

Cosmic Cooling as Admissibility Pathway: Expansion, Durable Structure, and the Selective Power of Computational Terminal States

Ian Staley

Independent Researcher

ORCID: 0009-0000-8592-3186

E-mail: ian.t.staley@gmail.com

Correspondence: ian.t.staley@gmail.com

Abstract

Standard cosmology treats the hot-to-cold thermal evolution of the universe as a thermodynamic consequence of metric expansion, and standard quantum mechanics treats indeterminacy as ontologically primitive. A prior framework developed by the present author proposed that a final boundary condition can constrain the admissible-history ensemble of quantum cosmology, with a tiering of terminal states—trivial, structured, and computational—in which computational terminal states exert the strongest selective pruning. That earlier work left open a cosmological question: how does thermal history relate to endpoint-conditioned history selection? The present paper advances the thesis that cosmic cooling is not itself selective, but that it establishes the intermediate regime within which a computationally organized terminal condition can exert non-trivial pruning on the admissible-history ensemble. Under this reading, quantum indeterminacy functions not as the ultimate explanatory terminus of physical theory, but as the permissive admissibility substrate through which endpoint-conditioned selection operates. This is an interpretive claim about explanatory role, not a revision of the quantum formalism or a denial of its fundamentality. The paper formalizes cooling as an intermediate-regime condition, distinguishes forced from selectable features of thermal history, addresses the trivialization, teleology, anthropic, underdetermination, and fundamentality objections, and offers three candidate differential predictions concerning error-correcting substrate prevalence, the selective fine-tuning profile of constants, and the relative timing of cosmogenic transitions. The proposal is situated against Wheeler’s participatory universe, the two-state vector formalism, the transactional interpretation, Price’s time-symmetric program, Tegmark’s mathematical universe hypothesis, and Smolin’s cosmological natural selection.

Keywords: final-state constraint, cosmic cooling, admissibility pathway, computational terminal state, endpoint-conditioned selection

1. Introduction

Two explanatory habits have become naturalized in modern physical cosmology, and they are usually held apart. The first is that the thermal trajectory of the observable universe—from the quark-gluon plasma at temperatures above 10^{12} K to the present-day cosmic microwave background temperature of $T_\gamma \approx 2.7255$ K—is a straightforward consequence of metric expansion: as the scale factor $a(t)$ grows, adiabatic cooling follows, and the successive epochs of nucleosynthesis, recombination, photon decoupling, and structure formation unfold as thermodynamic corollaries of this expansion [1]. The second is that quantum indeterminacy is ontologically primitive: whatever interpretive stance one takes toward the wavefunction, the branching of histories or the indeterminacy of measurement outcomes sits at the foundation of the theory and is not itself explained by something deeper [2], [3].

These two habits together leave a conceptual gap. If the universe cools as a matter of thermodynamic necessity and quantum branching is primitive, then the question of why the branching landscape is the *particular* landscape we observe—one that contains histories in which durable nuclei form, atoms bond, chemistry stabilizes, biological memory persists, and explicit computation eventually occurs—is answered only by appeal either to initial-condition fine-tuning, to a multiverse selection argument, or to the anthropic principle [4], [5]. None of these is without cost.

A third option is available. In a prior paper, the present author developed a framework in which the admissible-history ensemble of quantum cosmology is constrained by a final boundary condition, and in which the selective pruning exerted by different classes of terminal state can be ordered: trivial terminal states (e.g., generic thermal equilibrium) exert weak selection; structured terminal states (e.g., gravitationally bound configurations) exert moderate selection; and computational terminal states—those that favor histories in which information-processing architectures, self-replicating informational systems, and explicit computation emerge—exert the strongest selection [6]. The earlier paper did not, however, address cosmological thermal history directly. It gave the formal structure of terminal-state selection and argued on description-length grounds for the special selective power of computational endpoints, but it left open how cooling-driven cosmological evolution relates to endpoint reachability.

The thesis of the present paper is precise. Cosmic cooling is not itself the selector. Cooling is forced by expansion; it is a thermodynamic consequence and not a choice among alternatives. What cooling does is establish the intermediate regime—stable nuclei, durable chemistry, long-coherence memory substrates, low-thermal-noise conditions—within which a computationally organized terminal condition can exert non-trivial pruning. On this reading, quantum indeterminacy is not the ultimate explanatory terminus of the theory but the permissive admissibility substrate through which endpoint-conditioned history selection operates. The wavefunction formalism is unchanged; what shifts is what we take indeterminacy to be *for* in the explanatory order.

Three clarifications are essential at the outset. First, this is an interpretive claim about explanatory role, not a modification of the quantum formalism. No equation of motion is

rewritten; no new term is added to the Hamiltonian; no observational prediction is altered at the level of standard laboratory quantum mechanics. Second, the claim is not that quantum mechanics is non-fundamental in an ontological sense; it is that indeterminacy may serve as the admissibility structure over which temporally global constraints act, rather than as the final explanatory stopping point. Third, the endpoint in question is defined structurally—by the presence of computationally organized information-processing and self-replicating informational systems—not by reference to human observers. This last point matters for preempting the anthropic objection.

The remainder of the paper proceeds as follows. Section 2 summarizes the final-state framework at the level needed to make the present paper self-contained and identifies the gap that thermal history addresses. Section 3 develops the positive thesis: cosmic cooling as intermediate regime. Section 4 reframes quantum indeterminacy as permissive substrate. Section 5 addresses the slide from computational endpoint to agentic endpoint. Section 6 considers objections. Section 7 offers candidate differential predictions. Section 8 situates the view against adjacent programs. Section 9 concludes.

2. Background: Final-State Constraints and Informational Pruning

2.1 The admissible-history ensemble

The framework the present paper extends is most easily introduced through the decoherent histories formulation of quantum mechanics developed by Griffiths, Omnès, and Gell-Mann and Hartle [7], [8], [9]. In this formalism, a closed quantum system—ultimately the universe—is described not by a single evolving state vector but by an ensemble \mathcal{H} of coarse-grained histories, each labeled by a sequence of projections at successive times. A set of histories is *consistent* or *decoherent* when its decoherence functional is diagonal, permitting probabilities to be assigned obeying classical additivity. The Gell-Mann–Hartle decoherence functional

$$D(h, h') = \text{Tr}[C_{h'}^\dagger \rho_i C_h]$$

governs probabilities in the standard case, where C_h is the class operator associated with history h and ρ_i is the initial density matrix. A history h has probability $p(h) = D(h, h)$ when $\text{Re } D(h, h') \approx 0$ for $h \neq h'$.

2.2 Boundary-conditioned probability

Standard cosmology supplies only an initial boundary condition, encoded in ρ_i , and allows the unitary dynamics to determine the rest. Time-symmetric and boundary-conditioned extensions of quantum mechanics reconsider this asymmetry. The Aharonov–Bergmann–Lebowitz (ABL) rule, introduced in 1964, gives probabilities for intermediate measurement outcomes conditional on both pre- and post-selection [2]. For an intermediate measurement of an observable A with eigenstates $|a_n\rangle$, the ABL probability of obtaining outcome a_n given pre-selection of $|\psi_i\rangle$ at time t_i and post-selection of $|\phi_f\rangle$ at time t_f is

$$P_{\text{ABL}}(a_n) = \frac{|\langle \phi_f | U(t_f, t_n) | a_n \rangle \langle a_n | U(t_n, t_i) | \psi_i \rangle|^2}{\sum_m |\langle \phi_f | U(t_f, t_n) | a_m \rangle \langle a_m | U(t_n, t_i) | \psi_i \rangle|^2}.$$

The two-state vector formalism (TSVF) developed by Aharonov and Vaidman generalizes this by describing a system at time t through a forward-evolving ket $|\psi\rangle$ and a backward-evolving bra $\langle\phi|$, both defined by measurement outcomes at their respective time boundaries [10]. The TSVF is empirically equivalent to standard quantum mechanics when both boundaries are unknown; it becomes distinctive when the post-selection is physically specified.

In the final-state-constrained cosmological setting, the post-selection is not an experimental choice but a proposed property of the universe as a whole. Writing $|\Omega_f\rangle$ for a late-time boundary state (or more generally a density operator ρ_f), the admissible-history ensemble is restricted:

$$\mathcal{H}_{\text{adm}} = \{h \in \mathcal{H} : \langle \Omega_f | C_h | \psi_i \rangle \neq 0\},$$

and the probability measure on histories becomes

$$p(h \mid \rho_i, \rho_f) = \frac{\text{Tr}[\rho_f C_h \rho_i C_h^\dagger]}{\sum_{h'} \text{Tr}[\rho_f C_{h'} \rho_i C_{h'}^\dagger]}.$$

Equation (4) is the generalization to cosmology of the ABL rule. It inherits from ABL the property that the post-selection acts as a filter on the ensemble: histories compatible with ρ_f are retained; histories incompatible with ρ_f are excluded [11]. This is the core pruning operation.

2.3 Description length and the tiering of terminal states

The prior paper argued that the *kind* of ρ_f matters. Different classes of terminal state exert different selective pressure on \mathcal{H}_{adm} . The argument proceeded through the lens of algorithmic information theory, specifically the Kolmogorov complexity $K(x)$ of a history, defined as the length of the shortest program on a reference universal Turing machine that outputs x [12], [13]. The minimum description length (MDL) of the realized history ensemble under a boundary-conditioned measure is

$$L(\mathcal{H}_{\text{adm}}) = \sum_{h \in \mathcal{H}_{\text{adm}}} p(h \mid \rho_i, \rho_f) K(h \mid \rho_f),$$

where $K(h \mid \rho_f)$ is the conditional Kolmogorov complexity of history h given specification of the terminal condition [14]. A terminal state ρ_f is *selectively potent* to the extent that it reduces $L(\mathcal{H}_{\text{adm}})$ relative to the unconstrained ensemble; equivalently, to the extent that specifying ρ_f shortens the description of the histories that survive the filter.

The pruning power of ρ_f can be made quantitative. Let $\mathcal{I}(\rho_f)$ denote the mutual algorithmic information between the terminal state and a typical admissible history, averaged over the conditioned measure:

$$\mathcal{J}(\rho_f) = \mathbb{E}_{h \sim p(\cdot | \rho_i, \rho_f)} [K(h) - K(h | \rho_f)].$$

$\mathcal{J}(\rho_f)$ measures how many bits of the history description are recovered by knowing ρ_f . A terminal state that fails to constrain the ensemble has $\mathcal{J} \approx 0$; a terminal state that fully determines the history has $\mathcal{J} \approx K(h)$. The tiering of terminal states can then be stated precisely: a class \mathcal{F}_1 of terminal states is more selectively potent than a class \mathcal{F}_2 when $\mathbb{E}_{\rho_f \in \mathcal{F}_1} [\mathcal{J}(\rho_f)] > \mathbb{E}_{\rho_f \in \mathcal{F}_2} [\mathcal{J}(\rho_f)]$.

Three tiers of terminal state were identified on this basis:

Trivial terminal states (e.g., uniform high-entropy thermal states, generic heat-death configurations) impose almost no informational constraint. By construction, a high-entropy thermal state at temperature T is compatible with an exponentially large set of microstates, so specifying it as ρ_f constrains almost nothing about the antecedent history. Formally, $\mathcal{J}(\rho_f^{\text{triv}}) = O(\log \dim \mathcal{H}_{\text{Hilbert}})$, which is negligible relative to $K(h)$ for cosmologically extended histories.

Structured terminal states (e.g., configurations with specified gravitationally bound mass distributions, long-lived clustered remnants) exert moderate selection. They filter out histories in which structure fails to form, but the structural specification is comparatively compact and does not constrain the microstate degrees of freedom beyond the gross structural pattern. Typical $\mathcal{J}(\rho_f^{\text{struct}})$ scales with the specification complexity of the mass distribution, which for a clustered universe of N bound structures is $O(N \log N)$ —non-trivial but far from saturating $K(h)$.

Computational terminal states—configurations characterized by information-processing architectures, self-replicating informational systems, error-correcting memory substrates, and explicit computation in the Turing-relevant sense—exert the strongest selection, and the argument for this is essentially the *computational compressibility* of such configurations. A computational architecture is, by definition, a system whose microstate is correlated across large numbers of degrees of freedom by a comparatively short program (the code it is running, plus the code specifying its architecture). Conditioning on the presence of such an architecture is therefore conditioning on a regime of low algorithmic complexity relative to microstate count. More formally, if N_{micro} is the number of microstate degrees of freedom in a computational subsystem and K_{arch} is the algorithmic complexity of its architectural specification, then for architectures satisfying $K_{\text{arch}} \ll N_{\text{micro}} \log 2$ (the generic case for non-trivial computation), conditioning on the architecture recovers $\theta(N_{\text{micro}})$ bits of history information—scaling linearly with the size of the computational subsystem rather than logarithmically. This is the sense in which computational endpoints exert quantitatively stronger pruning than structured ones.

Table I summarizes the tiering.

Table I. *Selective pruning by terminal-state class.*

<i>Terminal-state class</i>	<i>Characteristic content</i>	<i>Selective pruning on \mathcal{H}</i>	<i>Reduction in L</i>
-----------------------------	-------------------------------	--	------------------------------------

Trivial	Generic thermal equilibrium; heat-death	Weak	Negligible
Structured	Gravitationally bound, clustered, long-lived remnants	Moderate	Partial
Computational	Information-processing architectures; self-replicating informational systems; explicit computation	Strong	Substantial

2.4 What the earlier paper did not address

The earlier framework was formal. It established that *if* a computational terminal state obtains, *then* the admissible-history ensemble is sharply pruned in the direction of histories that generate the required informational architecture. It did not address how such histories are made dynamically reachable in cosmological evolution. A boundary condition, however selective, cannot prune into existence a history that the intermediate dynamics cannot realize. The pruning operates on an ensemble already delimited by the laws of physics and by cosmological evolution.

This leaves a concrete gap. For a computational terminal state to exert its maximum selective pressure, there must be histories in the ensemble capable of producing the intermediate conditions under which computation is realizable: persistent matter, stable chemistry, durable memory, low thermal noise, error-correcting substrates. The cosmological thermal history of our universe *does* deliver these conditions, but the prior paper did not say why this matters for the pruning argument or how the two couple. Filling this gap is the task of Section 3.

3. The Cooling Universe as Intermediate Regime

3.1 The thermal history of the observable universe

The hot-to-cold trajectory of the observable universe is among the most firmly established structures in modern cosmology [15]. Beginning from the quark-gluon plasma at temperatures $T > 150$ MeV (roughly 10^{12} K, corresponding to cosmic time $t \sim 10^{-5}$ s), the universe passed through a rapid sequence of transitions: the quark-hadron phase transition at $T \sim 100$ MeV; neutrino decoupling at $T \sim 1$ MeV ($t \sim 1$ s); electron-positron annihilation; Big Bang nucleosynthesis at $T \sim 0.1$ MeV ($t \sim 3$ min), during which the light nuclei— ^2H , ^3He , ^4He , ^7Li —were synthesized; and radiation-matter equality at $T \sim 1$ eV [15]. At $T \approx 0.3$ eV (approximately 3000 K, $t \sim 380,000$ yr), hydrogen recombination occurred, free electrons bound to protons, and photons decoupled, giving rise to the cosmic microwave background now observed at $T_\gamma = 2.7255$ K [16]. Structure formation, star formation, and chemical enrichment followed over the subsequent ~ 10 Gyr.

The temperature history can be summarized by the scaling relation

$$T(a) \propto \frac{1}{a(t)}$$

for radiation, and more complicated relations for matter in the post-decoupling regime. The key observation is that cooling is thermodynamically forced: given the metric expansion described by the Friedmann equations

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2} + \frac{\Lambda}{3},$$

temperature decreases are not selectable among alternative histories compatible with expansion. Every history in which the Friedmann equations obtain and the universe expands features monotonic cooling at the radiation level.

Table II summarizes the cosmological epochs and the structural capacities each permits.

Table II. *Cosmological epochs and enabled structural capacities (approximate figures).*

<i>Epoch</i>	<i>Time</i>	<i>Temperature</i>	<i>Structural capacity enabled</i>
Quark-gluon plasma	$< 10^{-5}$ s	$> 10^{12}$ K	None beyond quark-gluon
Hadron formation	$\sim 10^{-5}$ s	$\sim 10^{12}$ K	Protons, neutrons
Nucleosynthesis	~ 3 min	$\sim 10^9$ K	Light nuclei (H, He, Li)
Recombination	$\sim 380,000$ yr	~ 3000 K	Neutral atoms; photon decoupling
Dark ages / first stars	$\sim 10^8$ yr	$\sim 10^2$ K	Molecular hydrogen; Population III stars
Galaxy formation	$\sim 10^9$ yr	~ 10 K	Metals, heavy elements
Present	$\sim 1.4 \times 10^{10}$ yr	~ 2.7 K	Chemistry, biology, computation

3.2 Forced versus selectable features

A critical conceptual distinction must be drawn before proceeding. Some features of the thermal trajectory are *forced* by expansion and the Friedmann dynamics: monotonic cooling, the approximate timing of the major transitions, the blackbody spectrum of the CMB. Other features are *selectable* in the sense that they depend on parameter choices and dynamical contingencies that differ across the ensemble of histories compatible with the laws of physics: the fine-structure constant α , the ratio of the proton to electron mass, the cosmological constant Λ , the baryon-to-photon ratio η , the precise magnitude of primordial density perturbations, and the parameters of chemistry and nuclear binding [4], [5], [17].

The claim of this paper is not that cooling per se is selected. That would collapse into the trivial assertion that universes that cool produce the structures cooling permits. The claim is rather that cooling constitutes the *pathway* through which selectable parameter configurations become consequential: cooling without appropriate parameter values yields only thin chemistry, unstable matter, or non-compounding structure; cooling with appropriate parameter values yields the durable structure a computational terminal state selectively favors. Selection operates at the level of which cooling trajectories (with which parameter values and which dynamical contingencies) survive the endpoint-conditioned filter of Eq. (4).

Figure I. *Cosmic cooling as admissibility pathway.*

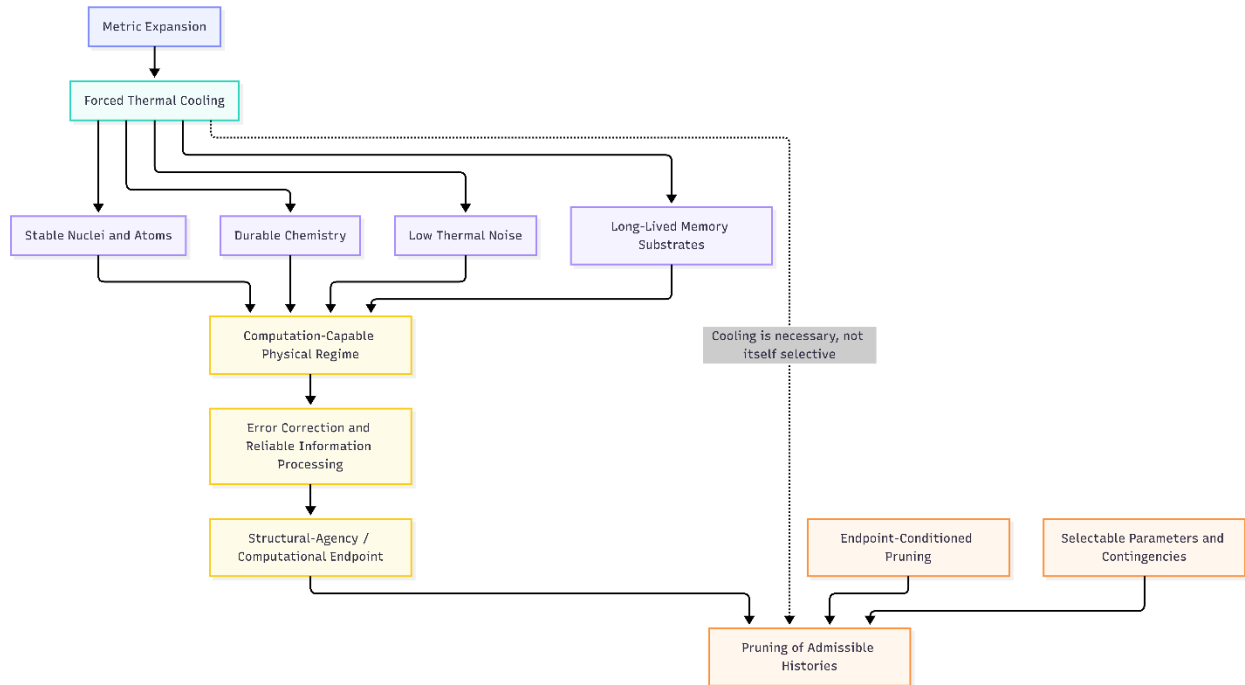


Figure I. *Cosmic cooling as admissibility pathway rather than selector. Metric expansion forces thermal cooling, which successively enables stable nuclei and atoms, durable chemistry, low thermal noise, and long-lived memory substrates. These conditions open a computation-capable physical regime in which error correction and reliable information processing become possible. Under the framework advanced here, pruning acts not on cooling itself, but on histories in which cooling together with suitable parameters and contingencies yields the intermediate regime required for a structural-agency or computational endpoint.*

3.3 Sub-claim (a): Stable nuclei, atoms, and chemistry require the cooling trajectory

The first sub-claim is that durable matter of the kind presupposed by a computational terminal state is available only in the post-cooling regime. The binding energy of the hydrogen atom is 13.6 eV; the binding energy of the deuteron is 2.22 MeV. Below these thresholds, bound atoms and nuclei are stable against thermal dissociation; above, they are not. The timing of nucleosynthesis is fixed by the condition that T drop below the deuteron binding energy while the neutron abundance has not yet entirely decayed; the timing of recombination is fixed by T dropping below hydrogen's ionization energy [15], [16]. Chemistry, as a generalization of atomic bonding, requires T to be well below typical molecular binding energies (~ 1 -10 eV), which is satisfied only in the post-decoupling universe.

None of this is in dispute; it is standard cosmology. The point for the present argument is that durable chemistry is a *necessary intermediate condition* for the content of a computational terminal state. Without stable atoms and molecules, there is no substrate for information-processing architectures of the kind the endpoint selects. Cooling is the pathway by which this substrate becomes available.

3.4 Sub-claim (b): Durable memory and error-correction require low thermal noise

The second sub-claim concerns the physical requirements of memory. Landauer's principle places a lower bound on the energy dissipated per irreversible bit operation [18], [19]:

$$E_{\text{erase}} \geq k_B T \ln 2.$$

At cosmological background temperatures of $T \sim 10^{12}$ K, $E_{\text{erase}} \sim 10^{-8}$ erg per bit—prohibitive relative to any plausible source of free energy for early-universe information processing. At $T \sim 300$ K, $E_{\text{erase}} \sim 3 \times 10^{-21}$ J per bit, accessible to chemical and biological systems.

More significantly, the durability of stored information depends on the ratio of the bistable energy barrier ΔE protecting a memory state to the thermal energy scale $k_B T$. The mean lifetime of a stored bit scales approximately as

$$\tau \sim \tau_0 \exp\left(\frac{\Delta E}{k_B T}\right),$$

which is an Arrhenius-type expression for thermally activated escape from a bistable well. For a fixed ΔE set by the chemistry of the memory substrate, τ grows exponentially as T falls. A memory that is reliably stable over timescales long enough for biological evolution or technological computation requires $T \ll \Delta E/k_B$. For covalent-bond-stabilized informational polymers (e.g., DNA, with relevant barriers in the eV range), this requires T in the hundreds of kelvins—far below the CMB temperature at recombination, but well within the post-cooling regime.

Error correction imposes additional constraints. The thermodynamic cost of an error-correcting code depends on the background noise temperature, and codes that can reliably preserve information over long periods require that thermal noise not dominate the signal [20]. Cooling reduces $k_B T$ and thereby makes error-correcting architectures thermodynamically accessible. A hot universe cannot sustain durable error-corrected memory; a sufficiently cold universe can.

3.5 Sub-claim (c): Turing-relevant computation requires persistent low-entropy substrate

The third sub-claim is the strongest. Explicit computation in the Turing-relevant sense—the controlled, reliable, resettable manipulation of symbols in accordance with a specified program—requires more than durable memory. It requires *persistent low-entropy substrate conditions* in which the computational architecture is dynamically stable, in which clock cycles can be defined, and in which the probability of thermally induced errors is low enough that computation proceeds as intended rather than collapsing into noise [21].

Seth Lloyd's bound on the total computational capacity of the observable universe,

$$N_{\text{ops}}(t) \lesssim \frac{2 \rho(t) c^5 t^4}{\pi \hbar},$$

gives the maximum number of elementary logical operations the universe could have performed up to cosmic time t [21], [22]. For the present-day universe, $N_{\text{ops}} \sim 10^{120}$. This bound depends critically on the low effective temperature of the present universe: a universe locked at $T \gg 10^6$ K would expend vastly more energy per bit-flip via the Margolus–Levitin bound, and would be dominated by thermal fluctuations rather than organized computation.

More concretely, Turing-relevant computation requires that the Kolmogorov complexity of the computational trajectory be substantially less than the complexity of the underlying microphysical trajectory. Computation occurs, by definition, when the program state evolves along a much lower-dimensional manifold than the full phase space permits. This separation—between the program state and the thermal bath—requires the thermal bath to be sufficiently cold that its fluctuations do not spill over into the program manifold. A hot universe does not permit such a separation; a cold universe does.

3.6 Cooling as admissibility pathway

Drawing the three sub-claims together: the cooling trajectory of the universe establishes, successively, the intermediate regimes in which stable chemistry, durable memory, and explicit computation become physically realizable. Cooling itself is forced; whether cooling *together with a given parameter configuration and dynamical contingency* produces the intermediate conditions a computational terminal state requires is not forced.

Selection operates on the conjunction. Let $\pi \in \Pi$ denote a point in the space of cosmological parameters and initial-condition contingencies compatible with a cooling trajectory—including the fine-structure constant α , the proton-to-electron mass ratio m_p/m_e , the cosmological constant Λ , the baryon-to-photon ratio η , nuclear-binding parameters, and the amplitude and spectrum of primordial density perturbations. Write $\rho_i(\pi)$ for the parameter-dependent initial density matrix and ρ_f^c for a computational terminal state. Eq. (4) becomes

$$p(h \mid \rho_i(\pi), \rho_f^c) \propto \text{Tr}[\rho_f^c C_h \rho_i(\pi) C_h^\dagger].$$

Integrating over histories while retaining π as a label, the induced measure on parameter space is

$$\mu(\pi \mid \rho_f^c) = \frac{\sum_h \text{Tr}[\rho_f^c C_h \rho_i(\pi) C_h^\dagger] \mu_0(\pi)}{\int_{\Pi} d\pi' \sum_h \text{Tr}[\rho_f^c C_h \rho_i(\pi') C_h^\dagger] \mu_0(\pi')},$$

where $\mu_0(\pi)$ is a prior measure on parameter space (taken uniform or minimally informative in the absence of further constraints). Eq. (13) is the induced posterior over parameters given the endpoint condition. It expresses precisely what is meant by selection operating *through* cooling rather than *on* cooling: cooling is the dynamical trajectory along which each π is tested against ρ_f^c , and $\mu(\pi \mid \rho_f^c)$ peaks at values of π for which the cooling trajectory delivers the intermediate regimes required by the endpoint.

This gives the framework a quantitative posture it otherwise lacks. The admissible parameter values are not those that make the universe exist, nor those that make the universe structured, but those for which the conjunction of cooling dynamics and parameter configuration produces histories compatible with a computational endpoint. Parameter configurations with appropriate α but wrong Λ (cooling proceeds but structure formation is cut off), or appropriate Λ but wrong α (structure forms but chemistry is thin), are exponentially suppressed in $\mu(\pi \mid \rho_f^C)$ relative to configurations that satisfy both.

On this reading, the hot-to-cold evolution of the observable universe is not proof of teleology nor merely thermal accounting. It is the intermediate regime through which endpoint-conditioned pruning becomes non-trivial on parameter space. The formal content of the framework is not that cooling is explained by the endpoint, but that parameter configurations selecting for histories with appropriate cooling pathways are over-represented in $\mu(\pi \mid \rho_f^C)$ relative to $\mu_0(\pi)$.

4. Quantum Indeterminacy as Permissive Substrate

4.1 Indeterminacy under a boundary-conditioned reading

In the standard reading of quantum mechanics, indeterminacy is taken as ontologically primitive: the Born rule returns probabilities that are not further reducible to hidden variables (subject to the usual no-go constraints on locality and non-contextuality) [23]. Under a boundary-conditioned reading of the cosmological ensemble, a different explanatory role becomes available without modifying the formalism.

Figure II. *Boundary-conditioned pruning of the admissible-history ensemble.*

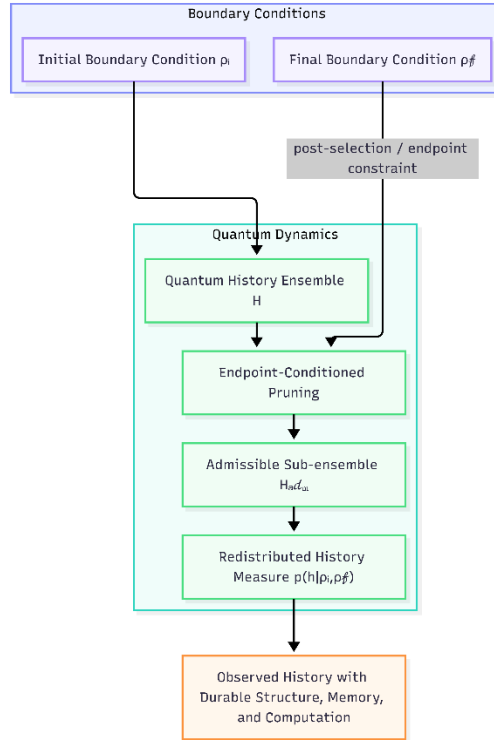


Figure II. Boundary-conditioned pruning of the admissible-history ensemble. An initial boundary condition generates a quantum history ensemble under the dynamics, while a final boundary condition imposes an endpoint-conditioned pruning on that ensemble. The result is an admissible sub-ensemble and a redistributed measure over histories, from which observed histories containing durable structure, memory, and computation emerge.

The standard formalism generates an ensemble \mathcal{H} of possible histories. Without a final boundary condition, probabilities on this ensemble are given by Eq. (1) and are interpreted as irreducible. With a final boundary condition, the ensemble is restricted to \mathcal{H}_{adm} by Eq. (3), and probabilities are redistributed by Eq. (4). The formalism is identical; what changes is the conceptual role of \mathcal{H} .

In the standard reading, \mathcal{H} is the complete explanatory ledger: every possible branching, however unproductive, is part of what the theory posits. In the boundary-conditioned reading, \mathcal{H} is the *admissibility structure* over which endpoint-conditioned selection operates. The branching is real, but it is not explanatorily terminal. What is explanatorily terminal is the conjunction of initial and final boundary conditions; \mathcal{H} is the space on which they act.

4.2 “Admissibility structure” versus “filtering medium”

The sharper rhetorical framing of this view is that quantum indeterminacy is a “filtering medium” rather than “the foundation.” This is catchy but overreaches. A filtering medium suggests that something sits *above* the quantum formalism and does the filtering from outside. That is not the claim. The claim is that the formalism generates an ensemble, and that a temporally global constraint—specifically a final boundary condition with non-trivial

informational content—selects a sub-ensemble. The selection is internal to the formalism as augmented by the boundary condition; nothing sits above.

The better framing is that indeterminacy functions as the *admissibility structure* on which the endpoint condition acts. This language preserves the fundamentality of the formalism: without quantum branching, there is no ensemble for the endpoint condition to operate on, and the selection has no content. Indeterminacy is a *necessary* feature of the theory under this reading; it is not a feature that could be explained away by a deterministic substitute. But it is not the final explanatory stop either. The final explanatory stop is the *conjunction*: branching structure plus boundary condition plus selective pruning.

4.3 What does and does not change

Several things do not change under this reading:

- (i) The laboratory predictions of quantum mechanics are unaffected. The ABL rule, the weak-value formalism, the Born rule in the standard single-measurement setting, and all the empirical content of quantum theory continue to obtain [10], [24].
- (ii) The formalism is not modified. No new terms enter the Hamiltonian; no new constants enter the theory; no equation of motion is rewritten.
- (iii) The ontological status of the wavefunction is left open. The boundary-conditioned reading is compatible with multiple positions on whether $|\psi\rangle$ is ontic or epistemic [25], [26].

What does change is the explanatory role of indeterminacy. In the standard reading, asking “why is there indeterminacy?” has no further answer; the Born rule is primitive. In the boundary-conditioned reading, the answer becomes: indeterminacy is the admissibility structure over which temporally global constraints act, and those constraints—specifically a computational terminal condition—determine which branches are realized with non-negligible measure.

This is analogous, but only analogous, to the Lagrangian formulation of classical mechanics. In Newtonian mechanics, forces are primitive and trajectories follow. In the Lagrangian formulation, trajectories are determined by extremization of an action functional over all possible paths. The path ensemble is not physical in the Newtonian reading, but it plays an indispensable formal role in the Lagrangian reading. Neither reading is “more fundamental” in a metaphysical sense; they differ in explanatory organization.

The boundary-conditioned reading of quantum indeterminacy is similarly a reorganization of explanatory priorities. It does not demote the quantum formalism to secondary status, and it does not elevate the boundary condition above the formalism. It says that once a boundary condition is included, the natural explanatory role of indeterminacy is as the admissibility substrate—not because indeterminacy has been explained by something deeper, but because the theoretical work of selection can now be located elsewhere in the formalism.

4.4 Non-empirical theoretical virtues

The boundary-conditioned reading is empirically equivalent to standard quantum mechanics at the laboratory level. The justification for adopting it therefore must rest on non-empirical theoretical virtues rather than crucial experiments. Three such virtues are worth making explicit.

Explanatory unification. Standard cosmology supplies an initial boundary condition (or, in the no-boundary proposal, a specific substitute) and takes the late-time structure of the universe as a derived consequence of initial data plus dynamics. Separately, standard quantum mechanics takes indeterminacy as primitive and the Born rule as fundamental. These two explanatory bases are not unified: initial-condition specification and Born-rule primitiveness do no common work. Under the boundary-conditioned reading, a single formal move—the inclusion of ρ_f alongside ρ_i in Eq. (4)—couples the treatment of cosmological initial data to the treatment of quantum probability. Both boundary conditions enter symmetrically into the measure on histories; the Born rule becomes a consequence of the fully specified boundary-conditioned measure rather than a separate primitive. This is explanatory unification in the standard sense: one formal structure does work that was previously done by two disconnected ones.

Reduction in arbitrary boundary-data specification. Standard cosmology faces a recurring awkwardness about initial data: the Weyl curvature hypothesis, the low-entropy past state, and the inflationary initial conditions are all in various ways attempts to motivate specific initial data on grounds other than arbitrary stipulation [27], [28]. Penrose's estimate that the Weyl curvature hypothesis requires a specification precision of order 1 part in $10^{10^{123}}$ illustrates how severe the arbitrariness can become when only initial boundary data is available to carry the explanatory load [27]. The boundary-conditioned reading distributes explanatory work across two boundary conditions rather than concentrating it in one. A computational terminal ρ_f^A is not a replacement for ρ_i , but it offloads part of the specification burden: some features of the realized histories are accounted for by the endpoint condition rather than by fine-tuned initial data. The arbitrariness of the initial specification is reduced, at the cost of specifying a late-time condition that is, arguably, independently motivated by the observed presence of structural agency in our cosmic history.

Parameter-space content. Perhaps most concretely, the induced measure on parameter space defined by Eq. (13) is not present in the standard framework at all. Standard cosmology treats the fundamental constants as fixed inputs to the theory; fine-tuning puzzles are noted but not explanatorily resolved by the theory itself. The boundary-conditioned reading generates $\mu(\pi \mid \rho_f^C)$ as a theoretical object—a posterior distribution on parameter space given the endpoint condition—that can in principle be studied, compared with observed parameter values, and used to make the differential predictions developed in Section 7. Whether or not the predictions turn out to be sharp, the existence of $\mu(\pi \mid \rho_f^C)$ as a theoretical object is a non-trivial content gain over standard frameworks that take parameters as fixed without further comment.

These three virtues—unification, reduction in arbitrary specification, and parameter-space content—are not empirical claims and do not by themselves settle whether the framework is correct. They do explain why adopting the reading is a non-trivial theoretical choice and not merely a linguistic rearrangement.

5. Why Agency, Not Just Computation

5.1 The slide to avoid

Informal restatements of the final-state framework tend to slide from “computational terminal state” to “agentic terminal state,” and from there to “endpoint characterized by observers” or “endpoint characterized by intelligent life.” Each step loosens the structural content of the claim and pulls the argument closer to anthropocentrism. It is necessary to say what structural agency *is*, and why the endpoint is better characterized in terms of it than in terms of generic computation.

A purely computational terminal state, strictly construed, would select for histories containing Turing-complete machinery—whether or not that machinery did anything interesting. A patch of rock in which sufficiently complex crystal defects can be read as a Turing-complete system, were anyone to read them that way, is computational by a liberal definition and yet contributes nothing distinctive to the endpoint. The selective character of a computational terminal condition must therefore be sharpened.

5.2 Structural features of agency

The features that matter for selective pruning, and that together constitute *structural* agency, are three:

Recursive self-modeling. A system has recursive self-modeling when it maintains an internal representation of its own state that is updated by its own processes and that is used to guide those processes. This is the core feature of predictive-processing and free-energy-principle accounts of biological organization [29], [30]. Structurally, recursive self-modeling requires persistent memory, updatable internal state, and feedback from outputs to inputs.

Goal-directed informational closure. A system exhibits goal-directed closure when its dynamics preferentially maintain a bounded set of internal states against environmental perturbation. Autopoietic and autonomy-based accounts in biology emphasize this feature [31], [32]. Structurally, closure requires boundary-crossing flows of matter and energy combined with internal regulation that stabilizes state variables.

Self-replicating informational organization. A system is self-replicating in the relevant sense when its internal organization specifies procedures for producing further systems of the same organizational type. Kauffman’s autocatalytic sets, Eigen’s hypercycles, and more recent work by England on dissipative adaptation all bear on this feature [33], [34], [35]. Structurally, self-replication requires copyable informational substrate, the machinery to

read and copy it, and environmental conditions permitting the copying to produce reliable descendants.

A computational terminal state that selects for all three features is an *agentic* terminal state in the structural sense. Agency here does not mean person-like agency. It does not imply consciousness, intentionality as understood in philosophy of mind, or moral status. It means the structural cluster of recursive self-modeling, goal-directed closure, and self-replicating organization. Integrated information theory's Φ measure [36], [37], while contested as a measure of consciousness, provides one candidate quantitative handle on the first feature, and it has been shown to correlate with fitness in simulated evolution [38], [39].

5.3 A formal handle on structural agency

The three features can be given operational form as a composite functional. For a subsystem S of a history h at time t , let $\Phi(S, t)$ denote the integrated information of S (an IIT-style measure of recursive self-modeling capacity) [36], [37], let $F(S, t)$ denote the variational free energy of S relative to its environment (an FEP-style measure of goal-directed closure) [29], [30], and let $R(S, t)$ denote the replication fidelity of S (the mutual information between S and its descendants under the system's own reproduction dynamics) [33], [35]. A structural-agency functional can be defined as

$$\mathcal{A}(h) = \max_{S \subset h, t} [\Phi(S, t) \cdot \exp(-\beta F(S, t)) \cdot R(S, t)],$$

with β a positive inverse-temperature-like parameter setting the weight on free-energy minimization. $\mathcal{A}(h)$ is positive only when all three features are simultaneously present above threshold; it is large when the features interlock in a stable subsystem. A structural-agency terminal state ρ_f^A is then defined as one whose support is restricted to final microstates $|x\rangle$ such that $\mathcal{A}(h_x) \geq \mathcal{A}_0$ for some threshold $\mathcal{A}_0 > 0$, where h_x is a history terminating in $|x\rangle$:

$$\text{supp}(\rho_f^A) \subseteq \{|x\rangle: \mathcal{A}(h_x) \geq \mathcal{A}_0\}.$$

This is schematic—each of Φ , F , and R remains contested in its own literature, and the multiplicative form is one choice among several possible compositions. But the formalization does real work. It states precisely what it would mean for a microstate to lie in the support of ρ_f^A , it makes explicit the three-way dependency, and it fixes the sense in which ρ_f^A is structural rather than observer-relative: \mathcal{A} is defined over subsystems and histories without reference to any privileged external observer.

5.4 Why the sharper endpoint is more, not less, selective

One might worry that narrowing the endpoint from generic computation to structural agency weakens the selective power of the framework. The opposite is true. A generic computational terminal state is satisfied by any history containing Turing-complete machinery, even machinery that does not do anything persistent or organized. A structural-agency terminal state is satisfied only by histories in which the three features above arise,

persist, and interlock. This is a narrower class of histories, which corresponds to a stronger filter in Eq. (4) and a larger reduction in L in Eq. (5).

Formally, if ρ_f^C denotes the generic computational terminal state and ρ_f^A denotes the agentic (structural-agency) terminal state with $\rho_f^A \preceq \rho_f^C$ in the partial order of informational content, then

$$L(\mathcal{H}_{\text{adm}} \mid \rho_f^A) \leq L(\mathcal{H}_{\text{adm}} \mid \rho_f^C),$$

with equality only when the agentic refinement adds no independent informational content. The selective pressure of the agentic endpoint dominates that of the merely computational endpoint.

5.5 Preempting the anthropic objection

Defining the endpoint structurally is essential for preempting the anthropic objection. If the endpoint were defined as “universes containing human observers,” the framework would collapse into a form of strong anthropic reasoning, vulnerable to all the standard critiques: that it trivially explains what it ought to predict, that it is unfalsifiable in principle, and that it smuggles in teleology under the cover of structural language [40], [41].

Structural agency is defined without reference to humans, carbon, biology, or consciousness. A universe in which structural agency is realized by silicon-based self-replicating systems, by exotic plasma structures sustaining recursive self-modeling, or by distributed computational architectures not resembling any earth biology would satisfy ρ_f^A equally well. The endpoint is defined by the cluster of organizational features, not by their physical realization in any particular substrate.

This distinction matters both philosophically and technically. Philosophically, it answers the charge that the framework is merely anthropocentrism in informational dress. Technically, it allows Eq. (12) to apply to histories that are very different from ours in chemistry and biology yet still admissible under the endpoint condition.

6. Objections

6.1 The trivialization objection

Objection. Cooling is forced by expansion. The Friedmann dynamics permit no alternative. Therefore cooling cannot be part of the selection story; it is thermodynamic inevitability. The present proposal collapses into the claim that universes that cool contain structures cooling permits—which is true but empty.

Reply. The objection conflates cooling per se with the conjunction of cooling and parameter configuration. Cooling is forced; the particular parameter values (fine-structure constant, baryon-to-photon ratio, cosmological constant, nuclear binding coefficients) that determine whether cooling yields durable chemistry, long-coherence memory, and computational substrate are *not* forced. These parameters sit at the boundary of what the standard model

takes as free. Selection operates at the level of the *conjunction*—not “cooling was selected” but “cooling-with-parameter-configuration- π was selected.” The framework is not trivial once this distinction is made. It predicts that parameter configurations permitting durable computational substrate are over-represented in the admissible ensemble relative to configurations that do not. This is non-trivial.

6.2 The teleology objection

Objection. A final boundary condition is final cause in disguise. The proposal reintroduces teleology, against which modern science has set its face.

Reply. Boundary conditions at late times are no more teleological than boundary conditions at early times. Standard cosmology routinely posits initial-condition boundary data—Penrose’s Weyl curvature hypothesis for the initial singularity is a well-known example [27], [28]—without anyone objecting that the initial condition is “first cause in disguise.” The asymmetry between our comfort with initial boundary data and our discomfort with final boundary data is not a feature of the physics; it is a feature of our psychology and our confidence in a certain direction of causal explanation.

Formally, Eq. (4) is time-symmetric in its treatment of ρ_i and ρ_f ; the ABL rule is explicitly a time-symmetric generalization of ordinary conditional probability [2], [10]. The TSVF and transactional interpretations have for decades argued that time-symmetry in the formalism is not teleology [42], [43]. What distinguishes teleology from boundary-conditioned physics is whether the endpoint is stipulated with ghostly causal power or is described structurally by a density operator. In the present framework, ρ_f^A is a density operator with specifiable informational content. It does not pull histories toward itself by occult means. It enters Eq. (4) as one of two boundary conditions and is treated formally on par with ρ_i .

6.3 The anthropic objection

Objection. This is simply the anthropic principle in new clothes. Observers select for universes in which observers exist.

Reply. The anthropic principle, in both weak and strong forms, refers to observer selection: we observe the universe we observe because we exist to observe it, and universes incompatible with our existence are not observed [4], [40]. The present framework does not invoke observer selection. It invokes an endpoint condition defined structurally in terms of agency features—recursive self-modeling, goal-directed closure, self-replicating organization—that are specified without reference to observers, consciousness, or human-like intelligence.

The framework is compatible with counterfactual universes in which structural agency is realized by non-human, non-biological, non-carbon-based systems. It is compatible with universes in which no observer ever contemplates the endpoint condition. What it is not compatible with is universes in which no structural agency is realized at all; those universes are pruned out. This is a selection principle, but it is a structural-organizational selection principle, not an observer-selection principle.

The distinction matters because observer selection is notoriously vulnerable to charges of triviality and unfalsifiability, while structural selection can in principle be refined to the point of making differential predictions (Section 7).

6.4 The underdetermination objection

Objection. The framework makes no empirical predictions different from standard cosmology plus standard quantum mechanics. Any observation consistent with the standard framework is consistent with the boundary-conditioned reading. Underdetermination renders the view empirically empty.

Reply. The objection is partly correct and partly overstated, and the distinction matters. It is correct that at the level of laboratory quantum mechanics and at the level of fixed-parameter cosmological prediction, the boundary-conditioned reading returns what standard theory returns. The formalism is identical; the measure on histories given fixed boundary conditions is identical. At this level, empirical discrimination is not available.

The objection is overstated in three ways.

First, the framework makes content commitments at the parameter-space level that standard theory does not. Eq. (13) defines a posterior measure $\mu(\pi \mid \rho_f^C)$ on the space of cosmological parameters; standard theory takes parameters as fixed inputs and does not generate such a measure. The differential predictions of Section 7 are derived from this posterior. Whether they can be sharpened depends on further work on the measure, but their existence as non-trivial content is not in dispute.

Second, empirical equivalence at the laboratory level is not unique to the present framework. It is shared by every interpretive program in the foundations of quantum mechanics: Copenhagen, Many-Worlds, Bohmian mechanics, the transactional interpretation, the TSVF, and Price's retrocausal program all return the same laboratory predictions. Dismissing the present framework on grounds of laboratory-level empirical equivalence would proscribe the entire interpretive enterprise, including views widely regarded as serious theoretical proposals. The appropriate standard for a framework of this type is non-empirical theoretical virtue (Section 4.4) plus differential content at the level accessible to the framework (Section 7), not laboratory-crucial experiment.

Third, the framework is falsifiable in a structural sense even where it is not empirically discriminating at the laboratory level. If the observed pattern of cosmological parameters were demonstrably *not* over-represented in regions of Π favoring durable computational substrate—if, for instance, α were located at a value far from the narrow range permitting stable chemistry, or if the cosmological constant were at a value incompatible with structure formation—the framework would be falsified. That we do not observe such values is not a trivial consequence of observer selection, because the framework's endpoint is defined structurally rather than observationally (Section 5.5). The absence of such falsifying observations is therefore a non-trivial confirmation rather than a tautology.

6.5 The Boltzmann brain objection

Objection. Any framework that selects for histories containing observers or agents faces the Boltzmann brain problem: if structural agency can arise through rare thermal fluctuations in a high-entropy equilibrium state, then such fluctuation-produced agents should dominate the admissible ensemble under ρ_f^A , since they are vastly more numerous than evolution-produced agents in the infinite-time limit. The framework would then predict that a typical observer is a fluctuation, contrary to observation.

Reply. The Boltzmann brain problem arises for any framework in which structural agency is the selection target and in which thermal fluctuations in long-lived equilibrium states can produce such agency. It is a serious objection, and the reply must be specific.

The structural-agency functional $\mathcal{A}(h)$ defined in Eq. (15) is not satisfied by isolated thermal fluctuations. All three features—integrated information Φ , goal-directed closure F , and replication fidelity R —require extended spatiotemporal organization with non-trivial causal history. A fluctuation-produced brain-like configuration has $\Phi > 0$ at one instant but has no causal history satisfying F (no stable boundary-crossing flows) and no replication dynamics satisfying R (no mechanism for producing descendants with preserved organization). The multiplicative form of Eq. (15) means that any factor at zero suppresses the whole: isolated fluctuations fail all three conditions and are therefore not in the support of ρ_f^A under the definition Eq. (16).

More formally, fluctuation-produced configurations satisfy $R \approx 0$ because they have no reliable self-reproduction mechanism; they satisfy $F \approx 0$ because they do not maintain bounded internal states against environmental perturbation over non-trivial timescales; and their Φ is typically evaluated at a single instant rather than persisting across cause-effect intervals long enough for the integrated-information measure to be meaningful. The structural-agency endpoint is therefore not the observer endpoint assumed in standard Boltzmann brain arguments. It is narrower, and the narrowing is precisely what excludes fluctuation agents.

This is a concrete advantage of defining the endpoint structurally rather than in terms of observer-existence. An endpoint defined as “universes containing observers” is vulnerable to Boltzmann brain collapse; an endpoint defined as “universes containing structurally agentic subsystems with persistent causal history and reliable replication” is not. The framework’s insistence on all three structural features is not aesthetic—it is what prevents the Boltzmann brain pathology.

6.6 The fundamentality objection

Objection. Calling quantum mechanics “non-ultimate” overreaches. The formalism is the most successful physical theory ever produced; demoting it to an admissibility substrate is unmotivated and gratuitously provocative.

Reply. Conceded, in the following sense. The claim is not that quantum mechanics is non-fundamental in an ontological sense. The claim is that indeterminacy may be non-ultimate in an *explanatory* sense under a boundary-conditioned reading. These are different claims.

The formalism retains its full explanatory and empirical power; what changes is whether indeterminacy functions as the final stop in the explanatory order or as the admissibility structure over which a temporally global constraint acts.

Language here is important. The earlier formulation of this view, using the phrase “filtering medium,” invited exactly the objection under consideration, because it suggested something sitting above the formalism doing the filtering. The preferred language of “admissibility structure” avoids this. Quantum indeterminacy is not replaced; it is reconceived in its explanatory role.

7. Testability and Differential Predictions

This is the section on which the paper’s empirical claims most directly rest, and honesty demands that its limits be acknowledged alongside its content. The differential predictions below are offered not as crucial experiments but as expectation-value differences under which the boundary-conditioned framework with a structural-agency endpoint predicts different tendencies than standard cosmology plus a trivial or merely structured endpoint.

7.1 Prediction 1: Prevalence of error-correcting substrate

Under a structural-agency terminal state, the admissible ensemble is weighted toward histories in which error-correcting informational substrates emerge and persist. Error-correction is a prerequisite for durable memory against thermal noise (Eqs. (9)-(10)) and is a prerequisite for self-replicating organization that transmits information across generations without unbounded drift.

Differential expectation. Histories in the admissible ensemble under ρ_f^A should exhibit a higher prevalence of error-correcting informational architectures—biological (DNA repair, proofreading polymerases), pre-biological (autocatalytic networks with self-stabilizing composition), and eventually technological (error-correcting codes in silicon and quantum computing)—than histories in which the endpoint is merely structured. The quantitative statement is that the measure on \mathcal{H}_{adm} under ρ_f^A should assign higher probability to histories with error-correcting informational density above some threshold than the corresponding measure under ρ_f^S (structured-only).

Limitation. This is an expectation-value prediction, not a sharp test. Assessing it requires a comparison across histories, which observationally we cannot make directly. The best empirical handle is the ubiquity and convergence of error-correction in terrestrial biology and the early emergence of proofreading mechanisms, which is qualitatively consistent with the prediction but cannot rule out standard explanations.

7.2 Prediction 2: Selective fine-tuning profile

Fine-tuning arguments in cosmology identify a suite of constants whose values appear finely tuned for the emergence of complex structure [4], [5], [17]. The boundary-conditioned framework with ρ_f^A predicts a specific pattern: the fine-tuning should be

particularly sharp for parameters that enable *durable computational substrate*, as distinct from parameters that merely enable structure.

Differential expectation. Compare two classes of constants. Class \mathcal{C}_1 are those whose values control whether structure forms at all—gravitational constant G , cosmological constant Λ , primordial density fluctuation amplitude $\delta\rho/\rho$. Class \mathcal{C}_2 are those whose values control whether the structures that form can host durable computation—fine-structure constant α , proton-electron mass ratio, nuclear binding details, neutron-proton mass difference, chemistry-relevant parameters. Under a structured-but-non-computational endpoint, \mathcal{C}_1 would be finely tuned and \mathcal{C}_2 would not. Under ρ_f^A , both \mathcal{C}_1 and \mathcal{C}_2 should be finely tuned, and \mathcal{C}_2 should be tuned with comparable or greater sharpness than \mathcal{C}_1 .

Quantitative content. Define the tuning sharpness of a parameter π_k relative to endpoint condition ρ_f as

$$\sigma_k(\rho_f) = \frac{\Delta\pi_k^{\text{viable}}(\rho_f)}{\pi_k^{\text{obs}}},$$

where $\Delta\pi_k^{\text{viable}}(\rho_f)$ is the width of the interval around the observed value π_k^{obs} within which the endpoint remains realizable and other parameters are held fixed. A smaller σ_k indicates sharper tuning. Existing fine-tuning estimates in the literature provide indicative numbers. For the cosmological constant, Weinberg’s anthropic bound gives $\sigma_\Lambda \sim 10^{-120}$ relative to Planck-scale natural values [44], though the interpretation of this ratio is contested [5], [17]. For the fine-structure constant relative to stable carbon nucleosynthesis via the Hoyle resonance, Oberhummer, Cs  t  , and Schlattl’s calculations give $\sigma_\alpha \lesssim 10^{-3}$ for a viable carbon yield within a factor-of-two change [45]. More recent lattice QCD calculations by Epelbaum et al. refine this to give $\sigma \sim 2\text{-}3\%$ for the light quark mass variation consistent with carbon-based chemistry [46]. For the proton-neutron mass difference relative to stable hydrogen chemistry, $\sigma_{\Delta m} \sim 10^{-2}$ [17], [47]. The observation relevant to the prediction is whether σ for \mathcal{C}_2 parameters is of comparable order to σ for \mathcal{C}_1 parameters when both are normalized on the same footing. The extant literature suggests that they are: the sharpness of the α tuning for carbon nucleosynthesis is of the same order as the sharpness of the baryon-photon ratio tuning for structure formation, which is a non-trivial coincidence under a merely structured endpoint but a predicted feature under ρ_f^A .

Falsification condition. The prediction would be falsified if future calculations established that σ for \mathcal{C}_2 constants was systematically *larger* (looser tuning) than σ for \mathcal{C}_1 constants—that is, if the constants controlling durable computational substrate turned out to have generous viable ranges while the constants controlling bare structure formation were sharply tuned. The framework predicts the reverse: comparable or sharper tuning of \mathcal{C}_2 . A systematic reversal in the sharpness hierarchy would constitute evidence against the structural-agency endpoint relative to a merely structured one.

Limitation. Fine-tuning arguments are notoriously sensitive to the choice of measure on parameter space. Without an agreed-upon measure, quantitative comparisons between \mathcal{C}_1 and \mathcal{C}_2 tuning degrees are contestable. The numbers above are indicative rather than

decisive, and sharpening them to the point where Eq. (17) gives rigorous discriminations across parameter classes is a concrete direction for follow-on work.

7.3 Prediction 3: Relative timing of cosmogenic transitions

Under ρ_f^A , the admissible ensemble should be weighted toward histories in which the relative timing of cosmogenic transitions—matter-radiation equality, recombination, first-star formation, chemical enrichment, planetary formation, biogenesis, the emergence of information-processing architectures—is well-ordered for the downstream emergence of structural agency.

Differential expectation. The timing sequence in the admissible ensemble should disproportionately feature histories in which each transition occurs soon enough relative to the others that the downstream transition has time to proceed before the universe has expanded and cooled past the point where it is dynamically favorable. For example, a history in which star formation is delayed so long that metal enrichment cannot reach the required levels before the dark-energy-dominated expansion shuts down star formation is disfavored under ρ_f^A but not under a merely structured endpoint.

Empirical handle. The observation that our universe’s transition timings appear, qualitatively, well-ordered for the emergence of complex chemistry and biology is consistent with the prediction. The fact that star formation peaked at $z \sim 2$ and has since declined, and that planetary formation with heavy-element abundances comparable to Earth’s became possible only after several Gyr of metal enrichment, fits the pattern.

Limitation. Like the previous prediction, this is a distributional claim, and direct testing would require comparative observation across multiple universes or a principled measure on cosmological parameter space.

7.4 Summary: empirical status

Table III summarizes the three predictions against the alternative of a structured-but-not-computational endpoint.

Table III. *Differential predictions under ρ_f^A versus ρ_f^S .*

<i>Prediction</i>	<i>Expected under ρ_f^A</i>	<i>Expected under ρ_f^S</i>	<i>Empirical handle</i>
Error-correcting substrate prevalence	High; convergent	Moderate; contingent	Biological error correction ubiquity
Fine-tuning of \mathcal{C}_2 constants	Sharp, comparable to \mathcal{C}_1	Not specifically sharp	Observed fine-tuning of α , nuclear parameters
Cosmogenic transition timing	Well-ordered for downstream agency	Well-ordered for structure only	Star formation history; metal enrichment timing

Honesty about the empirical status: these are soft predictions. They locate the framework in the interpretive/explanatory tier of proposals in philosophy of physics rather than in the direct-empirical tier. They are nevertheless non-trivial: they distinguish the structural-agency endpoint from merely structured endpoints in ways that, in principle, could be sharpened by future work on cosmological parameter measures and algorithmic-information measures on the admissible-history ensemble.

8. Relation to Adjacent Programs

The view advanced here has clear affinities and equally clear distinctions with several adjacent programs in the foundations of physics. Locating it precisely among these helps clarify what is new.

8.1 Wheeler’s participatory universe

John Archibald Wheeler’s “it from bit” and the associated participatory anthropic principle (PAP) propose that observer-participancy plays a constitutive role in bringing physical reality into being, with the delayed-choice experiment as the paradigmatic illustration [48], [49]. The present framework shares Wheeler’s emphasis on information as a fundamental organizing principle in physics, and it inherits Wheeler’s comfort with treating the universe as an informational whole subject to non-standard explanatory organization.

Distinction. Wheeler’s view is observer-centric in a strong sense: observer-participancy is the mechanism by which phenomena become real. The present framework is agency-centric in a structural sense: structural agency is the informational cluster selected by the endpoint, but it does not have constitutive or participatory causal power over the formalism. On Wheeler’s view, observation makes reality; on the present view, structural agency is what the admissible ensemble is selected for, but it does not make the ensemble. Wheeler’s view is also more radical ontologically; the present framework is deliberately minimalist about ontology.

8.2 Two-state vector formalism (TSVF)

Aharonov, Bergmann, Lebowitz, and Vaidman’s TSVF is the formal ancestor of much boundary-conditioned quantum mechanics, and Eqs. (2) and (4) are direct generalizations of the ABL rule [2], [10]. The present framework uses TSVF-style boundary-conditioned probability as its formal core.

Distinction. TSVF, as developed by its originators, is a time-symmetric formalism for laboratory quantum mechanics, where both pre- and post-selection are experimentally set. The present framework extends this to cosmology, where the post-selection is not an experimental choice but a proposed property of the universe as a whole, and the post-selection has *informational structure* (Eq. 5) rather than merely specifying a particular late-time measurement outcome. The TSVF is interpretively neutral about whether the post-selection has physical significance beyond its operational role; the present framework commits to a specific informational content for the endpoint.

8.3 Transactional interpretation

John Cramer’s transactional interpretation and its later elaboration by Ruth Kastner as the possibilist and relativistic transactional interpretation treat quantum events as handshakes between retarded “offer waves” and advanced “confirmation waves,” with time-symmetric causation built into the formalism [42], [43], [50].

Distinction. The transactional interpretation is a specific mechanism for wavefunction collapse: individual quantum events occur through the formation of transactions. The present framework is silent on the mechanism of individual collapse. It operates at a higher level: the ensemble of histories, not individual quantum events. Transactional interpretations can coexist with the present framework as an account of what happens at the level of individual events while the framework itself specifies how the history ensemble is pruned by a cosmological boundary condition.

8.4 Price’s time-symmetric program

Huw Price has argued, in work partly with Ken Wharton and others, that time-symmetry in quantum mechanics naturally leads to retrocausality once certain assumptions about ontology are granted [51], [52], [53]. The Price-Wharton-Leifer-Pusey literature has refined these arguments considerably.

Distinction. Price’s program is primarily concerned with locality, no-go theorems, and the foundations of quantum mechanics at the micro level; the question is whether retrocausality offers a natural escape from Bell-type non-locality. The present framework does not rely on retrocausation at the level of individual quantum events. It uses time-symmetry only as a property of the boundary-conditioned probability measure in Eq. (4), not as a commitment to signaling from future to past. In this sense, the present framework is *less* committed than Price’s program, and it could be adopted by physicists who are skeptical about micro-retrocausality.

8.5 Tegmark’s mathematical universe hypothesis

Max Tegmark’s mathematical universe hypothesis (MUH) proposes that physical reality is a mathematical structure, and that all computable mathematical structures exist physically in some sense, giving rise to a level-IV multiverse [54], [55].

Distinction. The MUH is an ontological claim about the nature of reality; it is maximally expansive. The present framework is a selective claim about the admissible-history ensemble of a single universe (or, if extended, a single cosmological ensemble), conditioned on an endpoint. These are not incompatible—one could hold a mathematical-structure ontology and still accept boundary-conditioned selection on the computable structures realized—but they are very different in ambition. The MUH says everything mathematically possible exists; the present framework says, within what exists, the admissible histories are those compatible with a structural-agency endpoint. The framework is agnostic about whether there is a multiverse, agnostic about whether reality is mathematical in the Tegmarkian sense, and committed only to the selection claim.

8.6 Smolin’s cosmological natural selection

Lee Smolin’s cosmological natural selection (CNS) proposes that black-hole formation in one universe seeds the formation of daughter universes with slightly varied constants, and that over many generations the distribution of universes is biased toward those producing many black holes [56], [57].

Distinction. CNS is a reproduction-with-variation account modeled on biological evolution, with black-hole fecundity as the fitness function. The present framework is not a reproduction account. It does not posit universe reproduction; it posits a single cosmological ensemble pruned by an endpoint condition. Interestingly, CNS has been critiqued on grounds that our universe appears better tuned for intelligent life than for black-hole production [58]; if one takes this critique seriously, one way to read the present framework is as a response that structural agency, not black-hole fecundity, is the relevant selection target. But the mechanism is different: there is no reproduction in the present framework, only endpoint-conditioned selection.

8.7 Penrose’s Weyl curvature hypothesis as initial-boundary complement

Of all the adjacent programs, Penrose’s Weyl curvature hypothesis (WCH) bears the most direct structural relationship to the present framework and deserves separate treatment [27], [28]. Penrose proposed that the initial singularity is characterized by vanishing (or near-vanishing) Weyl conformal curvature, $C_{abcd} \rightarrow 0$, which provides a low-entropy past boundary condition consistent with the observed smoothness and homogeneity of the early universe and with the thermodynamic arrow of time. The WCH is explicitly a *boundary condition*—an initial one—that constrains the admissible-history ensemble by fixing the gravitational entropy of the past.

The relationship to the present framework is direct and complementary. Where Penrose’s proposal constrains ρ_i in Eq. (4) by restricting the initial gravitational entropy, the present framework proposes a complementary constraint on ρ_f . The two proposals are not competitors; they address opposite temporal ends of the same boundary-conditioned formalism. A full cosmological theory in this style would include both: an initial condition specified in the spirit of the WCH (or its substitute in a quantum gravity theory) and a terminal condition in the spirit of a structurally agentic ρ_f^A .

Two observations make this pairing philosophically attractive. First, Penrose has estimated that the WCH requires a specification precision of order 1 part in $10^{10^{123}}$ for the observed universe [27], an astonishing figure that dramatizes how severe the initial-condition arbitrariness becomes when only initial data is available to do the explanatory work. Distributing some of the explanatory load to a non-trivial ρ_f does not eliminate this arbitrariness but does soften it: some features of the realized histories are no longer left entirely to initial-condition stipulation. Second, the WCH is itself not a theorem of the classical theory but a conjecture about what a successful quantum gravity theory will yield dynamically [28]. The parallel for the present framework is that ρ_f^A is also a conjecture, not a theorem: it proposes a boundary condition whose specific form would ideally be derived

from a deeper theory rather than stipulated, and whose structural content is motivated by observed features of the cosmic history we find ourselves in.

Distinction. Penrose’s WCH says nothing about whether the history ensemble is pruned toward a computational or agentic endpoint; it is concerned with the past. The present framework says nothing about the structure of the initial singularity; it is concerned with the future boundary. Taken together, they suggest a cosmological framework in which boundary conditions at both ends do coordinated explanatory work—a genuine two-boundary cosmology in the spirit Aharonov and Gruss explored for laboratory quantum mechanics [11].

8.8 Summary comparison

Table IV summarizes the adjacent programs across three dimensions: the temporal direction of the selection mechanism, the primary ontological commitment, and the empirical testability status.

Table IV. *Adjacent programs and the present framework.*

<i>Program</i>	<i>Temporal direction</i>	<i>Primary ontology</i>	<i>Testability</i>
Wheeler participatory	Present observation → past reality	Informational; observer-constitutive	Indirect (delayed-choice tests)
TSVF (Aharonov-Vaidman)	Pre- and post-selection	Time-symmetric formalism; ontology-neutral	Laboratory (weak values)
Transactional (Cramer, Kastner)	Retarded/advanced handshake	Real offer and confirmation waves	Interpretive; equivalent predictions
Price time-symmetric	Retrocausation at event level	Various; realist about λ	Indirect (Bell loopholes)
Tegmark MUH	None (timeless structure)	Mathematical monism	Contested; largely interpretive
Smolin CNS	Forward (evolutionary)	Multiverse with reproduction	Partial (neutron star mass predictions)
Penrose WCH	Initial boundary condition only	Low-entropy initial geometry	Indirect (CMB homogeneity; arrow of time)
Present framework	Endpoint-conditioned; single ensemble	Formalism-neutral; structural-agency endpoint	Soft differential predictions

The present framework occupies a distinctive position: it is time-symmetric in its formalism (inheriting from TSVF), single-universe in its ontological commitments (unlike MUH and CNS), complementary rather than competitive with WCH-style initial-boundary

proposals, agentic-endpoint in its selection target (where CNS uses black-hole fecundity and Wheeler uses observer-participancy), and structural in its characterization of agency (avoiding the anthropic vulnerability of observer-centric views).

9. Conclusion

This paper has advanced a specific and restricted thesis. The hot-to-cold thermal evolution of the observable universe is not itself selective—it is a thermodynamic consequence of metric expansion, forced by the Friedmann dynamics. But cooling establishes the intermediate regime in which a computationally organized, structurally agentic terminal condition can exert non-trivial pruning on the admissible-history ensemble. Quantum indeterminacy, under this boundary-conditioned reading, functions not as the ultimate explanatory terminus of the theory but as the permissive admissibility substrate through which endpoint-conditioned selection operates. The formalism is unchanged; the explanatory organization is.

What is genuinely new here, relative to the prior paper on final-state constraints and informational pruning [6], is not the final-state framework itself—that was established earlier—but the coupling between cosmological thermal history and endpoint-conditioned selection. The earlier paper established that computational terminal states prune the admissible ensemble; this paper establishes that cosmic cooling is the intermediate-regime condition through which such pruning becomes non-trivial. The cooling trajectory is neither a brute thermodynamic fact to be explained away nor a teleological direction imposed from outside. It is the pathway along which the selective content of the endpoint becomes physical.

Three commitments are worth restating at the close. First, the framework is interpretive, not revisionary: it does not modify the quantum formalism, does not introduce new terms into any equation of motion, and does not produce sharp empirical predictions that would distinguish it from standard theory at the laboratory level. Second, the endpoint is defined structurally—recursive self-modeling, goal-directed informational closure, self-replicating organization—without reference to human observers or carbon-based biology. This distinguishes the framework from anthropic reasoning. Third, the differential predictions offered in Section 7 are soft: expectation-value differences concerning error-correcting substrate prevalence, the fine-tuning profile of constants, and the relative timing of cosmogenic transitions. These predictions locate the framework in the interpretive tier of foundational programs rather than the direct-empirical tier.

Two lines of follow-on work are natural. The first is to formalize structural agency more precisely, perhaps by developing a principled measure on the space of informational architectures that instantiate the three features of recursive self-modeling, goal-directed closure, and self-replication. Integrated information theory's Φ measure, the free-energy-principle formulation of self-evidencing, and the autocatalytic-set formalism all offer partial candidates. The second is to sharpen the differential predictions, particularly the fine-tuning profile claim, by developing measures on cosmological parameter space that allow

quantitative comparison of the tuning of \mathcal{C}_1 versus \mathcal{C}_2 constants. Either line of work, if successful, would move the framework from soft-differential to sharp-differential predictions and strengthen its empirical standing.

A concise statement of the thesis, for reference: cosmic expansion-driven cooling is not itself selective, but it establishes the intermediate regime of durable chemistry, low-noise memory, and computational substrate within which a structurally agentic terminal condition can exert non-trivial pruning on the admissible-history ensemble; under this reading, quantum indeterminacy functions as the admissibility structure through which endpoint-conditioned selection operates, rather than as the ultimate explanatory terminus of the theory.

Use of Artificial Intelligence Tools

Large language model assistants (Anthropic Claude) were used during manuscript preparation for editorial tasks including prose refinement, citation cross-checking, structural review, and bibliographic formatting. All conceptual content, interpretive arguments, citation selections, and final wording are the author's own. The author reviewed and verified all AI-assisted output and takes full responsibility for the manuscript.

Data Availability Statement

This is a theoretical paper in the foundations of physics and cosmology; no new empirical data were generated or analyzed. All sources cited are publicly available through the references listed.

Funding

This research received no external funding.

Conflicts of Interest

The author declares no conflicts of interest. The author is employed by eCapital Corp in a strategy and architecture role unrelated to the subject matter of this paper; eCapital Corp had no role in the conception, preparation, or decision to publish this work.

References

- [1] S. Weinberg, *Cosmology*. Oxford, UK: Oxford University Press, 2008.
- [2] Y. Aharonov, P. G. Bergmann, and J. L. Lebowitz, "Time symmetry in the quantum process of measurement," *Physical Review*, vol. 134, no. 6B, pp. B1410–B1416, Jun. 1964, doi: <https://doi.org/10.1103/PhysRev.134.B1410>.
- [3] S. Weinberg, "Testing quantum mechanics," *Annals of Physics*, vol. 194, no. 2, pp. 336–386, 1989, doi: [https://doi.org/10.1016/0003-4916\(89\)90276-5](https://doi.org/10.1016/0003-4916(89)90276-5).
- [4] J. D. Barrow and F. J. Tipler, *The Anthropic Cosmological Principle*. Oxford, UK: Oxford University Press, 1986.
- [5] M. J. Rees, *Just Six Numbers: The Deep Forces That Shape the Universe*. New York, NY, USA: Basic Books, 2000.
- [6] I., "Final-state constraints and informational pruning in quantum histories," *International Journal of Quantum Foundations*, vol. 12, no. 2, p. 23, 2026. Available: <https://www.ijqf.org/>
- [7] R. B. Griffiths, "Consistent histories and the interpretation of quantum mechanics," *Journal of Statistical Physics*, vol. 36, no. 1-2, pp. 219–272, Jul. 1984, doi: <https://doi.org/10.1007/BF01015734>.
- [8] M. Gell-Mann and J. B. Hartle, "Classical equations for quantum systems," *Physical Review D*, vol. 47, no. 8, pp. 3345–3382, Apr. 1993, doi: <https://doi.org/10.1103/PhysRevD.47.3345>.
- [9] J. J. Halliwell, "A review of the decoherent histories approach to quantum mechanics," *Annals of the New York Academy of Sciences*, vol. 755, pp. 726–740, 1995, doi: <https://doi.org/10.1111/j.1749-6632.1995.tb39014.x>. Preprint: <https://arxiv.org/abs/gr-qc/9407040>.
- [10] Y. Aharonov and L. Vaidman, "The two-state vector formalism: An updated review," in *Time in Quantum Mechanics*, J. G. Muga, R. S. Mayato, and Í. L. Egusquiza, Eds. Berlin, Germany: Springer, 2008, pp. 399–447, doi: https://doi.org/10.1007/978-3-540-73473-4_13.
- [11] Y. Aharonov and E. Y. Gruss, "Two-time interpretation of quantum mechanics," arXiv preprint quant-ph/0507269, 2005. Available: <https://arxiv.org/abs/quant-ph/0507269>.
- [12] A. N. Kolmogorov, "Three approaches to the quantitative definition of information," *Problems of Information Transmission*, vol. 1, no. 1, pp. 1–7, 1965.
- [13] G. J. Chaitin, "On the length of programs for computing finite binary sequences," *Journal of the ACM*, vol. 13, no. 4, pp. 547–569, Oct. 1966, doi: <https://doi.org/10.1145/321356.321363>.

- [14] P. D. Grünwald and A. Vitányi, “Algorithmic information theory,” in *Handbook of the Philosophy of Science. Volume 8: Philosophy of Information*, P. Adriaans and J. van Benthem, Eds. Amsterdam, Netherlands: Elsevier, 2008, pp. 281–317. Preprint: <https://arxiv.org/abs/0809.2754>.
- [15] E. W. Kolb and M. S. Turner, *The Early Universe*. Redwood City, CA, USA: Addison-Wesley, 1990.
- [16] W. Hu and S. Dodelson, “Cosmic microwave background anisotropies,” *Annual Review of Astronomy and Astrophysics*, vol. 40, pp. 171–216, 2002, doi: <https://doi.org/10.1146/annurev.astro.40.060401.093926>.
- [17] F. C. Adams, “The degree of fine-tuning in our universe—and others,” *Physics Reports*, vol. 807, pp. 1–111, 2019, doi: <https://doi.org/10.1016/j.physrep.2019.02.001>.
- [18] R. Landauer, “Irreversibility and heat generation in the computing process,” *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183–191, Jul. 1961, doi: <https://doi.org/10.1147/rd.53.0183>.
- [19] C. H. Bennett, “Notes on Landauer’s principle, reversible computation, and Maxwell’s demon,” *Studies in History and Philosophy of Science Part B*, vol. 34, no. 3, pp. 501–510, Sep. 2003, doi: [https://doi.org/10.1016/S1355-2198\(03\)00039-X](https://doi.org/10.1016/S1355-2198(03)00039-X).
- [20] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, “Experimental verification of Landauer’s principle linking information and thermodynamics,” *Nature*, vol. 483, no. 7388, pp. 187–189, Mar. 2012, doi: <https://doi.org/10.1038/nature10872>.
- [21] S. Lloyd, “Ultimate physical limits to computation,” *Nature*, vol. 406, no. 6799, pp. 1047–1054, Aug. 2000, doi: <https://doi.org/10.1038/35023282>.
- [22] S. Lloyd, “Computational capacity of the universe,” *Physical Review Letters*, vol. 88, no. 23, p. 237901, May 2002, doi: <https://doi.org/10.1103/PhysRevLett.88.237901>.
- [23] J. S. Bell, “On the Einstein-Podolsky-Rosen paradox,” *Physics Physique Физика*, vol. 1, no. 3, pp. 195–200, 1964, doi: <https://doi.org/10.1103/PhysicsPhysiqueFizika.1.195>.
- [24] Y. Aharonov, D. Z. Albert, and L. Vaidman, “How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100,” *Physical Review Letters*, vol. 60, no. 14, pp. 1351–1354, Apr. 1988, doi: <https://doi.org/10.1103/PhysRevLett.60.1351>.
- [25] M. F. Pusey, J. Barrett, and T. Rudolph, “On the reality of the quantum state,” *Nature Physics*, vol. 8, no. 6, pp. 475–478, Jun. 2012, doi: <https://doi.org/10.1038/nphys2309>.
- [26] M. S. Leifer, “Is the quantum state real? An extended review of ψ -ontology theorems,” *Quanta*, vol. 3, no. 1, pp. 67–155, 2014, doi: <https://doi.org/10.12743/quanta.v3i1.22>.

- [27] R. Penrose, "Singularities and time-asymmetry," in *General Relativity: An Einstein Centenary Survey*, S. W. Hawking and W. Israel, Eds. Cambridge, UK: Cambridge University Press, 1979, pp. 581–638.
- [28] B.-L. Hu, "Weyl curvature hypothesis in light of quantum backreaction at cosmological singularities or bounces," *Universe*, vol. 7, no. 11, p. 424, Nov. 2021, doi: <https://doi.org/10.3390/universe7110424>.
- [29] K. Friston, "The free-energy principle: A unified brain theory?," *Nature Reviews Neuroscience*, vol. 11, no. 2, pp. 127–138, Feb. 2010, doi: <https://doi.org/10.1038/nrn2787>.
- [30] K. Friston, "Life as we know it," *Journal of the Royal Society Interface*, vol. 10, no. 86, p. 20130475, Sep. 2013, doi: <https://doi.org/10.1098/rsif.2013.0475>.
- [31] H. Maturana and F. Varela, *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht, Netherlands: D. Reidel, 1980.
- [32] A. Moreno and M. Mossio, *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Dordrecht, Netherlands: Springer, 2015, doi: <https://doi.org/10.1007/978-94-017-9837-2>.
- [33] S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*. New York, NY, USA: Oxford University Press, 1993.
- [34] M. Eigen and P. Schuster, *The Hypercycle: A Principle of Natural Self-Organization*. Berlin, Germany: Springer-Verlag, 1979, doi: <https://doi.org/10.1007/978-3-642-67247-7>.
- [35] J. L. England, "Statistical physics of self-replication," *Journal of Chemical Physics*, vol. 139, no. 12, p. 121923, Sep. 2013, doi: <https://doi.org/10.1063/1.4818538>.
- [36] G. Tononi, "Consciousness as integrated information: A provisional manifesto," *Biological Bulletin*, vol. 215, no. 3, pp. 216–242, Dec. 2008, doi: <https://doi.org/10.2307/25470707>.
- [37] G. Tononi, M. Boly, M. Massimini, and C. Koch, "Integrated information theory: From consciousness to its physical substrate," *Nature Reviews Neuroscience*, vol. 17, no. 7, pp. 450–461, Jul. 2016, doi: <https://doi.org/10.1038/nrn.2016.44>.
- [38] J. A. Edlund, N. Chaumont, A. Hintze, C. Koch, G. Tononi, and C. Adami, "Integrated information increases with fitness in the evolution of animats," *PLOS Computational Biology*, vol. 7, no. 10, p. e1002236, Oct. 2011, doi: <https://doi.org/10.1371/journal.pcbi.1002236>.
- [39] L. Albantakis, A. Hintze, C. Koch, C. Adami, and G. Tononi, "Evolution of integrated causal structures in animats exposed to environments of increasing complexity," *PLOS Computational Biology*, vol. 10, no. 12, p. e1003966, Dec. 2014, doi: <https://doi.org/10.1371/journal.pcbi.1003966>.
- [40] B. Carter, "Large number coincidences and the anthropic principle in cosmology," in *Confrontation of Cosmological Theories with Observational Data*, M. S. Longair, Ed.

Dordrecht, Netherlands: D. Reidel, 1974, pp. 291–298, doi: https://doi.org/10.1007/978-94-010-2220-0_25.

[41] N. Bostrom, *Anthropic Bias: Observation Selection Effects in Science and Philosophy*. New York, NY, USA: Routledge, 2002.

[42] J. G. Cramer, “The transactional interpretation of quantum mechanics,” *Reviews of Modern Physics*, vol. 58, no. 3, pp. 647–687, Jul. 1986, doi: <https://doi.org/10.1103/RevModPhys.58.647>.

[43] R. E. Kastner, *The Transactional Interpretation of Quantum Mechanics: The Reality of Possibility*. Cambridge, UK: Cambridge University Press, 2012, doi: <https://doi.org/10.1017/CBO9780511675768>.

[44] S. Weinberg, “Anthropic bound on the cosmological constant,” *Physical Review Letters*, vol. 59, no. 22, pp. 2607–2610, Nov. 1987, doi: <https://doi.org/10.1103/PhysRevLett.59.2607>.

[45] H. Oberhummer, A. Csótó, and H. Schlattl, “Stellar production rates of carbon and its abundance in the universe,” *Science*, vol. 289, no. 5476, pp. 88–90, Jul. 2000, doi: <https://doi.org/10.1126/science.289.5476.88>.

[46] E. Epelbaum, H. Krebs, T. A. Lähde, D. Lee, and U.-G. Meißner, “Viability of carbon-based life as a function of the light quark mass,” *Physical Review Letters*, vol. 110, no. 11, p. 112502, Mar. 2013, doi: <https://doi.org/10.1103/PhysRevLett.110.112502>.

[47] T. Damour and J. F. Donoghue, “Constraints on the variability of quark masses from nuclear binding,” *Physical Review D*, vol. 78, no. 1, p. 014014, Jul. 2008, doi: <https://doi.org/10.1103/PhysRevD.78.014014>.

[48] J. A. Wheeler, “Information, physics, quantum: The search for links,” in *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics*, Tokyo, Japan, 1989, pp. 354–368.

[49] J. A. Wheeler, “Law without law,” in *Quantum Theory and Measurement*, J. A. Wheeler and W. H. Zurek, Eds. Princeton, NJ, USA: Princeton University Press, 1983, pp. 182–213.

[50] R. E. Kastner, “The possibilist transactional interpretation and relativity,” *Foundations of Physics*, vol. 42, no. 8, pp. 1094–1113, Aug. 2012, doi: <https://doi.org/10.1007/s10701-012-9658-4>.

[51] H. Price, *Time’s Arrow and Archimedes’ Point: New Directions for the Physics of Time*. New York, NY, USA: Oxford University Press, 1996.

[52] H. Price, “Does time-symmetry imply retrocausality? How the quantum world says ‘maybe’?,” *Studies in History and Philosophy of Modern Physics*, vol. 43, no. 2, pp. 75–83, May 2012, doi: <https://doi.org/10.1016/j.shpsb.2011.12.003>.

[53] M. S. Leifer and M. F. Pusey, "Is a time symmetric interpretation of quantum theory possible without retrocausality?," *Proceedings of the Royal Society A*, vol. 473, no. 2202, p. 20160607, Jun. 2017, doi: <https://doi.org/10.1098/rspa.2016.0607>.

[54] M. Tegmark, "The mathematical universe," *Foundations of Physics*, vol. 38, no. 2, pp. 101–150, Feb. 2008, doi: <https://doi.org/10.1007/s10701-007-9186-9>.

[55] M. Tegmark, *Our Mathematical Universe: My Quest for the Ultimate Nature of Reality*. New York, NY, USA: Knopf, 2014.

[56] L. Smolin, *The Life of the Cosmos*. New York, NY, USA: Oxford University Press, 1997.

[57] L. Smolin, "The status of cosmological natural selection," arXiv preprint hep-th/0612185, Dec. 2006. Available: <https://arxiv.org/abs/hep-th/0612185>.

[58] H. Price, "Cosmology, time's arrow, and that old double standard," in *Time's Arrows Today: Recent Physical and Philosophical Work on the Direction of Time*, S. F. Savitt, Ed. Cambridge, UK: Cambridge University Press, 1995, pp. 66–94, doi: <https://doi.org/10.1017/CB09780511622861.005>.