The local quantum vacuum as the Past Hypothesis

David Wallace

October 2, 2023

Abstract

The 'Past Hypothesis', as advocated by David Albert and Barry Loewer, is the hypothesis that the world came into being in whatever particular low-entropy highly-condensed big-bang sort of macrocondition it is that the normal inferential procedures of cosmology will eventually present to us. I consider some hypotheseses about that that macrocondition is likely to be given what cosmology has already presented to us, and explore the consequences of these hypotheses for the broader ('Mentaculus') project of grounding physics and the special sciences in the Past Hypothesis. My main conclusion is that current cosmology suggests a unique, pure quantum state (the local quantum vacuum, or 'Bunch-Davies vacuum') for the initial state of the Universe, in which case statistical-mechanical probabilities emerge from quantum probabilities without any need for an intervening statistical postulate.

1 Introduction

The laws of microscopic physics make no distinction between past and future[1], but manifestly there is such a distinction — in the emergent physics of complex systems, in the observed world, in the very nature of causation and inference and epistemology. We could try to reconcile these facts by introducing new physics, or by supposing that aspects of causation and the flow of time introduce fundamental time asymmetry even without it showing up in the microphysical equations, or by rejecting the very idea that large-scale phenomena supervene on microphysics. But if we want to ground higher-level concepts in physics, and to leave the apparent laws of the microworld unchanged, on pain of contradiction we must break the symmetry by introducing some kind of time-asymmetric boundary conditions.

The most popular strategy for doing so, at least in the recent literature, is to introduce a so-called 'Past Hypothesis' (PH), which constrains the state

[1] It is often pointed out (see, e. g. , (Price 1996, p.18), (Maudlin 2007, pp.117-121), (Roberts 2022, ch.7).), that time reversal symmetry is broken in modern particle physics. But (as is generally accepted) that symmetry is not broken in a way that distinguishes past and future in a way that could ground the sorts of macroscopic asymmetries that are relevant here.

of the very early Universe in such a way as to guarantee (or at least make very plausible) that large-scale dynamics futureward of that constraint show the observationally-required asymmetries. The idea is far from new (it goes back to Boltzmann; see Sklar (1993, ch.8) and references therein for details) but its prominence in the recent philosophical literature is largely due to David Albert's defense of the idea in his influential book *Time and Chance*.[2]

Albert's version of PH (Albert 2000, p.96)) has two components: a hypothesis about the past *macrostate* (call this PHM), which is that "the world came into being in whatever particular low-entropy highly-condensed big-bang sort of macrocondition it is that the normal inferential procedures of cosmology will eventually present to us", and a hypothesis about the past *microstate* (call this PHμ), which is that the initial microstate is selected uniformly with respect to Liouville measure from those compatible with the macrostate given by PHM. (In fact Albert uses 'Past Hypothesis' to refer only to PHM, and states $PH\mu$ in terms of the probability measure over *present-day* microstates, but his presentation is equivalent to the one I give here and it is helpful to understand both PHM and PHμ as aspects of an overall hypothesis about the past. See (Winsberg 2004) and (Wallace 2023) for more on this point.)

A probability distribution over *initial* microstates of the Universe, together with precise (deterministic or stochastic) dynamics for those microstates, entails a probability distribution over all future microstates of the Universe, and indeed over all histories of the Universe. So a PH in Albert's sense does more than simply provide a grounding for the asymmetries of microphysics:[3] it provides, in principle, a complete reductive base for all of science. The point is not always apparent in *Time and Chance* but becomes explicit in Albert's subsequent work,much of it joint with Barry Loewer; cf (Loewer 2007; Loewer 2023; Albert 2015), in which they call the combination of PH and dynamical laws the 'Mentaculus', following a movie character's use of that term to refer to 'the probability map of the Universe'.

The literature engaging with these ideas is large and philosophically rich (see, for instance, the various papers in (Loewer, Weslake, and Winsberg 2023)), but one question has received surprisingly little attention[4]: if it is a foundation stone of all of science that "the world came into being in whatever particular low-entropy highly-condensed big-bang sort of macrocondition it is that the normal inferential procedures of cosmology will eventually present to us", which particular state is that? Presumably we cannot hope for a final answer yet, as cosmology is ongoing, but presumably its 'normal inferential procedures' have made some progress, and so might provide at least a preliminary insight into just what, specifically, PH actually says, and whether it does indeed have the

[2]So far as I can determine, the *term* 'Past Hypothesis' is due to Albert: he cites no source for it, and it does not appear in influential discussions of a cosmological boundary condition by Penrose (1989, ch.7), Lebowitz (1993), Sklar (1993, ch.8), or Price (1996, ch.2).

[3]As Wallace (2023) argues, and Albert (2023) is happy to grant, such a grounding apparently requires only some form of PHμ.

[4]An important, if partial, exception is Callender (2009), who considers the interplay between the Past Hypothesis and Newtonian gravity, though for the most part outside the cosmological context.

right form to serve the purposes for which it is being used.

My aim in this paper is to attempt to answer that question, at least to the degree that modern cosmology allows. I proceed in steps: after reviewing the basic logic of Albert's approach (section 2) I consider successively more realistic cosmological models: a toy Newtonian cosmology (sections 3-5); modern classical cosmology (sections 6-7); inflationary quantum cosmology (sections 9-10). By the end of the story, we will see (section 11) that while the core ideas of Albert and Loewer's Mentaculus remain unscathed, the details are altered in important ways.

While I focus on Albert's (and Loewer's) approach to the Past Hypothesis, this is largely for expository clarity. The main observations of the paper should continue to hold in most approaches to philosophy of statistical mechanics I know that use one or other form of the Past Hypothesis.

2 The logic of Albert's Past Hypothesis

To get clear on just how Albert's framework operates, let me review its original statement (Albert 2000, p.96), which forms the centerpiece of *Time and Chance* (there he refers to it as the 'Newtonian statistical-mechanical contraption for making inferences about the world', but it essentially coincides with what he later calls the Mentaculus). Albert describes this as consisting 'in its entirety, of three laws and one contingent empirical fact'. The contingent empirical fact is an agent's current (usually very coarse-grained) information about the present state of the world. Two of the laws are, essentially, PHM (recall: the specification of a 'particular low-entropy highly-condensed big-bang sort of macrocondition') and PHμ (the specification of the microcondition via a Liouville-uniform measure over microstates compatible with that macrocondition), and the third is the Newtonian law of motion (described by Albert as literally $F = ma$, but in context I think also including the details of the particular Newtonian force laws).

I shall make several observations about this framework. Firstly, and of greatest importance to this paper, PHM is not simply the hypothesis that the universe begins in some low-entropy state or other, though it is not infrequently[5] so characterized in the literature. It is better thought of as a placeholder: there is some macrocondition that will eventually be given to us by the workings out of cosmology, which (by the way) we expect will be a low-entropy, highly-condensed big-bang sort of macrocondition, and that macrocondition is the one we should use in statistical mechanics. If cosmology tells us that the initial macrocondition is P, then the Newtonian statistical-mechanical contraption consists of the present-day empirical data, the laws of Newtonian physics, the Statistical Postulate, and the postulate that P is indeed the initial macrocondition. (Hence the current paper's question: what is P, and does it work the way it is supposed to?) As such, PH itself requires no commitment to *entropy* even

[5]Cohen and Callender (2009, p.9) do so in their discussion of laws of nature; so does Earman (2006) in at least some parts of his wideranging critique of PH.

being well-defined for the early Universe, and indeed entropy plays a heuristic role at most in Albert's reconstruction of non-equilibrium statistical mechanics. So criticisms of PH as being inherently committed to high-level concepts like entropy miss their mark.

But (my second observation) there is still something emergent and high-level about PHM. For it requires that modern cosmology will give us a *macro*condition, which for Albert is a cell in some partition of phase space into many macroconditions. The idea of a macrocondition (or, equivalently, a macrostate) is that they provide a basis to write down autonomous high-level dynamics in which only a system's macrocondition enters into its dynamical equations — something which in general is possible in physics only given additional statistical-mechanical assumptions. But we can see that there is something (perhaps not viciously) circular about PHM if 'macrocondition' is so understood: PHM makes essential reference to higher-level dynamical ideas which are themselves legitimated only by a framework of which PHM is a part.

This is tied (my third observation) to the division of labor in Albert's system between PHM and PHμ. The two together provide a probability distribution over initial *microstates* of the Universe: it is the distribution uniform (PHμ) over the cosmology-provided macrocondition (PHM). PHM provides macroscopic information about the beginning of time; PHμ fills in the microscopic details; but it is only the joint product of the two that is plugged into the laws of physics and the present-day macrocondition in order to furnish information about past and future.

There is a division of labor in another way (my fourth observation): only PHμ introduces *probabilities* to statistical mechanics. The Newtonian law of motion is deterministic, and PHM and the 'contingent empirical fact' are categorical (non-probabilistic). Where physics is indeterministic at the macro level, in Albert's framework it can arise only from explicit introduction of probabilities in statistical mechanics. That in turn raises the question of how these probabilities are to be understood: in *Time and Chance* Albert is largely noncommittal but in subsequent work (Albert 2015) he explicitly endorses the view that they are Humean chancy laws in the fashion of (Lewis 1980), encoding that description of the initial state that best balances simplicity and strength. Others (e. g. ,Winsberg (2008), Demerast (2016), Myrvold (2019)) have challenged the desirability and/or coherence of this view of classical statistical-mechanical probabilities. Still others (e. g. ,Volchan (2007), Goldstein (2012), Lazarovici and Reichert (2015)) argue that this quantitative notion of probability can be weakened to a qualitative notion of 'typicality'.

But there is something a little odd about this quest for an interpretation of classical probabilities, for our world is quantum, not classical, and there are good reasons (Wallace 2016; see also Albert's own discussion in chapter 7 of *Time and Chance*) to doubt that a classical interpretation of the Statistical Postulate will transfer to quantum mechanics unscathed. And this brings me to my final observation: there is a tension between PHM, as Albert describes it, and the whole idea that we seek a *Newtonian* statistical-mechanical contraption. For if PHM is supposed to be that macrodescription of the world that 'the normal inferential procedures of cosmology will eventually present to us', we can right now be pretty confident that these methods will not present us with anything straightforwardly described by non-relativistic, non-quantum particle mechanics. This is not simply the observation that Newtonian mechanics is *false* (which Albert understands perfectly well); it is that at face value it is incompatible with PH as Albert states it.

I can see two ways of understanding Albert's framework in the light of this tension. The first is to insist that we are after all considering *classical* statistical mechanics, and replace real cosmology in PH with some fictitious analog or perhaps with pure stipulation. Statistical mechanics so understood will be a toy model, a training exercise to understand our world, or perhaps a simplified model of some subsystem of our world. The second — which I think is closer to Albert's own interpretation of his project (cf his remarks at (Albert 2023, pp.336-7)) — is to take 'the Newtonian law of motion' as itself just a placeholder for the true underlying laws of physics, and to argue or simply speculate that we have reason to think the conceptual ideas of his Newtonian statistical-mechanical contraption carry over to real physics *mutatis mutandis*.

The second half of this paper explores the second possibility; first, though, let us see what kind of toy model of cosmology we can build inside Newtonian physics.

3 Newtonian cosmology in a box: setup

Consider a fictitious cosmos in which space is flat and all of matter consists of point particles (all, for simplicity, with the same mass m), interacting under two-body forces between each pair of particles and defined by a two-body potential like

$$V(r) = -G\frac{m^2}{r} + \lambda(\mathrm{e}^{r_0/r} - 1) \tag{1}$$

where λ is positive but the sign of G is unspecified, and where r_0 is some fundamental lengthscale. This describes an inverse-square force between particles at distances large compared to r_0, going over to a strongly repulsive force at distances small compared to r_0; it might be thought of as approximating a world of fairly stiff spheres of radius $\sim r_0$. (This short-distance force regulates what would otherwise be singular behavior in the physics).

To model a cosmology, we might naturally consider an infinite space and an infinite number of these particles. But Newtonian mechanics is hostile to this conception of cosmology (Malament 1995; Norton 1999; Wallace 2017) and so for simplicity I consider something more flatly fictitious: there are finitely (but enormously) many of these particles, and they are contained within a box of finite (but enormous) size. (We will not need exact parameters here, but it is essential that the box is large enough that, were the particles to be uniformly distributed, the distance between nearest neighbors would be $\gg r_0$.) This is now a familiar statistical-mechanical problem, that of a dilute gas of classical point particles.

We now have in place the first lawlike component of the Mentaculus: the Newtonian kinematics and dynamics of the system. For the second component, PHM, let us assume a partition of the box into macrostates, following the usual rules for classical dilute gases: that is, we divide the classical phase space of a single particle into boxes of equal size, and individuate the macrostates of the whole system by the number of particles in each single-particle box. Two microstates within the same macrostate will define the same coarse-grained density of particles in phase space, differing only in the details. We will not need it much in what follows, but with this notion of macrostate available we now have access to the Boltzmannian definition of entropy: the entropy of a microstate, up to an arbitrary choice of scale, is the logarithm of the phase-space volume of the unique microstate in which it finds itself.

In our fictitious cosmos, I state PHM thus: the initil macrostate is that unique macrostate where the coarse-grained particle density is constant across space, the coarse-grained particle kinetic energy density is constant across space, and the coarse-grained particle momentum density is zero. At the beginning of time, in this cosmos, everything is uniform; the specific choice of macrostate is parameterised by the average mass density ρ and the average kinetic energy density k. (ρ is of course fixed by the total number N of particles, their individual mass m, and the total box volume V: $\rho = Nm/V$.)

To be sure, the conceptual status of the macrostate decomposition — and the version of PHM I stated through it — is unclear at best. I have said nothing about higher-level dynamics, and so the justification of this specific decomposition into macrostates has little to recommend it beyond historical tradition and intuitive plausibility. And even taking the decomposition as given, PHM cannot plausibly be the product of this universe's cosmologists, since its simple dynamics almost certainly do not support intelligent life at all, let alone intelligence sufficient to conduct cosmology research. These concerns are significant; still, I set them aside for now.

With PHM in place, completing the Newtonian Mentaculus is simple enough (at least technically): we just add PHμ, the assumption that the initial microstate is chosen uniformly, with respect to Liouville measure, from all those within the spatially uniform state selected by PHM. Now let's ask what sort of future it predicts.

4 Newtonian cosmology in a box: dynamics

Suppose for the moment that $G < 0$, so that the inverse-square force between pairs of particles is repulsive. Then our cosmology is boring in the extreme: the system pretty much remains in, or close to, its PHM-decreed macrostate for eternity. For suppose the system actually makes its way to a macrostate where the particle distribution is a bit more clumpy. The particles in the clumps mutually repel each other, and so the clumps dissipate, and the system moves back to uniformity. Deviations from uniformity are self-suppressing and so cannot build up. This means that the PHM macrostate is also the equilibrium macrostate of the system: our toy universe began in thermal equilibrium and will remain there — not literally forever, but for an unfathomably large time.

We can put essentially the same argument in terms of energy and entropy: allowing particles to clump increases their potential energy; that energy needs to come from somewhere, and so the kinetic energy must decrease. But phase-space volume increases very quickly with kinetic energy density, so (given PHμ) the system is highly unlikely to follow a trajectory leading to significant clumping; if it does, it is then highly likely to follow a trajectory that undoes the clumping. (This is just the standard argument for why, in ordinary statistical mechanics of dilute gases, the equilibrium state is uniform.)

If we simply reverse the sign of G, so that our potential simulates gravitational physics with a universal attractive force, the behavior of the model changes dramatically. Since the particles attract each other, the gravitational force induces small clumps to get larger, not to break up, and so once the system starts clumping it will just get more clumpy. Indeed, even microscopic levels of clumping (as will be present in the vast majority of even uniform macrostates) will build up over time and eventually reach macroscopic scale. A uniform, gravitationally-attracting gas is in general unstable.[6]

To be slightly more precise: the *Jeans instability* (see (Binney and Tremaine 2008, ch.5)) states that a clump of size λ in a self-gravitating gas will grow exponentially in scale with time provided that λ is large compared to the *Jeans wavelength* λ_J, given by

$$\lambda_J^2 = \frac{k}{G\rho} \tag{2}$$

Fluctuations over scales smaller than this are damped: the pressure in the gas from kinetic energy will overcome gravitational clumping and the clumps will smooth out before they have a chance to grow. Fluctuations over scales larger than this, however small initially, will grow exponentially as the clumps feed on themselves.

Since the vast majority of initial states will not be *completely* uniform on any scale, we can predict that our model will grow more and more inhomogeneous with time as the initially minute long-range fluctuations in the gas grow. Interestingly from the point of view of Albert's project, the development of inhomogeneties on a macroscopic scale will depend on the details of the initial fluctuations, which in our model are microscopic: that is, our model contains exactly the amplification of microscopic probability up to the macroscopic scale that Albert's project requires.

(Given the flatly unphysical model we are considering, the reader might be forgiven some skepticism that these technical results actually follow from the somewhat handwaving arguments I have given. But in fact systems of self-gravitating point particles have been very widely studied in physics, not because

[6]Technical note: this claim is true for a finite cloud of gas in empty space but is complicated by the box around our finite system. If the gas is sufficiently dense and/or hot, thermal motion can break up the clumps. I am assuming that our model is sufficiently cold and diffuse that this does not occur; to be exact, I am assuming it satisfies the conditions for gravothermal catastrophe (cf (Binney and Tremaine 2008, section 7.3) and references therein).

physicists care about toy models of point particles in themselves but because those models do a fairly good job of describing galactic dynamics, taking the 'points' to be stars (of course, the short-distance repulsive potential is replaced as necessary by a model of the actual physics of stellar collisions). Here one indeed finds — and observationally confirms — very much the physics we are describing here: the Jeans instability means that initial fluctuations in a model get magnified, and this lies at the core of galactic dynamics. A standard reference is (Binney and Tremaine 2008).)

5 Newtonian cosmology in a box: development

The PH in our Newtonian model is vaguely stated in two ways. One is familiar and largely harmless: the 'fundamental' parameters k and ρ are real numbers and most such numbers would not be exactly stateable in any finite way. (In a more realistic physics, we might also say that no observational evidence could ever fix the *exact* value of either.) Any physical law requiring real-valued parameters to state is vague in this way; this vagueness might tell us something interesting about laws *in general*, but it is not specific to Albert's project. The other arises from the use of uniform *macrostates* to state PHM: the division of the one-particle phase space into these macrostates was vaguely stated (indeed, pretty much unspecified!) in my account, and this vagueness thus infects the PH. And since the justification of one macrostate partition over another arises, if at all, from its suitability to ground higher-level autonomous dynamics with respect to that partition, this seems to be a vagueness of a fundamentally novel kind, as argued in detail by Chen (2022).

This vagueness in the specification of PH is not simply a philosophical problem; it has technical ramifications. Recall that in the case of an attractive potential, the macroscopic evolution of a system involves exponential growth of initially-tiny fluctuations from uniformity on lengthscales large compared to the Jeans wavelength. But the exact size of these fluctuations depends on how small the original partition of phase space, since that is what determines exactly how uniform are the individual microstates comprising the the 'uniform macrostate'. Different partitions will lead to quite different timescales for the system to develop inhomogeneities. and fairly clearly this is a matter of substantive physics, not simply of convention.

We could take this as evidence for an urgent need for more careful conceptual work on just what the macrostate partition is and how it is to be understood; but I want to suggest instead that the vagueness here is an artifact of dividing our hypothesis about the initial state into separate macro- and microscopic constraints (as mentioned in section 2). Suppose instead that we replace the combination of PHM and PHμ with

Newtonian initial-state hypothesis (uniform version): each particle is distributed at random with the probability distribution (over 1-particle phase space with position $\mathbf{q}$ and momentum $\mathbf{p}$

$$\Pr(\mathbf{q}, \mathbf{p}) = \mathcal{N} e^{-\mathbf{p}\cdot\mathbf{p}/2\sigma^2} \chi_B(\mathbf{q}) \quad (3)$$

where χ_B, the characteristic function of the box, is 1 when $\mathbf{q}$ is inside the box and 0 when it is outside, σ is a free parameter, and $\mathcal{N}$ is a normalization constant.

This distribution, given a large enough number of particles, will with very high probability have nearly uniform mass density and kinetic energy density (with $mk = 3\sigma^2/2$) and zero momentum density; in other words, with very high probability it entails some version of PHM (exactly which one is vague and depends on how demanding we are as to what constitutes 'very high probability'). And its microscopic probability distribution is just as suited for Albert's purposes as PHμ (indeed, it will very well approximate it) so it serves just as well as (PHM+PHμ) as a basis for statistical mechanics in our fictitious model. Yet it is microphysically stated, with no mention anywhere of 'macrostate' and no vagueness beyond that forced by the free parameter σ and the total particle number (essentially equivalent to the kinetic energy density and mass density in PHM). The alternative then looks more conceptually straightforward than Albert's original proposal (as well as seeming broadly as compatible with his Humean approach to probability and law).As we shall now see, it also appears a better fit to physics practice when we consider more realistic models. To begin doing so, I next provide a very brief summary of classical cosmology.

6 Classical cosmology: outline

There is a class of solutions to the equations of classical general relativity called the *Friedman-Lemaitre-Robertson-Walker* (FLRW) solutions, which describe a homogeneous, isotropic, expanding universe, filled with a uniform fluid. If the 'equation of state' of that fluid (which says how the fluid's temperature and energy density covary with its density as it expands uniformly) is specified, the resultant system can be solved to determine just how quickly this universe expands. Models of this kind, with the equation of state calculated using what we know from high-energy physics, have proved extremely effective at modelling the early Universe. As a concrete example which will be important to us later: the fluid is opaque to light above a temperature of about 3000 K (the temperature above which hydrogen atoms break up) but transparent below, so the light that initially permeated the fluid escaped once it reached that temperature c.380,000 years after the Big Bang and is still visible today with the right instruments. This 'cosmic microwave background radiation' is a literal image of the early Universe: its observation by Penzias and Wilson in 1965 was the decisive point at which the Big Bang theory of cosmology became widely accepted. (For a review of this material, see, e. g. , (Weinberg 2008, section 2.1) and references therein.)

This subject is often called ‘classical cosmology’, but in this context that name needs to be used with caution. The primordial fluid is thoroughly quantum-mechanical: no classical understanding of its physics suffices to calculate its equation of state. The early universe is quantum matter dynamically coupled to classical spacetime, or *semiclassical gravity* as it is normally called. (See (Wallace 2022) and references therein for more discussion of this point). Nevertheless, from about the time at which the universe becomes transparent, through to the point at which stars (which also require quantum mechanics) begin to form, we can to a pretty good approximation treat the expanding Universe as made up of classical particles, interacting only through gravity: that is, we can treat it as a general-relativistic version of our Newtonian model. (And indeed the ‘general-relativistic’ part is mostly needed only to see how Newtonian physics needs to be modified to treat an expanding universe.) The point at which the ‘Past Hypothesis’ would need to be stated is roughly the moment of transparency; all cosmology futurewards of that (again, at least until stars form) should be determined by the combination of classical gravitational physics and the Past Hypothesis.

Of course, our subject matter here — the physics of the Universe only with respect to certain degrees of freedom, only approximately, and only within certain time bounds — is more modest and circumscribed than Albert’s Mentaculus. Still, it is about as wide-ranging and cosmological as we can get within real physics while remaining in the domain of (microphysically-reversible) classical physics, so if we want insight into a physically realistic Past Hypothesis then it seems a good place to look.

7 Classical cosmology: structure formation

Post-transparency cosmology plays out quite similarly to our Newtonian toy cosmology with the addition of the expansion of the Universe.[7] The *average* density of matter constantly decreases as the Universe expands, but local fluctuations in that density reinforce and grow exponentially, eventually forming the fractal patterns of galaxies, clusters of galaxies, and superclusters of clusters that form the large-scale structure of the present-day Universe. Expansion changes the quantitative details of Jeans instability but not its basic nature, and the large-scale physics[8] of this epoch is reasonably well understood.

Prima facie, one might imagine that the small-scale seeds that grow through Jeans instability into large-scale structure are *microscopic*, statistical-mechanical

[7] This is well-known physics and I omit original sources: (Weinberg 2008) is a standard reference.

[8] The microscopic physics, by contrast, is hid in darkness — literally and figuratively. The dominant component to the mass density of the Universe is so-called ‘dark energy’, (aka the ‘cosmological constant’, which for various reasons is mysterious; most of the actual matter is so-called ‘dark matter’, which does not interact with light and so is very underconstrained by observational data. However, these puzzles mostly do not affect the large-scale physics we are concerned with here — indeed, the insensitivity of large-scale cosmology to the microphysics of dark matter is precisely why the latter is so hard to pin down observationally.

fluctuations in the early universe, so that the appropriate Past Hypothesis posits a macroscopically uniform early Universe — just as in our Newtonian model. However, it is clear in classical cosmology that for the Universe to form the structures we observe today, at the point of transparency it must already have developed *macroscopic* fluctuations, much larger than those that would be present randomly in (say) a uniform Boltzmannian macrostate. We know this from calculation — but we also know it directly from observation. Recall that the microwave background radiation is, literally, a picture of the Universe at the onset of transparency. Penzias and Wilson did not possess the tools to look at that picture very precisely — but the tools were developed, and in the 1990s the Cosmic Background Explorer Satellite (COBE) imaged the cosmic background radiation at a resolution sufficient to demonstrate that the early Universe did indeed show small, but clearly present, anistropies (the WMAP and Planck satellite missions in the 2000s and 2010s went on to greatly enhance the precision of these measurements).

What form does this anisotropy take? Alas, it's not a message from the Creator, or a smiley face: as best we can tell, the fluctuations are *random*, and indeed the probability measure from which they are drawn is fairly strongly constrained by the observational data and quite well known (in philosophical terms, we are assuming that the pattern of fluctuations is fairly typical among all those given by the probability measure and that the probability measure itself is isotropic; the correlation length in the fluctuations is then small enough compared to the size of the night sky that these requirements pretty much fixes the probability measure from which it is drawn). But note that this probability measure cannot be identified with PHμ in Albert's version of the Past Hypothesis: these are macroscopic fluctuations, involving variations in the temperature and density of the early Universe over scales that even then would have spanned hundreds of thousands of light years.

So what does the Past Hypothesis for classical cosmology look like, given these cosmological facts? Something like this:

PHM The macrostate of the Universe at the beginning of the transparency era is an *almost* uniform thermal state, selected at random from a certain probability distribution over almost-uniform thermal states (the precise mathematical form is not relevant here).

PHμ : The microstate of the Universe at that time is selected at random from all those compatible with the initial macrostate.

This form of PH actually makes the micro/macro distinction rather more natural than it was in our Newtonian model. But it does so because our theory really has two distinct sources of probability: one from the choice of overall fluctuation pattern over macrostates (which induces a probability distribution over large-scale structure in the Universe at later times) and one from the distribution over microstates compatible with the initial macrostate (which would play a foundational role in later statistical mechanics, at least if we continue to indulge the fiction that the latter can be treated classically).

This is, frankly, a somewhat awkward way to state a fundamental law. And indeed, no-one in cosmology thinks that the probability distribution over macrostates at the moment of transparency should be understood as a Past Hypothesis, rather than as a consequence of physical processes happening earlier in time — not least because a Past Hypothesis concerns the very beginning of the Universe, whereas the transparency epoch begins 380,000 years after the Big Bang and the physics of most of those 380,000 years is pretty well understood. However, it is the closest we can get to a genuine Past Hypothesis within cosmology while we stay in the domain of classical mechanics. To go further, we will have to consider how the Mentaculus program is modified when quantum physics is taken into account.

8 Quantum Boltzmannian Statistical Mechanics

There is a wide and longstanding consensus (see, e. g. , (Gell-Mann and Hartle 1989, Goldstein *et al* 2010) on how to extend at least the Boltzmannian macrostate framework to quantum theory. Namely: instead of macrostates being a set-theoretic partition of classical phase space, they are a decomposition of quantum-mechanical Hilbert space into orthogonal subspaces; instead of Boltzmann entropy being the logarithm of a macrostate's volume, it is the logarithm of a macrostate's dimension.

Given this setup, there is a fairly natural way to take the Mentaculus idea across from classical to quantum mechanics. Quantum dynamics (the Schrödinger equation, at least in the first instance) replaces classical dynamics; PHM becomes the hypothesis that the Universe's quantum state begins in whatever quantum macrostate is given to us by the normal inferential practices of cosmology; PHμ becomes the hypothesis that the initial microstate is selected from a uniform probability measure over microstates compatible with that macrostate. (The appropriate notion of 'uniform probability measure' is normally taken to be the Haar measure defined by the action of the unitary group on projective Hilbert space; the details will not be needed here.) Chen (2023) gives a more detailed statement of this framework.

On closer investigation, though, the resemblance to classical statistical mechanics proves superficial (here I summarize observations made in more depth in (Wallace 2023) and (Wallace 2016)). A classical microstate corresponds to a unique macrostate, but a quantum state may be a superposition of states lying in one, two, or many macrostates, and indeed generic states will not be contained in any single macrostate. Furthermore, even if a state begins in a single macrostate, evolution under the Schrödinger equation will generally evolve it into a macroscopic superposition.

This is of course the quantum measurement problem in statistical-mechanical guise, and how these macroscopic superpositions are to be understood depends on how it is to be solved: dynamical-collapse theories will need to modify the Schrödinger equation so that the state collapses stochastically into a definite macrostate; hidden-variable theories will need the classical reality described by the variables to have macroscopic features corresponding to one macrostate or another; the Everett interpretation accepts the simultaneous reality of all the macrostates and indexes definiteness to observers.[9] But all of these approaches have in common that the unitarily-evolving quantum state defines not a single macrostate but a probability distribution over macrostates. How these quantum probabilities are to be understood — as the result of stochastic dynamics, as arising from classical probabilities over hidden variables, as quantifying indexical uncertainty — is interpretation-dependent, but on any interpretation they need to be understood as additional to the purely statistical-mechanical probabilities that arise from the PHμ probability distribution over quantum states themselves.

This doubling up of probabilities — quantum and statistical — is awkward, and is not readily found in the mainstream quantum statistical mechanics literature (cf (Wallace 2016)), and it is tempting to ask whether the quantum probabilities might replace and not simply supplement the statistical-mechanical ones. Two strategies for doing so can be found in the literature. Firstly, and most flat-footedly, we might observe that the formalism of quantum mechanics admits mixed states as well as pure states, and that the predictions obtained from the mixed state corresponding to a normalized projection onto a macrostate are identical in every respect from those obtained by the uniform distribution over pure states compatible with that macrostate that PHμ. So we might consider replacing PHμ with

PHρ: The initial microstate of the Universe is the unique mixed state given by the normalized projection onto the macrostate specified by PHM.

PHρ is mentioned briefly in (Wallace 2023) and developed in depth by Chen (2020, 2021) under the name 'Wentaculus'. ('PHρ' alludes to the normal use of ρ in the physics literature to denote a mixed quantum state. In some of the foundational and mathematical-physics literature W is used instead — hence Chen's terminology. He would presumably prefer PHW to my PHρ!)

Of course, PHρ eliminates statistical probabilities only if mixed states can themselves be interpreted as non-statistical, and this is contentious, especially given that we are discussing the Universe as a whole. (When a subsystem is in a mixed state, there is always the possibility of interpreting it non-statistically as the result of entanglement with another system.) The second strategy holds on to the idea that the state of the Universe is (or at any rate might be) pure and attempts to derive statistical-mechanical probabilities from the dynamical magnification of quantum randomness to macroscopic indeterminism. The

[9] These strategies are not all alike. Recall that 'macrostates' are picked out by their fit to emergent and high-level dynamics. This fits naturally into the Everett interpretation, where branching structure is likewise emergent and high-level, and indeed picked out by the same process that selects macrostates themselves, but dynamical-collapse and hidden-variables theories are generally taken to modify the formalism at a fundamental level, and yet need to do so in a way compatible with the macrostate structure on pain of empirical inadequacy. This is a special case of a general concern with hidden-variable and dynamical-collapse approaches, developed in detail in (Wallace 2020). But for the purposes of this paper, let's stipulate that the concern can somehow be assuaged.

idea is that even if a (sufficiently complicated) system begins in a known pure state, its dynamics may carry it rapidly into a superposition of macrostates with the right probability distribution to underpin statistical mechanics. If so, PHμ might be replaced by some non-statistical specification of a unique microstate compatible with the PHM macrostate. This too is mentioned briefly in (Wallace 2023); reasons to find its dynamical conjecture plausible are given by Wallace (2016) and Albrecht and Phillips (2014). As a too-brief summary of this literature: on the one hand, statistical-mechanical systems display exactly the chaotic magnification of uncertainty that would cause quantum states to develop into macroscopic superpositions; on the other, the scope of statistical mechanics is vast and includes many systems, such as the stars in the Galaxy, where probabilities are not reducible to quantum uncertainties in at least a straightforward sense.

But for the purposes of this paper, both proposals are too abstract. They are concerned with what the Past Hypothesis logically might look like, when it is supposed to be given to us *a posteriori* by 'the normal inferential procedures of cosmology'. Let us see what they have delivered so far, now that we have the tools to follow them into the quantum regime.

9 Quantum cosmology: overview

The physics of the early Universe is reasonably well understood far earlier than the moment of transparency (at 380,000 years after the Big Bang, recall) — indeed, cosmologists are fairly confident in it all the way back to the 'electroweak epoch', which finishes around 10^{-12} seconds, at which time energy densities match those probed empirically in the Large Hadron Collider at CERN (and the Standard Model of particle physics can be extrapolated back much earlier than that, albeit that extrapolation is well beyond what we currently can test on Earth). We could imagine trying to state a Past Hypothesis at these early times that would suffice to underpin all later physics, but it would not differ interestingly from those we have considered before: in particular, it would still require a statistical-mechanical probability distribution over fluctuations, such as to evolve into the pattern of fluctuations observed later in the microwave background radiation and, ultimately, in the pattern of galaxies.

Beyond this, things get more speculative, but the strong majority view among cosmologists is that we can distinguish two more epochs of the early universe:

The Inflationary epoch: prior to the electroweak epoch, the physics of the universe is dominated by the presence of the so-called 'inflaton field', which drives a period of extremely rapid, exponential expansion of the Universe.

The Planck epoch: prior to the inflationary epoch, we reach a point where the energy density is so great that the very idea of physics playing out on a determinate spacetime just breaks down.

Little or nothing is known of the Planck epoch. It marks the breakdown of quantum field theory, the standard workhorse of modern theoretical physics;[10] in the Planck epoch, time itself becomes ill-defined and what passes for consensus physics breaks down into rival programs — string theory, loop quantum gravity, and more — which are presently incomplete and underconstrained by observation. It is not clear how to state a 'past hypothesis' in this regime, or even make sense of what an 'initial state' might be.

The inflationary epoch is much better understood[11] and much friendlier to the Mentaculus project: the details of the physics are unknown, but the general framework is quantum field theory, and the relevant notion of dynamics is the familiar determination of a later quantum state by an earlier one in accordance with the Schrödinger equation. (The quantum state encodes both the matter fields and the fluctuations of the spacetime metric around an isotropic background spacetime.) A past hypothesis in this context would be a specification of the earliest state in the inflationary epoch; it would perhaps underdetermine physics in the mysterious Planck regime but would still underpin practically all of science: not quite as sweeping a goal as the official Mentaculus, but not far short of it.

I should acknowledge frankly that serious scientists[12] have challenged the ideas of inflation, and that there are rival approaches; however, for the purposes of this paper it makes sense to put those challenges aside and ask what the Mentaculus would look like if this most-popular account of early-universe quantum physics were correct.

10 Quantum cosmology: origins of structure formation

In normal (that is: non-cosmological) quantum field theory, a special role is played by the 'vacuum' state, which is the lowest-energy quantum state. The vacuum is not 'nothingness' in any simple way: it is a complex state encoding a great deal of local and nonlocal structure. However, it is spatially and temporally uniform. It contains no particles, and never will; particle states, mathematically speaking, are 'created' from the vacuum by the action of certain operators, but this 'creation' is a formal process and does not correspond to anything physical. To get any interesting time evolution, one has to consider not the vacuum but some particle state, of which there are uncountably many, corresponding to the uncountably many ways to distribute particles across space.

Things are otherwise in an expanding universe. Indeed, in that context strictly speaking the notion of 'vacuum state' makes no sense: the vacuum state is by definition time-invariant, but the expansion process makes any state time-dependent. (In somewhat more technical terms: the vacuum is the lowest-energy

[10]More precisely, it marks the breakdown of what Wallace (2022) calls 'low-energy quantum gravity', the quantum field theory of general relativity coupled to matter fields.

[11]Once again see (Weinberg 2008), and references therein, for details.

[12]See, e. g. , (Ijjas, Loeb, and Steinhardt 2013).

state of the quantum Hamiltonian, but the Hamiltonian is the generator of time translation symmetries and an expanding universe has no such symmetries.[13]) But it is in general possible to define, in approximate but still mathematically well-stated terms, a *local vacuum* (sometimes called a *Bunch-Davies vacuum*[14]): a quantum state which at a certain instant in time has the same quantum structure as the true vacuum does.

Suppose that at a certain time the Universe is well-described by an expanding spacetime and the local quantum vacuum on that spacetime. How does it evolve forwards in time? The expansion of the Universe means that the local vacuum does not remain a local vacuum: instead, interesting physics occurs. As a simple example, particles spontaneously appear, and the local vacuum evolves into a multiparticle state. But more complicated evolutions than this — evolutions that cannot be described in a straightforward language of 'particles' — can also occur, and indeed the study of cosmic inflation is largely taken up with these somewhat alien ideas. The technical ideas do not matter here: what does matter is that a state as simple as the local vacuum will spontaneously evolve into a complex and energetic state as the Universe expands.

Or rather: it evolves into a superposition of such states. The evolution of the local quantum vacuum in an expanding universe displays exactly those features we saw as permitting a non-statistical version of PHμ. As it evolves, it becomes a superposition of macroscopically distinct, non-isotropic, non-homogeneous, distributions of mass-energy, even as the overall superposition remains both homogeneous and isotropic. In the admittedly-controversial language of the Everett interpretatation, the local vacuum evolves into a superposition of individually-anisotropic branches.

And in modern cosmology, it is this superposition — and the probability distribution over its terms defined by the quantum-mechanical (Born) probability rule — that is taken to provide a quantum-mechanical origin for the later, classical, fluctuations that are detectable in the microwave background radiation and that provide the seed for structure formation, and ultimately for the anisotropies that lead to stars and planets. This is not pure speculation: the quantum physics of inflation makes quantitative (albeit not terribly precise) predictions for the actual form of the probability distribution over microwave-background fluctuations, and that prediction can be compared to the actual, observed, pattern of fluctuations: the fit is pretty good, well within observational error.

So: the 'normal evidential procedures of cosmology' have led us to a Past Hypothesis that involves no macrostates, and no statistical probabilities, but only a single, categorical, statement about the quantum state of the universe at, if not the earliest moment there is, then the earliest moment we can yet study, and perhaps the earliest moment at which 'early' and 'moment' make

[13]Even this oversimplifies. The strictly-correct statement is that a sector of low-energy quantum gravity has a vacuum state only if it is an expansion around an extremum of the quantum effective action which has a timelike Killing symmetry.

[14]Strictly speaking, the Bunch-Davies vacuum is defined as the state which approaches the vacuum in the asymptotic past.

sense at all.

How secure is this result? It rests on inflationary cosmology, which as I have noted is a strong majority view but not without its detractors. But even if we stipulate the correctness of inflation, there is room to question whether the quantum state used — the local quantum vacuum — is the correct state. It is probably fair to say that it was originally adopted as much for convenience as for any other reason (it is by some way the simplest choice).

That said, there is no evidence against it either, no data better explained by a different state. (Physicists have explored other states — such as finite-temperature states, assumed to arise from pre-inflationary thermal physics — but to date such proposals have failed to improve on the local quantum vacuum.[15]) In other areas of physics — such as early classical cosmology — the choice of a simple, homogeneous, isotropic initial state is done for convenience and recognized as unrealistic. But the reason such states are unrealistic is that the observed world is neither homogeneous nor isotropic, and in a deterministic (and non-branching) universe with rotational and translational symmetries that means the initial state cannot be homogeneous or isotropic either. These arguments have no purchase in quantum theory. Of course there are small and localized perturbations that we could impose on the local quantum vacuum that would in all probability be unobservable today, and no evidence rules them out. But likewise, no evidence rules out local violations of the assumed laws of physics in distant and unassuming corners of the cosmos; that does not mean that it is scientifically reasonable to believe in them.

At any rate, if the form of the Past Hypothesis given by standard inflationary cosmology is not scientifically beyond question, certainly it is — by far — more scientifically evidenced than a Past Hypothesis based upon classical physics and a classical decomposition into macrostates, which is simply inadequate to the evidence if we take it to describe true cosmology, and is questionable even if we concoct fictional classical cosmologies once the details are thought through. In the conclusion of this paper, I will see what the Mentaculus project looks like if we substitute this quantum-cosmology version of the Past Hypothesis.

11 Conclusion

The quantum-field-theoretic contraption for making inferences about the world, outside the baroque and inaccessible physics applicable at Planckian scales, consists, in its entirety, of two laws and one contingent empirical fact.

The empirical fact is the one about what the macrocondition of the world happens to be,[16] and the laws are:

1. The dynamical laws of the quantum field theory comprised by general relativity and the Standard Model of particle physics, supplemented by

[15] For a route into this literature see (Kundu 2012) and references therein.

[16] Albert's own statement of this clause starts this way, but adds qualifications; I omit them for simplicity.

additional high-energy physics (dark matter and the inflaton field, most notably), and understood as an effective field theory valid at energies well below the Planck scale.

2. The Past-Hypothesis, which is that the world came into being (or at least coalesced out of Planck-scale physics) in the local quantum vacuum state for a homogenous, isotropic, inflationary spacetime.

The resemblance to Albert's contraption — to the Mentaculus — should be clear. The most important difference is that this contraption has no placeholder for the Past Hypothesis: where Albert's contraption had a space to be filled in by cosmologists, here we have looked at the cosmologists' handiwork and filled in the space accordingly. Otherwise, though, we have very much what Albert and Loewer seek: a reductive basis from which all of our large-scale, late-time physics and special sciences should follow.

Indeed, we have it in a format simpler than Albert guessed. There is no macrostate partition (at least in the lawlike part), and hence no perhaps-problematic blurring of macroscopic and fundamental; there is no nomic vagueness; there is no need for a Statistical Postulate, and so no puzzle about how the probabilities of that Postulate should be understood.

Albert and Loewer's Mentaculus is conceived of as an explicitly Humean project: to use a metaphor often applied to the Humean approach to laws, it is the best description God could give us of the Universe subject to the constraint that it fits onto a T-shirt. The version we have here almost certainly satisfies that constraint: particle physics PhD students the world over demonstrate that the Standard Model fits on a T-shirt (albeit in quite small font) and the additional space required to fit dark matter, the inflaton, any other bits of particle physics we missed, and a sharp description of the local quantum vacuum looks managable.

More seriously, the Mentaculus is normally described as best optimizing a tradeoff between the simplicity and strength of a description of the Universe, and in particular this is taken to be the rationale for the Statistical Postulate: yes, God *could* tell us the exact microstate of the early Universe, particle by particle, but who has that kind of time? (Albert 2015, section 1.4) So a probabilistic description is supposed to be the best we can do. But the version of the Past Hypothesis we extract from modern cosmology has no such tradeoff: in any of the main realist approaches to quantum mechanics, 'the quantum state at such-and-such time is ψ' is a categorical statement, not a partial summary of information that could be given in more detail if only we had the patience.

This raises the question: if the inflationary-cosmology version of the Mentaculus does not leave out any information about the initial quantum state, what *does* it leave out? — what details are left undescribed in exchange for simplicity? The answer here depends on our approach to understanding quantum mechanics.

In the de Broglie-Bohm theory, or similar hidden-variable theories, what is left out is the initial state of the hidden variables — where the Bohmian particles (or their supposed field-theoretic surrogates) actually are. The ultimate origin of probability in theories of this kind[17] is the probability distribution over hidden variables, and in the Mentaculus program this is to be understood the same way that the classical-statistical-mechanical probabilities in the Statistical Postulate were to be understood. In the end, then, the de Broglie-Bohm mentaculus is close to the classical one, only with the statistical postulate placed over hidden-variable microstates and determined by the initial quantum state.

In dynamical-collapse theories as they are usually understood, the laws are stochastic, and what is left out is the actual results of the various stochastic processes: the dynamical laws place a probability distribution over histories but do not tell us which actual history obtains. Indeed, for the Humean this probability distribution is one more instance of God's simplicity/strength tradeoff: telling us which history actually obtained would take implausibly long, and the best alternative is to summarize it partially through a set of stochastic differential equations.

And in the Everett interpretation, there is no simplicity/strength tradeoff at all. Probability in the Everett interpretation at the objective level (insofar as it can be understood at all, which is of course contested) is a categorical property of branches; at the personal level agents use it as a guide to where they can expect to be in the emergent multiverse, but not as a guide to what that multiverse is like. The Universe as a whole has neither an unknowable initial state nor stochastic dynamics: as long as we continue to set aside the unknown physics of the Planckian regime, the quantum-field-theoretic dynamics and the intial quantum state are a *complete* description of the Universe, and the complexity and anisotropy and inhomogeneity we observe is indexical.[18] The God of the Everettian Universe has no interest in compromises or tradeoffs: the description of the Universe which She puts on the T-shirt is no partial summary but Her own personal description, complete in every respect.

Acknowledgements

I am grateful to David Albert and Eddy Keming Chen for helpful conversations.

References

Albert, D. Z. (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.

Albert, D. Z. (2015). *After Physics*. Boston: Harvard University Press.

Albert, D. Z. (2023). Conclusion. In B. Loewer, B. Weslake, and E. Winsberg (Eds.), *The Probability Map of the Universe: Essays on David Al-*

[17] Some proposed relativistic versions of the de Broglie-Bohm theory (e.g. (Dürr, Goldstein, Tumulka, and Zanghí 2004)), and some other interpretations that broadly fit in the hidden-variable category, also explicitly add stochastic dynamics, which provides another source of probability; for simplicity I do not discuss this hybrid situation.

[18] For arguments along similar lines, see (Tegmark 1996).

bert's Time and Chance, pp. 351–374. Cambridge, MA: Harvard University Press.

Albrecht, A. and D. Phillips (2014). Origin of probabilities and their application to the multiverse. *Physical Review D 90*, 123514.

Binney, J. and S. Tremaine (2008). *Galactic Dynamics* (2nd ed.). Princeton: Princeton University Press.

Callender, C. (2009). The past hypothesis meets gravity. In G. Ernst and A. Hütteman (Eds.), *Time, Chance and Reduction: Philosophical Aspects of Statistical Mechanics*, Cambridge. Cambridge University Press. Available online at http://philsci-archive.pitt.edu/archive/00004261.

Chen, E. K. (2020). Time's arrow in a quantum universe: On the status of statistical mechanical probabilities. In V. Allori (Ed.), *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature*. New York: World Scientific.

Chen, E. K. (2021). Quantum mechanics in a time-asymmetric universe: on the nature of the initial quantum state. *British Journal for the Philosophy of Science 72*, 1155–1183.

Chen, E. K. (2022). Fundamental nomic vagueness. *Philosophical Review 131*, 1–49.

Chen, E. K. (2023). The past hypothesis and the nature of physical laws. In B. Loewer, B. Weslake, and E. Winsberg (Eds.), *The Probability Map of the Universe: Essays on David Albert's* Time and Chance, pp. 204–248. Cambridge, MA: Harvard University Press.

Cohen, J. and C. Callender (2009). A better best system account of lawhood. *Philosophical Studies 145*, 1–34.

Demerast, H. (2016). The universe had one chance. *Philosophy of Science 83*, 248–264.

Dürr, D., S. Goldstein, R. Tumulka, and N. Zanghí (2004). Bohmian mechanics and quantum field theory. *Physical Review Letters 93*, 090402.

Earman, J. (2006). The 'past hypothesis': Not even false. *Studies in the History and Philosophy of Modern Physics 37*, 399–430.

Gell-Mann, M. and J. B. Hartle (1989). Quantum Mechanics in the Light of Quantum Cosmology. In W. H. Zurek (Ed.), *Complexity, Entropy and the Physics of Information*, pp. 425–459. Redwood City, California: Westview Press.

Goldstein, S. (2012). Typicality and notions of probability in physics. In Y. Ben-Menahem and M. Hemmo (Eds.), *Probability in Physics*, pp. 59–71. Springer.

Goldstein, S., J. L. Lebowitz, C. Mastrodonato, R. Tumulka, and N. Zanghi (2010). Approach to thermal equilibrium of macroscopic quantum systems. *Physical Review E 81*, 011109.

Ijjas, A., A. Loeb, and P. Steinhardt (2013). Inflationary paradigm in trouble after Planck 2013. *Physics Letters B 723*, 261–266.

Kundu, S. (2012). Inflation with general initial conditions for scalar perturbations. *Journal of Cosmology and Astroparticle Physics 02*, 005.

Lazarovici, D. and P. Reichert (2015). Typicality, irreversibility and the status of macroscopic laws. *Erkenntnis 80*, 689–716.

Lebowitz, J. L. (1993). Boltzmann's entropy and time's arrow. *Physics Today 46*, 32–38.

Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability, volume II.* Berkeley: University of California Press. Reprinted, with postscripts, in David Lewis, *Philosophical Papers*, Volume II (Oxford University Press, Oxford, 1986); page numbers refer to this version.

Loewer, B. (2007). Counterfactuals and the Second Law. In H. Price and R. Corry (Eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited.* Cambridge: Cambridge University Press.

Loewer, B. (2023). The mentaculus: A probability map of the universe. In B. Loewer, B. Weslake, and E. Winsberg (Eds.), *The Probability Map of the Universe: Essays on David Albert's* Time and Chance, pp. 13–53. Cambridge, MA: Harvard University Press.

Loewer, B., B. Weslake, and E. Winsberg (Eds.) (2023). *The Probability Map of the Universe: Essays on David Albert's* Time and Chance. Cambridge, MA: Harvard University Press.

Malament, D. B. (1995). Is Newtonian cosmology really inconsistent? *Philosophy of Science 62*, 489–510.

Maudlin, T. (2007). *The Metaphysics within Physics.* Oxford: Oxford University Press.

Myrvold, W. C. (2019). Beyond chance and credence. Forthcoming.

Norton, J. (1999). The cosmological woes of Newtonian gravitation theory. In *The Expanding Worlds of General Relativity: Einstein Studies, Volume 7.* Boston: Birkhauser.

Penrose, R. (1989). *The Emperor's New Mind: concerning computers, brains and the laws of physics.* Oxford: Oxford University Press.

Price, H. (1996). *Time's Arrow and Archimedes' Point.* Oxford: Oxford University Press.

Roberts, B. W. (2022). *Reversing the Arrow of Time.* Cambridge: Cambridge University Press.

Sklar, L. (1993). *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics.* Cambridge: Cambridge University Press.

Tegmark, M. (1996). Does the universe in fact contain almost no information? *Foundations of Physics Letters 9*, 25–42.

Volchan, S. B. (2007). Probability as typicality. *Studies in the History and Philosophy of Modern Physics 38*, 801–814.

Wallace, D. (2016). Probability and irreversibility in modern statistical mechanics: Classical and quantum. To appear in D. Bedingham, O. Maroney and C. Timpson (eds.), *Quantum Foundations of Statistical Mechanics* (Oxford University Press, forthcoming). Preprint at https://arxiv.org/abs/2104.11223.

Wallace, D. (2017). More problems for Newtonian cosmology. *Studies in the History and Philosophy of Modern Physics 57*, 35–40.

Wallace, D. (2020). On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem. In J. Saatsi and S. French (Eds.), *Scientific Realism and the Quantum*, pp. 78–102. Oxford: Oxford University Press.

Wallace, D. (2022). Quantum gravity at low energies. *Studies in History and Philosophy of Science 94*, 31–46.

Wallace, D. (2023). The logic of the past hypothesis. In B. Loewer, B. Weslake, and E. Winsberg (Eds.), *The Probability Map of the Universe: Essays on David Albert's* Time and Chance, pp. 76–109. Cambridge, MA: Harvard University Press. Originally published online in 2011 (http://philsci-archive.pitt.edu/8894/).

Weinberg, S. (2008). *Cosmology*. Oxford: Oxford University Press.

Winsberg, E. (2004). Can conditionalizing on the 'past hypothesis' militate against the reversibility objections? *Philosophy of Science 71*, 489–504.

Winsberg, E. (2008). Laws and chances in statistical mechanics. *Studies in the History and Philosophy of Modern Physics 39*, 872–888.