Facts, Admissibility, and the Physical Inadmissibility of the Last Thursday Hypothesis

Keith Taylor VERSF Theoretical Physics Program

Abstract

The Last Thursday hypothesis claims that the universe could have been created arbitrarily recently, complete with fabricated memories and records, while remaining observationally indistinguishable from an older universe. Traditionally dismissed on methodological grounds alone, the hypothesis has never been ruled out physically. We show, under independently motivated physical constraints on irreversible commitment, that it is incompatible with any framework in which facts correspond to irreversible commitments with finite entropy cost — a characterization independently motivated by thermodynamic, information-theoretic, and computational considerations, developed here within the Void Energy–Regulated Space Framework (VERSF).

We derive a quantitative *fact capacity bound* on the number of irreversible commitments any finite event can produce, grounded in Landauer's principle, the Margolus–Levitin quantum speed limit, and the Bekenstein bound. For any creation event short relative to the age of the universe, this bound falls orders of magnitude below the structured content the universe exhibits. The hypothesis is therefore not merely unfalsifiable but physically inadmissible: it demands structure that cannot exist without the processes that generate it. This result extends to Boltzmann brain scenarios, retroactive simulation arguments, and all cosmologies that postulate structured pasts without entropy cost.

Contents

---

1. What the Last Thursday Hypothesis Asserts

The Last Thursday hypothesis proposes that the universe was created very recently — last Thursday, five minutes ago, or at any arbitrary moment — fully equipped with apparent memories, historical records, fossils, isotope ratios, cosmic microwave background anisotropies, and every other trace of a deep past already in place. Under this hypothesis, no observational experiment can distinguish such a universe from one that genuinely evolved over 13.8 billion years.

The hypothesis is logically consistent and empirically irrefutable by construction. Its traditional role in philosophy of science is not to be believed, but to expose the limits of evidence-based inference and to highlight the dependence of scientific reasoning on non-empirical constraints — simplicity, explanatory power, and resistance to ad hoc construction. The hypothesis also has real-world antecedents: Philip Gosse's *Omphalos* (1857) argued that God created the world with the appearance of a prior history, including fossils and geological strata, as a necessary feature of a functioning creation. The philosophical structure of the Last Thursday hypothesis is identical to the Omphalos argument, stripped of its theological specifics.

What the hypothesis specifically asserts is the conjunction of three claims:

(C1) All observable records are internally consistent with an apparent deep history.

(C2) These records need not correspond to any actual past — they may have been instantiated wholesale at the moment of creation.

(C3) There is no physical distinction between a universe with a real history and one instantiated with a fabricated history that produces identical present-state observations.

It is the third claim that we will show to be false. We note that (C2) is not merely a philosophical gambit — it is a covert claim about the universe's ontological status. A physical system whose structured content was not generated by its own dynamics is equivalent to what contemporary simulation discussions call a *shortcut simulation*: a data structure whose apparent history is injected rather than computed. The Last Thursday hypothesis is the limiting case of any cosmology that postulates structured pasts without commensurate irreversible process, and the argument developed here applies to all such cosmologies (Section 8).

---

2. Why the Hypothesis Has Survived

The Last Thursday hypothesis has persisted because it is designed to be immune to empirical refutation. Every piece of evidence one might cite against it — a fossil, a memory, a distant galaxy — is, by construction, part of the fabrication. The standard philosophical response is methodological: the hypothesis is maximally ad hoc, explains nothing, predicts nothing novel, and merely duplicates apparent structure without generative mechanism. Occam's razor, Bayesian parsimony, and explanatory economy all militate against it.

But these are principles of theory selection, not physical laws. They tell us which theory to *prefer*, not which states of affairs are *possible*. The hypothesis survives because physics, as conventionally framed, treats information as descriptive rather than constitutive. If the state of the universe is just a point in phase space, then any point is in principle accessible, and the question of how it got there is a matter of history, not law. We argue that this framing is incomplete.

The idea that irreversibility is constitutive rather than merely emergent has precedents in several research traditions. Prigogine's work on dissipative structures [11] argued that irreversible processes are not approximations to underlying reversible dynamics but generators of macroscopic order, with the arrow of time as a fundamental feature of nature rather than a statistical artifact. More recently, Deutsch and Marletto's constructor theory [12] reformulates physics in terms of which transformations are possible and which are impossible, treating the distinction between possible and impossible tasks as fundamental rather than derived. VERSF shares with both programmes the conviction that irreversibility is physically constitutive, but differs in grounding this conviction specifically in entropy-bearing commitments with finite distinguishability cost — a formulation that yields the quantitative bounds developed in Section 7.

---

3. Why This Is Not a Philosophical Argument

The standard philosophical dismissal of the Last Thursday hypothesis is methodological: the hypothesis is unfalsifiable and ad hoc, so we should not accept it. This leaves the hypothesis standing as a logical possibility — embarrassing, perhaps, but not defeated.

The argument developed in the following sections is physical, not philosophical. It does not say that the Last Thursday hypothesis is a bad explanation — it says that the state of affairs the hypothesis describes cannot be physically instantiated under any framework in which facts are irreversible commitments with finite cost. This is a resource argument, directly analogous to thermodynamic constraints on computation. Throughout this paper, "physically impossible" means inadmissible under the resource and transition constraints imposed by the physical laws governing entropy, computation, and state evolution — not logically contradictory. The Last Thursday hypothesis remains a logical possibility in the narrow sense that it violates no rule of formal logic. What it violates are the physical constraints that determine which states of affairs can be instantiated by any finite process. The impossibility is of the same kind as a perpetual motion machine: not a contradiction in terms, but a violation of the laws that govern what physical systems can do. Landauer's principle establishes that erasing one bit of information requires a minimum energy dissipation of $k_BT \ln 2$ — a physical law, not a practical limitation, confirmed experimentally [4]. In the same way, minimum distinguishability cost establishes that creating a fact requires a minimum entropy expenditure. A universe full of facts that were never physically created is as impossible as a computation that erased information without generating heat.

---

4. The VERSF Premise: What a Fact Is

The resolution begins with a precise physical definition of *fact*.

In the Void Energy–Regulated Space Framework, a fact is not a proposition, not a description, and not a state label. A fact is an irreversible physical commitment: a transition that consumes finite distinguishability, generates entropy, and persists across subsequent time ordering. By *finite distinguishability* we mean the principle that every physical distinction — every transition from one macroscopically distinguishable state to another — carries a minimum entropy cost, so that creating, maintaining, or erasing a record requires thermodynamic work that cannot be reduced below a positive lower bound. (In Landauer's formulation, this bound is $k_BT \ln 2$ per bit erased; the VERSF term "finite distinguishability" generalizes this to all irreversible commitments.) We use *minimum distinguishability cost* when invoking this as a general physical principle independent of VERSF, and reserve *finite distinguishability* for the full VERSF formalization. The distinction matters because the quantitative bounds of Section 7 require only the general principle, while the constitutive claims of Section 5.4 draw on the stronger VERSF framework. Facts are constitutive elements of physical reality, not annotations upon it.

To give this definition operational content that does not presuppose VERSF:

Definition (Fact). A fact is a physical record that is (i) redundantly encoded across many degrees of freedom, (ii) robust under coarse-graining and local perturbation, and (iii) irreversible in the sense that its erasure requires macroscopic entropy export to the environment. In operational terms: a fact is what remains invariant across repeated re-preparations of the local environment and serves as a stable constraint on future evolution.

This operational definition connects naturally to well-established physics. Condition (i) reflects the redundancy requirement of quantum Darwinism — information becomes a "fact" about a quantum system precisely when it has been proliferated into many environmental degrees of freedom [7]. Condition (ii) captures the thermodynamic robustness that distinguishes macroscopic records from microscopic fluctuations. Condition (iii) is a direct consequence of Landauer's principle: a record whose erasure costs entropy is a record that was created by an irreversible process. Together, these conditions identify "fact" with the class of physical records that cannot be undone cheaply — the class whose existence requires, and testifies to, irreversible thermodynamic work.

Before developing the consequences of this definition within VERSF, we note that it is independently motivated by three convergent lines of reasoning:

Thermodynamic motivation. Landauer's principle establishes that logically irreversible operations — those that destroy information — require a minimum entropy dissipation of $k_BT \ln 2$ per bit [1]. Any physical process that resolves an alternative — that commits the universe to one outcome rather than another — is logically irreversible and therefore thermodynamically costly. A "fact" in the sense used here is precisely such a resolution. The claim that facts have entropy cost is therefore not a VERSF postulate but a consequence of established thermodynamics, extended to its natural generality.

Information-theoretic motivation. A physical system's state can be described by its Shannon entropy, which measures the number of distinguishable configurations available to it. Each irreversible commitment reduces the number of accessible configurations and thereby increases the thermodynamic entropy of the environment. The creation of physical specificity — the transition from "could be A or B" to "is A" — is an information-theoretic act with physical consequences. VERSF formalizes this by treating distinguishability as a finite, conserved resource.

Computational motivation. The theory of computation imposes resource constraints on physical processes. A Turing machine that produces an output of Kolmogorov complexity K must execute at least K steps. A physical system that exhibits structure of complexity K must have undergone at least K irreversible transitions to produce that structure — or it must have received that structure as input from a system that did. There is no free lunch: structure requires process. This perspective aligns with Wheeler's "it from bit" programme [10], which holds that the physical world derives its existence from information-theoretic acts of registration. VERSF shares Wheeler's conviction that information is physically constitutive rather than merely descriptive, but departs from the participatory observer framework by grounding constitution in thermodynamic irreversibility rather than measurement: a fact becomes real not because an observer registers it, but because an irreversible transition has committed the universe to it at finite entropy cost.

VERSF synthesizes these motivations into a unified physical framework. The definition of fact carries three immediate consequences:

Facts have entropy cost. Every committed fact corresponds to an irreversible transition that dissipates distinguishability. A universe containing N committed facts must have paid an entropy cost proportional to N. This cost cannot be waived, deferred, or circumvented — it is the price of physical specificity.

Facts define time ordering. In VERSF, time is not a background parameter along which events are arranged. Time is the accumulated record of irreversible commitments. Each fact extends the causal chain; the ordering of facts constitutes temporal structure. A state with no committed facts has no time depth — it exists at a single moment, regardless of what labels or records it appears to contain.

Facts are not reducible to state encodings. A present-state configuration that appears to encode historical information is not physically equivalent to a state that was generated by that history. A state can *encode* the pattern of a fact — it can contain a data structure that looks like a record — but encoding is not commitment. Commitment requires entropy expenditure; encoding requires only bit arrangement. The distinction is not between what a state *says* and what a state *means*, but between what a state *contains* and what a state *has done*. Only committed facts constrain future admissible transitions; mere encodings carry no such constraint.

---

5. Why Fabricated Histories Are Physically Inadmissible

With these definitions in place, the Last Thursday hypothesis becomes not merely unattractive but physically inadmissible. We state the central claim precisely:

Claim. In any framework where facts are physically realized as irreversible, entropy-bearing commitments with finite cost — as characterized by the operational definition above and formalized within VERSF — the Last Thursday hypothesis is not merely unfalsifiable but inadmissible: the state of affairs it describes cannot be physically instantiated.

The hypothesis requires the universe to satisfy four conditions simultaneously, and these conditions are mutually incompatible.

5.1 The Entropy Constraint

The present state of the universe contains an enormous quantity of committed structure: particle positions, chemical abundances, gravitational configurations, biological complexity, neural states encoding memories. Each of these represents resolved distinguishability — a point at which the universe transitioned from multiple accessible configurations to one specific configuration.

The total entropy cost of this structure is the integral of all irreversible commitments required to produce it. The observable universe has a current entropy of approximately $S \sim 10^{104}$ k_B, dominated by supermassive black holes, with the cosmic microwave background contributing approximately $S_{CMB} \sim 10^{89}$ k_B [2]. A careful accounting is needed here: black hole entropy represents gravitational commitment capacity rather than "facts" in the operational sense of Section 4, and not every unit of thermodynamic entropy corresponds to a redundantly encoded, robust record. Nevertheless, the present state encodes an astronomically large set of stable records at every scale — from primordial nucleosynthetic abundances to geological strata to neural configurations — and the entropy associated with these records alone vastly exceeds what any brief creation event could produce. Even restricting attention to diffuse baryonic structure and radiation ($\sim 10^{89}$ k_B), the number of irreversible commitments required is enormous.

Throughout this paper, thermodynamic entropy is used as a scale-setting proxy rather than a literal count of operational facts. To ground the comparison, we note that a conservative lower bound on genuinely operational facts — stable baryonic records ($\sim 10^{80}$ nucleon identities fixed during nucleosynthesis), radiative correlations ($\sim 10^{68}$ CMB photons carrying recombination-era information), and long-lived classical configurations ($\sim 10^{22}$ stellar objects recording fusion histories) — already exceeds $\sim 10^{68}$, even excluding gravitational entropy and black holes. This bound alone suffices for the resource arguments that follow. Where we later compare against $\sim 10^{89}$ or $\sim 10^{104}$, the comparison is illustrative rather than identificational — it sets the scale of the resource gap, not the exact number of facts.

The Last Thursday hypothesis asserts that this cost was never paid — that the structure was instantiated wholesale without the entropy-generating processes that would ordinarily produce it. Under minimum distinguishability cost, this is not possible. Committed structure without entropy expenditure is a physical contradiction, directly analogous to a completed computation that consumed no energy. Just as Landauer's principle forbids information erasure without heat dissipation, minimum distinguishability cost forbids fact creation without entropy generation.

5.2 The Causal Ordering Constraint

The facts encoded in the present state of the universe are not independent — they are causally ordered. Fossils presuppose sedimentation; isotope ratios presuppose nuclear decay chains; memories presuppose neural encoding events; the cosmic microwave background presupposes a surface of last scattering; the abundance of light elements presupposes primordial nucleosynthesis.

These causal dependencies are not merely descriptive patterns — they are not narratives that happen to fit the data. They reflect the order in which irreversible commitments were made. The ^{14}C ratio in an ancient bone is not a number that happens to be consistent with a 5,000-year decay history. It is the *result* of 5,000 years of irreversible nuclear transitions, each of which consumed distinguishability and extended the causal record.

The Last Thursday hypothesis requires all of these causal dependencies to hold as structural features of the present state while denying the causal processes that would establish them. A defender might respond that these are simply correlated present-state conditions — patterns in the microstate that require no causal explanation beyond the specification of the state itself. We argue that in a finite-distinguishability framework, correlation without generative commitment is not physically admissible: the correlations are not merely consistent with a causal history but constituted by one, because each correlation is underwritten by a chain of irreversible transitions whose entropy cost is part of the physical content of the state.

Under VERSF, causal ordering is not a narrative overlay — it is the physical content of time. A state that exhibits causal structure without causal history does not merely lack an explanation; it lacks physical coherence. It is a state that refers to commitments that were never made, entropy that was never generated, and a temporal depth that does not exist.

5.3 The Compression Constraint

The third constraint concerns computational depth. Bennett's logical depth [3] measures not the *size* of the shortest program that generates a state (that is Kolmogorov complexity), but the *time* that program requires to execute. A state has high logical depth if every short program that produces it must run for many steps. The distinction matters: a high-entropy thermal equilibrium state has low Kolmogorov complexity (a short statistical specification suffices) and low logical depth (it can be produced quickly by random processes). The present universe is different. It has moderate Kolmogorov complexity — the laws of physics plus initial conditions provide a compact generating description — but enormous logical depth, because that compact description must be iterated forward through an astronomically large number of sequential physical updates (of order the universe's age in fundamental time units), producing enormous logical depth.

The Last Thursday hypothesis requires a creation event that produces a state of enormous logical depth in negligible time. It must generate the output of a deep computation — the full structured content of the present universe, including every correlation, decay chain, and nucleosynthetic abundance — without executing the computation. This is precisely what logical depth forbids. A state's logical depth is a property of the state itself, not of the observer: it reflects the minimum computational work required to produce that state from any compact description. No shortcut exists — not because we lack ingenuity, but because the structure of the output encodes the computational history that produced it.

A critic might object: could the creation event be governed by a compact rule that produces the same output as 13.8 billion years of evolution, but faster? This collapses under examination. If such a rule exists and the creation event executes it, then the creation event must perform computational work proportional to the logical depth of the output. The compactness of the rule does not reduce the number of steps required to run it — a short program that produces a deep output still requires many steps.

A subtlety deserves attention here. Bennett's own framework distinguishes reversible from irreversible computation, and in principle any computation can be performed reversibly — at the cost of retaining all intermediate results rather than erasing them. If the creation event were a reversible computation, Landauer's principle would not apply to the computation itself; thermodynamic cost would arise only at the final stage of erasing the intermediate workspace. This means the thermodynamic and computational arguments are not fully independent: a reversible creation event could in principle defer all entropy costs to a final erasure step.

However, two constraints remain binding regardless of reversibility. First, the Margolus–Levitin bound (Section 7.3) limits the *rate* of state transitions irrespective of whether those transitions are reversible or irreversible. A reversible computation of depth D still requires at least D sequential steps, each taking minimum time $\pi\hbar/2E$. The speed limit is on transitions per unit time, not on entropy per transition. Second, the final erasure of intermediate results — necessary to produce a clean output state rather than a state entangled with a massive computational workspace — itself incurs Landauer costs proportional to the information discarded. A reversible computation that produces the present universe and then erases its workspace has merely *deferred* the entropy cost, not eliminated it.

A further subtlety: could the creator retain the workspace without erasing it, thereby avoiding Landauer costs entirely? This is logically possible, but self-defeating in two respects. First, a retained workspace is a complete computational record of every intermediate step in the universe's evolution — every state transition, every branching point, every irreversible commitment in sequence. If this record exists (in the creator's domain or on some external substrate), then the creation event has produced not just the universe's present state but a full history of how that state was reached. The "fabrication" includes its own documentation.

A Last Thursday proponent might respond: the workspace is in the creator's domain, not in our universe — so our universe still has no internal genuine past. But this response concedes too much. If the workspace faithfully records the computational history that produced our universe's present state, then our universe's present state *was* produced by a genuine computational process — it is the output of a deep computation that actually executed. The question of where the computation's intermediate records are stored is a question about the *location* of the history, not its *existence*. Our universe's present state bears the same relationship to the workspace that any computation's output bears to its execution trace: the output's structure was produced by the process recorded in the trace. The Last Thursday hypothesis does not merely claim that the historical records are somewhere else — it claims that no historical process occurred at all. A creator who retains a complete workspace has refuted that claim by performing the very process the hypothesis denies. (This connects to the treatment in Q4, where we argue that a creation event with genuine temporal depth in the creator's frame constitutes a genuine causal history, relocated rather than fabricated.)

The compression constraint therefore rests primarily on the Margolus–Levitin bound (which is indifferent to reversibility) and secondarily on Landauer's principle (which applies to the inevitable erasure stage, or is circumvented only by retaining a genuine computational history).

The connection between computational structure and thermodynamic cost has independent support. Zurek [13] established that physical entropy and algorithmic complexity are related quantities, with physical entropy admitting a decomposition into contributions from algorithmic randomness (the incompressible information content of the microstate) and other terms reflecting the observer's coarse-graining. This result matters for the present argument because it connects the operational definition of fact to information-theoretic cost: a committed record that is redundantly encoded and robust under coarse-graining (conditions (i) and (ii) of Section 4) necessarily contributes to the algorithmic complexity of the macrostate, and Zurek's decomposition shows that this contribution has irreducible thermodynamic weight. A universe full of operationally robust facts is a universe whose macrostate has high algorithmic complexity — and that complexity is physically costly. Combined with Bennett's logical depth — which measures the computational work required to *produce* a state from a compact description — the implication is that states of high logical depth require many sequential computational steps regardless of whether those steps are individually reversible or irreversible. Zurek provides the entropy-complexity link; Bennett provides the depth-work link; the Margolus–Levitin bound provides the time-step link. Together they ground the claim that computational depth has irreducible temporal cost, with thermodynamic cost arising inevitably at the erasure stage.

The compression constraint therefore does not rest on the claim that no short program can describe the present universe. It rests on the claim that no short *execution* can produce it. The Last Thursday hypothesis attempts to collapse logical depth into a single boundary condition — to obtain the output of a deep computation without performing the computation. Under finite distinguishability, this is precisely what cannot be done.

5.4 The Lineage Constraint (Conditional)

The preceding three constraints address the *cost* of fabricating a universe. A fourth constraint addresses the *coherence* of such a fabrication.

Any fabricated history must contain a first impostor fact: a structure that presents itself as a committed record with causal antecedents but was in reality placed into existence without any generative process. Under the admissibility view developed here, this first impostor fails to qualify as a fact in the strong sense relevant to time-ordering. It may satisfy the operational conditions of Section 4 — redundancy, robustness, erasure cost — but it lacks the generative commitment chain that constitutes factual depth. The operational definition identifies facts by their surface properties; admissibility adds a constitutive requirement — that these properties were *produced* by the irreversible processes whose signatures they bear.

This constitutive requirement goes beyond what the operational definition alone demands. We argue that it is not an independent axiom but a consequence of the operational definition when applied recursively at all scales: satisfying conditions (i)–(iii) for every record and every environmental correlate of every record requires physical work commensurate with genuine production, because the correlations *are constituted by* the thermodynamic traces of the irreversible transitions, and producing them requires displacing the same degrees of freedom.

We state this formally:

Conjecture 1 (Thermodynamic cost of correlation reproduction). *Let $\mathcal{C}$ be a correlation web — the full mutual information structure among N subsystems, specified at the level of all k-body correlations (not merely pairwise) — generated by a physical process at total thermodynamic cost $W(\mathcal{C})$. Then no alternative physical process can reproduce $\mathcal{C}$ at cost $W' < W(\mathcal{C})$. That is, $W(\mathcal{C})$ is a lower bound on the thermodynamic cost of any process that instantiates $\mathcal{C}$, not merely the cost of the historically realized process.*

A note on the level of description: $\mathcal{C}$ here means the full mutual information structure — all orders of correlation, not just pairwise correlation functions. Matching only pairwise correlations is a weaker requirement that could plausibly be achieved at lower cost; the conjecture targets the full structure because this is what the operational definition demands (condition (i) requires redundancy across *many* degrees of freedom, which encodes higher-order correlations).

A proof would close the operational-to-constitutive gap entirely. We regard the conjecture as very likely true on the basis of the scaling arguments of Section 7, but we acknowledge it is unproven. A caveat: the conjecture as stated may require restriction to correlation webs generated by local interactions. Non-local correlation generation — where subsystems A and B are independently coupled to a common source C rather than interacting directly — can in principle produce identical pairwise correlations at different thermodynamic costs depending on the topology of the generating process. For the cosmological case relevant here, where correlations are overwhelmingly generated by local causal processes propagating through spacetime, this subtlety does not arise. But a fully general proof would need to address the topology-dependence of correlation cost. The paper's headline claim — that the Last Thursday hypothesis is physically inadmissible — rests on the quantitative bound and the qualitative constraints of Sections 5.1–5.3, which do not require Conjecture 1. The lineage argument that follows is a conditional strengthening: if the conjecture holds, fabricated histories are not merely resource-prohibitive but incoherent.

The lineage constraint yields a dilemma:

(L1) If the first impostor facts are not genuine facts (because they lack causal lineage), then no subsequent structure built on them constitutes genuine committed fact either. The entire fabricated universe is a chain of encodings rooted in a non-fact — physically complete in its particle content but empty of factual depth. In such a universe, no genuine irreversible commitment was ever made, so no fact constrains future evolution in the way the second law requires. Since we observe a universe in which entropy increases, records persist, and committed structure resists erasure, we do not inhabit an (L1) universe.

(L2) If the creation event produces structures that *do* qualify as genuine facts — with real entropy cost, real causal depth, real irreversibility — then it has done real physical work. The "fabrication" is indistinguishable from genuine history not because the fabrication is perfect, but because it *is* genuine history. The creator has not faked a past; the creator has *made* a past.

Under (L1), the Last Thursday universe is not the universe we observe. Under (L2), the Last Thursday universe has a genuine past after all. Either way, the hypothesis fails — not because fabrication is difficult, but because the concept of a fabricated fact is incoherent. A fact that was faked is either not a fact or not faked.

The committed/placed distinction is therefore empirically accessible precisely when the fabrication is incomplete, and physically redundant when the fabrication is total — because a total fabrication is genuine production. This connects directly to the argument of Section 6, where we show that indistinguishability at the single-time-slice level does not entail physical equivalence at the system-plus-environment level.

As emphasized in Section 3, the impossibility here is physical, not logical — of the same kind as a perpetual motion machine.

---

6. Indistinguishability Does Not Imply Equivalence

The deepest error in the Last Thursday hypothesis is the assumption that observational indistinguishability entails physical equivalence — that if two states yield identical predictions for all possible measurements, they are physically the same. A thought experiment makes the problem concrete.

Suppose two identical laboratories each contain a sealed box of radioactive isotopes. In Laboratory A, the isotopes were produced by genuine nuclear processes over many years, and the surrounding environment contains the correlated decay products, scattered photons, and thermodynamic traces of those processes. In Laboratory B, the isotopes were placed into the box one second ago by a hypothetical Last Thursday mechanism, complete with apparent decay-chain signatures, but with no environmental correlations beyond the box walls. At the moment of comparison, the two boxes are indistinguishable by any measurement confined to the box. But open the system to its environment: Laboratory A's isotopes are correlated with external records — detector logs, environmental radiation fields, thermal signatures in the shielding — that were produced by the same irreversible chain. Laboratory B's isotopes have no such correlations. The difference is not in the box but in the web of committed records connecting the box to its causal past. A single-time-slice measurement of the box sees no difference; a measurement of the box-plus-environment system does.

The thought experiment illustrates a general principle. Under finite distinguishability, physical identity is not observational identity. Two states can produce identical measurement outcomes while differing in their entropy history, their causal depth, and their admissibility status. The distinction is not epistemic but ontic: it concerns what the universe *has done*, not merely what it *looks like*. More precisely, the distinction is between:

- **Equivalence under single-time observables**: two states yield identical expectation values for all operators defined on a single time-slice.
- **Inequivalence under admissible transition constraints**: the two states differ in which future transitions are physically permitted, because only one carries the committed structure that constrains subsequent evolution.

This claim requires defense, since it appears to violate a widely held principle in physics: that states with identical observable consequences are physically identical. We offer two responses.

First, the principle as stated conflates two senses of "observable." In the narrow sense, an observable is a Hermitian operator on a Hilbert space, and two states that yield identical expectation values for all observables are indeed the same state. But this definition already presupposes that the state space is the correct description of physical reality. Under VERSF, the state space is not fundamental — it is an effective description that emerges from underlying irreversible commitments. Two configurations that yield identical expectation values within the effective state space may differ in their commitment histories, just as two identical files can exist — one produced by a lengthy computation, one injected by direct copy — matching in content while differing entirely in the resource history that produced them.

Second, the distinction between committed and uncommitted structure has observable *consequences*, even if it has no observable *signatures* at a single time-slice. A state with genuine causal depth admits certain future transitions and forbids others, because its committed facts constrain subsequent evolution. A state with identical present-time measurements but no causal depth — a "printed" state — would not carry these constraints. The difference is physical and, in principle, detectable across time, even if it is invisible at any single moment.

The claim that admissible future transitions differ between committed and placed states is framework-dependent: it holds in theories where irreversibility is constitutive (such as VERSF), but not in formulations that treat the instantaneous state as physically complete. The force of the argument is therefore not that standard quantum mechanics is internally inconsistent, but that it is incomplete with respect to the physical origin of irreversibility — that the question "how did this state come to be?" has physical content that the state vector alone does not capture.

A toy model makes this concrete. Consider two identical spin chains, each in a state with alternating up-down magnetization. Chain A reached this state through a sequence of irreversible spin-flip transitions, each dissipating energy into a thermal bath; the bath now carries a correlated thermal record of each transition. Chain B was prepared instantaneously by an external field that imposed the same spin configuration, with no bath interaction. At $t = 0$, the two chains are identical under any measurement restricted to the chain itself. But under any dynamics that couples the chain to its bath — including the dynamics that govern subsequent decoherence, thermalization, and measurement — the two systems diverge. Chain A's bath contains correlated records that constrain future evolution (the spin-bath entanglement structure restricts which transitions are thermodynamically accessible). Chain B's bath is uncorrelated; the chain is free to thermalize along any pathway. The distinction is not metaphysical — it is encoded in the joint state of the system-plus-environment, and it has measurable consequences for the statistics of future transitions. The committed state carries constraints that the placed state does not.

This also addresses the most pointed version of the objection to premise (iii): that the committed/placed distinction is itself unfalsifiable — that we have introduced a metaphysical distinction dressed in thermodynamic language, mirroring exactly the defect we attribute to the Last Thursday hypothesis. The charge is serious and deserves a direct answer.

The distinction is not unfalsifiable; it is falsifiable at the system-plus-environment level. A committed fact, by the operational definition of Section 4, is redundantly encoded across many degrees of freedom. This redundancy is not a hidden property — it is a measurable feature of the joint state of the record and its environment. A placed encoding that lacks environmental correlations is, in principle, distinguishable from a committed fact by any measurement that probes the correlations between the record and its surroundings. The distinction cannot be detected by a measurement restricted to a single subsystem at a single time — but it can be detected by measurements that probe inter-subsystem correlations, which are ordinary physical observables.

The Last Thursday hypothesis survives only if the fabrication extends to *all* environmental correlations at *all* scales — at which point the fabrication is doing the full thermodynamic work of genuine history, and the distinction between "fabricated" and "genuine" collapses (as argued in Q1). This is consistent with, and indeed follows from, the recursive closure argument of Section 5.4: the operational definition applied at all scales entails that any process reproducing the full correlation web has performed work commensurate with genuine production. The committed/placed distinction is therefore empirically accessible precisely when the fabrication is incomplete, and physically redundant when the fabrication is total — because a total fabrication *is* genuine production. The distinction is not metaphysical; it is empirical, but at the level of global correlations rather than local observables.

A natural question arises: what about a fabrication that extends to *almost* all correlations? The transition from encoding to commitment is not binary. Physical correlations come in degrees — a record may be correlated with some environmental subsystems but not others, with some degree of redundancy but not full redundancy. This is a feature of the framework rather than a vulnerability. It means that the Last Thursday hypothesis admits of *degrees* of physical inadmissibility. A fabrication that reproduces 90% of the environmental correlations has done 90% of the thermodynamic work — and the remaining 10% constitutes a measurable deficit in the commitment structure, a region where the fabricated universe is less factually deep than the genuine one. The hypothesis does not fail at a sharp threshold; it fails progressively, with the degree of inadmissibility proportional to the gap between the correlations present and the correlations required for full commitment. Perfect fabrication requires full thermodynamic work; imperfect fabrication leaves detectable seams.

The point can be sharpened by considering why the analogy most favorable to the Last Thursday hypothesis fails at cosmological scale. Consider two books with identical text: one written by an author over many months, one printed instantaneously by a machine given the complete text as input. The books are observationally indistinguishable, and their future behavior as books is identical — a book's causal powers depend on its text, not its production history. This is the strongest intuition pump for the Last Thursday position: if the product is identical, why should the process matter?

The answer is that the universe is not like a book. A book's content is separable from its production process: the text is encoded in ink patterns whose causal history is irrelevant to their function. The universe's content is not separable from its production process. The "content" of the universe — the structured correlations between subsystems, the thermodynamic traces of irreversible transitions, the causal ordering encoded in every decay chain and nucleosynthetic yield — just *is* the accumulated record of the irreversible processes that produced it. At cosmological scales, there is no "text" apart from the "writing process." The book analogy fails precisely where the universe-scale argument succeeds: a book's future behavior is independent of its production history, while the universe's admissible future transitions are constrained by the commitment depth of its present state.

The Last Thursday hypothesis is the claim that the universe is a printed book. Finite distinguishability says that printed books of this complexity cannot exist without the authorship process that generates them — because the book's content, at this scale, *just is* the accumulated record of the process that wrote it.

---

7. The Fact Capacity Bound

The foregoing arguments can be made quantitative. Any physical system of finite extent and finite duration has a bounded capacity for fact creation. The total number of irreversible commitments a system can make is limited by its available distinguishability — its entropy budget.

Definition (Fact Capacity). The fact capacity $\mathscr{F}(\tau, E, T)$ of a physical event is the maximum number of irreversible binary commitments that can be produced in proper time τ by a system with available energy E at effective temperature T.

We derive bounds on $\mathscr{F}$ from three independent components, each grounded in established physics.

7.1 The Energy Bound

The simplest constraint comes directly from Landauer's principle. At temperature T, each irreversible binary commitment dissipates at least $k_BT \ln 2$ of energy. A system with total available energy E can therefore make at most

$$\mathscr{F}_E \leq E / (k_BT \ln 2)$$

irreversible commitments. This is an upper bound on the total fact capacity, independent of time. For the observable universe at CMB temperature ($E \sim 10^{70}$ J, $T \sim 2.7$ K):

$$\mathscr{F}_E \lesssim 10^{70} / [(1.38 \times 10^{-23})(2.7)(0.693)] \sim 10^{92}$$

This provides a conservative sanity bound on the number of Landauer-priced irreversible bit-erasures that can be supported at an effective bath temperature T. It is not a claim that all physical record formation in cosmology proceeds at a uniform temperature, nor that every unit of entropy corresponds one-to-one with operational "facts" as defined in Section 4. Early-universe processes occurred at far higher temperatures; black hole interiors are not thermal baths at 2.7 K;

and many forms of record creation may be logically reversible in principle. The bound assumes the most favorable conditions for the Last Thursday scenario — low temperature, maximal energy availability, ideal efficiency — and still illustrates the scale of the resource gap.

A reviewer may reasonably ask: what temperature should a creation event operate at? The choice of $T = 2.7$ K is deliberately generous to the Last Thursday proponent, since lower temperatures yield more bits per joule. If the creation event is a cosmological-scale process occurring at Planck temperature ($T \sim 10^{32}$ K), the energy bound tightens dramatically to $\mathscr{F}_E \sim 10^{51}$ — a far more severe constraint. If it occurs at nucleosynthesis temperatures ($T \sim 10^{10}$ K), $\mathscr{F}_E \sim 10^{70}$. The CMB temperature yields the loosest possible bound; any physically realistic creation scenario gives a tighter one.

We note that a more careful treatment would *strengthen* this bound. The Landauer cost per bit scales as $k_BT \ln 2$, so irreversible transitions at high temperature are thermodynamically more expensive per bit than those at low temperature. The universe's entropy was generated across its full thermal history, from $T \sim 10^{10}$ K during nucleosynthesis to the present $T \sim 2.7$ K. An integrated bound $\mathscr{F}_{int} = \int dE / (k_BT(t) \ln 2)$, tracking the available energy and temperature at each epoch, would typically yield a smaller total under realistic cosmological thermal histories than our constant-T estimate — because most of the universe's irreversible work was done at far higher temperatures, where each joule purchases fewer bits. By evaluating the bound at the present CMB temperature, we have given the Last Thursday proponent the most generous possible accounting. The resource gap under a realistic thermal history is larger, not smaller.

For scale: $\mathscr{F}_E \sim 10^{92}$ exceeds the diffuse baryonic and radiation entropy ($\sim 10^{89}$ k_B) but falls well short of the total cosmological entropy including black holes ($\sim 10^{104}$ k_B [2]). These quantities are not identical to operational "facts" as defined in Section 4, but they serve as rough benchmarks for the scale of irreversible structure. Even at the most generous temperature, the energy bound cannot account for the full entropy budget of the observable universe.

7.2 The Information Bound

An independent constraint comes from the Bekenstein bound [9], which limits the total information content of any finite physical region. Interpreted as an upper bound on entropy (or equivalently, information) for given energy E and radius R, the maximum number of distinguishable quantum states is bounded by:

$$S_{Bek} \leq 2\pi RE / (\hbar c \ln 2)$$

This bound is temperature-independent and provides a ceiling on the total information that *can exist* within a given region, regardless of how it was produced. For the observable universe ($R \sim 4.4 \times 10^{26}$ m, $E \sim 10^{70}$ J):

$$S_{Bek} \lesssim 10^{123} \text{ bits}$$

The Bekenstein bound does not directly limit fact *creation* — it limits fact *storage*. But it provides an important cross-check: the total structured information content of the observable universe cannot exceed $\sim 10^{123}$ bits, and any creation event must produce this structure through physical transitions, each of which is separately constrained by the energy and operations bounds. The Bekenstein bound confirms that the universe's information content is enormous but finite, and that any account of how that content came to exist must respect the resource constraints governing physical transitions.

7.3 The Operations Bound

A separate constraint limits how many operations — irreversible or otherwise — a physical system can perform per unit time. The Margolus–Levitin theorem [8] establishes that a quantum system with average energy E above its ground state can transition between orthogonal states at most

$$N_{ops} \leq 2E\tau / (\pi\hbar)$$

times in proper time τ. This is a fundamental quantum speed limit, independent of the system's architecture or the specific Hamiltonian. Since each fact-creating transition produces a macroscopically distinguishable state change (a resolved alternative that satisfies the operational definition of Section 4), the orthogonality condition required by the Margolus–Levitin theorem is satisfied. For a creation event of duration τ with energy E, the number of distinguishable operations — and therefore the number of potential fact-creating transitions — is bounded by N_{ops}.

For a creation event lasting one Planck time ($\tau = t_P \approx 5.4 \times 10^{-44}$ s) with the total energy of the observable universe:

$$N_{ops} \lesssim (2 \times 10^{70} \times 5.4 \times 10^{-44}) / (\pi \times 1.05 \times 10^{-34}) \sim 10^{60}$$

This calculation generously grants the creation event the full energy of the observable universe as computational fuel. This is itself problematic: a creation event that uses the energy of the universe it is creating as its own resource is circular — the energy must exist before it can be used as fuel, but the hypothesis claims the energy was created by the event. We set this circularity aside and grant the assumption, since the bound is devastating even on these generous terms. The bound further assumes all operations are independent, parallelizable, and fact-creating. Since operational facts require redundancy and stabilization (Section 4), they cannot be produced at a rate exceeding the available orthogonalization capacity; thus N_{ops} is a generous upper bound on any fact-creation mechanism. Even so, it falls many orders of magnitude short of standard cosmological entropy scales, underscoring the mismatch between the resources available to an instantaneous creation event and the structured content observed in the universe.

7.4 The Combined Bound

The energy bound and the operations bound constrain different resources: the first limits the total number of irreversible commitments a system can *pay for*, the second limits how many transitions a system can *execute* in a given time. The effective fact capacity of a creation event is constrained by whichever bound is tighter:

$$\mathcal{F}(\tau, E, T) \leq \min(E / (k_B T \ln 2), 2E\tau / (\pi\hbar))$$

The Bekenstein bound (§7.2) does not appear in this expression because it plays a different role. The energy and operations bounds constrain fact *creation* — the rate and cost of producing irreversible commitments. The Bekenstein bound constrains fact *storage* — the total information content that can exist within a finite region. It serves as an independent ceiling: even if a creation event could somehow circumvent the Landauer and Margolus–Levitin constraints, the resulting structure could not exceed $S_{Bek} \sim 10^{123}$ bits within the observable universe. The three bounds are complementary: the Bekenstein bound limits what can exist, the energy bound limits what can be paid for, and the operations bound limits what can be executed in time.

What must the bound be compared against? The fact capacity bound provides a well-defined ceiling on what a creation event can produce. To complete the argument, we need a floor — a lower bound on the structured content the universe actually contains. Cosmological entropy estimates ($\sim 10^{89}$ to $\sim 10^{104}$ k_B) provide rough benchmarks, but thermodynamic entropy and operational facts (as defined in Section 4) are related but not identical. A sharper approach is direct enumeration: counting only the most robust, unambiguous examples of committed records that satisfy all three operational conditions (redundancy, robustness, erasure cost).

Consider: the $\sim 10^{80}$ baryons in the observable universe, each of whose identities (proton vs. neutron) was fixed by irreversible weak-interaction processes during Big Bang nucleosynthesis; the $\sim 10^{68}$ photons in the CMB, each carrying polarization and frequency information imprinted during recombination; the $\sim 10^{22}$ stellar-mass objects whose internal compositions record billions of years of nuclear fusion; the $\sim 10^{9}$ independent multipole modes in the CMB angular power spectrum, each encoding a distinct density fluctuation from the early universe. Even this absurdly conservative inventory — ignoring all geological, chemical, and biological records — yields a lower bound well above 10^{68} independent committed records. Each is redundantly encoded (correlated with surrounding structure), robust (stable against local perturbation), and costly to erase (reversing nucleosynthesis or recombination requires macroscopic entropy export).

The quantitative bound is therefore not merely illustrative. Its ceiling (the fact capacity) is rigorously derived; its floor (the number of operational facts) is conservatively estimated but robustly large.

The dilemma of duration. The Last Thursday hypothesis is silent on the mechanism and duration of the creation event, so a critic could posit a creation event lasting far longer than a Planck time. The operations bound scales linearly with τ. But any creation event faces a structural dilemma: short creation events lack the resources, and long creation events produce the very history they are supposed to fake. A creation event of duration τ necessarily produces τ of genuine irreversible history. Any creation event long enough to meet the operations bound is long enough to have generated genuine committed structure, at which point the "fabrication" is partially genuine. The hypothesis requires fabrication *without* genuine history, not fabrication that incidentally produces it.

The following table illustrates the scaling, comparing the operations bound against the conservative enumeration floor ($\sim 10^{68}$ committed records) and cosmological entropy benchmarks:

Creation duration τ	N_ops (upper bound)	vs. 10^{68} (enumerated)	vs. 10^{89} (diffuse)	vs. 10^{104} (total)
$t_P \approx 10^{-43}$ s	$\sim 10^{60}$	deficit $\sim 10^{8}$	deficit $\sim 10^{29}$	deficit $\sim 10^{44}$
10^{-15} s (femtosecond)	$\sim 10^{88}$	surplus	borderline	deficit $\sim 10^{16}$
1 second	$\sim 10^{104}$	surplus	surplus $\sim 10^{15}$	borderline
1 year ($\sim 3 \times 10^{7}$ s)	$\sim 10^{111}$	surplus	large surplus	surplus

Every "surplus" entry is subject to the dilemma: a one-second creation event that meets the 10^{104} threshold has, by doing so, generated a genuine one-second history of irreversible commitments. The surplus entries show that the operations bound is met, not that fabrication has occurred. For the energy bound, which is time-independent, $\mathscr{F}_E \sim 10^{92}$ falls short of the 10^{104} total regardless of duration.

The gap between ceiling and floor spans tens of orders of magnitude for any creation event short relative to the age of the universe. The primary argumentative weight is carried by the qualitative constraints of Sections 5.1–5.4, which establish that fabricated histories are physically inadmissible on structural grounds. The fact capacity bound supplements these qualitative arguments by showing that the resource gap is not marginal but enormous under any reasonable accounting.

7.5 The Bound as a Physical Constraint

The fact capacity bound is not an artifact of VERSF. It follows from any framework that accepts three premises: (i) irreversible transitions have minimum energy cost (Landauer), (ii) state transitions have minimum rate (Margolus–Levitin), and (iii) the structured content of the universe is constituted by irreversible transitions rather than being a freely specifiable boundary condition. Premises (i) and (ii) are established physics. Premise (iii) is the substantive claim, which we have argued in Section 4 is independently motivated by thermodynamic, information-theoretic, and computational considerations — though we acknowledge it is not universally accepted.

A reader who grants (i) and (ii) but rejects (iii) is making the following claim: that the universe's present microstate can be treated as a free initial condition with no thermodynamic price — that the specificity of the present is not owed to any process. This position faces two independent difficulties, neither of which presupposes VERSF.

The specification problem. Classical phase space is continuous. The microstate of the observable universe is a point in a phase space of approximately 10^{80} particle degrees of freedom — a space of dimensionality on the order of 10^{81}. Specifying a single point in a continuous space of this dimensionality requires infinite information: any finite specification determines at best a region of nonzero measure, not a point. A freely specifiable boundary condition that selects the *particular* microstate we observe — with its specific particle positions, momenta, chemical abundances, isotope ratios, and correlation structure — requires infinite precision and therefore infinite information content.

This is not an artifact of classical mechanics. Quantum mechanics discretizes the state space, but the number of distinguishable quantum states within the observable universe is bounded above by the Bekenstein limit at $\sim 10^{123}$. A specification that selects one state from 10^{123} alternatives carries approximately 10^{123} bits of information — a finite but astronomically large quantity that must be supplied *ab initio* as a single boundary condition, with no physical process to generate or justify it.

A produced state does not face this cost. Sequential irreversible transitions progressively narrow the accessible region of phase space through a finite number of finite-cost steps, each resolving one distinction and paying the corresponding Landauer price. The final microstate is not specified in advance — it is *selected* incrementally by the physical processes that built it. The information content of the present state is accumulated across 13.8 billion years of irreversible evolution, not front-loaded into a boundary condition. The initial condition need carry only low specificity — the information content of the laws of physics plus a low-entropy starting state — with all subsequent specificity generated by the forward evolution.

The free-boundary-condition framework therefore faces a dilemma. Either the specification of the present microstate carries infinite information (in the classical case) or $\sim 10^{123}$ bits of information (in the quantum case), supplied as a brute stipulation with no generative account — or there exists some principle constraining which boundary conditions are physically admissible, in which case one has reintroduced precisely the kind of constitutive constraint that premise (iii) asserts. The choice is not between "facts require process" and "facts are free." It is between "facts require process" and "facts require an inexplicable specification of comparable or greater information cost." Premise (iii) is not an additional assumption imposed on physics; it is the only alternative that does not replace one explanatory burden with a strictly larger one.

The genericity problem. A second, independent difficulty reinforces the first. The space of possible microstates is overwhelmingly dominated by thermal equilibrium configurations. This is the standard counting argument from statistical mechanics: high-entropy states vastly outnumber low-entropy states. Absent an additional selection principle (such as a Past Hypothesis), a freely specified initial condition has no reason to be atypical. Without a mechanism that selects for low-entropy, highly correlated states, a free specification overwhelmingly produces a generic state: one near thermal equilibrium, with no persistent causal ordering, no accumulating structure, and no arrow of time. A generic state evolves generically — thermal fluctuations around equilibrium, with no directed behavior.

We observe the opposite: persistent directionality, accumulating structure, causal ordering that extends coherently across 10^{60} Planck times. This is not what a freely specified state looks like. It is what a *produced* state looks like — one whose directionality was inherited from the processes that built it. The present state's non-generic character is evidence that it was produced rather than specified. The causal ancestry is not a metaphysical add-on; it is the explanation for why the universe is non-generic.

The fundamentalist must respond by invoking the past hypothesis — the brute postulate that the universe began in a low-entropy state. The Humean response — that the past hypothesis is itself a fundamental law requiring no further explanation — is available, but it carries the cost of treating maximal specificity as primitive, an unexplained axiom with no physical mechanism behind it. The present framework offers an alternative: this microstate because it was *produced* by directed irreversible processes from a simpler initial condition, and produced states have non-generic correlations that freely specified states do not.

The combined force of the two arguments. The specification problem and the genericity problem attack from opposite directions. The specification problem says: selecting the observed microstate from the space of all possible microstates requires enormous information, which must come from somewhere. The genericity problem says: absent a selection mechanism, the probability of landing on a non-generic microstate is vanishingly small. Together they establish that premise (iii) is not an optional interpretive commitment but a near-forced conclusion: the structured content of the present universe requires a generative account, because the alternatives — brute stipulation of enormous information content, or fantastically improbable selection from an overwhelmingly generic space — are each more physically demanding than the premise they are invoked to avoid.

This also sharpens the asymmetry between the Big Bang and the Last Thursday hypothesis. The Big Bang's initial condition is low-entropy and low-specificity — it carries minimal information content (a few hundred bits: initial density, expansion rate, a nearly scale-invariant perturbation spectrum with amplitude $\sim 10^{-5}$) and requires no fine selection from a vast space of alternatives. It is the *least* demanding initial condition: a simple starting point from which all subsequent specificity is generated by forward evolution. The Last Thursday initial condition, by contrast, demands maximal specificity — the full directed correlation structure of 13.8 billion years of evolution, requiring either $\sim 10^{123}$ bits of brute specification or an equivalently improbable selection — without any process to produce it. The asymmetry is not arbitrary; it is quantitative. The Big Bang is the one initial condition that does not need inherited specificity, because it has none to inherit. The Last Thursday state is the one initial condition that demands all of it up front.

Indeed, the Last Thursday hypothesis contains an internal contradiction at this level — one that the specification problem makes precise. In the classical case, specifying the present microstate requires infinite information. But a finite creation event — a beginning on Thursday — is a finite physical process with finite energy, finite duration, and therefore finite information capacity. A finite source cannot supply infinite information. The classical version of the hypothesis is therefore self-refuting: it demands a microstate whose specification requires more information than any finite origin can provide. In the quantum case, where the state space is discrete, the information cost is finite but astronomically large ($\sim 10^{123}$ bits). This avoids logical self-contradiction but does not avoid physical impossibility: the finite creation event must somehow front-load $\sim 10^{123}$ bits of structured information into its initial condition with no generative process to produce them.

Both cases reveal the same structural problem. A finite beginning, physically understood, is a state of minimal commitment: a point from which specificity can accumulate through irreversible processes. A state that already carries the full information content of 13.8 billion years of structured evolution is not a beginning — it is a summary. The hypothesis asks for a universe that simultaneously begins and has already happened. Under finite distinguishability, these are incompatible: a universe cannot have a finite origin and contain information that requires temporal depth to produce, because the information *is* the temporal depth.

The fact capacity bound quantifies the gap between these two positions: it measures the distance, in irreversible transitions, between a low-specificity initial condition and the high-specificity structured state we actually observe. That distance is the work the universe has done — and the work the Last Thursday hypothesis claims was never performed.

The bound also provides a principled resolution to computability questions. Computation, in this framework, is the forward creation of facts within finite resource limits. A computed state is one that was reached by a sequence of irreversible commitments from a well-defined initial condition. A stipulated state — one placed into existence without process — is not computed and, under finite distinguishability, is not physically admissible. The universe is necessarily a computation, not a stipulation.

---

8. Broader Implications

The Last Thursday hypothesis can be understood as the limiting case of a broader class of scenarios in which structured information is stipulated rather than generated. In contemporary discussions of simulation [6], it is common to distinguish between simulations that compute physical dynamics forward from initial conditions and those that shortcut the computation by injecting a precomputed or preselected state. The former generate their informational content through their own dynamics (albeit on a different physical substrate); the latter do not. The Last Thursday hypothesis belongs to this second category: it is equivalent to a simulation-without-computation — a data structure whose apparent history is injected rather than produced.

This reframing unifies the three applications that follow. Boltzmann brains, shortcut simulations, and radical skepticism are all instances of the same structural claim: structured content without commensurate irreversible process. Each postulates a present state whose information content was not produced by the dynamics that ostensibly generated it. And each faces the same resource dilemma established in Sections 5–7: either the process that produced the structured content did the full thermodynamic work (in which case it *is* genuine history, wherever the computation was physically located), or it did not (in which case the content fails the operational definition of committed fact). There is no middle ground — no way to obtain the output of a deep computation without performing the computation.

Throughout this section, "simulation" refers only to the shortcut case unless otherwise specified: simulations that bypass the physical computation of history by directly stipulating a final state. Simulations that genuinely compute physical evolution forward are not challenged by the present argument; they instantiate real causal depth on a different substrate and therefore constitute genuine history rather than fabrication.

8.1 Boltzmann Brains

The standard Boltzmann brain problem concerns the statistical likelihood of thermal fluctuations producing isolated observers with fabricated memories. In conventional statistical mechanics, such fluctuations are enormously improbable but not physically forbidden — and in an eternal de Sitter universe, they may even dominate the observer population, leading to paradoxes for cosmological measure theory [5].

The present argument introduces a distinct and more fundamental objection. Even granting the thermal fluctuation, the resulting brain-shaped object does not contain *genuine* committed facts. A Boltzmann fluctuation is a random excursion from thermal equilibrium — it produces a low-entropy configuration, but it does so without executing the irreversible commitment chains that would constitute genuine memories, genuine causal records, or genuine temporal depth.

The distinction is between two claims:

(B1) A thermal fluctuation can produce a physical configuration identical to a brain with memories. *(Granted.)*

(B2) A thermal fluctuation can produce a brain that *has* memories — that contains committed facts with causal depth. *(Denied.)*

Under finite distinguishability, (B1) does not entail (B2). The Boltzmann brain has the *pattern* of memory without the *substance* of memory. To see why, recall the operational definition of fact from Section 4: a fact is a record that is redundantly encoded across many degrees of freedom, robust under perturbation, and irreversible to erase. A genuine memory satisfies all three conditions — it is encoded in synaptic weights, correlated with environmental records, and erasable only at thermodynamic cost.

A precision is needed here: a Boltzmann brain that persists long enough to instantiate neural activity does have *local* redundancy (synaptic weights, ion gradients) and *local* erasure cost. At the scale of the brain itself, the internal states satisfy the operational definition — they are physically real configurations with genuine thermodynamic properties. What fails is condition (i) at the *global* level: the memories reference events (a childhood, a conversation, a sunset) that left no traces anywhere else in the universe. The redundant encoding that constitutes a genuine memory extends beyond the brain into the environment — footprints, documents, other brains, photon scattering histories — and the fluctuation did not produce these external correlates. The Boltzmann brain's memories are therefore facts about the brain's internal state (the synaptic pattern really is what it is) but not facts about the events they purport to represent (no sunset occurred, no conversation happened). The distinction is between the local physical reality of the brain state and the referential content of its memories. The brain is real; its autobiography is not.

A physicist might object: the fluctuation that assembles the brain *also* assembles the correlated physical substrate — the synaptic weights, the ion gradients, the local molecular environment — so why isn't that sufficient for genuine memory? The answer is that genuine redundancy extends beyond the brain: a real memory is correlated with environmental records (footprints, documents, other brains, photon scattering histories) that the fluctuation did not produce. The fluctuation assembles a self-consistent local encoding, but it does not assemble the web of redundant external records that constitutes a committed fact in the operational sense. A Boltzmann brain's "memories" are locally encoded but globally unanchored — they reference a causal history that left no traces anywhere else in the universe. In particular, the correlations that constitute genuine redundancy are absent beyond the fluctuation volume, whereas committed facts necessarily involve correlations extending into the wider environment.

This is not a statistical argument about fluctuation probability — it is a structural argument about what fluctuations can and cannot produce. A Boltzmann fluctuation can assemble the encoding of a fact, but it cannot commit a fact, because commitment requires irreversible entropy expenditure distributed across many degrees of freedom, and fluctuations are, by their nature, reversible excursions that do not generate net entropy.

A refinement is needed here, because the standard Boltzmann brain problem is not restricted to small fluctuations. In principle, a thermal fluctuation could encompass the entire observable universe — producing not just an isolated brain but a complete universe-scale configuration with full environmental correlations. The standard problem is precisely that *any* configuration, including one with all the correlations we observe, has some nonzero fluctuation probability in an eternal thermal bath.

Our argument applies differently to this case. A universe-scale fluctuation that produces the full web of environmental correlations — every redundant encoding, every inter-system correlation, every thermodynamic trace — has, by doing so, satisfied the operational definition of fact for every record it contains. The resulting state is indistinguishable from a genuinely evolved universe not because the distinction has been erased but because the fluctuation has done the equivalent work. This is the Boltzmann brain analogue of the (L2) horn of the lineage dilemma (Section 5.4): a fluctuation large enough and structured enough to produce genuine committed facts is a fluctuation that has performed genuine irreversible work, and its products *are* genuine facts. The "fluctuation" label does not diminish their factual status.

What our argument rules out is the intermediate case that drives the standard paradox: a fluctuation that produces an observer with the *appearance* of living in a structured universe but without the full environmental correlations that constitute genuine commitment. Such a fluctuation produces encodings without facts — a Boltzmann brain, not a Boltzmann universe.

It is important to be precise about what this result does and does not accomplish. The standard Boltzmann brain paradox has two components: a *metaphysical* question (does a Boltzmann brain have genuine memories and experiences?) and a *cosmological measure* question (which fluctuation scale dominates the observer count in eternal de Sitter space?). Our argument addresses the first question decisively: a Boltzmann brain's internal states are genuine local facts about the brain's configuration, but its "memories" are not referential facts about external events — they lack the global correlation web that would anchor them to the events they purport to represent. The brain is real; its autobiography is not. Our argument does not address the second question, which concerns the relative statistical weight of brain-scale vs. universe-scale fluctuations and the correct measure over observer-moments in eternal inflation. The measure problem survives as a technical challenge in cosmology. What our argument removes is the metaphysical anxiety that motivates it — the worry that you might *be* a Boltzmann brain with fabricated memories. Under finite distinguishability, this worry is not merely unlikely but incoherent.

8.2 Retroactive Simulation

The simulation hypothesis, in its strongest form, suggests that our universe could be a simulation running on some external substrate [6]. As noted above, the critical distinction is between simulations that compute and simulations that shortcut.

Under finite distinguishability, the simulation argument bifurcates:

If the simulation computes the history genuinely — executing each irreversible step, paying each entropy cost — then the simulated facts are physically real. The simulation has done real work, and the distinction between "simulated" and "physical" collapses. The simulated universe has genuine causal depth because the simulation process *provided* that depth through its own irreversible computation.

If the simulation shortcuts the history — stipulating the present state without computing the intervening steps — then the simulated universe lacks genuine committed structure and is subject to the same inadmissibility constraints as the Last Thursday creation. The simulation has produced an encoding without commitment, and the resulting state is not a universe with a past but a data structure that references a past that never occurred.

A simulation cannot have it both ways: it cannot shortcut the entropy cost of fact creation while producing genuine facts. The simulation hypothesis is therefore not refuted but *constrained*. Simulations that do real work produce real facts. Simulations that cheat produce states that are, under finite distinguishability, physically incomplete.

This constraint is stronger than it first appears, because it applies to the simulation's *starting state*. A simulation that begins at any point after the Big Bang — say, one billion years in — must stipulate the state at that moment. But that state already contains enormous structured content: nucleosynthetic abundances, CMB photons released from the surface of last scattering, proto-galactic density perturbations, all the causal ordering accumulated over the first billion years. Under the present framework, those are facts that were injected, not computed. The simulation has shortcut the first billion years of irreversible commitments.

The only starting state that carries no stipulated facts is the Big Bang initial condition itself — low entropy, low specificity, a few hundred bits (Section 7.5). Everything else is a state that already contains structured information requiring a generative account. The only non-shortcut simulation is therefore one that begins at (or near) the Big Bang and computes forward through the full history. Every other simulation is a shortcut simulation to some degree, with the degree of shortcutting proportional to the structured content of its stipulated starting state.

One nuance is worth preserving: a simulation could compute its starting state in a preliminary pass — run the first billion years, save the output, then use that as input for a second-stage simulation. The total computation covers the full history even though no single run does. The criterion is not single continuous execution from the Big Bang but total computational work accounting for all structured content. This qualification reinforces rather than weakens the constraint: what matters is whether the facts were computed *somewhere*, not whether they were computed in one pass. The simulation constraint therefore reduces to: the only simulation that doesn't cheat is one whose total computational history is equivalent to the universe's actual history — a result that leaves the simulation hypothesis formally intact but physically empty.

A natural intermediate case deserves attention: partial simulation. A simulator might compute some aspects of the universe genuinely while interpolating or approximating others — running full physics for stellar interiors, say, while using statistical shortcuts for intergalactic gas dynamics. Under the framework of degrees of commitment developed in Section 6, this produces a universe with uneven factual depth. The genuinely computed regions contain committed facts with full causal lineage. The approximated regions contain encodings with partial or absent commitment — structures that exhibit the *signatures* of irreversible production but were produced by a cheaper process that did not generate the full environmental correlation web. The result is a universe with "seams" at the boundaries between computed and approximated regions: measurable deficits in the inter-subsystem correlations that characterize genuine commitment. A perfect partial simulation — one in which the approximated regions are indistinguishable from genuinely computed ones at all scales — would require the approximation to reproduce all correlations, at which point it has done the full thermodynamic work and the "approximation" label is misleading. Partial simulation, like partial fabrication, is constrained by the same dilemma: shortcuts that save thermodynamic work leave detectable deficits, and shortcuts that don't save work aren't shortcuts.

8.3 Simulated vs. Self-Grounding Universes

The compute-vs-shortcut distinction addresses whether a simulation's *content* is genuine. A prior question concerns whether a simulation can be ontologically self-sufficient — whether it can ground its own existence in the way a real universe does.

A simulation is parasitic in a way a real universe is not. It requires three things that must come from outside:

Rules. The simulation's physics must be specified — coded, implemented, imposed from without. In the real universe, the laws of physics are not injected; they are constitutive features of how irreversible commitments propagate. In a simulation, they are encodings of someone else's committed facts about how physics works. The simulation's dynamics are borrowed structure, derived from irreversible work done in the host universe: the work of discovering, formalizing, and implementing those rules.

Initial conditions. Even a Big-Bang-start simulation receives its initial state as input. That input — the low-entropy starting configuration, the perturbation spectrum, the coupling constants — is information generated by processes in the host universe. The simulation does not generate its own starting point; it inherits one.

Substrate. The computer running the simulation is itself a product of irreversible commitments in the host universe. Silicon wafers, quantum processors, whatever the physical implementation — they exist because the host universe did the thermodynamic work of producing them.

This dependency structure creates an asymmetry that goes beyond the compute-vs-shortcut distinction. The real universe is self-grounding: its facts, its rules, and its substrate are the same thing — the accumulated record of irreversible commitments. A simulation is externally grounded: its facts may be computed, but its rules, initial conditions, and substrate are imported from a prior layer of committed reality. The simulation cannot bootstrap itself.

Under VERSF, this makes the simulation hypothesis's regress problem precise. If our universe is a simulation, it runs on a substrate in a host universe. That host universe either is real (self-grounding, generating its own facts through its own dynamics) or is itself a simulation running on a deeper host. The regress must terminate in a self-grounding universe — one whose rules, initial conditions, and substrate are constituted by its own irreversible commitments rather than imported from outside. The simulation stack bottoms out at reality.

This does not refute the simulation hypothesis — it constrains its ontological structure. A simulated universe can contain genuine committed facts (if it computes them). But it cannot be the fundamental layer. There must exist a non-simulated ground: a universe whose physical laws are not encoded but constitutive, whose initial conditions are not inherited but primitive, and whose substrate is not manufactured but self-identical with the irreversible processes it supports. VERSF identifies what that ground consists of: irreversible commitments with finite distinguishability cost, propagating under their own dynamics without external specification.

8.4 Radical Skepticism

More broadly, any skeptical scenario that postulates structured experience without structured history — brains in vats, Cartesian demons, spontaneous creation — faces the same objection. Structure requires process. Process requires entropy. Entropy requires time. You cannot have the structure without the time it takes to build it.

The contribution here is to reframe what these skeptical scenarios actually require. The brain-in-a-vat scenario is typically presented as a question about *knowledge*: how do you know you're not a brain in a vat? Under finite distinguishability, it becomes a question about *resources*: what would it take to produce a brain-in-a-vat with the structured content of genuine experience? The answer is: the same irreversible work as producing the experience genuinely. A vat that produces all the correlations, all the redundant encodings, all the environmental traces that constitute genuine committed facts has done the thermodynamic work of a genuine world. A vat that produces only the local encodings (the brain's internal states) without the environmental correlations has produced a Boltzmann brain — an encoding without commitment, subject to the same structural deficiencies discussed in Section 8.1.

This does not refute skepticism as a logical exercise. One can always doubt whether the physical principles invoked here are themselves part of the fabrication. But this is an infinite regress that applies to *any* physical argument about *anything*. The regress is not specific to our argument; it is a feature of the relationship between physical theory and radical doubt. Within the domain of physical theory — the domain in which questions about the universe's structure can be meaningfully posed — the Last Thursday hypothesis is ruled out, and ruled out not by preference or parsimony, but by resource constraints as fundamental as the conservation of energy.

---

9. Objections and Replies

Q1: Can't entropic history be faked like anything else?

If the creator is powerful enough to fabricate fossils and memories, why can't it also fabricate entropy trails?

This is the most natural objection, and it rests on a misunderstanding of what the argument claims. We are not saying that a creator could not arrange particles into a configuration that *looks like* it has an entropy trail. Of course it could — that is precisely what claim (C1) asserts, and we grant (C1) without reservation. The point is that arranging particles into such a configuration *is itself* an entropy-generating process. The "faking" is not free.

To fabricate an entropy trail for one fact, you must commit at least one new fact — the act of fabrication. To fabricate the entropy trail for the entire universe, you must perform work at least commensurate with the entropy content of the universe. The fabrication does not escape the constraint; it is subject to it. You cannot fake a receipt for entropy expenditure without expending entropy to produce the fake. The fraud is as expensive as the honest transaction.

This is directly analogous to counterfeiting. One can print a banknote that is observationally indistinguishable from a genuine one. But printing the counterfeit requires ink, paper, machinery, and time — real resources. The claim is not that perfect counterfeits are detectable; it is that producing them is not cheaper than earning the money honestly. A creator that fabricates the entire entropic history of the universe has done work equivalent to actually running that history — at which point the "fabrication" and the "genuine history" are the same thing.

Q2: Doesn't this argument assume VERSF is correct?

The argument depends on facts being irreversible commitments. What if that's wrong?

The argument is strongest within VERSF, but it does not require VERSF specifically. It requires three premises: (i) irreversible transitions have minimum energy cost (Landauer's principle — established physics), (ii) state transitions have minimum duration (quantum speed limits — established physics), and (iii) the structured content of the universe is constituted by irreversible transitions rather than being a freely specifiable initial condition.

Premise (iii) is the substantive one, and we acknowledge it is not universally accepted. But it is not a VERSF peculiarity. Any framework that takes thermodynamic irreversibility seriously as a constitutive feature of reality — rather than an emergent approximation over fundamentally reversible dynamics — will reach the same conclusion. VERSF provides the most explicit formulation, but the argument applies to any physics in which the second law is not merely statistical but structural.

The real question is whether premise (iii) is independently defensible. We believe it is, for the reasons given in Section 4: Landauer's principle, Shannon entropy, and computational complexity theory all converge on the conclusion that physical specificity has irreducible cost. A reader who rejects (iii) is not disagreeing with VERSF — they are claiming that the universe's structure is a free parameter with no thermodynamic price. That is a position one can hold, but it is a strong and non-obvious claim that bears its own burden of justification.

Q3: Isn't this just Occam's razor dressed up in physics language?

You're still just saying the simpler explanation is better. What's actually new here?

Occam's razor says: prefer the hypothesis with fewer assumptions. The present argument says: the state of affairs described by the Last Thursday hypothesis cannot be physically instantiated. These are categorically different claims.

Occam's razor leaves the Last Thursday universe as a logical possibility that we choose not to believe. The fact capacity bound says it is not a logical possibility at all — it requires more irreversible commitments than can be produced in the time available. This is the difference between saying "we have no reason to think this bridge will collapse" and saying "this bridge cannot support its own weight." The first is an epistemic judgment; the second is an engineering constraint.

The novelty is the transition from epistemology to physics. Every previous argument against the Last Thursday hypothesis has been about our reasons for belief. This argument is about the universe's capacity for structure.

Q4: What if time itself was created along with everything else?

If the creator made time, then there's no "duration" constraint on creation — the creator operates outside of time.

This is a more sophisticated version of the objection, and it requires a careful answer.

If the creator operates outside of time entirely, then the creator performs no irreversible transitions — because irreversibility requires a before and an after, which requires temporal ordering. A creator with no temporal extension commits no facts, generates no entropy, and produces no structure. An atemporal creation event has zero fact capacity. It can produce a void — a state with no committed structure — but it cannot produce a universe with 10^{104} bits of committed specificity.

If the creator operates within some external time but creates our time as part of the fabrication, then the constraint shifts but does not vanish. The creator must still perform irreversible work in its own temporal framework, and that work must be sufficient to produce the committed structure of our universe. The entropy cost is paid in the creator's time rather than ours, but it is still paid. The creator's act of fabrication is a genuine physical process with genuine temporal depth — it has simply been relocated to an external frame. Under these conditions, the Last Thursday hypothesis is not refuted but *reinterpreted*: our universe has a genuine causal history, but that history is embedded in the creator's timeline rather than our own.

A further variant: what if the creator operates within time but under different physical laws — laws where Landauer's principle does not hold, where information can be manipulated without thermodynamic cost? In that case, the resource constraints of Sections 5 and 7 would not apply in the creator's domain. But the output must still satisfy the constraints of *our* physics to be a universe with genuine facts as we define them. A universe whose internal structure exhibits committed facts — redundantly encoded, robust, costly to erase — must contain the thermodynamic signatures of those commitments regardless of how the creator's physics works. If the output lacks those signatures, it is not a universe with genuine facts. If it has them, the creator has produced genuine committed structure, and the "fabrication" is genuine production by another name. The constraint is on what the output must be, not on how the creator operates.

This reinterpretation deserves emphasis, because it reveals the self-undermining structure of the objection. The hypothesis was introduced to deny the reality of the past — to claim that our apparent history is a fiction. But the only way to create that fiction, under finite distinguishability, is to perform real work with real temporal depth. The creator's act of fabrication *is* a past — a sequence of irreversible commitments that constitutes genuine causal history. The Last Thursday hypothesis, pushed to its logical conclusion, does not eliminate the past; it relocates it. The past still happened. It just happened somewhere else. And if the creation process has genuine temporal depth in the creator's frame, then the question "is our past real?" reduces to "does it matter whose clock the work was done on?" — a question about labeling, not about physics.

Q5: Doesn't quantum mechanics allow something from nothing? What about vacuum fluctuations?

Particles appear from the vacuum all the time. Why can't the universe?

Virtual particle pairs are not "something from nothing." They are temporary excursions within quantum field theory that conserve energy over time intervals bounded by the uncertainty principle. They do not constitute committed facts — they are uncommitted fluctuations that resolve back to the vacuum state. No irreversible transition occurs; no distinguishability is consumed; no entropy is generated.

Real particle creation — as in Hawking radiation or pair production in strong fields — does involve genuine irreversible transitions and does generate entropy. But these processes take time, consume energy, and are subject to exactly the resource constraints we describe. The universe's vacuum can produce real particles, but it cannot produce 10^{104} bits of committed structure instantaneously, for the same reason a thermal reservoir cannot spontaneously organize itself into a library.

The cosmological case is sometimes phrased as "the universe is a free lunch — it emerged from a quantum fluctuation of zero total energy." This picture has a serious literature behind it (notably Tryon's vacuum fluctuation proposal [14] and Vilenkin's tunneling-from-nothing model [15]). Even granting these scenarios, the initial quantum event was followed by 13.8 billion years of inflationary expansion, cooling, nucleosynthesis, gravitational collapse, and structure formation — each step involving irreversible commitments that built up the present entropy budget. In these models, the initial fluctuation or tunneling event is not the creation of the structured universe; it is the starting condition for a 13.8-billion-year computation. The structured content we observe is produced by the subsequent inflationary and post-inflationary evolution, not by the nucleation event itself. The Last Thursday hypothesis claims the computation can be skipped. It cannot.

Q6: If the argument works, doesn't it prove too much? Doesn't it rule out divine creation entirely?

If no recent creation event can produce the universe, are you ruling out all creation accounts?

No. The argument constrains *when* and *how* creation can work, not *whether* creation occurred.

A creation event that establishes initial conditions and then lets physics run forward for 13.8 billion years is perfectly consistent with the fact capacity bound. The committed structure is generated by the subsequent evolution, not by the creation event itself. The bound only excludes creation events that attempt to front-load the entire structured content of the present universe into a single moment.

Theological creation accounts that posit a beginning followed by natural development are not touched by this argument. The specific claim that is ruled out is the conjunction of recent creation and apparent deep history — the claim that the universe looks old but isn't, that the evidence of age is fabricated rather than genuine. The argument says: if the evidence of age is present, the age is real, because the evidence *consists in* the accumulated cost of the processes that produced it.

Q7: What would falsify this argument?

Every good physical argument should be falsifiable. What would prove you wrong?

The argument would be falsified by any of the following:

A demonstration that Landauer's principle fails. If information can be erased without entropy cost, then fact creation has no minimum energy price, and the entropy constraint (Section 5.1) collapses. Current experimental evidence strongly supports Landauer's bound [4], but a reproducible violation would undermine this argument.

A demonstration that computational depth can be produced without temporal depth. If a physical system can generate a state of high logical depth (in Bennett's sense) without executing the corresponding number of irreversible steps, then the compression constraint (Section 5.3) fails.

This would require a mechanism for "shortcutting" computation that violates known limits in computational complexity theory.

A physical system exhibiting committed structure without entropy history. If a laboratory experiment could produce a macroscopic object with internal causal structure — correlations, decay chains, layered records — without the corresponding entropy expenditure, the core claim of the paper would be directly refuted.

A viable alternative account of facthood. If facts can be coherently defined without reference to irreversible commitment — if physical specificity can exist without thermodynamic cost — then the entire framework of the argument dissolves. We have argued that this is difficult to sustain (Section 4), but we acknowledge it as the most direct route to overturning the conclusion.

The argument is falsifiable. It makes specific claims about the relationship between structure, entropy, and time, and these claims can in principle be tested.

Q8: Can a fake fact generate real facts downstream? What happens after creation?

Suppose the universe was created last Thursday with impostor facts. Couldn't real physics take over from that point and generate genuine facts going forward?

This is the question the lineage constraint (Section 5.4) addresses directly. The short answer is: if real physics takes over and generates genuine irreversible commitments, then those new commitments are real facts — but they do not retroactively validate the impostor facts they were built on.

Consider the analogy of a forged document used to establish a line of legal precedent. Subsequent court decisions that cite the forged document are genuine legal acts — real judges made real rulings. But the entire chain of precedent is compromised at the root. The later rulings are real events with real consequences, but their *justificatory lineage* is broken. They refer to an authority that does not exist.

The physical version is more severe. A genuine fact is not merely an irreversible transition — it is an irreversible transition whose causal antecedents are themselves committed facts (Section 4, condition (i): redundant encoding across many degrees of freedom, including environmental correlations with prior facts). A post-creation fact that refers to an impostor antecedent inherits a structural deficiency: it is correlated with a record that *encodes* a causal history but does not *instantiate* one. The correlation pattern is present, but the physical substance underwriting it — the chain of entropy expenditures linking each record to its cause — is missing at the root.

This creates a detectable asymmetry, at least in principle. Facts generated after the creation moment have genuine causal depth extending back to the creation moment. Facts that were placed at the creation moment have zero causal depth — they are the floor on which everything else rests, but the floor itself rests on nothing. In a genuine universe, there is no such floor: the causal chain extends continuously to initial conditions. In a fabricated universe, there is a sharp discontinuity — a boundary between placed structure and generated structure — that constitutes a physical seam in the fabric of committed facts.

The Last Thursday hypothesis requires this seam to be undetectable. But under finite distinguishability, the seam is not merely difficult to detect — it is a structural feature of the state. Facts on one side of it have causal lineage; facts on the other side do not. The universe cannot be "half faked" without the boundary between real and fake constituting a physical discontinuity in the commitment record.

---

10. Conclusion

The Last Thursday hypothesis has long served as a philosophical curiosity — a demonstration that evidence alone cannot rule out radical alternatives to our standard picture of the universe. We have shown that this status was a consequence of treating information as descriptive rather than constitutive. Once facts are recognized as irreversible physical commitments with finite entropy cost, the hypothesis is not merely implausible but physically impossible.

The argument proceeds on four independent fronts: thermodynamic (committed structure requires entropy expenditure), causal-structural (causal ordering is physical content, not narrative overlay), information-theoretic (computational depth requires temporal depth), and — conditionally, pending Conjecture 1 — genealogical (a fabricated fact is either not a fact or not fabricated). The first three converge on a single conclusion without additional assumptions: the present state of the universe is not a configuration that *could* have been instantiated — it is a record that *had* to be written.

The fact capacity bound makes this quantitative: any creation event short relative to the age of the universe falls orders of magnitude short of the committed structure the universe contains. Without the lineage constraint, the Last Thursday hypothesis is ruled out on resource grounds; with it, the very notion of a fabricated fact dissolves. The deep past is not inferred from evidence — it is required by the information content of the present.

The past is not a story we tell about the present. It is the accumulated cost of everything the universe has done. You cannot have the facts without paying for them, and you cannot pay for them without the time it takes to do so.

References

[1] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM J. Res. Dev.* **5**, 183–191 (1961).

[2] C. Egan and C. Lineweaver, "A larger estimate of the entropy of the universe," *Astrophys. J.* **710**, 1825–1834 (2010).

[3] C. H. Bennett, "Logical depth and physical complexity," in *The Universal Turing Machine: A Half-Century Survey*, ed. R. Herken (Oxford University Press, 1988), pp. 227–257.

[4] A. Bérut *et al.*, "Experimental verification of Landauer's principle linking information and thermodynamics," *Nature* **483**, 187–189 (2012).

[5] R. Bousso and B. Freivogel, "A paradox in the global description of the multiverse," *J. High Energy Phys.* **2007**, 018 (2007).

[6] N. Bostrom, "Are you living in a computer simulation?" *Philos. Q.* **53**, 243–255 (2003).

[7] W. H. Zurek, "Quantum Darwinism," *Nat. Phys.* **5**, 181–188 (2009).

[8] N. Margolus and L. B. Levitin, "The maximum speed of dynamical evolution," *Physica D* **120**, 188–195 (1998).

[9] J. D. Bekenstein, "Universal upper bound on the entropy-to-energy ratio for bounded systems," *Phys. Rev. D* **23**, 287–298 (1981).

[10] J. A. Wheeler, "Information, physics, quantum: the search for links," in *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics* (Tokyo, 1989), pp. 354–368.

[11] I. Prigogine, "Time, structure, and fluctuations," *Science* **201**, 777–785 (1978).

[12] D. Deutsch and C. Marletto, "Constructor theory of information," *Proc. R. Soc. A* **471**, 20140540 (2015).

[13] W. H. Zurek, "Algorithmic randomness and physical entropy," *Phys. Rev. A* **40**, 4731–4751 (1989).

[14] E. P. Tryon, "Is the universe a vacuum fluctuation?" *Nature* **246**, 396–397 (1973).

[15] A. Vilenkin, "Creation of universes from nothing," *Phys. Lett. B* **117**, 25–28 (1982).