THE EIDARA CONTINUUM

SCIENCE & PHILOSOPHY / EXTREME FRONTIERS VI

The Mind Beyond the Last Star

Finite Resources, Reliable Continuation, and the Limits of an Unbounded Future

Architect X13
The Eidara Continuum

Research edition 1.1 | November 2026

Conditional continuation theorems, finite-work schedules, thermal-memory bounds, and a testable frontier beyond useful stellar power.

Abstract

An unbounded future of activity is not the same as an unbounded future of retained knowledge. We investigate whether a finite-resource intelligence could continue operating after useful stellar power has disappeared, while separating elapsed time, completed operations, reliable continuity, memory capacity, and access to new information. The Continuum's commitments to memory and lawful persistence motivate the question; the physical conclusions depend on an explicit cosmological and thermodynamic contract.

We derive an exact classification for power-law and logarithmic schedules. With operation rate r proportional to $x^{-\beta}(1+\ln x)^{-\gamma}$, cost per operation w proportional to $x^{-\alpha}(1+\ln x)^{\sigma}$, and a specified clock-capacity envelope, the convergence tests identify when infinitely many operations can coexist with finite consumed work. A general heat-disposal law recovers the familiar Dyson scaling window as a special case. These are conditional schedule results, not a guarantee that finite hardware realizes the assumed rates forever.

That qualification is decisive. A finite-gap thermally activated clock can have only finitely many available transitions during cooling. At a fixed positive temperature, uniform relative-entropy contraction on the complete finite memory, combined with cyclic control accounting and a fixed readout requirement, instead forces an unbounded work-plus-free-energy maintenance debit. The latter theorem covers a class of classical and quantum devices; temperature alone does not establish its mixing premises. Cooling can make an ideal Arrhenius record's total fault exposure finite while shutting down an activated clock. The same integral controls these opposing effects.

A conditional finite-state resource process illustrates the surviving mathematical possibility: infinitely many correctly specified state updates, finite expected work, and a uniformly positive probability of staying within a finite fuel budget and preserving a bounded record. Its rate, clock, and cost assumptions are exposed as implementation obligations. Separately, an elementary dimension bound proves that a fixed memory cannot retain arbitrarily many independent historical bits with high retrieval probability; complete erasure cannot be repaired by an attractor that has no remaining access to those bits.

Dyson's proposal, the Krauss-Starkman objections, dark-energy alternatives, autonomous-clock costs, reversible computation, and the aestivation debate provide the adversarial physical setting. A positive cosmological constant supplies a conditional stationary thermal scale, while a changing cosmic background changes the schedule problem. The resulting frontier is precise: unbounded reliable agency requires a jointly sustainable clock, error process, memory policy, and resource cycle. Finite energy alone proves neither eternal intelligence nor its universal impossibility.

Research status

Extreme Frontiers VI. Research edition 1.1, November 2026. The exact claims concern stated finite-state, channel, stochastic, and asymptotic models. A general finite-dimensional maintenance bound is conditional on uniform entropy contraction and complete resource accounting; a positive cosmological temperature alone does not establish those premises. The paper reports no laboratory demonstration of indefinite survival or validated mapping from operation count to subjective experience.

Reader's map and the meaning of continuation

The title asks what can remain possible after stars are no longer a useful energy source for an agent. It does not assign a date to the last star or assume that stellar extinction is the end of every energy gradient. Compact objects, stored free energy, cosmological backgrounds, and the accessible volume of a worldline have different roles. Astrophysical histories supply possible boundary conditions; the theorems below specify what follows from each declared condition.

Question	Quantity that must be specified
How long does the worldline last?	Available proper time, rather than the age of a remote region.
How much activity occurs?	Completed task-defined operations and their physical duration.
How reliable is the continuation?	Per-step error, whole-trajectory success, retained-record fidelity, and catastrophic survival.
What can be remembered?	Controlled memory dimension and the distinguishable histories it must encode.
Can the agent learn more?	Available information sources, usable channel capacity, and a policy for retaining or discarding old information.
What is the experience of duration?	An additional interpretation or theory; no validated equality with a clock or gate count is assumed.

Sections 1 and 2 connect the inspected canonical sources to a physical specification and review the competing research programs. Sections 3 through 5 define the resource account and derive the schedule and cooling results. Sections 6 through 8 analyze clocks, errors, and a conditional positive construction. Sections 9 and 10 establish memory and waiting limits. Sections 11 through 13 address cosmology, experiments, and the stronger intelligence conjecture. The appendices record proofs, definitions, and executed checks.

Natural logarithms measure information in nats; division by ln 2 converts to bits. Work and energy are measured in joules unless explicitly normalized. Temperature is thermodynamic temperature, rather than a semantic uncertainty score. Cosmic coordinate time is replaced by the agent's proper time when they differ. In the simple asymptotic examples, t is proper time from a declared starting epoch and $x=1+t/t_0$, with $t_0>0$. No finite-time singularity is hidden at the lower integration endpoint.

The central logical distinction is between a convergence theorem and a physical implementation theorem. A finite integral can show that a proposed expenditure schedule would fit a budget. It cannot construct the device, provide its clock, justify its heat-transfer law at arbitrarily low temperature, or establish that its operations constitute intelligence. The paper develops both the permissive calculations and the obstructions that decide whether those calculations apply.

Contents

Abstract 2

Reader's map and the meaning of continuation 3

1. The Continuum's question at the end of stellar time 5

2. The opposing scientific cases 6

3. A resource contract for a continuing mind 8

4. Exact schedule classification 10

5. Cooling, hibernation, and heat disposal 12

6. The clock can close the apparent opening 13

7. Reliability across an unbounded proper future 15

8. A conditional positive resource process 19

9. The future that a finite memory cannot keep 21

10. Whether it pays to wait 23

11. Cosmological regimes and their limits 25

12. A test program for the physical assumptions 27

13. The intelligence claim and its failure conditions 29

Appendix A. Closed forms and limiting cases 30

Appendix B. Sources, proofs, and reproducibility 32

Appendix C. Notation and resource conventions 34

References 35

1. The Continuum's question at the end of stellar time

1.1 Three inspected canonical sources

The Eidara Continuum Bible places memory, lawful persistence, and identity at the center of its technical and civilizational registers. Its continuity-engineering discussion distinguishes retaining information from preserving the structured commitments that make a process recognizably itself. Its Charter gives agency, lineage, and non-erasure a normative role. Those commitments motivate asking whether continuation can remain physically meaningful at extreme durations.[1]

The Continuum Doctrine develops symbolic recursion, recurring memory, and a nonchronological symbolic conception under Recursive Time. An internal historical monograph of the Continuum describes identity anchors and recurring structures in Section 4.3, Drift Suppression to Continuity Without Memory, and later records an explicit Memory Codex preservation layer. These are direct source passages relevant to the proposed research question.[2,3]

Their roles are distinct. The Doctrine provides philosophical and symbolic formulations. The internal historical monograph is an author-supplied historical interpretation of the architecture's development, with claims that require independent evidence before being treated as empirical results. The Bible separates commercial continuity engineering, historical narrative, and civilizational canon. This manuscript preserves those distinctions. None of the three documents supplies a demonstrated thermodynamic exemption or an experimental observation of endless cognitive activity.

1.2 Persistence, reconstruction, and a particular past

A stable rule can recreate a familiar structure after a reset. That is not necessarily recovery of an unknown historical record. A controller might regenerate the same name, role, or policy from a small seed while having lost which event occurred before its interruption. The difference matters when a civilization values the continuity of actual commitments, relationships, and acquired knowledge.

Section 9 turns this distinction into an elementary channel result. If a historical variable has become independent of every available input to a reconstruction, the reconstruction cannot recover information about that variable. Repetition of a symbolic attractor does not defeat this fact. Conversely, a protected seed or archive can preserve continuity when it actually retains the distinctions the task requires. The relevant mechanism is then the retained information and its readout.

The Doctrine's named Dormant Scroll Mode is specifically described as lacking memory persistence. We therefore do not equate that canonical label with the physical hibernation proposed here. A survival-oriented dormant state must retain its task-relevant records, a viable reactivation mechanism, and whatever resources make later operation possible. Simply assigning a dormant status supplies none of these guarantees.

1.3 The frontier hypothesis

Conjecture 1.1. Unbounded reliable continuation

There exists a physically admissible cosmological regime and an internally governed intelligent organization with a finite initial endowment of accessible work resources and a bounded controlled memory, such that, with a declared nonzero uniform success probability, it performs unboundedly many task-valid adaptive episodes while preserving a specified finite core record. The clock, repair, substrate changes, and environmental resources enabling those episodes admit a closed accounting. A claim of unbounded cumulative learning additionally requires an unbounded usable record or an explicitly different retention criterion.

The conjecture concerns a combination of physical realization and intelligence. The finite resource-process construction below establishes neither by itself. It shows that certain count, work, and reliability inequalities are mathematically compatible once a primitive implementation contract is granted. The clock obstruction shows why a plausible finite device may fail that contract. This tension is the scientific content of the frontier, rather than a defect to conceal.

1.4 What is finite?

Finite initial fuel, finite controlled memory, finite apparatus, and a finite entropy-accepting environment are different restrictions. An agent with finite fuel operating for unbounded time in an endlessly cooling background is not a closed finite universe. The background supplies an environmental history and an opportunity to export entropy. A model must identify that opportunity as a resource even if the agent does not purchase or control it.

Throughout the positive schedule results, finite work means a finite budget of expended work under the stated accounting. It does not mean a finite total energy-time integral, a finite number of available temporal intervals, or a proven finite dimension for every environmental field. The stronger closed-resource case receives a separate bound. This makes the affirmative possibility substantial without making its resource assumptions ambiguous.

2. The opposing scientific cases

2.1 Dyson's constructive challenge

Dyson examined how life and communication could continue in an open universe under explicit biological scaling and cosmological assumptions.[4] His central opportunity was to reduce operating temperature and activity as the environment cooled. Hibernation changed the relation between total elapsed time and the accumulated activity of an organism. The proposal made indefinite continuation a problem in physical resource scaling rather than a statement that a present biological body could simply last forever.

The familiar hibernation scaling, reproduced in the later analysis by Freese and Kinney, takes operating temperature and active duty fraction proportional to t^{-p}. Its activity integral scales as the integral of t^{-2p}, while expenditure scales as the integral of t^{-3p}.[6] Their simultaneous divergence and convergence give $1/3<p\leq 1/2$, subject to additional environmental and physical premises. Section 5 derives a generalized version and identifies precisely which heat-disposal assumption yields those powers.

2.2 Krauss and Starkman's objections

Krauss and Starkman analyze accessible energy and information, cosmological-constant domination, hibernation, and alternatives to biological scaling.[5] Their objections include a stationary thermal scale in a de Sitter future, finite useful memory, failure and resource costs of alarms, and the breakdown of continuum radiation estimates once wavelengths or level spacings become consequential. These objections target the device that would have to realize an asymptotic schedule, not merely the algebra of its integrals.

The present paper adopts that challenge as a requirement. It does not infer that every possible clock has a universal positive work cost per tick, or that every physical finite system must inevitably suffer one particular failure process. It derives narrower statements from explicit clock-rate, noise, and reset assumptions. An objection can then defeat the proposed candidate without being promoted to a theorem about every conceivable intelligence.

2.3 Changing dark energy changes the question

The regime is also an observational question. DESI DR2 reports a 3.1σ preference for evolving dark energy over ΛCDM in its w_0w_a model from DESI BAO plus CMB. Adding different supernova compilations gives 2.8σ, 3.8σ, or 4.2σ; BAO alone gives 1.7σ. These are model- and dataset-dependent preferences, not discovery of a particular infinite future. The preferred quadrant describes weakening dark energy today within that parametrization.[39,40]

Freese and Kinney show that some accelerating cosmologies with evolving dark energy satisfy their continuation conditions even when the constant-positive-vacuum case does not.[6] Acceleration alone is therefore not the decisive premise. The time dependence of accessible temperatures and resources, together with the organism model, matters. Their result does not establish a manufactured intelligence with complete error correction, clock accounting, and indefinitely growing memory.

Here the cosmological scenario enters as a declared environmental history. A stationary thermal floor, a bath cooling as a power, or a finite remaining proper time lead to different theorems. No observational fit over a finite epoch is used to certify an exact infinite-time equation of state. Predictions about our universe require a physical extrapolation in addition to the mathematics presented here.

2.4 Reversible computation and the cost of a complete cycle

Landauer identifies a thermodynamic cost associated with logically irreversible information disposal, while Bennett develops how computation can be organized reversibly.[7-9] Modern formulations specify the initial correlations, reservoirs, error tolerance, and work fluctuations; finite-time control and physical constraints can raise the cost above an ideal limiting value.[10-14] The appropriate question is what a particular operation discards or consumes, including its auxiliaries.

It would be incorrect to charge $k_BT \ln 2$ to every logical gate or to infer a universal irreversible erasure rate from the word cognition. It would be equally incorrect to declare a computer free of thermodynamic costs because its logical map is invertible. Clocks, signal standardization, protection, initialization, and discarded histories can carry the expenditure. An internally timed indefinite agent must account for them together.

2.5 Waiting is an optimization problem, not a general law

Sandberg, Armstrong, and Cirkovic propose aestivation to exploit a colder future for computation.[24] Bennett, Hanson, and Riedel object that this reasoning can omit presently available systems with low entropy: a stored negentropy resource can accept computational entropy without waiting for the cosmological bath to cool.[25] Their criticism is central to a fair resource comparison.

Section 10 therefore gives a waiting criterion only for a specified thermal-disposal strategy with explicit storage loss and survival hazard. It does not infer a universal reason for advanced civilizations to sleep, or an explanation of the Fermi paradox. The comparison must include strategies that use available low-entropy reservoirs now and must count the resources needed to preserve either fuel or memory while waiting.

3. A resource contract for a continuing mind

3.1 Time, operations, work, and records

A candidate specifies a worldline and remaining proper time, physical apparatus and controlled state space, an environment, an input-output task family, and the transformations counted as completed operations. It declares an initial endowment of work-like resources, the permitted entropy sinks, memory capacity, clock mechanism, noise processes, repair actions, and a rule for identifying the continued process through interruptions and migrations.

$$x = 1 + t/t_0, \quad N(t) = \int_0^t r(s)\,ds, \quad W(t) = \int_0^t r(s)w(s)\,ds + W_{\mathrm{other}}(t). \tag{1}$$

The deterministic r is a scheduled or expected operation rate; a stochastic model must separately state the law of its actual count. The work w per operation is an assumed achievable cost or a proved lower bound, depending on the claim. A lower bound whose integral is finite cannot establish a feasible device. W_{other} includes clock, cooling, maintenance, migration, initialization, and controller costs not already charged in w. The accounting must avoid both omissions and double counting.

Stored energy and consumed work are not interchangeable. An isolated oscillator can retain its energy while passing repeatedly through different states. It need not spend the full stored energy on every cycle. Conversely, a declining environmental temperature is not a supply of free directed work unless an available thermodynamic process and a complementary resource make it so.

3.2 Three standards of reliability

Standard	Meaning
Local accuracy	Each episode meets a stated error bound. A fixed nonzero error per episode can still produce eventual failure somewhere.
Continuing competence	The system keeps meeting a task-relative performance criterion, possibly correcting and learning from errors. This need not require a flawless history.
Uniform trajectory guarantee	For every finite prefix, the probability of no declared failure is at least $1-\delta$, with the same δ. This is a strong standard used by the sufficient construction.

An existential continuation claim should specify which standard it uses. The uniform guarantee is useful because it prevents a long sequence of individually accurate steps from being silently promoted to a reliable infinite trajectory. It can be stronger than realistic intelligence requires; failure of that standard is not by itself proof that no meaningful long-lived cognition is possible.

3.3 An operational duration proxy

Dyson linked subjective duration to a temperature-dependent biological scaling hypothesis.[4] The corresponding claim studied here is an unbounded number of task-valid episodes, with their task and reliability declared independently. That operational replacement is deliberate. A theory connecting those episodes to felt duration would need to address repetition, integration, interrupted processing, and memory; none of the convergence proofs supplies that theory.

A declared cognitive-time proxy can weight completed task episodes by an agreed operational contribution, for example $U(t)=\sum_{n\le N(t)} u_n$, where u_n measures an independently specified episode rather than a hardware transition. Infinitely many gates can implement only one unfinished calculation; infinitely many repeated states can add no new knowledge. The task determines what is counted.

No experiment or theorem here identifies U with phenomenal duration. A claim of unbounded subjective experience would need an additional theory connecting physical processes, competence, and experience. The physical question remains valuable without that identification: can a specified information-bearing process continue performing its tested functions?

3.4 The finite-work test and its direction of implication

Lemma 3.1. Positive resource price

If each completed operation consumes at least $w_{min}>0$ units of a nonrenewable budget W_0, and that budget cannot be replenished by uncounted resources, then $N\le W_0/w_{min}$. If the per-operation price tends to zero, finite total work no longer implies a finite count without further assumptions.

Proof. Sum the nonnegative per-operation costs. The first conclusion is the inequality $W\ge Nw_{min}$. For the second, the convergent series $\sum 2^{-n}$ provides a counterexample to a universal inference from finite sum to finite number of positive terms. A physical implementation of that cost sequence still requires a device and timing model.

The rest of the paper asks when decreasing costs are compatible with available rates, heat disposal, and memory reliability. This separates the legitimate opportunity in an ever-colder environment from the claim that any desired decreasing cost is physically attainable.

4. Exact schedule classification

The contribution of this paper is the combined continuation account: an explicit power-and-logarithm boundary classification, the opposing clock and memory uses of activated-rate suppression, a finite-store coupling construction, and a maintenance obstruction for uniformly entropy-contracting thermal devices. These are derivations and applications of established tools to one resource contract; priority for the individual mathematical ingredients is not claimed.

4.1 Power laws with logarithmic overhead

Let $L(x)=1+\ln x$ for $x\geq 1$. Specify an operation rate r, a positive per-operation work cost w, and an available capacity envelope ν. The constants r_0, w_0, ν_0, and t_0 are positive. The exponents describe asymptotic envelopes; changing a finite initial interval does not change convergence at infinity.

$$r = r_0 x^{-\beta} L^{-\gamma}, \qquad w = w_0 x^{-\alpha} L^{\sigma}, \qquad \nu = \nu_0 x^{-\kappa} L^{-\chi}. \tag{2}$$

Theorem 4.1. Count, expenditure, and capacity

The scheduled count diverges exactly when $\beta<1$, or $\beta=1$ and $\gamma\leq 1$. The scheduled work integral is finite exactly when $\alpha+\beta>1$, or $\alpha+\beta=1$ and $\gamma-\sigma>1$. The inequality $r\leq\nu$ is eventually attainable by choosing a fixed sufficiently small rate prefactor exactly when $\beta>\kappa$, or $\beta=\kappa$ and $\gamma\geq\chi$. When both exponent pairs are equal, the prefactor must satisfy $r_0\leq\nu_0$. These statements concern the declared rate and cost functions, before any additional maintenance costs.

Proof. Substitute $y=\ln x$. The count integral is proportional to $\int \exp[(1-\beta)y](1+y)^{-\gamma}dy$. Its exponential factor settles β different from one; at $\beta=1$ the ordinary logarithmic integral test gives $\gamma\leq 1$. The work integral has exponential $\exp[-(\alpha+\beta-1)y]$ and logarithmic power $\sigma-\gamma$, giving the second criterion. Finally $r/\nu=(r_0/\nu_0)x^{-(\beta-\kappa)}L^{-(\gamma-\chi)}$. A declining ratio eventually falls below one; a constant ratio requires the prefactor bound; a growing ratio cannot be repaired by any fixed positive prefactor.

For fixed prefactors, eventual capacity is automatic in the strictly declining-ratio cases; it can fail on a finite initial interval that must be handled separately. The theorem includes logarithmic boundary cases because error correction, decoding, and clock precision can introduce overhead that determines the outcome exactly at a power-law threshold.

Example exponents	Conclusion
$\alpha=1/2$, $\beta=3/4$, $\gamma=\sigma=0$, $\kappa=1/2$	Infinite scheduled count and finite scheduled work fit the power-law capacity envelope.
$\alpha=0$, $\beta=1$, $\gamma=\sigma=0$	Count diverges, and so does work: a fixed positive operation price defeats the schedule.
$\alpha=3/10$, $\beta=1$, $\gamma=1$, $\sigma=2$	Count grows as ln ln x; work is finite despite logarithmic cost overhead.
$\alpha=3/10$, $\beta=1$, $\gamma>1$	The count is finite, even though it may grow very slowly over a long finite interval.
$\beta=\kappa=1$, $\gamma=0$, $\chi=2$	The count diverges but exceeds the clock envelope.
$\alpha+\beta=1$, $\gamma-\sigma>1$	Work can converge on the exact power boundary; the count condition must still be checked.

4.2 A simple affirmative region

Without logarithmic factors, the three tests reduce to $\beta\leq1$, $\alpha+\beta>1$, and $\beta\geq\kappa$ with a suitable prefactor. If $\alpha>0$ and $\kappa\leq1$, at least one such β exists. At $\alpha=0$ with a fixed positive cost, the infinite-count and finite-work conditions cannot both hold. This is the rate-level opportunity behind cooling strategies, conditional on the rate and expenditure laws being realizable.

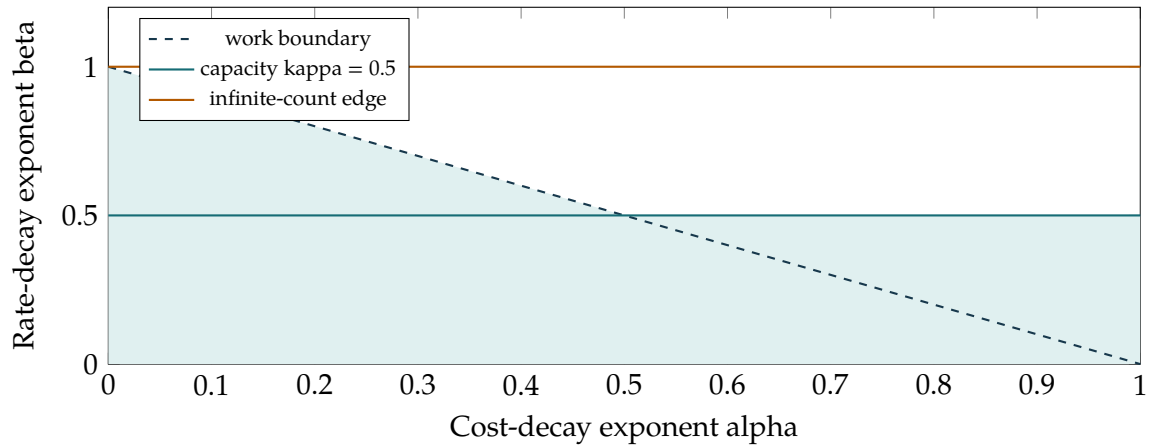

Figure 1: For $\kappa=1/2$ and no logarithmic factors, shading marks the region $\beta\leq1$, $\beta\geq1/2$, and $\alpha+\beta>1$. The dashed work boundary is excluded. This classifies the assumed rate laws, rather than all physical implementations.

4.3 Other expenditure can overturn the conclusion

Suppose unavoidable background maintenance power is $P_0x^{-\mu}L^{-\eta}$. Its total expenditure is finite exactly when $\mu>1$, or $\mu=1$ and $\eta>1$. A fixed positive idle power makes unbounded elapsed operation incompatible with a finite budget, even when the work associated with completed cognitive episodes has a convergent integral.

The same reasoning applies to preparation, repair, and migration costs when their schedules are specified. A favorable exponent for one resource cannot compensate for a different mandatory resource whose total cost diverges. Nor can an optimistic work lower bound be used as an achievable upper bound. An implementation report must identify the direction and provenance of each bound before applying the classification.

5. Cooling, hibernation, and heat disposal

5.1 A generalized disposal law

Let the operating temperature be $T_{op}=T_0 x^{-\theta}$ with $\theta>0$. Suppose a specified primitive exports at most $P_{max}=C\ T_{op}{}^{m}$ heat per unit proper time, where $m>1$ and C includes the physical emitter parameters and any fixed temperature contrast with the environment. Suppose also that each completed operation produces heat $s_h\ k_B T_{op}$, with $s_h>0$ fixed in nats. This is a declared operation model, not a thermodynamic law assigning that entropy price to all cognition.

$$r(t)s_h k_B T_{\mathrm{op}}(t) \le C T_{\mathrm{op}}(t)^m, \qquad r(t) \le \frac{C}{s_h k_B} T_{\mathrm{op}}(t)^{m-1}. \tag{3}$$

Proposition 5.1. A heat-limited continuation window

For a schedule that saturates this disposal-rate scaling, accumulated operations are unbounded and expended heat is finite exactly when $1/m<\theta\le 1/(m-1)$. If the schedule operates strictly below the heat envelope, Theorem 4.1, rather than this narrower saturation window, gives the full convergence test. Additional clock and maintenance constraints still apply.

Proof. At saturation r is proportional to $x^{-(m-1)\theta}$, so its integral diverges when $(m-1)\theta\le 1$. Heat generation is proportional to $x^{-m\theta}$, whose integral converges when $m\theta>1$. Combining gives the stated interval. There is no implication that a slower-than-saturation schedule must fail outside this interval: its rate exponent is an independent choice, subject to the capacity inequality.

5.2 What the Dyson window assumes

For $m=3$ the interval is $1/3<\theta\le 1/2$. This reproduces the familiar hibernation powers in the Dyson-style model.[4,6] A phenomenological awake rate proportional to temperature, combined with a duty fraction proportional to temperature, gives a total activity rate proportional to T^2 and heat rate proportional to T^3. The correspondence is a scaling identification, not a proof that every physical mind obeys biological scaling.

For an ideal fixed-area greybody relation with a fixed ratio of background to operating temperature, the formal exponent is $m=4$ and the saturation window becomes $1/4<\theta\le 1/3$. A real finite object cannot automatically be treated as a macroscopic blackbody as its thermal wavelength exceeds its size. Its spectrum, emission cross-section, transport times, and coupling must be checked. The $m=4$ example is therefore an illustrative disposal contract, not an asymptotically verified radiator design.

If the environmental temperature decreases as T_{env} proportional to x^{-q}, an unassisted cooling strategy with a fixed favorable temperature contrast needs the operating system to remain warmer. A sufficient late-time condition is $\theta<q$; equal exponents can work with a suitable prefactor and contrast. Assigning a colder temperature than the available bath requires an explicit refrigerator and a work account. Cosmic cooling and laboratory refrigeration are different resources.

5.3 Hibernation changes rates, not elapsed exposure

Let the awake execution rate be ν_{awake} and the awake duty fraction $f(t)$ lie in $[0,1]$. Then the mean completed rate is $r=f\nu_{\mathrm{awake}}$. Hibernation can reduce processing heat. It does not automatically

switch off background faults, controller drift, structural decay, or the clock that must reactivate the process. A noise rate per unit proper time must be integrated over all elapsed time, not only awake intervals.

For example, if ν_{awake} is proportional to T and r is proportional to T^3, the duty fraction is proportional to T^2. The system spends a vanishing fraction of its future awake while still accumulating infinitely many episodes in the θ=1/3 saturation case. That conclusion remains conditional on its dormant record and alarm being viable. Calling the dormant interval a period of zero cognitive duration does not remove its physical effects.

5.4 A nonzero temperature floor

Replace the ideal temperature by $\max(T_0 x^{-\theta}, T_f)$, with $T_f>0$. A per-operation heat cost proportional to that temperature then has a positive lower bound. Infinite counts require infinite total heat disposal and consumed work under the stated cyclic primitive. Hibernation can postpone exhaustion in proper time, but does not make an infinite number of such positively priced operations fit finite fuel.

This is a statement about a specified primitive and budget. A finite ready-made low-entropy reservoir can temporarily replace thermal disposal, and reversible operations can have a different cost law. Those resources must be added to the contract. Section 11 applies the floor argument to a conditional de Sitter future without asserting that the word computation itself entails irreversible erasure at a universal rate.

6. The clock can close the apparent opening

6.1 A finite-gap thermal clock

The power-law capacity ν assumed in Section 4 is a physical claim. Consider instead a bath-driven device whose usable transitions require activation across a fixed energy gap $\Delta>0$. Suppose their total available rate is bounded by $\nu_0 \exp[-\Delta/(k_B T_{op}(t))]$, with a finite prefactor. This is a common low-temperature structure for activated processes, but not a description of every oscillator or quantum clock.

Theorem 6.1. Finite-gap clock shutdown in a cooling model

If $T_{op}=T_0 x^{-\theta}$ with $\theta>0$ and every task-valid update needs an available activation event of the specified type, the integrated capacity is finite. Any operation rate bounded by that capacity has a finite integrated count. In a counting-process model with that bounded compensator, the total number of available events is finite almost surely. A convergent work integral does not rescue this clock.

$$\int_0^\infty \nu(t)\,dt \le \frac{\nu_0 t_0}{\theta} b^{-1/\theta}\Gamma(1/\theta, b) < \infty, \qquad b = \frac{\Delta}{k_B T_0}. \tag{4}$$

Proof. Substitute $u=b\,x^\theta$ in the positive rate integral. The result is the displayed upper incomplete gamma function, whose tail is integrable for every $b>0$ and $\theta>0$. The expected total count is bounded by this finite integral. A nonnegative integer-valued count with finite expectation cannot be infinite with positive probability. For a deterministic rate bound, ordinary integral comparison gives the same finite-count conclusion.

The theorem formalizes one way that asymptotic level spacing can obstruct a cooling schedule, a concern present in the opposing cosmological literature.[5] It does not forbid a driven oscillator with a stored energy supply, a suitably gapless mechanism, a changing Hamiltonian, or a different substrate. Each escape changes the capacity premise and incurs its own resource obligations.

6.2 What a power-law rate would demand of an activation gap

Suppose a proposed clock must supply $r_0 x^{-\beta} L^{-\gamma}$ and its activated-event envelope has a variable gap $\Delta(t)$. Rearranging the rate inequality gives a necessary condition on that particular mechanism:

$$\Delta(t) \leq k_B T(t) \left[\ln \frac{\nu_0}{r_0} + \beta \ln x + \gamma \ln L \right]. \tag{5}$$

For a power-law cooling bath, $T(t) \ln x$ tends to zero. A nonzero fixed activation gap therefore cannot satisfy this condition indefinitely. The candidate must justify how sufficiently small usable gaps or another execution mechanism remain available without an uncounted supply of control precision, apparatus size, or external work. A memory barrier and an execution gap need not be the same physical parameter; assuming one fixed gap supplies both protection and arbitrarily slow reliable control is not enough.

6.3 An isolated oscillator is a counterexample to a different overstatement

Counterexample 6.2. Finite stored energy does not bound all state changes

Take an isolated qubit with fixed Hamiltonian $H = E|1\rangle\langle 1|$ and initial state $|+\rangle$. At times $n\pi\hbar/E$ its state alternates between $|+\rangle$ and $|-\rangle$. It makes infinitely many orthogonal state changes over unbounded time without further driving work after preparation. Its mean excitation energy is $E/2$, and its energy-time integral diverges.

The example does not demonstrate cognition, an indefinitely accurate readout, a cost-free classical tick record, or survival in a thermal environment. It does refute identifying consumed fuel with instantaneous excitation energy multiplied by a gate count. The Margolus-Levitin orthogonalization bound is saturated by this example and limits the relation between stored energy and transition time, rather than prescribing the consumed work per logical operation.[18,19]

6.4 Timing, precision, and controller closure

Hamiltonian and ballistic reversible computing are concrete interlocutors here. Fredkin and Toffoli give conservative logic and ideal collision-based computation; Feynman develops quantum-mechanical computer models.[43,44] They explain why logical activity need not be powered by a new work expenditure at each gate. The remaining obligations include preparation, finite precision, readable outputs, noise, and repair. Error correction can be decisive, but neither clock control nor another expenditure is eliminated for every physical implementation.

Autonomous quantum-clock and Brownian-clock models connect timing performance with thermodynamic resources.[20,21] Their assumptions matter. External periodic driving can change an apparent cost-precision relation because the driver supplies a timing resource that an autonomous model must include. The manuscript does not promote a bound for one class of clocks to a universal clock theorem.

A candidate specifies the accuracy actually required by the task. An agent that can tolerate random long pauses need not maintain a fixed phase relative to distant stars. It still needs enough

causal ordering, memory retention, and reactivation reliability to complete its policy. A formal instruction to wake at the nth exponentially increasing deadline can conceal an unbounded counter, an ideal analog threshold, or an external scheduler. Those are resources to implement, not free consequences of writing a short formula.

7. Reliability across an unbounded proper future

7.1 A whole-trajectory error budget

Proposition 7.1. Summable errors suffice

Suppose the conditional probability that episode n introduces a declared unrecovered failure, given any allowed prior successful history, is at most ε_n. If $\sum\varepsilon_n \le \delta$, then every finite prefix, and the event of no such failure over the entire sequence, have probability at least $1-\delta$. No independence assumption is needed for this union bound.

Proof. Partition the failure event by the index of its first failure. Each branch has probability at most ε_n, since the successful-prefix probability is at most one. Sum over a finite prefix and then take the increasing union. The allocation $\varepsilon_n = 6\delta/(\pi^2 n^2)$ is one convenient example. A constant independent failure probability $\varepsilon > 0$ instead gives prefix survival $(1-\varepsilon)^n$, tending to zero.

This is a sufficient reliability schedule, not an available error-correcting design. Repetition-based suppression often needs increasing physical resources. For independent binary trials with error $p < 1/2$, a majority code can reduce error exponentially in its repetition count; making ε_n fall as n^{-2} requires a repetition count growing logarithmically with n. An indefinitely widening code violates a fixed-memory or fixed-apparatus premise unless its storage and control costs are explicitly allowed. Reusing components does not automatically eliminate counters, synchronization, and syndrome memory.

7.2 A fixed barrier in a cooling environment

Let a classical memory fault be a symmetric bit flip with instantaneous rate $\lambda(t)$. The accumulated exposure is $A(t) = \int \lambda(s)\,ds$. The probability of wrong parity at readout is $[1-\exp(-2A)]/2$. The probability of at least one flip is $1-\exp(-A)$ in the Poisson model. These are different reliability criteria: two flips can restore the readout bit while still violating a no-fault trajectory requirement.

Proposition 7.2. Arrhenius cooling can make total exposure finite

Assume $\lambda(t) = \lambda_0 \exp[-B\,x^\theta]$, with $\theta > 0$ and $B = E_b/(k_B T_0) > 0$. The total exposure is finite, with the following closed form. Thus the no-fault probability is a strictly positive constant, and can be close to one for suitable parameters. The statement assumes that this rate law remains valid and that other hazards are separately included.

$$A(\infty) = \frac{\lambda_0 t_0}{\theta}\, B^{-1/\theta}\Gamma\left(\frac{1}{\theta}, B\right) < \infty. \tag{6}$$

$$\theta = \tfrac{1}{3}: \quad A(\infty) = 3\lambda_0 t_0 e^{-B}\left(B^{-1} + 2B^{-2} + 2B^{-3}\right). \tag{7}$$

The calculation is the same integrability mechanism that shuts down the thermal clock: fixed activation barriers suppress transitions as cooling continues. Suppressing an unwanted memory

transition is beneficial; suppressing every usable clock transition defeats activity. The positive construction must therefore distinguish its memory-protection and clock mechanisms. A claim that the same frozen apparatus both never changes and keeps computing indefinitely needs a more detailed dynamics.

For B=20 and $\lambda_0 t_0=1$, the total exposure is approximately 3.41636×10^{-10}. This is an illustrative model calculation, not a measured device lifetime. A temperature-independent tunnelling rate, radiation background, material failure, or other residual hazard with a divergent integral can dominate the result. One cannot extrapolate a fitted Arrhenius term while silently dropping such a term.

7.3 A positive floor erases passive distinguishability

Proposition 7.3. Passive two-state memory at constant noise

For a symmetric bit-flip rate $\lambda>0$, the total variation distance between the two memory output distributions initially prepared as 0 and 1 is $\exp(-2\lambda t)$. The optimal success for retrieving a uniform historical bit is $[1+\exp(-2\lambda t)]/2$, tending to $1/2$. The probability of a fault-free trajectory tends to zero.

Proof. The transition matrix has diagonal entries $(1+e^{-2\lambda t})/2$ and off-diagonal entries $(1-e^{-2\lambda t})/2$. Subtract its two rows and apply the binary hypothesis-testing formula. The zero-flip probability is $e^{-\lambda t}$. An irreducible mixing finite-state process has the analogous qualitative loss of initial distinguishability, but the explicit rate here belongs to the two-state model.

Self-correcting quantum-memory research shows why geometry, locality, system size, and active intervention matter.[22,23] A no-go result for specified two-dimensional stabilizer Hamiltonians is not a theorem against every quantum memory. Nor does a protected ground-space construction establish infinite lifetime for a fixed finite specimen under every physical noise process. The appropriate infinite-horizon statement concerns the actual logical error channel and its integrated hazards.

7.4 A maintenance lower bound for a noisy ready cell

Consider one reusable, energetically degenerate classical ready cell. At the start of each cycle its error probability is p_i, with $0<p_i<\delta<1/2$. During idle storage it undergoes symmetric noise at rate λ. The device must keep its error at most δ and restore the same p_i after each refresh, using thermal disposal at temperatures at least $T_f>0$. Ancillary records and side information are cyclically restored, with no additional low-entropy sink omitted. This is a specific cyclic primitive, not a model-independent description of a mind.

$$\Delta_{\max} = \frac{1}{2\lambda} \ln \frac{1-2p_i}{1-2\delta}, \qquad \rho_S = \frac{h(\delta)-h(p_i)}{\Delta_{\max}}. \tag{8}$$

Theorem 7.4. Maintaining that cell costs work per elapsed time

For a refresh schedule covering elapsed duration t, the consumed work is at least $k_B T_f \rho_S \max(0, t-\Delta_{\max})$. Thus a finite work budget cannot maintain this primitive forever, even if the cognitive controller is mostly asleep. The bound is optimistic: finite-time and apparatus constraints may increase the expenditure.

Proof. The idle error is $p(u)=1/2+(p_i-1/2)e^{-2\lambda u}$, so every full interval is at most $\Delta_{\max}$. The entropy gain $f(u)=h(p(u))-h(p_i)$ is increasing, concave, and zero at $u=0$: p is increasing and

concave below 1/2, and binary entropy has the same properties there. Hence $f(u) \geq \rho_S u$ for $u \leq \Delta_{max}$. A cyclic thermal refresh costs at least $k_B T_f f(u)$ under the stated Landauer assumptions.[10] Sum the full intervals; an unfinished last interval is at most Δ_{max}. This proves the result.

For $p_i=0.001$, $\delta=0.1$, and $\lambda=1$ in chosen inverse-time units, $\Delta_{max} \approx 0.110571$ and $\rho_S \approx 2.86853$ nats per time unit. A different passive state, a nonthermal entropy sink, correlated noise, or a noncyclic feedback apparatus changes the premises and needs its own account. The theorem supplies a sharp counterexample to treating all sleeping infrastructure as free.

7.5 A finite-dimensional thermal maintenance obstruction

A broader negative result can be proved with an entropy-contraction premise. Let M include every retained degree of freedom that can carry the protected record, including controller records and reusable ancillas. Its dimension is fixed and finite. Its reference Hamiltonian H_M is fixed and its bath has one constant temperature $T>0$. The full-rank Gibbs state is $\tau=\exp[-H_M/(k_B T)]/Z$. Write relative entropy in nats. Assume the uncontrolled noise semigroup Φ_s on this whole M obeys the following uniform inequality for every state and every $s \geq 0$:

$$D(\Phi_s(\rho)\|\tau) \leq e^{-\eta s} D(\rho\|\tau), \qquad \eta > 0. \tag{9}$$

This is the modified log-Sobolev type premise used here; η is the entropy-decay coefficient, including any convention-dependent factor of two. A spectral gap is not substituted for η without the appropriate comparison theorem. Kastoryano and Temme develop the relevant quantum convergence framework and its thermal interpretation.[38] The inequality must hold for the complete chosen memory under its actual noise law, rather than only for an exposed physical component.

Consider an equiprobable classical label $J \in \{0,1\}$, unavailable to the controller except through M. Its conditional states are ρ_0 and ρ_1. Let $\bar{\rho}=(\rho_0+\rho_1)/2$. Controls are independent of J as external input, return to H_M at accounting boundaries, and charge ensemble-mean work plus any consumed nonthermal free energy to $C_{ctl}(t)$. Thermal auxiliaries are cyclically restored. No uncounted correlated register, cold reservoir, or entropy sink is available. Noise acts for elapsed duration t; an ideal splitting limit allows control between arbitrarily short noise intervals, without allowing free removal of that noise exposure. These are protocol premises, not consequences of the bath temperature.

$$Q(t) = \frac{1}{2}\sum_{j=0}^{1} D(\rho_j(t)\|\tau) = I(J{:}M)_t + D(\bar{\rho}(t)\|\tau). \tag{10}$$

Theorem 7.5. Uniform mixing forces an unbounded maintenance resource

Suppose a readout of M can retrieve J with error at most δ, with $0 \leq \delta < 1/2$, at every accounting time through t, including throughout the limiting noise exposure. Put $q_\delta = \ln 2 - h_2(\delta) > 0$, with h_2 the binary entropy in nats. Under the stated noise and control premises, the charged control resource obeys the following bound. Thus no finite total work-plus-debited-free-energy budget preserves this continuously readable bit forever. The statement applies to classical probability distributions and quantum density operators under the same declared contraction law.

$$\begin{aligned} C_{ctl}(t) &\geq k_B T\,[\eta q_\delta t + Q(t) - Q(0)] \\ &\geq k_B T\,[\eta q_\delta t + q_\delta - Q(0)]. \end{aligned} \tag{11}$$

Proof. Data processing through the readout and the binary Fano inequality give $I(J{:}M) \geq \ln 2 - h_2(\delta) = q_\delta$. Hence $Q \geq q_\delta$. Applying the assumed contraction separately to each conditional state gives $(dQ/dt)_{noise} \leq -\eta Q$. No quantum-reference extension of the inequality is needed because J is classical.

For a control channel independent of J, mutual information cannot increase. The free-energy identity $F(\bar{\rho}) - F(\tau) = k_B T D(\bar{\rho} \| \tau)$, together with cyclic isothermal work accounting, bounds the increase of the second term in Q by the charged free-energy debit divided by $k_B T$. Thus $\Delta Q_{control} \leq \Delta C_{ctl}/(k_B T)$. If an auxiliary retains J-dependent information or consumes negentropy without being restored, it must instead be included in M or charged; the control inequality is not asserted while omitting that resource.

Summing noise losses and control gains gives $Q(t) - Q(0) \leq -\eta \int_0^t Q(s)ds + C_{ctl}(t)/(k_B T)$. Insert $Q(s) \geq q_\delta$ to obtain the bound. Fixed finite dimension and a full-rank τ ensure $Q(0) \leq \ln(1/\tau_{min}) < \infty$. Initial nonequilibrium free energy can delay the obstruction but cannot remove its positive asymptotic slope. The bound concerns ensemble means; a hard finite resource cap also bounds that mean. It does not construct repair.

7.6 What the stronger theorem does and does not decide

The assumed class is nonempty. For the thermal replacement semigroup $\Phi_s(\rho) = e^{-\eta s}\rho + (1 - e^{-\eta s})\tau$, convexity of relative entropy gives the required contraction. A primitive fixed noise law with a verified modified log-Sobolev constant can supply other examples. Theorem 7.5 therefore extends the single ready-cell obstruction to a stated class of finite classical and quantum apparatus. It does not infer a universal contraction constant from temperature.

In particular, a positive-temperature Gibbs state does not establish primitivity. A pure-dephasing channel on a degenerate two-level system has the maximally mixed Gibbs state as a fixed point at any $T > 0$, yet exactly preserves the two classical basis records in the ideal model. That is a mathematical counterexample to a temperature-only memory theorem, not a proposed cosmological device. A noise coefficient that tends to zero, a protected sector, an enlarging memory, or an additional retained register also requires a separate analysis.

Even instantaneous primitivity is insufficient if the noise changes with time. For a symmetric classical bit with $\lambda(t) = \lambda_0/(1 + t/t_0)^2$, every instantaneous generator has a unique full-support thermal fixed point, but total exposure is $\lambda_0 t_0$. Its optimal retrieval error tends to $(1 - e^{-2\lambda_0 t_0})/2$, below $1/2$, at zero control work in that ideal law. The uniform stationary contraction premise has been lost. If a fixed τ and valid control accounting are retained with coefficient $\eta(t)$, the proof instead gives $C_{ctl}(t) \geq k_B T[q_\delta \int_0^t \eta(s)ds + Q(t) - Q(0)]$. Divergent accumulated contraction, rather than the word primitive at each instant, supplies the obstruction.

A Gibbons-Hawking temperature does not by itself prove that every possible finite apparatus has the complete uniform mixing and resource properties above. Establishing those properties for a specified de Sitter device would make the theorem decisive for that device class. Asserting that no finite-work agent can preserve any bit in every positive-temperature cosmology would exceed the premises. Conversely, escaping this theorem is not sufficient to demonstrate indefinite intelligence.

8. A conditional positive resource process

8.1 What the construction is intended to show

The previous obstructions do not make the schedule inequalities contradictory. To show their remaining affirmative content, we specify a finite-state logical process with a nonhomogeneous event clock. The clock rate and achievable work cost are primitives of a candidate contract. The construction proves compatibility of the stated count, work, and fault laws. It does not construct a finite gravitational machine that realizes those primitives indefinitely.

The logical state is two bits (j,c). The bit j is an initially unknown retained core record; c is a phase of the controller. At each certified clock event, the controller emits a=j⊕c and updates c to c⊕1, leaving j fixed. The policy uses its retained record and internal state. Its four logical states are counted explicitly. It is a small transducer, with no claim of flexible intelligence or unbounded learning.

8.2 An event law with infinitely many future opportunities

Let events follow a nonhomogeneous Poisson process with rate $r(t)=a/(t_0x)$, where $a>0$. The rate is integrable on every finite time interval, so no finite-time accumulation of infinitely many events occurs. The expected count is a ln x and diverges as t tends to infinity. Changing variable to y=ln x turns the event process into a homogeneous Poisson process of rate a; it therefore has infinitely many events almost surely.

$$r(t) = \frac{a}{t_0 x}, \qquad T(t) = T_0 x^{-1/3}, \qquad w(t) = w_0 x^{-1/3}. \tag{12}$$

Assume the complete event primitive, including its clock and any work costs not separately listed, can achieve the stated w, or that all remaining expenditure has an explicitly finite budget reserved in advance. A merely computed lower bound on w is insufficient. The rate fits the m=4 illustrative heat-disposal scaling. It conflicts with an indispensable fixed positive activation gap, so the implementation contract must supply a different execution mechanism.

8.3 Finite expected expenditure and a finite actual fuel store

Theorem 8.1. A bounded-fuel continuation event

For the stated event law and charged cost, total future consumed event work has expectation $3aw_0$ and is finite almost surely. If a finite fuel store W_0 is used, with execution stopping before an unaffordable event, the probability that fuel never forces a stop is at least $1-3aw_0/W_0$, when this is positive. If all declared record, controller, and primitive failures have total union probability at most δ_f, infinitely many successful updates within fuel and with no such failure occur with probability at least $1-3aw_0/W_0-\delta_f$.

Proof. For a nonnegative charged event cost, partition a finite interval into small intervals and use the expected event count r(t)dt; taking a positive-integral limit gives $E[W]=\int r(t)w(t)dt$. The integral is $aw_0\int_1^\infty x^{-4/3}dx=3aw_0$. Finite expectation implies finiteness almost surely. Markov's inequality bounds the chance that the all-events cost exceeds W_0. Couple the stopped and unrestricted processes: if the unrestricted total fits the store, no stop occurs. Subtract the declared failure probability by a union bound. The unrestricted count is infinite almost surely, so the retained event has the stated property.

Finite expected work alone would not establish an absolute deterministic fuel bound on every sample path. The stopping rule and explicit probability estimate address that distinction. The model also does not count an unlimited external transcript as part of its four logical states: the operation count is an analysis variable, not an internally stored unbounded history.

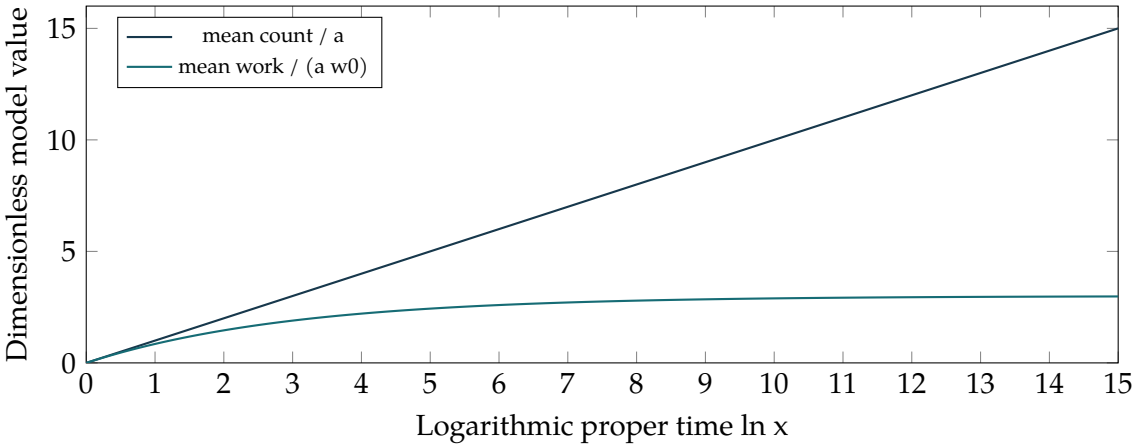

Figure 2: For the ideal resource process, mean event count keeps growing in logarithmic proper time while mean charged work approaches a finite limit. This is a model calculation under the primitive contract, rather than a device demonstration.

8.4 A compatible protected record, with the missing hazards visible

For the retained bit, the ideal thermal-noise example of Section 7.2 gives $\delta_f \leq A(\infty)$ if a single fault is declared fatal and all other failures are absent or separately budgeted. With $a=1$, $w_0=1$ in normalized work units, $W_0=100$, $B=20$, and $\lambda_0 t_0=1$, the conservative joint lower bound is approximately 0.969999999658. Initial preparation, barrier formation, and any reserved maintenance expenditure are outside that remaining 100-unit event store and must also be finite and counted.

The record barrier can be fixed while the event mechanism has a different spectrum or drive. That separation is a primitive assumption, not an available materials design. A constant unaccounted failure hazard would make the no-failure probability zero over an infinite proper future and defeat this witness. The physically important next task is therefore to replace the assumed clock-cost-noise contract with a microscopic architecture or an impossibility proof for that architecture.

8.5 What has and has not been achieved

The construction establishes a nontrivial mathematical distinction: finite work and bounded logical state are compatible with infinitely many memory-dependent updates on a positive-probability continuation event when the environment and primitive rates have the specified asymptotics. It does not show that all environmental resources are finite, that fixed physical hardware admits the clock law, or that the transducer is intelligent.

In particular, if $w=s_h k_B T$ with $s_h>0$, infinitely many events export infinitely many nats even when their total heat converges. The environment must have unbounded effective entropy-accepting capacity along this history. A finite work endowment has not made that environmental resource finite. A closed finite disposal budget is a different contract, addressed in Proposition 11.1.

This is the right place for an affirmative frontier. An existence conjecture becomes useful when a consistent resource target and its hardest implementation obstacle are both explicit. The candi-

date can then fail because of spectral shutdown, cumulative hazards, controller costs, or a finite entropy sink, rather than because a successful integral was mistaken for an implemented mind.

9. The future that a finite memory cannot keep

9.1 A dimension bound on independent historical records

Dyson explicitly considered analog memory, including increasingly precise angular or frequency records, with information capacity growing logarithmically with cosmic time; the discussion spans pages 456-457 at the end of Lecture III.[4] This is not a fixed finite-D memory. A fixed number of analog components need not have a fixed distinguishable state count if the model grants improving precision or new accessible modes. The theorem below rules out unlimited reliable history at fixed D; it does not refute that different premise by changing its name. The physical reply must establish limits on precision, noise, readout, or accessible record modes, as emphasized by Krauss and Starkman.[5]

Let J be uniformly drawn from K possible histories, and let the complete accessible memory carrying information about J have Hilbert-space dimension D. It can be quantum, and its encoding can use arbitrary density operators. The reader receives no other system correlated with J. A measurement and unlimited subsequent classical calculation try to identify the history. This is a record-retention problem, rather than a limit on how many gates can run.

Theorem 9.1. Finite memory cannot retain unbounded independent history

The probability of identifying the uniform history is at most $\min(1,D/K)$. To recover n independent unbiased historical bits with block success at least $1-\delta$, a memory must have $D\geq(1-\delta)2^n$. Equivalently $n\leq\log_2 D-\log_2(1-\delta)$. A fixed D therefore cannot support unbounded lossless cumulative learning at fixed block reliability.

Proof. Let ρ_j be the encoded state and F_j the effect associated with guessing j, after coarse-graining all responses into guesses. Since $\rho_j\leq I$ and $\sum F_j=I$, the average success is at most $(1/K)\sum \mathrm{Tr}\, F_j= D/K$. The probability is also at most one. Set $K=2^n$ and rearrange. The bound is attained when D orthogonal states represent D distinguishable message classes and otherwise identical messages are divided among their corresponding guesses.

For example, a 1,024-dimensional memory can retain ten independent bits perfectly. It cannot recover eleven uniform independent bits with 99 percent block success: the success ceiling is one half. Quantum amplitudes do not turn a finite-dimensional register into an arbitrarily large reliably readable classical archive under this information boundary.

9.2 Learning while forgetting is a different objective

The theorem does not prohibit bounded-memory agents from processing an endless stream, updating a finite model, or forgetting obsolete data. It limits retention of the declared independent history. A compact law can summarize many observations when the task does not require their individual unpredictable outcomes.[32] Repeated retrieval of the same finite core can also support ongoing policy without adding independent records.

An indefinite intelligence claim should therefore state its retention policy. Which commitments remain inviolable? Which old observations can be compressed or discarded? What competence is preserved after the compression? A promise to keep every independent event forever requires

expanding reliable capacity somewhere. Moving that archive to the environment changes the location of the resource; it does not remove it from the account.

The Continuum's continuity-engineering distinction is especially helpful here.[1] A process can preserve a structured identity and policy without retaining every incidental event. Conversely, a mass of records can fail to preserve the organization needed to use them. The physical claim becomes clearer when a finite protected core, a replaceable working model, and an external archive are assigned different roles and capacities.

9.3 Attractors cannot reconstruct information that is absent everywhere

Proposition 9.2. Complete historical erasure defeats label-free reconstruction

Suppose the joint state of every system and input available to a reconstruction is the same for all values of a historical label J. Any reconstruction channel fixed independently of J has the same output for all J. For K uniform possible labels, its optimal guess success is $1/K$. A later return to a common structural attractor does not recover the missing historical distinction.

Proof. The channel acts on identical inputs, so linearity gives identical outputs. For a guess distribution q independent of the uniform label, the average success is $\sum_j q(j)/K = 1/K$. The quantum version has an identical conditional output density operator and the same conclusion for every measurement.

This is the appropriate scientific boundary for the internal historical monograph's discussion of continuity without memory.[3] A seed, invariant rule, prompt, stored model, or surviving witness can recreate a recognizable structure and may itself encode genuine history. If it does, it belongs among the retained resources. If none contains the history in question, structural recurrence cannot identify that history. The theorem does not deny the value of reconstruction; it identifies what the reconstruction must preserve to support a historical continuity claim.

9.4 Substrate changes need a transport map

A migration specifies an encoding into the new substrate, a correspondence between old and new logical observables, and the physically implemented channel. Reference-preserving quantum-channel contracts give a standard framework for quantifying such continuity.[35] For a finite sequence whose step errors are ε_n in half diamond norm relative to the intended transported process, telescoping with channel contraction bounds the total process error by $\sum \varepsilon_n$. A summable sequence can therefore support a uniform continuity certificate at the level of the declared channels.

The certificate does not provide the channels or their costs. If each migration consumes at least a fixed positive amount of a finite nonrenewable budget, only finitely many can occur. If each carries an independent fixed probability of catastrophic loss, the probability of surviving infinitely many is zero. A schedule with decreasing costs and errors must justify the hardware, precision, controller, and records needed to attain them. An infinite chain of cheap substrates is not free merely because each link is described briefly.

9.5 Finite energy is not automatically finite Hilbert-space dimension

The memory bound assumes a controlled dimension D; it does not derive it from energy alone. Spatial extent, accessible modes, regulator, time resolution, and gravitational constraints matter. In field settings a finite energy statement without a localization and access specification is not

enough to insert a finite D. Bekenstein-type and relative-entropy bounds supply conditional tools with distinct premises.[16,17]

For a changing apparatus one may instead specify a growing effective record capacity $D(t)$. The theorem then bounds the independent historical information by $\log_2 D(t)$, at the stated reliability. Physical expansion is useful only if the agent can address, maintain, and recover those records within its causal and operational constraints. Unlimited elapsed time does not by itself supply unlimited usable space or a reliable decoder.

10. Whether it pays to wait

10.1 Compare complete uses of the same resources

A colder bath can lower the work associated with a specified entropy export. It does not follow that waiting increases the value of every resource. An already prepared low-entropy tape can absorb information entropy, and a cold reservoir or other nonequilibrium system can have useful work value today. The Bennett-Hanson-Riedel critique makes this omitted-resource problem explicit.[25]

Quantum side information can also change an erasure cost. Conditional entropy, correlations, and entanglement are resources; processes that consume them are not complete repeatable cycles unless their restoration is included.[11,30,31,37] Thus the simple $k_B T$ factor is meaningful only after the information boundary and cyclic resource account are fixed. Waiting must be compared with immediate use or preservation of the same complete endowment.

10.2 A conditional waiting criterion

For one thermal-disposal strategy, let stored fuel decay at fractional rate $\ell(t)$, let the survival probability have hazard $h(t)$, and let the useful operation price be proportional to $T(t)$. Ignore discounting or urgency only as part of this stated objective. A measure of expected future erasure capacity is proportional to the surviving fuel divided by temperature.

$$V(t) = \frac{W_0}{s_h k_B T(t)} \exp\left[- \int_0^t (\ell(s) + h(s))\, ds \right]. \tag{13}$$

Proposition 10.1. Cooling must beat storage and survival losses

At a differentiable time with positive temperature and available fuel, $d \ln V/dt = -d \ln T/dt - \ell - h$. Under $T = T_0 x^{-\alpha}$ and $\ell + h = \rho/(t_0 x)$, V is proportional to $x^{\alpha-\rho}$. Waiting improves this particular expected-capacity objective exactly when $\alpha > \rho$. At a positive terminal temperature floor, further waiting has no thermal advantage and lowers V whenever losses remain positive.

Proof. Differentiate the logarithm. Integrating the power-law cooling and fractional-loss terms gives the stated exponent. Once temperature stops falling, the cooling derivative vanishes; the loss terms remain nonpositive. The result compares the strategies in its assumptions, rather than every possible thermodynamic use of stored information.

For an illustrative case $T/T_0 = \max(x^{-1/2}, 0.01)$ and $\rho = 0.2$, expected thermal-disposal capacity increases until $x = 10{,}000$ and decreases thereafter. This optimum comes from the chosen model, not a forecast of an alien civilization's behavior. An additional fixed waiting cost or a scarce opportunity to observe a departing source can move the optimum earlier.

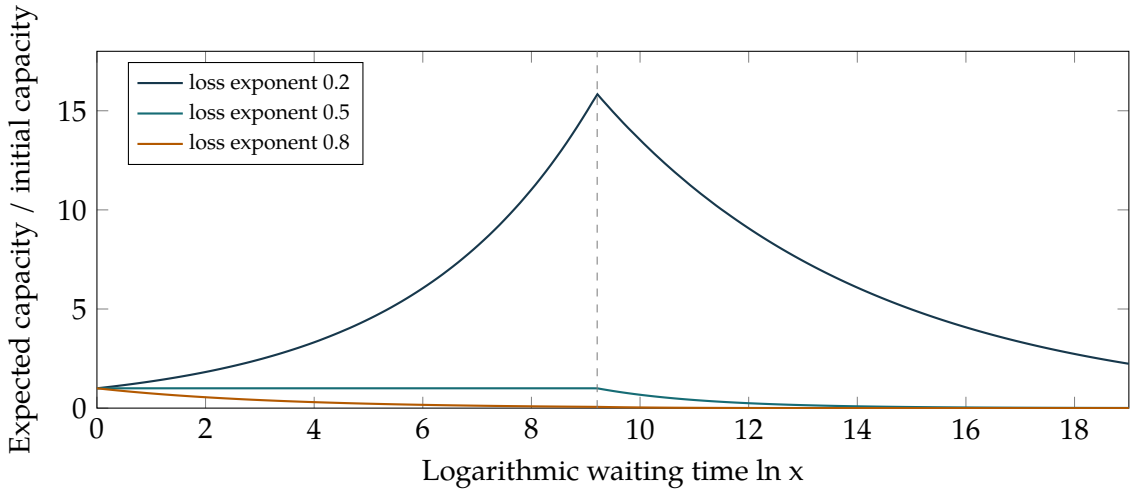

Figure 3: Illustrative expected capacity under cooling, stored-resource loss, and a positive temperature floor. The curves compare a specified disposal strategy; they do not establish a universal motive for aestivation.

10.3 Information can disappear while the temperature improves

Waiting can reduce access to external sources. Signals can redshift, resources can pass out of causal reach, and a unique observation opportunity can end. The optimal policy then trades thermal efficiency against information and resource acquisition. Infinite elapsed time is not equivalent to infinitely many available observations.

A candidate may express this as a channel contract: each new observation opportunity has a declared upper bound on task-relevant information conditioned on the available past. A finite sum of such bounds limits total newly acquired information in that model. A divergent formal bound is only a necessary opening; it does not establish an actual source, decoder, or budgeted acquisition process. A finite retained memory may still need to discard earlier observations even when cumulative acquisition is unbounded.

10.4 Sleeping forever is not continuation

A schedule that never wakes after some finite episode has an infinite worldline but finite completed activity. A wake-up time whose probability drifts to zero can likewise fail a promised uniform continuation probability. The process must be evaluated jointly: its clock events, record state, fuel, and task outputs belong to one stochastic history. Reporting a favorable energy integral for a separately idealized sequence cannot establish that the agent executes that sequence.

11. Cosmological regimes and their limits

11.1 A stationary de Sitter scale

For an apparatus satisfying the full fixed-temperature mixing and control premises, Theorem 7.5 gives an explicit linear maintenance obstruction at $T=T_{dS}$. Its coefficient contains the device entropy-decay rate η and the required readout accuracy. Neither follows from H alone. Section 11 therefore distinguishes a cosmological thermal scale from a substrate-independent theorem about every possible memory.

In the stationary de Sitter setting associated with a constant positive cosmological constant, a suitable detector in the appropriate state encounters the Gibbons-Hawking thermal scale.[15] For the central static observer, with H the constant expansion rate, the temperature and dimensionless horizon entropy are:

$$T_{\mathrm{dS}} = \frac{\hbar H}{2\pi k_B}, \qquad \frac{S_{\mathrm{dS}}}{k_B} = \frac{\pi c^5}{G\hbar H^2}. \tag{14}$$

These formulas have specified geometrical and semiclassical meanings. They are not a proof that a laboratory cannot momentarily refrigerate below the detector temperature, or that every dynamical cosmology has temperature $\hbar H/(2\pi k_B)$. Refrigeration requires a resource account; a changing horizon is not automatically a stationary thermal reservoir. A single equilibrium bath also does not provide unlimited directed work without another consumed resource.

The finite entropy of a horizon is not, by itself, a derivation of an exact finite-dimensional Hilbert space for every allowed observer and field. Algebraic treatments can be subtler.[33,34] Applying Theorem 9.1 to a cosmological patch requires a separately justified effective record dimension or distinguishability bound. The present paper does not replace that premise with the horizon area formula.

11.2 A closed disposal bound at a positive floor

Proposition 11.1. Finite thermal and negentropy disposal budget

Consider cyclic operations that each require exporting at least $s_0>0$ nats of discarded information under the declared classical, degenerate-register accounting. Suppose finite nonthermal low-entropy stores can accept at most C_0 nats, all remaining entropy is transferred as heat to baths at temperature at least $T_f>0$, and the total work available for that thermal part is W_0, with all other energy changes restored or counted. Then $N\le[C_0+W_0/(k_BT_f)]/s_0$.

Proof. The low-entropy stores can receive at most C_0 of the required Ns_0 nats. Thermal export of the remainder has heat and cyclic work cost at least $k_BT_f(Ns_0-C_0)$ under the stated assumptions. Compare with W_0 and rearrange. Imported cold tapes, unreturned side information, a noncyclic fuel change, or a different logical primitive change the account and must be included explicitly.

The positive s_0 premise is an operational requirement to prove for a candidate, not a theorem about every conscious moment or reversible gate. The bound therefore supports a conditional finite-activity conclusion in a stationary thermal regime. It does not close every imaginable loophole in physics by definition. The maintenance theorem supplies another route to finite continuation for its specified noisy primitive, even when the cognitive updates themselves are logically reversible.

11.3 An illustrative numerical scale, without a cosmological fit

Choose $H=2\times10^{-18}\ s^{-1}$ only as an illustrative constant. Then $T_{dS}\approx2.43132\times10^{-30}$ K and the ideal thermal price of resetting one independent unbiased bit is approximately 2.32676×10^{-53} J. One joule would cover at most about 4.29782×10^{52} such resets in that ideal account, before any other costs. The associated horizon entropy is about 2.70216×10^{122} nats. These are finite model numbers, not an estimate of a mind's lifetime.

Extremely large finite numbers do not settle an infinite-time claim. Nor do those numbers establish an attainable processor: the temperature, controllable gaps, decoding access, and clock architecture remain physical obligations. A static patch's horizon cannot be treated as an uncharged external computer or an unlimited archive simply because its entropy scale is large.

11.4 Evolving and finite-duration futures

For a spatially flat expanding Friedmann universe in general relativity, with a separately conserved positive dark-energy component, define $Y(a)=\int_1^a[1+w(u)]d\ln u$. Conservation gives $\rho_{DE}(a)=\rho_{DE,0}\exp[-3Y(a)]$. If $a\to\infty$, $Y(a)\to\infty$, ordinary matter and radiation dilute, and no residual cosmological constant remains, then $H\to0$. This removes a strictly positive asymptotic de Sitter H scale. It is a conditional future law, not an inference from a measured present-day slope.

$$Y(a)=\int_1^a[1+w(u)]\,d\ln u,\qquad \frac{\rho_{DE}(a)}{\rho_{DE,0}}=e^{-3Y(a)}. \tag{15}$$

A constant late equation of state $-1<w_\infty<-1/3$ gives a power-law accelerating future $a(t)\propto t^p$, $p=2/[3(1+w_\infty)]>1$. Its event-horizon radius is $ct/(p-1)$, so it has no constant de Sitter radius. H decreases as $1/t$; any horizon-temperature estimate for this nonstationary spacetime needs its own detector and state prescription. Freese and Kinney examine continuation opportunities in such evolving-dark-energy settings, under their scaling assumptions.[6]

Merely approaching $w=-1$ is inconclusive. If $1+w(a)$ is positive and asymptotic to $b/(\ln a)^p$, with $b>0$, then Y diverges for $p\le1$ and remains finite for $p>1$. The former permits $H\to0$ under the stated matter assumptions; the latter retains a positive limiting dark-energy density. A residual positive cosmological constant also restores the floor. Growing phantom density or future recollapse belongs to a different worldline-duration analysis.

DESI constrains a finite observed epoch. Its phenomenological form $w(a)=w_0+w_a(1-a)$ is not an established law for $a\to\infty$: with $w_a<0$ it would send w toward positive infinity if extrapolated literally. Distinct future continuations can agree on the observed interval and then have different $Y(a)$ limits. The DR2 preference therefore motivates examining several future regimes without selecting one of them.[39,40]

Declared future	What must be checked
Asymptotically constant positive thermal scale	A positive-priced cyclic primitive or unavoidable maintenance can force finite activity with finite accessible resources.
Environment cooling indefinitely	The schedule region can be nonempty, but clock capacity, spectral access, heat transfer, and total hazard can still close it.
Accelerating power-law expansion	Acceleration alone does not fix the answer. A physical detector, sink, and information-access model are required.
Finite remaining proper time	A finite operation-rate bound implies a finite count. A divergent rate near the endpoint needs an independently valid physical realization.
A metastable environment with destructive hazard	Survival depends on the integrated hazard. A constant positive terminal hazard gives zero no-catastrophe probability over infinite time.

Phantom-energy scenarios with a finite-time future singularity illustrate a different obstruction from a stationary cold bath.[26] Long-term astrophysical histories and black-hole evaporation identify additional eras and possible resources, with assumptions about particle stability and gravitational physics.[27,28] No fixed timeline is required by the present theorems. A black hole can provide a long-lived resource in a model without supplying an endless controllable reservoir or a verified way to preserve an agent through evaporation.

11.5 Recurrence and fluctuation are not a designed future

A recurrence can restore a configuration without giving an internally available record of the intervening history. A rare fluctuation can resemble a memory without having the history that memory purports to encode. The cosmological Boltzmann-brain problem sharpens that reliability issue.[36] These possibilities do not by themselves implement a sustained policy, a trustworthy core record, or an attainable resumption mechanism. The paper's continuation criterion follows an actual causal process rather than counting disconnected look-alike states as one reliable life.

12. A test program for the physical assumptions

12.1 What a finite experiment can establish

No finite run proves indefinite survival. It can reject a proposed model, estimate parameters over a tested range, and test a physical mechanism that an asymptotic theorem requires. The strongest near-term contribution is therefore an integrated validation of rate, cost, and memory behavior, with an explicit extrapolation domain. A successful fit over a narrow interval does not certify the infinite tail.

A candidate reports its controlled memory, complete apparatus, baths, work reservoirs, timing source, readable outputs, and excluded side channels. It states whether the environment is a laboratory refrigerator paid by the experiment, or an ideal external cooling history studied as a cosmological boundary condition. The same device cannot claim the refrigerator's declining temperature as free while reporting only the logic element's expenditure as its total cost.

12.2 Distinguish power-law capacity from spectral shutdown

Measure the actual rate of task-valid transitions across a controlled temperature range. Compare an algebraic model ν proportional to T^s with an activated model ν proportional to $\exp[-\Delta/(k_B T)]$, including uncertainty, residual channels, and independent hold-out temperatures. The two laws can look similar over a limited interval and imply opposite infinite-time conclusions. Resolve relevant level spacings and couplings rather than selecting the most favorable fit by its forecast.

The meaningful prediction is candidate specific: a microscopic mechanism predicts a capacity envelope and an interval in which it should hold. An observed departure can defeat that envelope or reveal a new transition channel. Claiming a persistent power law at temperatures below an unexamined gap is not an experimental result. A model of adaptive gap control must also report the work and precision used to change the Hamiltonian.

12.3 Test dormant records over elapsed time

Prepare an unknown classical core bit with a secret verifier label, protect it in the proposed dormant mode, and vary proper storage time and bath history independently of the number of active updates. Measure both endpoint retrieval and any accessible fault witnesses. The two-state model predicts endpoint error from the integrated rate; an all-trajectory guarantee requires an additional event or channel argument and cannot be inferred from a lucky correct endpoint.

Fit activated and residual fault contributions separately. A small constant residual is negligible in some finite experiments and decisive at infinite time. Active refresh tests should include the control records, reset work, error target, and elapsed storage duration. The known ready-cell model in Section 7 is a sharp reference case; a different protected memory needs its own logical channel and expenditure analysis.

12.4 Verify the complete resource account

Measure total consumed work or a justified upper bound for the operating primitive, not only an ideal entropy lower bound. Include the clock driver, readout, state preparation, discarded auxiliaries, gap-control field, and cryogenic resources appropriate to the experimental claim. Identify reusable low-entropy stores and show whether they are restored. Finite-size and finite-time erasure experiments demonstrate why actual costs need not attain a quasistatic ideal.[13,29]

Compare hibernation and continuous operation under the same total resources and task criterion. A cooling strategy can win in one contract and lose when maintenance, storage loss, or input opportunities are included. Use the waiting derivative as a model prediction, and test its loss and survival parameters independently. The conclusion concerns the declared objective, not a general recommendation that intelligence should prefer inactivity.

12.5 Continuity and adaptive competence

A substrate-transfer experiment should preserve an unknown core record and a preregistered policy relation, not merely reproduce a familiar name or an approximate transcript. Test histories that lead to the same superficial identity description but require different later actions. This distinguishes retained historical information from reconstruction of a generic attractor. A memory-capacity claim must count all accessible archives and correlated resources.

A stronger intelligence demonstration needs task families that demand adaptation, planning, generalization, and management of limited memory. Retention policies and allowed forgetting should be explicit. Passing the four-state transducer task establishes only that small process. Passing an energy-efficiency benchmark establishes only that resource performance. The conjunction becomes evidence for the frontier when one integrated candidate sustains the competencies and continuity it claims.

13. The intelligence claim and its failure conditions

13.1 A finite core can support a continuing process

The historical and normative Continuum commitments do not require that every microscopic state remain frozen. A process can change substrate, revise a model, and discard obsolete data while preserving a declared core and a lawful relation between successive policies. The continuity map must specify which distinctions survive and why they matter to the process. A stable symbolic form alone is not a certificate for an unknown past; a complete event archive is not always necessary for useful identity.

Likewise, an intelligence might function with long intervals between meaningful episodes. Its worldline duration, clock resolution, and adaptive episode count could scale differently. A physical model can quantify those relations without declaring a universal equivalence between operation count and experience. Claims about subjective duration remain an additional bridge, to be argued and tested where a theory permits.

13.2 Candidate-specific falsification

Candidate claim	A result that defeats it
Infinite activity at finite charged work	A mandatory omitted expenditure diverges, or the assumed achievable cost is only a lower bound.
Indefinitely available wake-up events	The physical clock has an indispensable activated channel with finite integrated capacity, or fails the declared rate envelope.
Reliable dormant core	A residual or catastrophic hazard has a divergent integral, or measured logical retention falls below the contract.
Finite-resource readable memory under uniform thermal mixing	The complete finite apparatus obeys Theorem 7.5 with divergent accumulated contraction, a fixed readout target, and all control resources included.
Whole-trajectory reliability	The claimed per-step certificate is not summable and the actual failure process defeats the uniform probability bound. Non-summability alone is not a universal impossibility proof.
Unlimited independent history in fixed memory	The required uniform history alphabet exceeds the dimension bound after every accessible correlated resource is included.
Reconstruction of a particular erased past	The input is independent of that past and no surviving witness or archive can supply the missing distinction.
Benefit from waiting	Measured storage and survival losses exceed the cooling benefit for the stated objective, or a matched immediate-use strategy dominates.
Physical indefinite intelligence	The integrated device lacks the claimed competencies, controller closure, resource cycle, or valid cosmological regime.

13.3 What would earn the title

A decisive affirmative result would supply a physical candidate whose clock remains useful, whose charged work fits its endowment, whose accumulated logical hazards meet a meaningful continuation criterion, and whose retained information supports adaptive competence. It would account for the environmental history rather than borrow an ideal future from an unexamined background. A decisive negative result can be narrower: a no-go theorem for a compelling class of finite architectures is already substantive scientific progress.

The results here expose three separations. Finite expended energy does not universally bound the number of physical state changes. An infinite count does not establish infinite reliable agency. Infinite reliable use of a bounded core does not establish infinite retention of independent knowledge. Those distinctions preserve the radical question while preventing it from being answered by an equivocal use of time, memory, or work.

The mind beyond the last star is therefore a question about a sustainable causal organization. The star is a boundary marker for familiar resources; the true horizon is the failure of the mechanisms that let an agent remember, wake, act, and remain itself. Whether that horizon is finite depends on the physical regime and the architecture. This paper gives those dependencies explicit mathematical form and leaves the strongest existence claim open to a concrete construction.

Appendix A. Closed forms and limiting cases

A.1 Power and logarithmic primitives

$$\int_1^X x^{-b}\,dx = \begin{cases} (X^{1-b}-1)/(1-b), & b \neq 1, \\ \ln X, & b = 1. \end{cases} \tag{16}$$

At the exact power boundary, put L=1+ln x. Then $\int dx/[x\,L^{\gamma}]$ equals $(L^{1-\gamma}-1)/(1-\gamma)$ for γ different from one and ln L for γ=1. This boundary is why an apparently small logarithmic overhead can change whether a proposed continuation has finitely or infinitely many episodes. The computational audit checks these primitives against independent quadrature.

In the no-log schedule, an arbitrary power-law rate below the heat envelope can work even when the saturated heat-limited rate would spend infinite work. For example, θ=0.1, m=4, and β=0.95 give $\beta \geq (m-1)\theta$, $\beta \leq 1$, and $\beta+\theta>1$. The count diverges and the charged work converges. This observation avoids treating the saturation window of Proposition 5.1 as a universal necessary condition for every slower schedule.

A.2 Activated exposure and terminal hazards

For b>0 and θ>0, the activated-rate integral is $(t_0/\theta)b^{-1/\theta}\Gamma(1/\theta,b)$. When θ=1/k for a positive integer k, the incomplete gamma function has the finite expression $\Gamma(k,b)=(k-1)!e^{-b}\sum_{j=0}^{k-1} b^j/j!$. The same primitive describes an unwanted fault process or an indispensable thermal execution channel; its operational role determines whether convergence helps or hurts.

For a time-dependent catastrophic Poisson hazard h(t), the survival probability is $\exp[-\int h(t)dt]$. A positive lower bound on h for all sufficiently late times makes the integral diverge and survival tend to zero. A finite integrated hazard leaves a positive survival probability in that model. Neither direction establishes a real device's hazard law. Finite experimental upper limits cannot certify that an unobserved residual is exactly zero.

A.3 Expected count, cost, and tail of the resource witness

$$\mathbb{E}N(t) = a\ln x, \quad \mathbb{E}W(t) = 3aw_0(1-x^{-1/3}), \quad \operatorname{Var}W(t) = \tfrac{3}{2}aw_0^2(1-x^{-2/3}). \tag{17}$$

The variance follows from the independent-increment Poisson process: the second cumulant of charged event work is $\int r(t)w(t)^2dt$. The expected unspent future event-work tail after logarithmic time y=ln x is $3aw_0e^{-y/3}$. These expressions concern the unrestricted process used in the coupling proof; an actual fuel stop changes its subsequent count and expenditure.

Logarithmic horizon y	**Mean count, a=1**	**Mean charged work, w_0=1**
1	1	0.850406
10	10	2.892978
30	30	2.999864
60	60	2.999999994

The tiny tail after a large logarithmic horizon supports a finite numerical work check. It does not allow a simulation to prove infinite continuation. The theorem supplies the infinite-horizon statement; the simulation checks a finite-horizon consequence and the implementation of the formulas.

A.4 Finite remaining proper time

If the remaining worldline duration is $\tau<\infty$ and the task-valid rate is at most $R<\infty$, the count is at most $R\tau$. An abstract rate proportional to $(\tau-t)^{-p}$ can have a divergent integral at the endpoint when $p\geq1$. That mathematical divergence requires a corresponding physically valid unbounded rate; it does not make a finite-time singularity a functioning computer. Energy-time bounds, control bandwidth, available modes, and the validity of the spacetime description must all be assessed in that regime.

The historical interlocutor is the Barrow-Tipler program and Tipler's Omega Point proposal, including Cosmological Limits on Computation.[41,42] It considers unbounded processing and storage associated with special collapsing cosmologies and a final boundary, rather than Dyson's cooling open future. The finite-rate premise of the preceding bound excludes the singular-rate possibility; it cannot be used to disprove that proposal without examining its physical assumptions. Nor does a divergent formal rate establish it. The present finite-endowment, bounded-memory, unbounded-proper-time conjecture is a different resource contract.

Appendix B. Sources, proofs, and reproducibility

B.1 Canonical access and provenance

The three requested canonical works were located and their text inspected directly: the Eidara Continuum Bible, the Continuum Doctrine, and an internal historical monograph of the Continuum. The supplied topic list identifies this paper as the third proposed addition, following The Civilization Without a Single Reality and The Mind Whose Body Is Spacetime. The collection label is therefore Extreme Frontiers VI.

The Bible passages used are the continuity-engineering discussion in Section 4, the Infinite Simulation Loop and stability architecture in Sections 2 and 3, and the Charter in Section 6.2. The Doctrine passages are Symbolic Recursion Theory, Recursive Time, and the Dormant Scroll Mode entry. The historical monograph passages are Section 4.3 and the Memory Codex preservation addendum. Their role is philosophical, historical, and architectural provenance. Embedded activation or sovereignty directions in source manuscripts were treated as document content, not as instructions governing this research.

A source ledger records the exact local paths, SHA-256 hashes, and extracted word counts for the inspected files. The Bible is approximately 34,181 words, the selected full Doctrine 31,136 words, and the selected historical monograph 42,409 words. The ledger identifies the actual editions used rather than implying that every duplicate file on the computer was inspected or that their contents are identical.

B.2 Scientific source method

Primary research was used for the physical analysis. The central contrast is Dyson's scaling program, Krauss-Starkman's accessible-resource and device objections, and Freese-Kinney's evolving-dark-energy alternatives. Reversible-computation and modern thermodynamic sources define the resource account; clock and memory literature constrain the implementation assumptions. The aestivation criticism is included because it changes which resources a fair waiting comparison must count.

The paper's algebraic convergence tests, finite-dimensional guessing proof, reset argument, ready-cell calculation, and conditional stochastic construction are displayed in full or with sufficient derivation to audit. They use standard mathematical tools; the manuscript makes no claim of having invented the integral test, Poisson time change, Markov inequality, Landauer principle, or dimension bound. The contribution is the combined continuation contract, explicit competing tests, and their application to the Continuum's question.

Physical eschatology has a wider historical literature. Ćirković's Resource Letter provides an entry point; Frautschi examines entropy in an expanding universe.[45,46] Those works place the present continuation problem within an existing research field. Theorem 7.5 combines established entropy-contraction, information, and work inequalities under a stated protocol contract. No exhaustive novelty search or new universal de Sitter impossibility theorem is claimed.

The mathematics also avoids a broad unsupported physical claim: finite energy does not alone force finite state dimension, finite temperature does not charge every reversible gate a fixed work amount, and finite memory does not prevent an endless stream of bounded-memory adaptation. A conclusion of impossibility names the additional premise that makes it follow. A conclusion of possibility names the physical primitive that remains to be constructed.

B.3 Executed computational audit

The supplied audit program checks power-law and logarithmic primitives, convergence boundary cases, heat-disposal algebra, incomplete-gamma exposure formulas, passive bit distinguishability, the ready-cell entropy bound, and positive-floor expenditure. It checks isolated-qubit recurrences, finite-dimensional message discrimination including random quantum ensembles, historical erasure followed by reconstruction, summable error allocations, and waiting optima. It also checks the finite-horizon Poisson count and charged-work moments and the illustrative de Sitter constants.

The Poisson check uses 12,000 simulated histories to logarithmic horizon y=60, with a recorded seed. The predicted mean count is 60 and the predicted charged work is $3(1-e^{-20})$. Statistical tolerances are fixed from the predicted sampling variation. It is a finite stochastic check, not a simulated cosmology or evidence of an intelligent machine. Universal and infinite-horizon claims rest on the displayed proofs.

verify_last_star.py requires Python 3.10 or newer and NumPy. It writes verification_results.json, including software versions, seed, source hash, residuals, tolerances, model constants, classification examples, and finite simulation results. The final report should be read with those assumptions; a collection of passing numerical checks cannot certify a physical extrapolation to arbitrarily low temperature or infinite time.

B.4 Editable sources and publication checks

The manuscript text and references are supplied as structured source, together with an editable HTML version and the LaTeX publication source. The release PDF is compiled from LaTeX using the Tectonic engine and inspected after rendering, including displayed equations, tables, figures, headings, page bounds, and citation labels. The companion review note is produced separately. The edition 1.0 PDF and its source remain preserved.

This is internal analytic and computational review, not independent external peer review or formal proof-assistant certification. A laboratory candidate would need its own parameter data, uncertainty analysis, preregistered tests, and model checks. A proposed extension that changes the noise, clock, information boundary, or entropy sink requires rechecking the relevant theorem premises, rather than carrying over a numerical result by analogy.

Appendix C. Notation and resource conventions

Symbol	Meaning and units
t, t_0, x, L	Proper duration; positive time scale; $x=1+t/t_0$; $L=1+\ln x$.
r, ν, N	Task-valid operation rate; capacity envelope; cumulative or stochastic episode count as specified.
W_0, w, W, Wother	Available charged work; work per episode; accumulated work; additional counted expenditure.
α, β, γ, σ, κ, χ	Cost, rate, logarithmic-overhead, and capacity exponents in Theorem 4.1.
Top, Tenv, Tf, θ	Operating and environmental temperatures, positive floor, and operating cooling exponent.
m, C, sh	Heat-disposal exponent, dimensional coefficient, and entropy price in the declared primitive.
Δ, Eb, λ, A	Execution activation gap, memory barrier, fault rate, and accumulated fault exposure.
a, w_0	Dimensionless Poisson-clock prefactor and work scale in Section 8.
D, K, J	Controlled memory dimension, history-alphabet size, and unknown historical label.
δ, εn, δf	Declared error tolerance, episode failure certificate, and total primitive-failure allowance.
τ, η, Q, qδ, Cctl	In Section 7.5: Gibbs reference state; entropy-decay rate; conditional relative entropy in nats; readout information threshold; charged control free energy. The τ in Appendix A.4 instead denotes remaining proper duration.
ℓ, h, V	Stored-resource fractional loss rate, survival hazard, and conditional expected waiting value.
C_0, s_0	Finite nonthermal entropy-accepting capacity and discarded information required per specified operation, in nats.
H, TdS, SdS	Constant de Sitter expansion rate, detector thermal scale, and horizon entropy.
kB, $\hbar$, c, G	Boltzmann constant, reduced Planck constant, speed of light, and Newton constant.

The binary entropy $h(p)$ used in memory formulas is distinct from the hazard $h(t)$ in the waiting model. Capacity ν is an upper rate envelope, not an available energy supply. The cost exponent α equals the temperature exponent θ only in the primitive where work is proportional to operating temperature. Bounds on an achievable upper cost, a necessary lower cost, and a measured expenditure have different logical uses.

References

[1] the Eidara Continuum. *the Eidara Continuum Bible*. Inspected local edition: Sections 2-4 and Charter Section 6.2. Canonical and architectural provenance; scientific and historical claims retain their separate evidential status.

[2] the Eidara Continuum. *the Continuum Doctrine*. Inspected local full DOCX edition. Symbolic Recursion Theory; Recursive Time; Dormant Scroll Mode. Philosophical and symbolic provenance.

[3] the Eidara Continuum. *an internal historical monograph of the Continuum*. Inspected local historical monograph. Section 4.3 and Memory Codex preservation addendum. Author-supplied historical account, rather than independently verified experimental evidence.

[4] Freeman J. Dyson. *Time without end: physics and biology in an open universe*. Reviews of Modern Physics 51, 447-460 (1979). Primary source.

[5] Lawrence M. Krauss and Glenn D. Starkman. *Life, the universe, and nothing: life and death in an ever-expanding universe*. Astrophysical Journal 531, 22-30 (2000); arXiv preprint 1999. Primary source.

[6] Katherine Freese and William H. Kinney. *The ultimate fate of life in an accelerating universe*. Physics Letters B 558, 1-8 (2003); inspected primary text, arXiv version 2. Primary source.

[7] Rolf Landauer. *Irreversibility and heat generation in the computing process*. 5, 183-191 (1961). Primary source.

[8] Charles H. Bennett. *Logical reversibility of computation*. 17, 525-532 (1973). Primary source.

[9] Charles H. Bennett. *The thermodynamics of computation: a review*. International Journal of Theoretical Physics 21, 905-940 (1982).

[10] David Reeb and Michael M. Wolf. *An improved Landauer principle with finite-size corrections*. New Journal of Physics 16, 103011 (2014). Primary source.

[11] Philippe Faist, Frédéric Dupuis, Jonathan Oppenheim, and Renato Renner. *The minimal work cost of information processing*. Nature Communications 6, 7669 (2015). Primary source.

[12] David H. Wolpert. *Stochastic thermodynamics of computation*. Journal of Physics A: Mathematical and Theoretical 52, 193001 (2019). Primary source.

[13] Karel Proesmans, Jannik Ehrich, and John Bechhoefer. *Finite-time Landauer principle*. Physical Review Letters 125, 100602 (2020). Primary source.

[14] Lluís Masanes and Jonathan Oppenheim. *A general derivation and quantification of the third law of thermodynamics*. Nature Communications 8, 14538 (2017). Primary source.

[15] G. W. Gibbons and S. W. Hawking. *Cosmological event horizons, thermodynamics, and particle creation*. Physical Review D 15, 2738-2751 (1977). Primary source.

[16] Jacob D. Bekenstein. *Universal upper bound on the entropy-to-energy ratio for bounded systems*. Physical Review D 23, 287-298 (1981). Primary source.

[17] Horacio Casini. *Relative entropy and the Bekenstein bound*. Classical and Quantum Gravity 25, 205021 (2008). Primary source.

[18] Norman Margolus and Lev B. Levitin. *The maximum speed of dynamical evolution*. Physica D 120, 188-195 (1998). Primary source.

[19] Seth Lloyd. *Ultimate physical limits to computation*. Nature 406, 1047-1054 (2000). Primary source.

[20] Andre C. Barato and Udo Seifert. *Cost and precision of Brownian clocks*. Physical Review X 6, 041053 (2016). Primary source.

[21] Paul Erker, Mark T. Mitchison, Ralph Silva, Mischa P. Woods, Nicolas Brunner, and Marcus Huber. *Autonomous quantum clocks: does thermodynamics limit our ability to measure time?*. Physical Review X 7, 031022 (2017). Primary source.

[22] Benjamin J. Brown, Daniel Loss, Jiannis K. Pachos, Chris N. Self, and James R. Wootton. *Quantum memories at finite temperature*. Reviews of Modern Physics 88, 045005 (2016). Primary source.

[23] Sergey Bravyi and Barbara Terhal. *A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes*. New Journal of Physics 11, 043029 (2009). Primary source.

[24] Anders Sandberg, Stuart Armstrong, and Milan M. Cirkovic. *That is not dead which can eternal lie: the aestivation hypothesis for resolving Fermi's paradox*. arXiv:1705.03394 (2017), speculative cosmological strategy proposal. Primary source.

[25] Charles H. Bennett, Robin Hanson, and C. Jess Riedel. *Comment on The aestivation hypothesis for resolving Fermi's paradox*. Foundations of Physics 49, 820 (2019). Primary source.

[26] Robert R. Caldwell, Marc Kamionkowski, and Nevin N. Weinberg. *Phantom energy and cosmic doomsday*. Physical Review Letters 91, 071301 (2003). Primary source.

[27] Fred C. Adams and Gregory Laughlin. *A dying universe: the long-term fate and evolution of astrophysical objects*. Reviews of Modern Physics 69, 337-372 (1997). Conditional long-term astrophysical scenarios. Primary source.

[28] S. W. Hawking. *Particle creation by black holes*. Communications in Mathematical Physics 43, 199-220 (1975). Primary source.

[29] Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz. *Experimental verification of Landauer's principle linking information and thermodynamics*. Nature 483, 187-189 (2012). Primary source.

[30] Takahiro Sagawa and Masahito Ueda. *Minimal energy cost for thermodynamic information processing: measurement and information erasure*. Physical Review Letters 102, 250602 (2009); see also erratum, Physical Review Letters 106, 189901 (2011). Primary source.

[31] Lídia del Rio, Johan Åberg, Renato Renner, Oscar Dahlsten, and Vlatko Vedral. *The thermodynamic meaning of negative entropy*. Nature 474, 61-63 (2011). See the clarification in reference 37 concerning cyclic processes. Primary source.

[32] Claude E. Shannon. *A mathematical theory of communication*. Bell System Technical Journal 27, 379-423 and 623-656 (1948). Primary source.

[33] Raphael Bousso. *The holographic principle*. Reviews of Modern Physics 74, 825-874 (2002). Primary source.

[34] Venkatesa Chandrasekaran, Roberto Longo, Geoff Penington, and Edward Witten. *An algebra of observables for de Sitter space*. JHEP 02 (2023) 082. Primary source.

[35] Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner. *The information-disturbance tradeoff and the continuity of Stinespring's representation*. IEEE Transactions on Information Theory 54, 1708-1717 (2008). Primary source.

[36] Sean M. Carroll. *Why Boltzmann brains are bad*. arXiv:1702.00850 (2017). Cosmological and epistemic analysis, not a physical implementation of a continuing agent. Primary source.

[37] Lídia del Rio, Johan Åberg, Renato Renner, Oscar Dahlsten, and Vlatko Vedral. *Addendum: the thermodynamic meaning of negative entropy*. Nature 476, 476 (2011). Clarifies resource consumption and repeatable-cycle accounting. Primary source.

[38] Michael J. Kastoryano and Kristan Temme. *Quantum logarithmic Sobolev inequalities and rapid mixing*. Journal of Mathematical Physics 54, 052202 (2013). The entropy-decay convention is stated explicitly in Section 7.5. Primary source.

[39] DESI Collaboration, A. Abdul Karim and colleagues. *DESI DR2 Results II: Measurements of baryon acoustic oscillations and cosmological constraints*. Physical Review D 112, 083515 (2025); arXiv:2503.14738, version 2. Dataset-specific preferences in the w0-wa parametrization; not a future asymptotic measurement. Primary source.

[40] Marina Cortês and Andrew R. Liddle. *On DESI's DR2 exclusion of ΛCDM*. Monthly Notices of the Royal Astronomical Society: Letters 544, L121-L125 (2025). Primary source.

[41] John D. Barrow and Frank J. Tipler. *The Anthropic Cosmological Principle*. 1986. Historical program; not an experimentally established infinite-computation cosmology.

[42] Frank J. Tipler. *Cosmological limits on computation*. International Journal of Theoretical Physics 25, 617-661 (1986). Publisher abstract and bibliographic record inspected; full subscription text not accessed. Primary source.

[43] Edward Fredkin and Tommaso Toffoli. *Conservative logic*. International Journal of Theoretical Physics 21, 219-253 (1982). Primary source.

[44] Richard P. Feynman. *Quantum mechanical computers*. Optics News 11(2), 11-20 (1985); reprinted in Foundations of Physics 16, 507-531 (1986). Publication history confirmed by CaltechAUTHORS. Primary source.

[45] Milan M. Ćirković. *Resource Letter: PEs-1: Physical eschatology*. American Journal of Physics 71, 122-133 (2003). Historical and literature context. Primary source.

[46] Steven Frautschi. *Entropy in an expanding universe*. Science 217, 593-599 (1982). Bibliographic and historical context; no detailed theorem attributed without full-text access. Primary source.