysics, and serve appropriately to contextualise the caricatures of history often swiftly presented in textbooks.

On which note, in addition to these contributions Myrvold's book serves as an excellent introduction to the foundations of statistical mechanics at the ad-

*katie.robertson@stir.ac.uk
†alexander.r.franklin@kcl.ac.uk
[1] All page numbers without further citation refer to Myrvold (2021).

vanced undergraduate/graduate student level. What's especially salutary about Myrvold's approach is that, unlike many of the philosophical tracts on this subject over the last few decades, he does not restrict focus to what's known as 'Boltzmannian statistical mechanics'. Rather, he recognises and respects the orthodoxy in the physics community that statistical mechanics developed with contributions from both Boltzmann and Gibbs, and that many of the objections levelled at the so-called Gibbsian formulation also apply to neo-Boltzmannian approaches.[2]

The book is organised as follows: chapter 1 concerns 'the puzzle of predictability', and why statistical considerations are brought into physics even if the world is ultimately deterministic. Chapter 2 charts the familiar distinction between chance and credence as two senses of the term 'probability', and chapter 3 offers an illuminating discussion of why indifference is insufficient for understanding probability, and the pitfalls of frequentism. Chapter 4 sets the stage for chapter 5's introduction of Myrvold's flagship concept of 'epistemic chances' by discussing the role dynamics play in determining our choice of measure – a choice which the previous chapters establish is needed. These ideas will be discussed extensively in what follows, along with the application of 'epistemic chances' to statistical mechanical probabilities in chapter 8. While this review won't cover the chapters on thermodynamics (chapter 6) and statistical mechanics (chapter 7) there is much to learn from these, as they offer a refreshing and rigorous introduction to thermal physics. Finally, the nature of probability is widely agreed to be radically different in the quantum domain, and Myrvold's chapter 9 considers how epistemic chances fit into different interpretations of quantum mechanics. Chapter 10 concludes.

In the remainder of this review we will focus on an exposition of Myrvold's primary philosophical contribution after which the book is titled: a new account of 'hybrid probabilities' that go beyond chance and credence. Much of the following will be expository, though we hope that covering this material will highlight to readers some of the many virtues of the work and will encourage their own exploration of the text. We follow this with a more critical analysis of his positive philosophical project, and draw some comparisons to related projects.

2 Exposition

Myrvold's aim is to go beyond objective chance and subjective credence with his halfway house that he terms 'epistemic chance'. Being sensitive about how the word 'objective' gets bandied around is a theme of recent work in philosophy of physics; for example, Guido Bacciagaluppi (2020) discusses how the epistemic/ontic distinction crosscuts the subjective/objective distinction. Jenann Ismael (2009) rejects the dichotomy between something's either being an

[2] Also see Wallace (2020) for a defence of the Gibbsian approach to the foundations of statistical mechanics.

epistemic state *or* a fact about the world. Myrvold's epistemic chances follow in the Ismael tradition: the epistemic chances in gambling games and statistical mechanics (and perhaps elsewhere) come about from our epistemic states *and* facts about the world – specifically dynamical facts.

How does this work? Myrvold builds upon Poincaré's method of arbitrary functions, and beautifully explicates how just a smidge of uncertainty about the state of the physical system, once evolved under the dynamics, leads to something that looks and dances like chance. The core idea is that certain dynamics wash out the differences between various probability distributions over initial conditions.

Imagine for instance that Alice thinks that the croupier is equally likely to impart a range of initial velocities v_1–v_2 to the roulette wheel, but Bob thinks that the croupier is most likely to impart a certain velocity v_3 with a Gaussian distribution around this velocity. The difference between them is washed out in the sense that both Alice and Bob will assign incredibly similar probabilities to later macrovariables, such as the probability of landing on a red slot.

2.1 The Importance of Dynamics

Why bring in dynamics at all? One might think that instead of engaging with the details of dynamics, something like a principle of indifference will do the job. This principle holds that, roughly, if we have no reason to distinguish the likelihood of several possibilities we should assign them equal probability. Myrvold masterfully charters the history of this principle in Chapter 3. The key problem is that telling us "to treat equals equally" leaves out "which alternatives are to be regarded as equal"? (p.53). As Myrvold puts it, if I don't know the colour of a book, should I treat its being 'red or not red' as equally likely or 'red, blue or green' as equally likely? This problem will be familiar to readers versed in the Bertrand paradox, and the history of the 'paradox' that Myrvold discusses is edifying.

Thus, there is a substantive choice over whether we set the probability distribution as uniform with respect to one or other variable (cf. Myrvold, Ch. 3 and Section 4.6, Ismael (2009, p. 96)). This is where the dynamics enter; they pick out a preferred set of variables and a measure over those variables *but* the justification for any particular measure's being classed as 'natural' or 'special', will not have its roots in the principle of indifference.

Myrvold illustrates the importance of dynamics using 'the parabola gadget'; this is a toy model where a ball starts on the diagonal and travels vertically until it reaches the parabola, at which point it travels horizontally until it reaches the diagonal, and this is iterated (see figure 1). We leave the gadget running and ask: are we more likely to find the ball in zone A or zone B on figure 2?

Bob answers: zone B, exclaiming "it's twice as big"! But astute Alice notices that Bob's reasoning could only be applicable at one particular time, $t = n$.

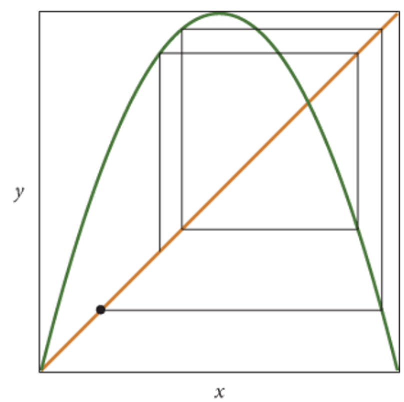

Figure 1: The parabola gadget's position (represented by the black dot) after four iterations From p. 79. Copyright Myrvold 2021, reproduced with permission.

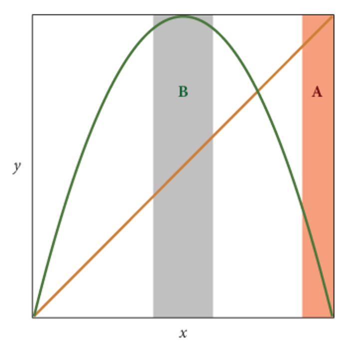

Figure 2: Are Alice and Bob more likely to find the ball in A or B? From p. 80. Copyright Myrvold 2021, reproduced with permission.

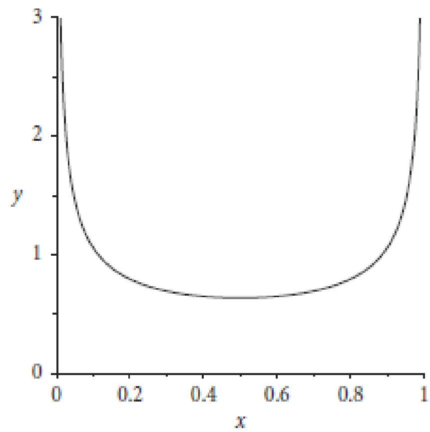

Figure 3: The U-shaped density function for the distribution upon which all not-too-wiggly probability distributions converge after a number of iterations. From p. 82. Copyright Myrvold 2021, reproduced with permission.

Even if we agree with Bob that equal intervals in variable x should be thought of as equally likely, all the points in area B at time n will find themselves in area A at $n+1$. Looking at the dynamics, there's a tendency for the ball to be funnelled towards the edges.

The consequence is that even if we start with the uniform distribution over initial positions of the ball (on the diagonal, i.e. variable x) evolving this distribution under the dynamics leads to a different distribution: the uniform distribution over x is not invariant under the dynamics. But there is a distribution μ (U-shaped – highly concentrated on the right and left hand sides of the gadget with low probability in the centre, figure 3) that *is* invariant, and Myrvold shows (p. 83) how even after only 5 iterations of the dynamics, a variety of different initial distributions (like the uniform and Gaussian ones discussed above) come to approximate this invariant distribution closely.

Myrvold proves that provided Alice and Bob don't believe the process that initially places the ball targets or avoids very tiny areas – so that their initial distribution rapidly varies or wiggles – then a bound can be set on how quickly their initial probability distribution will approximately reach (or perhaps better: approach) the invariant μ, called the 'attractor' distribution. This condition on reasonable initial credence functions we can call: 'not too wiggly'.

Alice and Bob don't need to be very ignorant about the mechanism by which the ball is placed for this argument to get off the ground. How little uncertainty is required? Just a nanometer shake in the hand placing the ball leads to a very small uncertainty in the position of the ball, and after only 32 iterations, the ball is already more likely to be in A than B.

After some time has passed, reasonable (i.e. not too wiggly) initial probability distributions will quickly approach the invariant distribution μ – so we can use this dynamically privileged distribution to calculate macroscopic variables, i.e. answers to questions such as 'how likely is the ball to be in a certain area?' In this way, μ can act as a surrogate for the actual evolved credence function.

Clearly, surrogacy – replacing the actual fine-grained distribution with the attractor distribution – is only a good strategy once some time has passed; it is a good measure to use *later* but there is, as Myrvold emphasises, no justification to use it as a distribution over initial conditions. Indeed, it might not be wise to: imagine that you have access to a high precision measuring device at the initial time and you wish to know the location of the ball after a single iteration.

To sum up: because of the nature of the dynamics, most initial distributions – credence functions – converge towards an attractor distribution provided that they are not too wiggly. Given even a very small amount of uncertainty (that is, even a sharply peaked initial distribution), we still end up with very similar probability distributions later. Thus, this attractor distribution can be called objective in two senses. Firstly, there can be later intersubjective agreement about the predictions of macrovariables. Secondly and relatedly, these predictions (after the washing out has happened) will hardly vary given different initial distributions; there is what is referred to as invariance between perspectives which entails objectivity by some lights (see e.g. John (2021), Nozick (2001), Reutlinger (2013), and Robertson and Prunkl (forthcoming), though see de Canson (2022) for a dissenting view).

2.2 Time-asymmetry

The dynamics are crucial to Myrvold's picture, but dynamics in physics are also often time-symmetric: how does this reasoning work towards the past? Answering this question is central to carrying the lessons learnt about the parabola gadget across to statistical mechanics.

In the parabola gadget case, we first need some time-reversal invariant dynamics; the parabola dynamics are not invertible[3] without an additional pointer variable z whose position – either near the top or the bottom of its scale – at $n+1$ is a consequence of whether the ball was on the right or the left hand side at n. If at time n the ball is on the left, the value of z at $n+1$ is $z/2$; but if the ball is on the right, the value of z at $n+1$ is $1-z/2$ (i.e. $z/2$ from the top rather than $z/2$ from the bottom). If you are certain that the ball starts on the left, then after two time steps, z will either be in the interval [0, 1/4] or [3/4, 1], and after three time steps, z will be in [0, 1/8] or [3/8, 5/8] or [7/8, 1]. Over time, the distribution for z gets fragmented into these disaggregated sets (see p. 87 for more details).

What is the attractor distribution for an initial distribution over (x, z)? Myrvold

[3] Because each location in x could have been reached either from the left or right hand side.

proves that the invariant attractor $\rho(x, z)$ is the μ (U-shaped) distribution for x, and the uniform distribution for z. Any reasonable – i.e. not too wiggly – credence function will converge towards this attractor, (cf. p. 89). But what about when we look to the past? The dynamics are invertible, so an attractor is an attractor to the past *and* the future. In this way, the convergence results do not care about temporal direction.

This seems strange: I could forwards-evolve an initial distribution that is very different from the attractor ρ, call it β; after some time τ, it will approximate ρ. But evolving ρ for the same time period backwards will also take me to ρ, not β. How is this resolved?

First, remember the conditions on the convergence results: from the previous section we know that "as long as the density function for the value of x at time t is not too wiggly, the probability that the state of the system being in a set A at later time $t+n$ is approximately equal to $\rho(A)$ for large n" (p. 89). And now for the variable z: "as long as the density function for the value of z at time t is not too wiggly, the probability that the state being in set A at earlier time $t-n$ is approximately equal to $\rho(A)$ for large n", (p. 89).

Now, consider what temporal evolution does to these variables: forwards-evolution makes the z variable very wiggly (but nonetheless makes it approximate the uniform distribution). If we back-evolve the wiggly distribution we get from forwards evolving β, we do *not* get the attractor, but instead the original β.[4]

The crux of the matter is *when* to apply surrogacy. For forwards evolution and predictions, we can replace the evolved distribution with the surrogate ρ; the information about the wiggliness of z won't matter for the future evolution. But for past evolution and retrodictions, replacing it with the surrogate means you've thrown away information about the past encoded in the wiggliness of z, which is going to be important for your credences about the past! So don't throw it away!

Note that the dynamics, and mathematical results about convergence, are time-symmetric. Say you heard from an oracle that the ball would be at a certain position in the future $(t+n)$ – this would be encoded in a wiggly distribution for the x variable at t. Thus it *would* matter for forwards evolution, and in this situation you wouldn't be justified in using ρ as a surrogate.

Convergence doesn't care about temporal direction. But our knowledge does: we don't have access to oracles, so we take our knowledge of the current state and dynamics to forward evolve. But we can do better than just the current state for the past – often we have more knowledge, such as our memory that the ball was on the left hand side 10 iterations ago.

[4]General features of the dynamics are as follows: forwards evolution makes distributions for x smooth out to U, and ones for z fragment but in such a way that 'any interval that is not too small will be half covered by this support'. Backwards evolution will make z distributions smoother and x distributions wigglier – but wiggly in a way that approximates U.

When forming credences about the past – making retrodictions using a distribution – it would be sheer madness to throw away information about the past (by using the smoothed out, i.e. coarse-grained surrogate distribution). But when making predictions we don't have any access to information about the future *other than* what we would find from forwards-evolving the current state. Had we access to oracles, then time-symmetry would be restored. In this way, the time-asymmetry of the parabola gadget is parasitic on the asymmetry of our knowledge.

2.3 Turning to Statistical Mechanics

This story carries smoothly across to statistical mechanics. Here, the attractor states – i.e. the states invariant under time evolution – are the equilibrium states familiar in statistical mechanics, such as the microcanonical ρ_{mc} or canonical ensemble ρ_c. These equilibrium states do not vary in time.

The dynamics – whether mixing or some other formal property that Myrvold discusses (p. 178) – will take a large class of initial distributions towards the attractor state. In this way, we can think of equilibration as akin to convergence (in Myrvold's technical sense).

A large class of initial distributions will converge to, and so approximate, the equilibrium state.[5] This equilibrium state can then be used as a surrogate: much like in the case above, the probability assigned to various macrovariables will be the same for either the evolved initial distributions or the equilibrium distribution.

But it is important to note that the microcanonical/canonical distributions are surrogate distributions rather than veridical ones. Why is that? Because they assign a high probability to something that we *know* to be impossible – for example, that the system was in equilibrium in the recent past, rather than having evolved from a non-equilibrium state. But there are also reasons for thinking that a system cannot reach equilibrium exactly *anyway*. If the dynamics are invertible, then the 'total variation distance' (the maximum amount the two differ over the probability of any proposition[6]) between the equilibrium distribution and the evolved from out-of-equilibrium distribution is unchanging in time.[7] That means that there will always be *some* propositions for which they differ. This strict unattainability of equilibrium is also familiar from quantum theory: if there's some initial time-dependence in quantum state, under unitary evolution, the state will not become time-independent.

But provided that there is a finite set of measurements, then it is possible that there will be a time at which the evolved initial distribution and the equilibrium distribution return answers as similar as one would like for measurements

[5] See Myrvold's section 8.2 for the nuance about different types of convergence.
[6] For two probability functions this is zero only if they agree on the probability of everything.
[7] See pp. 176-178 for discussion.

that are not infinitely precise (see Short (2011)). This is when the equilibrium distribution is an excellent surrogate.

Replacing the full details of the initial credence function evolved under the dynamics with the attractor distribution is akin to 'coarse-graining' – some details have been thrown away in the 'smoothing' of the probability distribution. Coarse-graining has been seen as controversial in statistical mechanics; see p. 168 for why it is a problem for both Boltzmannian and Gibbsian approaches, should the two be sharply distinguished. In particular, coarse-graining has been seen as 'anthropocentric' as some justify throwing away the details because *we* cannot track them (cf. K. Ridderbos (2002) and T. M. Ridderbos and Redhead (1998)). But our epistemic limitations don't seem to be invoked, on Myrvold's account, when replacing the evolved distribution with the surrogate (i.e. one type of 'smoothing'/coarse-graining), although the asymmetry of epistemic access is involved, as we saw in the previous section. The motivation for surrogacy is that this distribution will give the same answers for all the macrovariables: it is a *good* approximation as guaranteed by the convergence results. (Or in our preferred terms, we are justified in abstracting away from these details). While macrovariables for Myrvold are central to what creatures like us can measure or are interested in, as one of us (KR) has argued elsewhere exactly what we are interested in, and exactly what our measurement capacities are seems to play very little role in the detailed argument: see Robertson (2020).[8]

This coarse-graining is a key step in what Myrvold terms 'the Markovian recipe'. The Markovian recipe says that coarse-graining just once at the end of the time-evolution gives the same answer as coarse-graining at the beginning, evolving and then coarse-graining: on our reading of Myrvold, this bears remarkable similarities to what Wallace terms the 'Simple Dynamical conjecture' in his framework.[9]

But here we face the analogous problem from the parabola gadget case: should we expect the Markovian recipe to work towards the past? In Myrvold's words, should we expect the recipe to be temporally democratic?

Once again, the answer hinges on when it is a good idea to throw away information – and what information is available to throw away. Replacing the evolved distribution with the surrogate equilibrium distribution means that all the wiggles that encode information about the past are lost. The Markovian recipe – should it hold for certain dynamics – tells us that this is not a problem for predictions. In Myrvold's own words, "[w]henever partial equilibration washes out the past history of a system, for the purposes of prediction, the Markovianity condition will be satisfied for all ρ that represent reasonable credence functions.

[8]Ismael (2023) discusses how our position as agents *does* enter the picture.

[9]There are some differences: Wallace (2023) emphasises the construction of a coarse-grained dynamics where there is a repeated coarse-graining at each small time step. In contrast, Myrvold's Markovian recipe doesn't talk about this repeated coarse-graining, but we can think of Myrvold's case as a single time step in the Wallace picture; Wallace's view of statistical mechanics is discussed further in §3.2.

This will not be the case for backwards-in-time evolution – if there's any possibility that records of the past state of a system exist, then it is *not* the case that the current state of the system is all that matters to credences about its past" (p. 191).

Because of our different epistemic access to the future and past, surrogacy is only justified towards the future when setting credences.

3 Discussion

3.1 Bayesianism and Boltzmann Brains

Myrvold's analysis has very interesting consequences for how we should think about temporal asymmetry. David Albert's (2000) *Time and Chance* has been taken as the starting point in much of the philosophical literature;[10] Albert seeks to account both for the universe's manifest time asymmetry and the probabilities of *all(!)* events starting with a three-part package developed with Barry Loewer: a statistical postulate – uniform probability over all microstates compatible with the macrostate of the universe –, the Newtonian laws, and the Low Entropy Past Hypothesis.

One might think that the uniform probability distribution and the fundamental physical laws would be together sufficient, however this famously leads to disastrous retrodictions: just as it is overwhelmingly likely that the system will head towards equilibrium in the future (since there are so many more states corresponding to equilibrium), so it looks overwhelmingly likely that our current state is a mere fluctuation from equilibrium.

Hence, for Albert a crucial role of the Past Hypothesis is that this rules out the disastrous retrodictions and, consequently, radically minimises the chance that any individual who is considering the past is a so-called 'Boltzmann Brain' – that all of their memories and perceptions are false and they are simply a fluctuation from equilibrium with instantaneous consciousness and a vast array of misleading impressions to boot. The Past Hypothesis (almost) rules out these sceptical possibilities by insisting that the universe in fact started out approximately 13.7 billion years ago in a 'low entropy' macrostate and has been evolving according to the laws of physics since then.

Myrvold persuasively argues that the Past Hypothesis is only needed because of a faulty statistical postulate. (He is not denying that the claims of the Past Hypothesis are true, but the epistemological status needn't have the transcendental status that Albert and Loewer ascribe to it). Myrvold argues that we needn't rely on the Past Hypothesis to shore up our ordinary inferences and undermine the sceptic. His argument goes as follows:

- Call the slightly faded photo on your desk the evidence E.

[10] Some other starting points: Price (1997) and Reichenbach (1956).

- Hypothesis H_1: some time back it was a nice, shiny photo; H_2: it was a pile of dust.
- Whilst both are possible, E is *likely* given H_1 but wildly unlikely given H_2, since such a fluctuation is incredibly rare. Given these features of the likelihoods, the Bayesian analysis is:

$$\frac{Cr(H_1|E)}{Cr(H_2|E)} = \frac{Cr(E|H_1)}{Cr(E|H_2)} \times \frac{Cr(H_1)}{Cr(H_2)} \tag{1}$$

Thus, Myrvold emphasises even if your priors were considerably higher in the sceptical hypothesis, they would be swamped by the likelihoods.

But had we followed Myrvold's path instead of Albert's we wouldn't have needed the Past Hypothesis in the first place; positing a uniform probability distribution is something that should be done as a surrogate only when the system has been evolving under the right kind of dynamics, it should not be applied over the *initial conditions*. The sceptical scenario arises as a consequence of applying this distribution at the *wrong* time.

Advocates of the Albert and Loewer picture might respond that they want to be talking about chances rather than credences because credences are often infected with a too subjectivist gloss. We think a fair response is that Myrvold has shown how such credences may be suitably 'objectified'. Moreover, when evaluating sceptical hypotheses it's surely the agent-centric credences that are the subject of discussion.

As Myrvold colourfully puts it, the Past Hypothesis is a way to claw out of a hole we didn't need to leap into in the first place.

3.2 Asymmetries Built On Asymmetries

This is impressive and interesting. But it comes at a price: Myrvold has assumed rather than explained the asymmetry of epistemic access – that we know about the past but not the future.

Imagine running the Bayesian argument from the previous section but towards the future. The evidence E remains the same: you currently have a faded, crumpled photo in front of you. H_3 is the hypothesis that at some time *in the future* the photo is nice and shiny. H_4 is the hypothesis that at some time *in the future* the photo is a pile of dust.

Why shouldn't we conclude symmetrically that the pile of dust is very unlikely towards the future, just as we came to the conclusion that the pile of dust toward the past was unlikely? In other words, why should we think (as we should!) that $Cr(H_4/E) \gg Cr(H_3/E)$ while on the other hand $Cr(H_1/E) \gg Cr(H_2/E)$? The photo being shinier *later* strikes us as overwhelmingly unlikely, and that is because it would need to be on an entropy-decreasing trajectory – something that our past experience and records tell us is overwhelmingly unlikely.

But there's more to it than that: Myrvold shows through the analogy of the parabola gadget that equilibrium states are attractor states in the sense that after some time our probability distributions are highly peaked around such states (assuming they don't start out too wiggly!). So one should expect to observe dust (H_4) as the system evolves under its dynamics. That's why $Cr(H_4/E) \gg Cr(H_3/E)$. On the other hand $Cr(H_1/E) \gg Cr(H_2/E)$ because we know about the past, and we have records (at many length scales) that the photo was shinier in the past. Given the lack of oracles/records about the future we must rely on evolving our probability distribution in line with the known dynamics. It's thus that record and epistemic asymmetries lead to asymmetric expectations about the systems around us.

Where does the record asymmetry come from? Myrvold claims that this "can be traced to causal asymmetry; there are processes that reliably produce records of past events; there is nothing comparable when it comes to future events" (p. 113).

But some might want more: an explanation from physics about how records exist, why we have an asymmetry of epistemic access, and the source of the causal asymmetry. Albert and Loewer's ambitious plan was to explain all these asymmetries via the entropy gradient, such that the knowledge and causal asymmetries are thus explanatorily downstream of the entropic gradient. Much can be (and has been – see Loewer, Weslake, and e. Winsberg E. (2023) and references therein) said about whether or not that project succeeds and there are good reasons to think that, at least, the project is incomplete (see e.g. Fernandes (2017) and Ismael (n.d.)). However, their goals are yet grander: they purport to be able to account for the entropic asymmetry in terms of a constraint on the initial conditions. If the full reductive project is successful then 'past' will just be our label for whichever end of the universe has low entropy.

Huw Price (1997) criticises approaches where explaining temporal asymmetries nonetheless involves making reference to a future/past distinction. Pricean projects are those that locate and conceptually reduce the arrow of time to seemingly non-temporal features. The common analogy, going back to Boltzmann (1897/2003, p. 413) is to the reduction of our folk understanding of 'up' and 'down' to the directions with respect to the Earth's gravitational field; likewise our understanding of 'past' and 'future' stems from features of physics, such as towards and away from the low entropy initial condition for Albert (2000). An alternative to Albert's project that also satisfies this Pricean standard is developed in a number of places by David Wallace (2014, 2023); on this view the initial uniform probability distribution and low entropy past hypothesis are replaced by the constraint that the initial quantum state be appropriately 'simple', and the fundamental laws are those of quantum theory. As noted above, one advantage of Wallace's approach that he shares with Myrvold is that this is developed in terms that respect the practice of physics; that is, the use of both Boltzmannian and Gibbsian concepts in statistical physics.

For those engaged in the Pricean project, the challenge is to answer questions like ‘why shouldn’t the Markovian recipe work in the other direction?’ without saying the word ‘past’!

Myrvold’s account of the difference between the temporal directions in terms of features of our knowledge and the causal asymmetry is not aimed at satisfying Price’s demand. This certainly does not impugn Myrvold’s own project, but it’s worth emphasising that, as such, Myrvold and Albert have different goals and Myrvold’s critiques of Albert should be understood in that light.

3.3 Expanding on Explanation

Myrvold (p. 196) endorses the view that in statistical mechanics, the pertinent notion of ‘explanation’ is that of rational expectability: on this view, an event’s occurrence has been explained, if it has been shown that one should have rationally expected it.[11]

Using this account of explanation, Myrvold’s project is very successful. Rather than explaining – to take Albert’s example – why milk mixes with coffee, we rather explain why it is rational to expect this (cf. E. Winsberg (2008) for more on this). Myrvold’s account does so beautifully; as demonstrated in previous sections: as long as one has a not too wiggly initial distribution, then, after not too many time steps, under certain conditions one’s distribution will be very well approximated by a local equilibrium distribution.[12]

But this still leaves unanswered questions. Why does our world have these dynamics, why is it generic that things mix rather than unmix? Myrvold points out (p. 197) that were there devices that stirred milk and coffee without mixing them, then our rational expectations would be different. But this doesn’t offer an explanation of, for example, why we live in an entropy-increasing world, rather than a decreasing one and why we have knowledge of the past but not the future. Myrvold tells us why we should *expect* entropy to increase, but not why it does in fact increase.

We started this review with a brief précis of a familiar puzzle: how come the probabilities of statistical mechanics seem both to be epistemic, and yet describe the evolution of physical systems independently of any agents’ interactions?

As noted, Myrvold addresses this puzzle by formulating the notion of epistemic chances, and we agree that his appeal to the parabola gadget/method of arbitrary functions helps explain intersubjective agreement on the probabilities of the theory. However, we are concerned that Myrvold parses his explanatory achievements in rather subjectivist terms. He purports to show why we should have certain expectations – that is, why our credences are such that our personal probabilities match up with those of statistical physics. One might hanker for

[11] There are other places where he suggests that we should really be in the business of providing *causal* explanations, but tempers this with a justifiably agnostic view about whether concepts of causation are operative in this domain.

[12] Conditions will include, e.g., timescales for coffee and milk to mix but not for milky coffee to evaporate; see Ma (1985).

more objective forms of explanation that show why the world is as it is rather than just why rational agents would expect to see what we see, but one book cannot resolve all such worries!

3.4 Objectified Credences and Epistemic Chances

One final issue concerns the metaphysics of probability. As developed above we think that Myrvold's book provides an important contribution to the literature, especially in his demonstration that a vast range of initial probability distributions will converge on the same coarse-grained probability distributions through physical mechanisms, such as his parabola gadget, or equilibration. It would be quite reasonable to regard these as objectifying processes for they generate inter-subjective agreement on the probabilities for particular outcomes. But as his approach to asymmetry relies on the fact that for him the input probability distributions are credences, we think it would be preferable call these 'objectified credences'.

Why have a seemingly terminological squabble between 'objectified credences' and 'epistemic chances'? The latter name emphasises that the *source* of the probabilities is our credences, but nonetheless the resulting distribution can be seen to be objective.

Talk of 'chance' might lead to the anticipation of a more objective or worldly account of explanation than the one operative in Myrvold's work, as discussed in the previous section. Use of 'chance' thus seems to fit poorly with the (substantial) achievements of this project, as chances are often associated with Lewis's Best Systems Account or propensity views, both of which explicitly consider worldly explanations.

In contrast, foregrounding 'credence' emphasises the dialectical role that *credence* as the *source* of probability plays in his project. For example, others have sought to use similar formal tools but also to eliminate the epistemic component; see e.g. Strevens (2011). Myrvold is not optimistic for the prospects of such projects, comparing them to getting "soup from a stone" (p. 117). Yet if such projects were successful, notice that Myrvold's response to the Albertian past hypothesis would have to be recast, as it hinges on the source of statistical mechanical probabilities being credences – that's how the asymmetry of epistemic access is relevant – no matter how thin/unimportant the details of these credences are.

The point is that much of the positive upshot of Myrvold's approach and critique of Albert relies on probabilities viewed as credences, but it does not rely on our *particular* ignorance: as long as we are not a Laplacean demon, however small our initial error, the long term dynamics will objectify our credences.

4 Conclusions

For those who think credence is the ultimate source of probability in the world, we don't think you can do better than Myrvold's account.

We conclude with some observations that should serve as important morals of the book:

First, Myrvold emphasises a theme that the devil is in the details (a philosophy of physics tradition going back at least to Batterman (2002)). He claims that the higher-level statistical regularities are not *in spite of* the sensitivity to the lower-level details but in part *depend on* the details of that sensitivity. More complex stories are available than just averaging, and averaging often misleads! We think that it is not so much that *only* the sensitivity is important, but our slight lack of knowledge coupled with the dynamics gives us these macroscopic probabilities. Even the slightest epistemic disadvantage from Laplace's demon (as long as that's not fine-tuned, or wiggly) allows for remarkable macroscopic predictive power given the dynamics of our world: our incomplete knowledge turns out to be a feature not a bug.

Second, proving that systems exactly reach equilibrium is impossible, but that they are well approximated by the equilibrium distribution is essential and achievable.

Third, the principle of indifference is both irrelevant to understanding statistical mechanics and, in general, severely misleading even while in many presentations of the subject it's relied on either implicitly or explicitly.

Myrvold says that this is the book that he wishes had been available as a graduate student and junior researcher; we wholeheartedly endorse this observation and are sure that current and future generations will be glad of this book.

References

Albert, David Z. (2000). *Time and Chance*. Harvard University Press.

Bacciagaluppi, Guido (2020). "Unscrambling Subjective and Epistemic Probabilities". In: *Quantum, Probability, Logic: The Work and Influence of Itamar Pitowsky, eds. Meir Hemmo and Orly Shenker*. Springer.

Batterman, Robert W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press.

Boltzmann, Ludwig (1897/2003). "On Zermelo's Paper 'On the Mechanical Explanation of Irreversible Processes'". In: *The Kinetic Theory of Gases: An Anthology of Classic Papers with Historical Commentary, ed. S. Brush*. World Scientific, pp. 412–419.

de Canson, Chloé (2022). "Objectivity and the Method of Arbitrary Functions". In: *British Journal for the Philosophy of Science* 73.3, pp. 663–684. DOI: 10.1093/bjps/axaa001.

Engel, E (1992). “A road to randomness in physical systems”. In: URL: https://api.semanticscholar.org/CorpusID:122478987.

Fernandes, Alison Sutton (2017). “A deliberative approach to causation”. In: *Philosophy and Phenomenological Research* 95.3, pp. 686–708.

Ismael, Jenann (2009). “Probability in Deterministic Physics”. In: *Journal of Philosophy* 106.2, pp. 89–108. DOI: 10.5840/jphil2009106214.

— (2023). “Reflections on the asymmetry of causation”. In: *Interface Focus* 13.3, p. 20220081.

— (n.d.). *Do statistical mechanical probabilities need quantum shuffling? Albert’s Boltzmannian Story and GRW collapses*. URL: www.jenanni.com/papers/6Do_Statistical_Mechanical_Probabilities_Need_Quantum_Shuffling.pdf.

John, Stephen (2021). *Objectivity in Science*. Cambridge University Press.

Loewer, Barry, Brad Weslake, and eds. Winsberg Eric (2023). *The Probability Map of the Universe: Essays on David Albert’s Time and Chance*. Harvard University Press.

Ma, K. (1985). *Statistical Mechanics*. New York: World Scientific.

Myrvold, Wayne C (2021). *Beyond chance and credence: A theory of hybrid probabilities*. Oxford University Press, USA.

Nozick, Robert (2001). *Invariances: The structure of the objective world*. Harvard University Press.

Price, Huw (1997). *Time’s Arrow and Archimedes’ Point: New Directions for the Physics of Time*. Oxford University Press.

Reichenbach, Hans (1956). *The Direction of Time*. University of California Press.

Reutlinger, Alexander (2013). “Can interventionists be Neo-Russellians? Interventionism, the open systems argument, and the arrow of entropy”. In: *International Studies in the Philosophy of Science* 27.3, pp. 273–293.

Ridderbos, Katinka (2002). “The Coarse-Graining Approach to Statistical Mechanics: How Blissful is Our Ignorance?” In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 33.1, pp. 65–77. DOI: 10.1016/s1355-2198(01)00037-5.

Ridderbos, T. M. and M. L. G. Redhead (1998). “The Spin-Echo Experiments and the Second Law of Thermodynamics”. In: *Foundations of Physics* 28.8, pp. 1237–1270. DOI: 10.1023/a:1018870725369.

Robertson, Katie (2020). “Asymmetry, Abstraction, and Autonomy: Justifying Coarse-Graining in Statistical Mechanics”. In: *British Journal for the Philosophy of Science* 71.2, pp. 547–579. DOI: 10.1093/bjps/axy020.

Robertson, Katie and Carina Prunkl (forthcoming). “Is Thermodynamics Subjective?” In: *Philosophy of Science*, pp. 1–16. DOI: 10.1017/psa.2023.58.

Short, Anthony J (2011). “Equilibration of quantum systems and subsystems”. In: *New Journal of Physics* 13.5, p. 053009. DOI: 10.1088/1367-2630/13/5/053009. URL: https://dx.doi.org/10.1088/1367-2630/13/5/053009.

Strevens, Michael (2011). “Probability Out of Determinism”. In: *Probabilities in Physics*. Ed. by Claus Beisbart and Stephan Hartmann. Oxford University Press, pp. 339–364.

Wallace, David (2014). "Probability in Physics: Statistical, Stochastic, Quantum". In: *Chance and Temporal Asymmetry*. Ed. by Alastair Wilson. Oxford University Press, pp. 194–220.

— (2020). "The necessity of Gibbsian statistical mechanics". In: *Statistical mechanics and scientific explanation: Determinism, indeterminism and laws of nature*. Ed. by Valia Allori. World Scientific, pp. 583–616.

— (2023). "The Logic of the Past Hypothesis". In: *The Probability Map of the Universe: Essays on David Albert's Time and Chance*. Ed. by Barry Loewer, Brad Weslake, and Eric Winsberg. Harvard University Press. URL: http://philsci-archive.pitt.edu/8894/.

Winsberg, Eric (2008). "Laws and Chances in Statistical Mechanics". In: *Studies in History and Philosophy of Modern Physics* 39.4, pp. 872–888. DOI: 10.1016/j.shpsb.2008.05.005.