

Finite Distinguishability and Irreversible Commitment: A No-Go Result for Infinite-Precision Physics

Abstract

We establish that irreversible commitment—the process by which physical possibilities resolve into definite facts—is impossible in any system with infinite distinguishability. This result functions as a no-go theorem: no physical theory permitting unbounded state resolution in finite regions can accommodate irreversible measurement outcomes, monotonic entropy increase, persistent records, or asymmetric temporal ordering. Finite distinguishability is therefore not an empirical discovery or modeling convenience, but a necessary condition for fact-producing physics. We formalize this argument, address objections from quantum mechanics and eternalism, connect the result to established information-theoretic bounds, and identify the class of physical frameworks excluded by this constraint.

Plain Language Summary. When you flip a coin and it lands heads, that outcome becomes a fact—you can't "un-flip" it. But many mathematical models of physics are perfectly reversible: in principle, every process can be run backward. How can irreversible facts exist in a reversible universe? This paper shows that facts can only exist if there's a limit to how finely we can distinguish physical states. If we could make infinitely precise distinctions, any apparent "fact" could always be undone by accessing finer details. The existence of genuine facts—measurement outcomes, memories, records—therefore requires that physical distinguishability be finite. This isn't just a practical limitation; it's a logical requirement for facts to exist at all.

Contents:

Abstract	1
1. Introduction: The Problem of Facts	6
2. Formal Definitions and Framework.....	6
2.1 State Spaces and Distinguishability	6
2.2 Irreversible Commitment	8
2.3 Physical Realizability.....	9
2.4 Admissibility Versus Mathematical Possibility.....	9
3. The No-Go Theorem	10
3.0 Assumptions (Admissibility Layer)	10
3.1 Refinement Lemma.....	10
3.2 Statement.....	11
3.3 Proof.....	11
3.4 Contrapositive Formulation	13
3.5 Scope and Generality	14
4. Information-Theoretic Formalization	14
4.1 Setup and Notation.....	14
4.2 Reformulated Definitions.....	15
4.3 The Information-Theoretic No-Go Theorem	16
4.4 Connection to the Data Processing Inequality	17
4.5 Channel Capacity Interpretation	18
4.6 Worked Example: Bit Erasure.....	19
4.7 Measurement as Capacity Saturation	20
4.8 Entropy Increase as Information Compression.....	21
5. Addressing Potential Objections	22
5.1 Is the Argument Circular?	22
5.2 Does Quantum Mechanics Already Solve This?.....	23
5.2.1 Decoherence Does Not Create Commitment	24
5.3 What About Many-Worlds?	24
5.4 What About Block Universe Interpretations?	25

5.5 Why Not Just Accept Reversibility?	25
5.6 Is This Just Effective Field Theory?	26
5.7 What About Non-Local Correlations?	26
5.8 What About Bohmian Mechanics?	27
6. Information-Theoretic Connections.....	28
6.1 The Bekenstein Bound	28
6.2 The Holographic Principle	28
6.3 Conceptual Foundation	28
7. Entropy, Erasure, and Landauer's Principle	29
7.1 Entropy Requires Finite Resolution	29
7.2 Information Erasure	29
7.3 Landauer's Principle as Consequence	29
7.4 Entropy as Commitment Ledger	30
8. Measurement and Outcome Finality	30
8.1 Measurement as Commitment	30
8.2 Finite Outcome Sets	30
8.3 Collapse as Closure.....	31
8.4 The Emergence of Classicality	31
8.5 Selection Mechanism.....	31
9. Time, Causation, and Temporal Direction	32
9.1 Time as Commitment Ordering	32
9.2 Why Time Has a Direction.....	33
9.3 Causation as Constraint Propagation	33
9.4 Discrete Temporal Structure	34
10. Excluded Physical Frameworks	34
10.1 What Infinite Distinguishability Would Entail	34
10.2 Classification of Frameworks	34
10.3 Interpretation of Exclusions	35
10.4 Constraint-Based Physics.....	35
11. Empirical Content and Constraints	36
11.1 Is This Empirically Meaningful?	36

11.2 Relation to Existing Physics	36
12. Conclusion	37
12.1 Summary of Results.....	37
12.2 Broader Implications.....	37
12.3 Outlook	37
Appendix A: Formal Statement and Proof.....	38
A.1 Metric Space Formulation: Definitions.....	38
A.2 Theorem (Metric Formulation)	38
A.3 Proof.....	38
A.4 Corollary	39
A.5 Measure-Theoretic Formulation: σ -Algebras	39
A.6 Theorem (σ -Algebra Formulation)	40
A.7 Proof.....	40
A.8 Connection to Information Theory.....	41
Appendix B: Classical Illustration—Bat and Ball	41
B.1 Standard Causal Narrative.....	41
B.2 Pre-Impact: Admissible Histories.....	41
B.3 Impact as Commitment	42
B.4 Retrospective Fixation.....	42
B.5 No Backward Causation.....	42
B.6 Classical Intuition Explained	42
Appendix C: Glossary of Key Terms	43
Appendix D: Operational Lift, Admissibility Tightening, and Empirical Contours	43
D.1 Motivation and Scope	43
D.2 The Operational Lift Principle	44
D.3 Reframing Step 3 of §3.3	44
D.4 Lemma: No Erased-but-Accessible Distinctions	45
D.5 Environment and Nonlocal Correlations.....	45
D.7 Relation to Planck and Holographic Scales	46
D.8 Falsifiability Criterion.....	46
References.....	47

Reversibility and Irreversibility	47
Information and Thermodynamics	47
Information Bounds and Holography.....	47
Quantum Measurement and Decoherence	47
Measurement Theory.....	47
Foundations and Interpretation	48
Hidden Variables and Pilot-Wave Theory	48

1. Introduction: The Problem of Facts

Modern physics relies fundamentally on irreversible facts. Detector clicks occur, records persist, memories form, entropy increases, and causal chains produce outcomes that cannot be undone. Yet the mathematical structures most commonly used to model physical reality—continuous state spaces, infinite-precision variables, and reversible dynamics—quietly permit perfect recoverability in principle. This tension between irreversible physical commitment and reversible mathematical description has remained unresolved.

This paper addresses a single foundational question: *under what conditions is irreversible commitment physically possible at all?*

We establish a necessity result: irreversibility cannot arise in systems with infinite distinguishability. If arbitrarily fine distinctions between states are physically accessible, then any apparent many-to-one evolution can be refined into an information-preserving mapping, rendering commitment illusory. Finite physical distinguishability is therefore not an assumption imposed for convenience, but a structural requirement for the existence of facts.

Remark (Minimal empirical premise). The premise "facts exist" is used in the weakest operational sense: experiments yield stable, reproducible records that can be compared across observers and times. Any theory that denies this undermines the empirical practice by which the theory itself is justified. The argument therefore treats record existence as a minimal precondition for doing physics at all.

No assumption is made here about the metaphysical uniqueness of outcomes beyond the operational existence of stable records; the argument concerns the physical possibility of record-finality within an admissible domain. This framing is compatible with relational, perspectival, or even Everettian interpretations, provided they acknowledge that *within* an accessible domain, records must stabilize.

The argument operates at the level of *admissibility*—identifying which mathematical descriptions can correspond to fact-producing physics—rather than proposing specific dynamics. This generality is a strength: the result constrains all candidate theories regardless of their particular field equations or interpretive commitments.

2. Formal Definitions and Framework

2.1 State Spaces and Distinguishability

Let S denote a state space representing physically possible configurations of a bounded region.

Definition 1 (Distinguishability). Two states $s_1, s_2 \in S$ are *physically distinguishable* if there exists a physically realizable procedure M that reliably yields different outcomes for s_1 and s_2 using finite resources.

In plain terms: Two states are distinguishable if you can actually tell them apart using some physical measurement or procedure. This is different from being *mathematically* different—two numbers can differ in their trillionth decimal place, but no physical apparatus could detect that difference.

Definition 2 (Infinite Distinguishability). A state space S has *infinite distinguishability* if, for any $\epsilon > 0$ and any partition of S into cells of diameter ϵ , each cell contains states that are themselves physically distinguishable by some finer procedure.

Equivalently: S has infinite distinguishability if there is no minimum resolution below which distinctions become physically inaccessible.

Moreover, the refinement is *iteratable*: for any physically distinguishable pair $s_1 \neq s_2$, there exists a procedure distinguishing them at strictly finer resolution, and this can be repeated without terminal scale. This ensures the induction in the main proof (§3.3) is well-founded.

Definition 3 (Finite Distinguishability). A state space S has *finite distinguishability* if there exists some $\delta > 0$ such that states separated by less than δ (in an appropriate metric) are not physically distinguishable by any realizable procedure.

Finite distinguishability implies a maximum information content: the number of distinguishable states in any bounded region is finite.

In plain terms: Infinite distinguishability means you can always "zoom in" further and find finer distinctions that are physically real and accessible. Finite distinguishability means there's a bottom level—a finest grain—beyond which no physical procedure can detect differences. Think of it like pixels on a screen: below the pixel level, there's nothing finer to see.

Remark (Metric-independence). The ϵ/δ language is used only to express an operational idea: whether there exists a minimum physically resolvable scale of distinction. The argument does not depend on a particular metric choice; any operational distinguishability relation induces a topology of resolvable distinctions. "Infinite distinguishability" means there is no terminal scale at which refinements become physically inaccessible.

Remark (Non-vacuity). The infinite-distinguishability condition is intentionally strong because it targets the precise claim made—often implicitly—by infinite-precision physics: that arbitrarily fine distinctions are not merely mathematically labelable but physically recoverable in principle. Many idealized frameworks (e.g., exact classical phase-space realism with infinite precision, continuum field models taken as literally complete at arbitrarily fine scales, or block-universe pictures with unlimited microstate specification) effectively commit to this stance. The theorem therefore functions as a diagnostic: it identifies which idealizations cannot be treated as complete descriptions of fact-producing reality.

Remark (Continuous spaces with finite distinguishability). Finite distinguishability does not require discrete state spaces at the mathematical level. A continuous manifold with a minimum resolvable scale δ has finite distinguishability: the effective state count is V/δ^n where V is the volume and n the dimension. The mathematical description may employ real numbers; what matters is that no physical procedure can distinguish states separated by less than δ . The discreteness is operational, not necessarily structural. Physics may use \mathbb{R} for convenience while physical distinguishability remains finite.

2.2 Irreversible Commitment

Definition 4 (Irreversible Commitment). A physical process $P: S \rightarrow S$ exhibits *irreversible commitment* if:

- (i) There exist distinct states $s_1, s_2 \in S$ such that $P(s_1) = P(s_2)$, and
- (ii) There exists no physically realizable process $Q: S \rightarrow S$ such that $Q(P(s_i)) = s_i$ for all such s_i .

Condition (i) requires genuine many-to-one mapping: multiple prior states yield the same outcome. Condition (ii) requires that this mapping be non-invertible in principle, not merely in practice.

Clarification (Accessible domain). Irreversible commitment is non-invertibility *with respect to the physically accessible degrees of freedom and admissible operations on them*. This domain restriction is essential: the question is not whether some hypothetical super-observer with access to all correlations in the universe could invert the process, but whether inversion is possible using resources within the accessible domain. If information is "preserved" only in degrees of freedom outside this domain, it is not preserved in any operationally meaningful sense.

Irreversible commitment is thus distinguished from:

- *Reversible dynamics* (one-to-one mappings)
- *Epistemic coarse-graining* (information hidden but recoverable within accessible domain)
- *Practical irreversibility* (recovery difficult but possible in principle within accessible domain)

In plain terms: Irreversible commitment is what happens when multiple possibilities collapse into a single definite outcome that cannot be "uncollapsed." When a detector clicks, when you remember something, when a record is written—these are commitments. The key question is whether such commitments are genuine (the alternatives are truly gone from the accessible domain) or merely apparent (the alternatives are hidden somewhere accessible and could in principle be recovered).

2.3 Physical Realizability

The phrase "physically realizable" requires explicit definition to avoid circularity.

Definition 5 (Physical Realizability). A process is *physically realizable* if it can be implemented using:

- (i) Finite energy
- (ii) Finite time
- (iii) Finite spatial extent
- (iv) Operations that do not presuppose access to information that has been irreversibly committed elsewhere

Condition (iv) prevents definitional circularity: we cannot define recoverability by appealing to a "cosmic ledger" whose existence already presupposes irreversible commitment.

Clarification on condition (iv): This condition prevents a specific circularity: defining "recoverability" by appealing to a hypothetical cosmic ledger that records all information. Such a ledger would itself require irreversible commitment to exist—the ledger's records must be facts. The condition ensures we don't smuggle commitment in through the back door while asking whether commitment is possible. Operationally: a recovery procedure must work using resources available within the light cone of the process, not by consulting an external record whose existence presupposes what we're trying to establish.

2.4 Admissibility Versus Mathematical Possibility

Throughout this paper, we distinguish *mathematical possibility* from *physical admissibility*.

Mathematical structures may permit infinite state density, perfect reversibility, or unbounded precision without internal contradiction. Physical admissibility, however, requires that states be preparable, distinguishable, and evolvable using finite physical resources.

Admissibility thus functions as a constraint layer on formal theories. It determines which mathematical descriptions correspond to realizable physics. The central result of this paper operates entirely at this level.

3. The No-Go Theorem

3.0 Assumptions (Admissibility Layer)

The following assumptions define the admissibility framework within which the theorem operates. Making them explicit controls the logical terrain and clarifies what the theorem does and does not assume.

(A1) Operationality. Distinguishability is defined by physically realizable tests under bounded resources. Two states are distinguishable if and only if some admissible procedure can reliably separate them.

(A2) Refinement Accessibility. If arbitrarily fine distinctions are physically accessible (infinite distinguishability), then there exists an admissible refinement procedure that can preserve (track) those distinctions through any process that does not itself constitute commitment at that refined level.

(A3) Record Minimality. "Facts exist" means stable records exist that can be compared across observers and times. This is the minimal empirical premise required for physics to be possible.

(A4) Domain Restriction. Irreversible commitment is non-invertibility with respect to the physically accessible degrees of freedom and admissible operations on them—not with respect to hypothetical extensions beyond physical accessibility.

These assumptions are not hidden premises smuggled into the argument; they are the explicit conditions under which the theorem holds. Rejecting any of them has consequences: rejecting (A1) makes distinguishability undefined; rejecting (A2) defines finite distinguishability by fiat; rejecting (A3) abandons empirical grounding; rejecting (A4) allows "cosmic ledger" escapes that presuppose what they deny.

Remark (Transcendental status of A3). The premise "facts exist" is not a metaphysical assumption but a *transcendental condition* for the practice of physics. Any argument against this premise would itself require formulation, communication, and evaluation—all of which presuppose the existence of the facts being argued about (the statements made, the records of the argument, the outcome of evaluation). Denying A3 is not merely uncomfortable but performatively incoherent: the denial cannot be maintained as a fact.

The argument thus has the structure: *if physics is possible at all, then finite distinguishability holds*. This is not circular but conditional, and the condition is one that anyone engaging in physics has already accepted.

3.1 Refinement Lemma

Before stating the main theorem, we establish a key lemma that makes the role of refinement explicit.

Lemma (Refinement Lemma). If two states $s_1, s_2 \in S$ are physically distinguishable, then there exists: (i) An admissible refinement variable Φ , and (ii) An admissible measurement M producing values in Φ ,

such that $M(s_1) \neq M(s_2)$. That is, Φ separates s_1 and s_2 .

Proof. By Definition 1, distinguishability means there exists a physically realizable procedure that reliably yields different outcomes for s_1 and s_2 . Let M be such a procedure and Φ its outcome space. Then $M(s_1) \neq M(s_2)$ by construction. The procedure M is admissible by hypothesis (it is physically realizable under bounded resources). ■

This lemma is nearly tautological—it simply unpacks what "distinguishable" means operationally. Its importance lies in making explicit the step that critics often attack: the "appearance" of Φ is not an addition of structure but a naming of structure that distinguishability already commits us to.

3.2 Statement

Theorem (No-Go for Infinite-Precision Irreversibility). Let S be a state space with infinite distinguishability. Then no process $P: S \rightarrow S$ can exhibit irreversible commitment.

Equivalently: Irreversible commitment requires finite distinguishability.

What this means: If you can always find finer and finer physical distinctions, then nothing is ever truly lost. Any process that seems to erase information or collapse possibilities is actually preserving that information in finer details you haven't looked at yet. Real, permanent facts—like measurement outcomes or memories—can only exist if there's a limit to how finely the universe can be carved up.

3.3 Proof

Suppose S has infinite distinguishability and $P: S \rightarrow S$ appears to exhibit irreversible commitment, with $P(s_1) = P(s_2) = s^*$ for distinct s_1, s_2 .

Step 1 (Many-to-one implies distinguishable inputs): Since $P(s_1) = P(s_2)$ with $s_1 \neq s_2$, we have at least two distinct input states mapping to the same output.

Step 2 (Apply Refinement Lemma): By infinite distinguishability, s_1 and s_2 are physically distinguishable. By the Refinement Lemma (§3.1), there exists an admissible refinement variable Φ and measurement M such that $M(s_1) = \varphi_1 \neq \varphi_2 = M(s_2)$.

Step 3 (Construct lifted process): Define the extended state space $S' = S \times \Phi$. The key observation is that the process P , being physical, must act through some physical mechanism that engages the degrees of freedom in Φ . Define:

$$P'(s, \varphi) = (P(s), \varphi_{\text{out}}(s, \varphi))$$

where $\varphi_{\text{out}}(s, \varphi)$ is the *actual final state* of the Φ degrees of freedom after P acts on state (s, φ) . This is not a constructed mathematical function but a description of what physically happens to all accessible degrees of freedom.

The critical point: under infinite distinguishability, φ_{out} is itself physically accessible. Since s_1 and s_2 are distinguishable via Φ , and the process P cannot "erase" this distinction without engaging Φ (erasure being physical), the distinction persists in φ_{out} unless P explicitly destroys it—but explicit destruction would constitute irreversible commitment at the Φ level, contradicting the hypothesis that no irreversible commitment occurs under infinite distinguishability.

Step 4 (Lifted process is injective): By construction:

- P' projects to P under coarse-graining (ignoring Φ)
- P' is injective: $P'(s_1, \varphi_1) \neq P'(s_2, \varphi_2)$ since $\varphi_{\text{out}}(s_1, \varphi_1) \neq \varphi_{\text{out}}(s_2, \varphi_2)$

Step 5 (Invertibility at refined level): Since P' is injective and Φ is physically accessible (by assumption A2), the process is invertible at the refined level. Therefore P does not exhibit irreversible commitment—the apparent information loss is recoverable.

Step 6 (Induction): If apparent loss remains at level S' , repeat. By infinite distinguishability, refinement is always available. By induction, no information loss is ever final.

Conclusion: Under infinite distinguishability, all apparent irreversibility is an artifact of incomplete description. No process exhibits genuine irreversible commitment. ■

Remark (Mathematical vs. physical invertibility). A map can be mathematically invertible yet physically non-invertible because the inverse would require operations outside admissibility (infinite precision, infinite memory, infinite control). Conversely, a map can be mathematically many-to-one yet appear irreversible only because we haven't accessed sufficiently fine degrees of freedom. The no-go theorem establishes that under infinite distinguishability, every apparently irreversible map falls into the second category: the "irreversibility" is always an artifact of incomplete description, never a feature of the physics. This is the core claim. Finite distinguishability is what makes some maps genuinely non-invertible within the admissible domain.

Clarification (No hidden-register escape). The refinement argument is not an addition of hidden variables; it is an explicit statement of what infinite distinguishability means operationally. If arbitrarily fine distinctions are physically accessible, then any purported information loss can be tracked by physically accessible refinements. The lifted space $S \times \Phi$ is therefore not a metaphysical extension but a formal representation of accessible structure that the infinite-distinguishability hypothesis already commits us to. If such refinements are not physically accessible, then distinguishability is finite in the relevant sense.

The construction doesn't posit new physics. It makes explicit what infinite distinguishability means operationally. To claim s_1 and s_2 are physically distinguishable just IS to claim there exists some physically accessible degree of freedom that differs between them. The lifted space doesn't add Φ ; it names what distinguishability already commits us to. If no such Φ exists, then s_1 and s_2 weren't distinguishable in the first place—which is precisely finite distinguishability.

Remark (Recovery vs. non-disturbance). The no-go targets the existence of a physically realizable inverse in principle, not an intervention-free readout. If distinctions are physically accessible under infinite distinguishability, then there exists some admissible procedure that can recover them, even if recovery requires interaction. If no such recovery procedure exists in principle, then the distinctions were not physically accessible and distinguishability is finite in the relevant sense.

Clarification (What is and is not assumed about information preservation). The argument does not assume a separate axiom of "information conservation." It uses only the operational meaning of infinite distinguishability: if arbitrarily fine distinctions are physically accessible, then there exists an admissible procedure that can track those distinctions through any process that does not explicitly erase them. If a process truly destroys those distinctions while they remain physically accessible, then the process itself constitutes irreversible commitment—contradicting the hypothesis that irreversible commitment is impossible under infinite distinguishability. Thus the incompatibility is structural: either (a) distinctions remain accessible and therefore cannot be irreversibly collapsed without finite distinguishability, or (b) they are not accessible, which is precisely finite distinguishability.

Clarification (Local commitment vs. global unitarity). The no-go result does not require that information be destroyed "in the universe as a whole." It requires only that certain distinctions become inadmissible to recover within the physical resources and degrees of freedom that remain operationally accessible after commitment. A globally unitary completion may conserve information in an enlarged description, but if that enlarged description remains physically accessible in bounded regions, then commitment does not occur. If it is not physically accessible, then distinguishability is finite in the relevant operational sense.

3.4 Contrapositive Formulation

The contrapositive is equally important:

Corollary. If irreversible commitment occurs in a physical system, then that system has finite distinguishability.

Since we observe irreversible facts—measurement outcomes, records, memories, entropy increase—we conclude that physical reality has finite distinguishability. This is an empirical conclusion, not merely a definitional one.

Clarification (Chaos does not create commitment). Chaotic mixing and practical unpredictability can explain why inversion is *difficult*, but not why it is *inadmissible in principle*. The present result targets principled non-invertibility: if arbitrarily fine distinctions remain

physically accessible, then any apparent irreversibility can always be re-described as reversible information flow into finer structure. Chaos accelerates loss of practical trackability, but it does not, by itself, create irreversible commitment.

3.5 Scope and Generality

The theorem makes no assumptions about:

- Specific dynamics or field equations
- Quantum versus classical physics
- Continuous versus discrete underlying structure
- Determinism versus indeterminism

It applies to any candidate physical theory. The constraint is structural: infinite distinguishability and irreversible commitment are logically incompatible.

4. Information-Theoretic Formalization

The conceptual argument of §3 can be made mathematically precise using information theory. This formalization connects the no-go result to established theorems, provides quantitative bounds, and makes the argument auditable by mathematical physicists.

Why information theory? Information theory provides a rigorous language for talking about what can be known, transmitted, and lost. "Entropy" measures uncertainty or information content; "mutual information" measures how much knowing one thing tells you about another. By translating our definitions into this language, we can prove precise theorems and connect to established results like the Data Processing Inequality and channel capacity bounds.

4.1 Setup and Notation

Let X be a random variable representing the pre-process state, drawn from a state space S with probability distribution $p(x)$. The Shannon entropy is:

$$H(X) = -\sum_x p(x) \log p(x)$$

Intuition: Entropy measures "how much you don't know" or equivalently "how much information is contained." A fair coin has entropy 1 bit; a biased coin has less. A system with N equally likely states has entropy $\log N$.

For a physical process $P: S \rightarrow S$, let $Y = P(X)$ denote the post-process state. The mutual information between input and output is:

$$I(X; Y) = H(X) - H(X|Y) = H(Y) - H(Y|X)$$

where $H(X|Y)$ is the conditional entropy—the remaining uncertainty about X given knowledge of Y .

Intuition: Mutual information $I(X; Y)$ measures how much information about the input survives in the output. If $I(X; Y) = H(X)$, nothing was lost. If $I(X; Y) < H(X)$, some information was destroyed.

Remark (Link between geometric and entropic formulations). The information-theoretic formalization treats distinguishability as an induced hierarchy of coarse-grainings $\mathcal{P}_1 \prec \mathcal{P}_2 \prec \dots$ on the state space, where refinement corresponds to physically accessible discrimination procedures. "Infinite distinguishability" in §2 means there is no terminal partition; the entropic condition in §4 captures the same fact by asserting that accessible refinements can increase $H(X|\mathcal{P})$ without bound. The two formulations are equivalent descriptions of the same operational constraint.

Definition (Refinement Sequence). A *refinement sequence* on state space S is a chain of partitions $\mathcal{P}_1 \prec \mathcal{P}_2 \prec \dots$ where \mathcal{P}_{i+1} refines \mathcal{P}_i (every cell of \mathcal{P}_{i+1} is contained in some cell of \mathcal{P}_i), and each partition corresponds to a physically realizable discrimination procedure. S has *infinite distinguishability* if and only if every refinement sequence can be extended: for all \mathcal{P}_n , there exists \mathcal{P}_{n+1} with $H(X|\mathcal{P}_{n+1}) < H(X|\mathcal{P}_n)$ for some distribution over S .

This makes the "unbounded refinement" condition precise and connects the geometric and entropic formulations rigorously.

4.2 Reformulated Definitions

Irreversible Commitment (Information-Theoretic): A process P exhibits irreversible commitment if there exists an admissible input ensemble $p(x)$ such that:

$$I(X; Y) < H(X)$$

Equivalently: $H(X|Y) > 0$ for some admissible distribution. This formulation avoids trivial cases (e.g., delta distributions where $H(X) = 0$) by requiring that information loss occurs for at least one physically preparable input ensemble.

Finite Distinguishability (Information-Theoretic): A bounded region has finite distinguishability if and only if there exists a maximum entropy $H_{\max} < \infty$ such that for all preparable distributions:

$$H(X) \leq H_{\max}$$

This is equivalent to the state space having finite effective cardinality $N = 2^{H_{\max}}$.

Infinite Distinguishability (Information-Theoretic): A system has infinite distinguishability if, for any description at resolution level \mathcal{P} with entropy $H(X_{\mathcal{P}})$, there exists a finer resolution \mathcal{P}' with:

$$H(X_{\mathcal{P}'}) > H(X_{\mathcal{P}})$$

and all distinctions in \mathcal{P}' are physically accessible.

4.3 The Information-Theoretic No-Go Theorem

Theorem (Information-Theoretic No-Go). Let S have infinite distinguishability. For any process P with apparent information loss $H(X|Y) > 0$ at description level \mathcal{P} , there exists a refinement \mathcal{P}' such that:

$$I(X'; Y') = H(X')$$

where X' and Y' are the refined descriptions. All apparently lost information is recoverable at finer resolution.

Proof.

(1) Suppose at description level \mathcal{P} we have $H(X|Y) > 0$ (apparent information loss).

(2) The conditional entropy $H(X|Y) > 0$ means that multiple distinct input states map to the same output:

$$\exists x_1 \neq x_2 : P(x_1) = P(x_2) = y^*$$

(3) By infinite distinguishability, there exists a refinement \mathcal{P}' that resolves finer structure within the \mathcal{P} -equivalence classes. Let Φ denote these additional degrees of freedom.

(4) Define the extended output:

$$Y' = (Y, \Phi)$$

where Φ encodes the information distinguishing x_1 from x_2 .

(5) Since Φ is physically accessible (by infinite distinguishability) and tracks the input distinctions, we have:

$$H(X'|Y) = H(X'|Y, \Phi) = 0$$

(6) Therefore:

$$I(X'; Y') = H(X') - H(X'|Y') = H(X')$$

All information is preserved.

(7) If apparent loss remains at level \mathcal{P}' , repeat the refinement. By infinite distinguishability, refinement is always available.

(8) By induction, no information loss is ever final at any resolution level.

Conclusion: Under infinite distinguishability, all apparent irreversibility is an artifact of incomplete description. No process exhibits genuine irreversible commitment. ■

Remark (Sufficient statistics formulation). The information-theoretic result can be stated more precisely using the concept of sufficient statistics. Under infinite distinguishability, refinements can always be chosen so that Y' becomes a *sufficient statistic* for X' —that is, $H(X'|Y') = 0$, meaning Y' captures all information about X' with no residual uncertainty. In finite distinguishability, no such sufficient refinement exists beyond the terminal partition. The output cannot be refined into a sufficient statistic for the input because the refinement hierarchy terminates. This is the information-theoretic essence of irreversible commitment: commitment occurs precisely when no sufficient statistic for the input is physically accessible from the output.

4.4 Connection to the Data Processing Inequality

The Data Processing Inequality (DPI) states that for any Markov chain $X \rightarrow Y \rightarrow Z$:

$$I(X; Z) \leq I(X; Y)$$

Processing cannot create information about the source. However, DPI does not guarantee that information IS lost—it permits $I(X; Y) = H(X)$ (perfect preservation).

In plain terms: DPI says you can't make information appear out of nowhere by processing data. But it doesn't say information must be lost—you might preserve everything. The question is: when must information actually be destroyed?

The finite distinguishability constraint provides the complementary bound. Let N be the maximum number of distinguishable states. Then for any process with $|\text{input states}| > |\text{output states}|$:

$$H(X|Y) \geq \log(|\text{input states}|) - \log(N)$$

At some resolution, loss becomes genuine because no further refinement exists. The state space "bottoms out."

In plain terms: If you try to cram more information than a system can hold, some must be lost. Finite distinguishability sets a hard limit on how much information any region can contain. Exceed that limit, and information is genuinely destroyed—not just hidden.

Together, DPI and finite distinguishability yield:

Proposition. In a system with finite distinguishability (maximum N states), any process mapping $M > N$ equiprobable input states to $K < M$ output states exhibits irreversible commitment of at least:

$$H(X|Y) \geq \log(M/N)$$

Derivation: For an equiprobable ensemble over M distinguishable inputs and an admissible output domain of capacity N , basic counting bounds imply this inequality whenever $M > N$. The pigeonhole principle guarantees that at least $\lceil M/N \rceil$ inputs map to some common output, forcing conditional uncertainty.

This is a quantitative lower bound on information loss—not merely the assertion that loss occurs, but how much must occur.

4.5 Channel Capacity Interpretation

A physical process P can be viewed as a noisy channel with input X and output Y . The channel capacity is:

$$C = \max_{\{p(x)\}} I(X; Y)$$

For a deterministic process (no noise in the forward direction), Y is a function of X , so $H(Y|X) = 0$ and:

$$I(X; Y) = H(Y)$$

If P is many-to-one (irreversible commitment), then $|\text{range}(P)| < |\text{domain}(P)|$, so:

$$H(Y) < H(X)$$

and information is genuinely lost.

Under infinite distinguishability, any channel can be "upgraded" to a refinement with:

$$C' = H(X')$$

by accessing finer output degrees of freedom. The effective channel capacity is always equal to source entropy—no information is ever lost.

Under finite distinguishability, channel capacity is bounded:

$$C \leq H_{\text{max}} = \log N$$

This hard ceiling forces information loss when input entropy exceeds capacity.

4.6 Worked Example: Bit Erasure

Consider the simplest case: erasure of a single bit.

Why this matters: Erasing a bit—resetting a memory cell to zero, for instance—is the most elementary irreversible operation. If we can't genuinely erase one bit, we can't genuinely erase anything. This example shows exactly where the argument bites.

Setup: Input $X \in \{0, 1\}$ with $H(X) = 1$ bit. The erasure process maps both states to a single output:

$$P(0) = P(1) = \emptyset$$

At the coarse level:

- $H(Y) = 0$ (output is constant)
- $I(X; Y) = 0$
- $H(X|Y) = 1$ bit (complete loss)

Under infinite distinguishability:

The erasure process necessarily involves physical degrees of freedom—heat dissipation, material reconfiguration, field dynamics. If distinguishability is infinite, these carry a record:

$$Y' = (\emptyset, \phi) \text{ where } \phi \in \{\phi_0, \phi_1\}$$

The extended output distinguishes the two cases:

- $P'(0) = (\emptyset, \phi_0)$
- $P'(1) = (\emptyset, \phi_1)$

Now $H(Y') = 1$ bit and $I(X; Y') = 1$ bit. No information lost.

If further refinement is always possible, this escape is always available. Erasure never completes.

Under finite distinguishability:

Suppose the physical substrate has N distinguishable states. The erasure process must map both inputs to states within this finite set. If both inputs map to indistinguishable final states (as required for erasure), then:

$$H(X|Y') = 1 \text{ bit (genuinely lost)}$$

No refinement can recover the distinction because no finer physical structure exists.

Landauer's bound follows: The 1 bit of lost information must be compensated by entropy increase elsewhere (heat dissipation $\geq kT \ln 2$), since total information is conserved but local information is destroyed.

4.7 Measurement as Capacity Saturation

Quantum measurement can be understood as channel capacity saturation.

Setup: A quantum system in state $|\psi\rangle = \sum_i \alpha_i |i\rangle$ is measured in the $\{|i\rangle\}$ basis.

Pre-measurement: The quantum state encodes information in amplitudes $\{\alpha_i\}$ and relative phases. For a d -dimensional system:

$$H_{\text{quantum}} = \log d \text{ (maximum distinguishable outcomes)}$$

Measurement as channel: The measurement process M maps quantum states to classical outcomes:

$$M: |\psi\rangle \rightarrow i \text{ with probability } |\alpha_i|^2$$

Information accounting:

- Input entropy: $H(\psi)$ can be arbitrarily large if we consider continuous parameters
- Output entropy: $H(\text{outcome}) \leq \log d$
- Under finite distinguishability: $H_{\text{max}} = \log d$

The measurement channel saturates capacity. All information beyond $\log d$ bits is irreversibly lost. The recorded outcome i constitutes an irreversible commitment precisely because:

1. The finite-dimensional output space cannot encode finer distinctions
2. No physical refinement accesses additional degrees of freedom
3. $H(\text{input}|\text{output}) > 0$ genuinely and irrecoverably

Born rule compatibility: Finite-outcome commitment is naturally consistent with Born-rule statistics. The structure $|\alpha_i|^2$ is compatible with:

- Finite output distinguishability (d outcomes)
- Unitarity at the pre-measurement level
- Consistency across subsystems

This suggests a route by which Born-rule statistics can be viewed as compatible with finite-outcome commitment under unitary pre-measurement evolution; a full derivation would require additional assumptions (e.g., noncontextuality or symmetry constraints) beyond the scope of this paper.

4.8 Entropy Increase as Information Compression

The Second Law can be reformulated as forced information compression.

Consider an isolated system with N_{max} distinguishable macrostates. At time t_1 , the system occupies one of W_1 microstates consistent with macrostate M_1 :

$$S_1 = k_B \log W_1$$

At time t_2 , it occupies one of W_2 microstates consistent with macrostate M_2 :

$$S_2 = k_B \log W_2$$

Under infinite distinguishability: All W_1 microstates remain individually trackable. The system's microstate is always recoverable. Entropy is merely a coarse-grained description with no fundamental status:

$$H(\text{micro}|\text{macro}) = \log W$$

$$\text{but } H(\text{micro}_{t_2}|\text{micro}_{t_1}) = 0.$$

Under finite distinguishability: Microstate distinctions exceeding N_{max} are physically inaccessible. When dynamics mix states beyond the resolution limit:

$$H(\text{micro}_{t_2}|\text{micro}_{t_1}) > 0 \text{ (genuine information loss)}$$

Entropy increase is real, not merely epistemic. The Second Law holds because:

1. Dynamics tend to spread distributions across state space
2. Finite distinguishability prevents tracking of fine structure
3. Information is genuinely destroyed, not merely hidden

The entropy bound $S_{\text{max}} = k_B \log N_{\text{max}}$ is the thermodynamic reflection of finite distinguishability.

Summary of §4: The information-theoretic formalization shows that irreversibility isn't just a philosophical puzzle—it has precise mathematical content. Information loss requires finite distinguishability because infinite distinguishability always provides an "escape route" where supposedly lost information persists in finer details. Established results like the Data Processing Inequality and channel capacity bounds confirm this structure. The Second Law of thermodynamics, Landauer's bound on erasure, and the finality of quantum measurements all emerge as consequences of finite distinguishability.

5. Addressing Potential Objections

This section anticipates and responds to the most common objections to the argument. Readers already convinced may skip to §6.

5.1 Is the Argument Circular?

Objection: The definition of "physically realizable" presupposes finite resources, making the conclusion trivial.

Response: The argument can be run in two independent modes.

Empirical mode: We observe irreversible facts (detector clicks, persistent records, entropy increase). These observations falsify infinite distinguishability as a property of physical reality, regardless of how we define admissibility.

The empirical mode does not depend on any definition of admissibility. It runs as follows:

1. We observe irreversible facts—this is not a theoretical claim but the precondition of observation itself. Every measurement, every record, every memory constitutes an irreversible commitment.
2. The no-go theorem proves: if distinguishability is infinite, irreversible facts cannot occur.
3. By modus tollens: distinguishability is finite.

This is an empirical conclusion from observable phenomena, analogous to inferring finite speed of information propagation from relativistic observations. The argument does not assume finiteness; it derives it.

The Empirical Argument (Summary)

Independent of all definitions of admissibility:

1. We observe irreversible facts (measurements yield stable outcomes, records persist, entropy increases).
2. The no-go theorem proves: infinite distinguishability \Rightarrow no irreversible facts.
3. By contraposition: irreversible facts \Rightarrow finite distinguishability.
4. Therefore: physical reality has finite distinguishability.

This has the same logical status as inferring finite light speed from the observed failure of superluminal signaling, or inferring energy conservation from the observed failure of perpetual motion.

Structural mode: Even granting the definitional point, the argument has substantive content. It identifies which mathematical structures are compatible with fact-producing physics. Many candidate frameworks—exact classical mechanics, unitarily closed quantum mechanics, block

universes with infinite precision—are revealed as incomplete. Not because they violate known laws, but because they cannot accommodate the existence of facts.

Even granting unbounded resources, infinite distinguishability still blocks principled commitment, because any apparent many-to-one mapping can be refined into a one-to-one mapping whenever distinguishability remains unbounded.

The argument thus functions as a *compatibility filter*, not a derivation from nothing. The no-go is not merely a "finite resources" statement; it is a structural incompatibility claim: if arbitrarily fine distinctions remain physically accessible, then commitment cannot be final, regardless of available resources.

5.2 Does Quantum Mechanics Already Solve This?

Objection: Quantum mechanics already limits distinguishability through complementarity and uncertainty. The argument adds nothing new.

Response: Quantum mechanics limits *simultaneous* distinguishability of non-commuting observables, but does not obviously impose finite distinguishability in the relevant sense:

1. Hilbert spaces may be infinite-dimensional
2. Position and momentum eigenstates form continuous spectra
3. Unitary evolution is perfectly reversible

The measurement problem persists precisely because standard quantum mechanics lacks a native mechanism for irreversible commitment. The projection postulate is added by hand; decoherence displaces but does not destroy information; many-worlds avoids commitment entirely.

The present argument identifies what any *solution* to the measurement problem must provide: a finite bound on distinguishable outcomes that cannot be refined away by accessing finer degrees of freedom.

Complementarity limits joint sharpness of certain observables, but does not by itself guarantee a finite bound on recordable, irreversibly committed distinctions in bounded regions; a theory can have continuous spectra and unitary reversibility while still respecting uncertainty.

Quantum mechanics is *compatible* with finite distinguishability (POVMs have finite outcome sets), but does not *guarantee* it at the foundational level. The no-go theorem clarifies what is required. Finite distinguishability in this paper concerns the physically committable distinctions that can be stabilized as records, not the mathematical cardinality of spectra in the kinematic formalism.

The objection correctly notes that POVMs have finite outcome sets. But this observation relocates rather than resolves the problem. *Why* do measurements have finite outcomes? The projection postulate is added axiomatically; it doesn't follow from unitary dynamics. The question is whether finite outcomes reflect (a) a fundamental constraint on distinguishability, or

(b) a contingent feature of measurement apparatus that could in principle be refined. If (b), then measurement finality is illusory. If (a), then quantum mechanics implicitly assumes finite distinguishability—precisely the condition this paper identifies as necessary. The no-go theorem clarifies what any interpretation must provide: not merely finite outcomes in practice, but a principled account of why refinement terminates.

5.2.1 Decoherence Does Not Create Commitment

A common response holds that decoherence explains irreversibility: environmental entanglement makes recovery practically impossible by spreading information across enormous numbers of degrees of freedom. But practical impossibility is not principled inadmissibility.

The distinction between practical and principled irreversibility is not academic fastidiousness—it determines whether facts are real or apparent. If information persists in environmental correlations, then the outcome is not a fact but an indexical appearance relative to a subsystem description. The "collapse" is revealed as coarse-graining, not commitment.

This has consequences: if decoherence preserves information in principle, then quantum Darwinism describes apparent robustness of records, not genuine factuality. The bat-and-ball becomes the bat-and-ball-and-environment-and-everything-it-correlates-with, never settling into a definite event. The claim "this ball went there" becomes permanently provisional—always recoverable by a sufficiently powerful observer who can track environmental correlations.

More fundamentally: if practical difficulty sufficed for factuality, then factuality would be technology-dependent. What counts as a "fact" would change as measurement precision improves. This conflates epistemology with ontology. The present argument targets whether facts exist at all, not whether we can access them.

Decoherence explains why interference disappears from local observations. It does not explain why outcomes become facts. That requires finite distinguishability.

5.3 What About Many-Worlds?

Objection: In Everettian quantum mechanics, commitment is branch-relative. Information is globally conserved (unitarity preserved) while facts exist relative to branches. This satisfies both the "facts exist" premise and global reversibility.

Response: Many-worlds relocates commitment to branch-relative description without explaining it. Within each branch, observers record definite outcomes. But what makes branch-relative facts *facts*?

If branches are merely labels on a globally unitary evolution, and inter-branch coherence is in principle recoverable at finer scales, then branch-relative "facts" inherit the same problem as any other apparent commitment under infinite distinguishability. The branching structure would be refinable, not fundamental.

The Everettian faces a dilemma:

- (a) Accept that branches themselves have finite distinguishability—making the branching structure discrete and non-refinable. But then finite distinguishability is built into the interpretation, conceding the present result.
- (b) Acknowledge that branch-relative facts are not genuine commitments but indexical descriptions of a fully reversible whole. But this faces the self-undermining objection: the Everettian's own assertion of many-worlds is itself a branch-relative "fact" with no genuine factuality.

Option (a) is compatible with this paper's framework. Option (b) renders the interpretation self-undermining in the sense of §5.4.

5.4 What About Block Universe Interpretations?

Objection: Eternalism holds that past and future events are equally real. This doesn't require infinite distinguishability.

Response: Block-universe interpretations face a dilemma:

If the block has infinite precision: All microstates are fully specified with unlimited resolution. But then no irreversible commitment occurs—"commitment" becomes an indexical illusion rather than a physical process. The apparent arrow of time has no structural basis.

If the block has finite information density: The temporal asymmetry we observe reflects genuine structural features of the block. Certain directions accumulate commitments; others do not. The framework is then compatible with our result, but the block is not the causally inert manifold often imagined.

The present framework does not refute eternalism *per se*. It constrains which versions of eternalism remain physically admissible. Specifically, the result constrains eternalism to versions with finite information density if one wants irreversibility to be a physical feature of the block rather than merely an indexical description.

5.5 Why Not Just Accept Reversibility?

Objection: Perhaps reality is fundamentally reversible and irreversibility is merely apparent or perspectival.

Response: This position faces severe difficulties:

1. *Empirical:* Irreversibility is not merely observed but constitutive of observation. Without irreversible records, no measurement could ever be completed or remembered.

2. *Self-undermining*: The claim "irreversibility is illusory" is itself a factual claim that requires irreversible commitment to formulate, communicate, and record.
3. *Explanatory vacuum*: If all processes are reversible, why do we universally experience temporal asymmetry? Boundary conditions alone cannot explain this without presupposing the very asymmetry they purport to derive.

Accepting fundamental reversibility does not dissolve the problem; it makes it insoluble.

Clarification (Bullet-biting is incoherent). One might attempt to "bite the bullet" and consistently accept that all facts, including the assertion of universal reversibility, are illusory. But bullet-biting here is not merely uncomfortable—it is operationally incoherent. The claim "I accept universal reversibility" cannot be meant in any standard sense. To assert is to commit; to communicate is to create a record; to argue is to constrain future admissible responses. The consistent reversibilist cannot consistently *do* anything that constitutes assertion, since assertion requires their interlocutor to be unable to "un-receive" the message. The position is not wrong but unoccupiable—it cannot be maintained by any process that could convey it.

5.6 Is This Just Effective Field Theory?

Objection: All you've shown is that physics has UV cutoffs. This is standard effective field theory—new physics appears at shorter scales, and our current theories are effective descriptions valid above some minimum length. Nothing new here.

Response: The relationship to effective field theory deserves clarification. EFT holds that descriptions valid at scale L may break down at scales $\ll L$, where new physics appears. This is compatible with either:

(a) *Finite distinguishability*: Physics terminates at some fundamental scale. There is no "new physics" below some minimum; the scale is absolute.

(b) *Infinite distinguishability*: New physics appears at every scale, with no terminus. Each effective theory is replaced by a finer one, ad infinitum.

Standard EFT is agnostic between these options. The present argument establishes that option (b) is incompatible with facts. This is not a claim about any particular EFT's validity, but about the necessary structure of whatever fundamental theory underlies all effective descriptions. EFT tells us that our current theories break down at short scales; the no-go theorem tells us that this breakdown must eventually terminate in finite distinguishability, not continue indefinitely.

5.7 What About Non-Local Correlations?

Objection: Information about local alternatives might be preserved in non-local correlations, accessible only through measurements on distant systems. Entanglement could preserve information that appears locally lost.

Response: Non-local correlations do not alter the core argument. If information about a local commitment is preserved in correlations with distant systems, then either:

- (a) Those correlations are themselves physically accessible from the local region—in which case commitment hasn't occurred locally; the information is still available.
- (b) They are not physically accessible from the local region—in which case finite distinguishability holds locally, which is the relevant operational domain.

The no-go theorem applies to the operationally relevant domain. If facts are local (as they appear to be—my detector clicks here, my memory forms here), then distinguishability must be finite in bounded regions. Global preservation in non-local correlations would render local facts illusory, returning us to the self-undermining problems of §5.5.

5.8 What About Bohmian Mechanics?

Objection: Bohmian mechanics has definite particle positions at all times, deterministic dynamics, and produces definite measurement outcomes—yet it uses continuous configuration space. Doesn't this escape the no-go theorem?

Response: Bohmian mechanics presents an interesting case that illuminates rather than escapes the argument. Analysis reveals it faces the same dilemma:

- (a) *If particle positions have infinite precision* (continuous configuration space taken literally), then measurement outcomes are determined by infinitely precise initial conditions. But this precision is not physically accessible—no finite procedure can determine positions to arbitrary accuracy. The apparent definiteness of outcomes reflects our epistemic coarse-graining, not genuine commitment. The theory describes a universe of infinitely precise facts that are operationally indistinguishable from probabilistic outcomes. The "facts" are not operationally accessible facts in the sense of §2.
- (b) *If configuration space has finite effective precision* (positions distinguishable only to some δ), then Bohmian mechanics implicitly incorporates finite distinguishability, and the no-go theorem is satisfied. The continuous mathematics is a convenience; the physics respects finite resolution.

Bohmian mechanics is thus *compatible* with finite distinguishability but does not *escape* the theorem. Its apparent definiteness either reflects finite operational distinguishability (satisfying the theorem) or is merely mathematical rather than physical (the "facts" being inaccessible and hence not facts in the operational sense).

The same analysis applies to any hidden-variable theory: if the hidden variables have infinite precision, they cannot ground operational facts; if they have finite precision, finite distinguishability holds.

6. Information-Theoretic Connections

6.1 The Bekenstein Bound

The Bekenstein bound states that the maximum entropy of a region is proportional to its boundary area:

$$S \leq 2\pi k R A / (\hbar c)$$

where R is the region's radius and A its surface area. This implies finite information content in bounded regions—precisely the finite distinguishability condition derived here from admissibility requirements alone.

6.2 The Holographic Principle

The holographic principle generalizes this insight: the information content of any spatial volume is bounded by the information that can be encoded on its boundary, at roughly one bit per Planck area:

$$I_{\max} \approx A / (4l_P^2)$$

where $l_P \approx 1.6 \times 10^{-35}$ m is the Planck length.

This is not infinite. A region of radius R contains at most:

$$N_{\text{states}} \approx \exp(\pi R^2 / l_P^2)$$

distinguishable configurations—a vast but finite number.

6.3 Conceptual Foundation

The present argument provides a *conceptual foundation* for such bounds. They are not merely empirical regularities or quantum-gravitational predictions to be derived from more fundamental principles. They are *necessary conditions* for the existence of irreversible facts.

Scope note. The bounds cited here are used as consistent exemplars of finite-information principles already suggested in established physics. The no-go result does not depend on their validity. Rather, it provides a conceptual foundation for why any successful fundamental theory is expected to implement some finite-information constraint, whether or not it takes the precise holographic form.

Any theory violating holographic bounds would permit infinite distinguishability and therefore fail to accommodate commitment. The bounds are not negotiable features of specific models; they are admissibility requirements that any fact-producing physics must satisfy.

7. Entropy, Erasure, and Landauer's Principle

7.1 Entropy Requires Finite Resolution

Entropy is standardly interpreted as:

$$S = k_B \ln \Omega$$

where Ω is the number of microstates consistent with macroscopic constraints. This formula presupposes that Ω is finite or at least that state space has bounded effective resolution.

If distinguishability were infinite, entropy could never meaningfully increase. Any apparent entropy growth would correspond to redistribution of information across scales, not genuine loss. One could always define finer-grained entropies that remain constant.

Entropy monotonicity (the Second Law) therefore *presupposes* finite distinguishability. Without it, the Second Law becomes a bookkeeping convention rather than a physical constraint.

7.2 Information Erasure

Information erasure is the paradigmatic irreversible commitment. To erase a bit is to render the prior distinction (0 vs 1) physically inaccessible.

If arbitrarily fine degrees of freedom remained accessible, erasure would be impossible in principle. The "erased" information would persist in finer structure. True erasure requires a finite state space within which all distinctions can be exhausted.

7.3 Landauer's Principle as Consequence

Landauer's principle states that erasing one bit of information requires dissipating at least:

$$E \geq k_B T \ln 2$$

of energy as heat.

This principle is often taken as foundational, linking information and thermodynamics. However, the present analysis reveals it as *contingent on* finite distinguishability.

Only when erasure is a genuine physical operation—possible only with finite state spaces—can an energetic cost be meaningfully assigned. In systems with infinite distinguishability, no bit is ever fully erased, and Landauer's bound loses operational meaning.

Finite distinguishability is thus *logically prior* to Landauer's principle.

7.4 Entropy as Commitment Ledger

From this perspective, entropy increase tracks the accumulation of irreversible commitments. Each commitment eliminates prior alternatives from physical accessibility. Entropy is not a measure of ignorance but a ledger of lost possibilities.

This interpretation unifies thermodynamic irreversibility with measurement, memory, and causation. All arise from the same structural requirement: finite distinguishability enables commitment; entropy records its accumulation.

8. Measurement and Outcome Finality

8.1 Measurement as Commitment

Measurement is the paradigmatic instance of irreversible commitment. A measurement does not merely reveal a pre-existing value; it produces a definite outcome from a set of prior possibilities. Once registered, alternative outcomes are no longer physically accessible.

This finality distinguishes genuine measurement from reversible interaction. Any account preserving full recoverability fails to explain why outcomes ever become definite.

8.2 Finite Outcome Sets

In quantum mechanics, measurements are represented by positive-operator-valued measures (POVMs). A POVM partitions the state space into a finite or countable set of outcomes:

$$\{E_i\} \text{ where } \sum_i E_i = I, E_i \geq 0$$

Each outcome corresponds to a physically distinguishable result. This partitioning encodes the physical limits of distinguishability.

If outcome resolution were infinite, measurement would never complete. The system would remain indefinitely suspended among ever-finer alternatives. Finite outcome sets are required for closure.

8.3 Collapse as Closure

The apparent "collapse" of the quantum state has traditionally been interpreted as dynamical disturbance. Once finite distinguishability is acknowledged, collapse admits reinterpretation as *closure* rather than force.

Prior to measurement, the system encodes a space of admissible possibilities. Measurement selects one outcome and renders others inaccessible. No additional mechanism is required beyond admissible completion of the outcome set.

Collapse reflects the exhaustion of distinguishability at the measurement resolution, not a violation of unitary evolution at finer scales.

8.4 The Emergence of Classicality

Classical behavior emerges when irreversible commitments accumulate faster than quantum coherences can be sustained. Finite distinguishability ensures eventual domination by committed facts, producing:

- Stable measurement outcomes
- Persistent records
- Effective classicality at macroscopic scales

This is not decoherence in the standard sense (which merely delocalizes information). It is commitment-induced closure enabled by finite state resolution.

8.5 Selection Mechanism

Remark (What this framework does and does not explain). The present framework establishes that commitment *must* occur but does not specify the *selection mechanism*—how one outcome is selected from the admissible set. This is intentional. The no-go theorem is a necessity result (*commitment requires finite distinguishability*) not a dynamical proposal (*here is how commitment works*).

Candidate selection mechanisms compatible with finite distinguishability include:

- *Objective stochastic processes* (GRW-type spontaneous collapse)
- *Gravitationally-induced decoherence* (Penrose objective reduction)
- *Relational actualization* (perspectival but operationally non-trivial)
- *Entropic selection* (outcome probabilities from accessible state volumes)
- *Pilot-wave dynamics* (Bohmian selection via particle positions, with finite effective precision)

Each supplies dynamics; the present paper supplies the structural constraint any such dynamics must satisfy. The Born rule, in particular, may emerge from such dynamics under additional

assumptions (symmetry, noncontextuality, decision-theoretic constraints), but this derivation lies beyond our scope.

The framework is thus *compatible* with multiple dynamical proposals while being *neutral* among them. What it excludes is any account that preserves infinite distinguishability—including interpretations where "selection" is merely apparent or indexical without genuine commitment.

9. Time, Causation, and Temporal Direction

This section explores implications of the framework for understanding time and causation. These implications range from relatively secure (temporal asymmetry requires commitment accumulation) to more speculative (time as emergent from commitment ordering).

We distinguish between:

- *Constraint claims:* Any account of temporal asymmetry must invoke something with the structure of irreversible commitment. This follows from the no-go theorem.
- *Constitutive claims:* Time *is* commitment ordering; before/after are *created* by facts. This is a stronger interpretive proposal compatible with but not strictly entailed by the theorem.

Readers may accept the constraints while remaining agnostic about the constitutive claims. The former are results; the latter are suggestions.

9.1 Time as Commitment Ordering

Time is commonly modeled as a coordinate parameter. Such representations struggle to account for temporal directionality and irreversibility.

Within the present framework, time is reinterpreted as an *ordering relation over irreversible commitments*:

- Each commitment establishes a before/after distinction
- The accumulation of such distinctions generates a partial order
- This order *is* physical time

This is a constitutive claim. More conservatively, one might say: this order is *isomorphic* to physical time, or *grounds* our experience of time. The no-go theorem establishes only that such an order requires finite distinguishability; it does not settle whether time *is* this order or merely *correlates* with it.

In plain terms: We usually think of time as a river that events float down. But perhaps it's the reverse: "before" and "after" are created when possibilities become facts. Time doesn't contain events; events generate time.

Remark (Modal neutrality). The framework is neutral on the modal status of pre-commitment alternatives. Whether alternatives are "real possibilities" in a robust metaphysical sense, or merely formal placeholders in an admissibility structure, commitment still functions to eliminate them from physical accessibility. The no-go theorem applies regardless: infinite distinguishability prevents any transition from having the structure of elimination, whether what is eliminated is "real possibilities" or "admissible descriptions."

9.2 Why Time Has a Direction

Standard explanations of temporal asymmetry appeal to:

1. *Boundary conditions* (low-entropy past)
2. *Dynamical asymmetries* (CP violation)

Both face regress: Why those boundary conditions? Why that asymmetry?

Finite distinguishability offers a structural alternative. Temporal direction is *constituted* by irreversible commitment, not imposed on a pre-existing temporal manifold. The asymmetry requires no prior explanation because it is not derived from something more basic.

To ask "why does time flow forward?" is to ask why facts exist rather than not—a question that lies outside physics.

9.3 Causation as Constraint Propagation

Causation arises naturally as the propagation of constraints induced by commitments:

1. An outcome is committed
2. This commitment restricts the space of admissible future states
3. Later states are constrained by earlier commitments

Causes are identified with prior commitments that limit future possibilities. This requires no metaphysical notion of "influence" or "power." Causal structure is bookkeeping over eliminated alternatives.

Finite distinguishability ensures such eliminations are genuine and irreversible, grounding causation in physical structure.

9.4 Discrete Temporal Structure

Because commitments are discrete events, time inherits a counting structure. While macroscopic time appears continuous, it is underwritten by sequences of finite commitments.

Let $N(t_1, t_2)$ denote the number of irreversible commitments between events t_1 and t_2 . Physical duration corresponds to this count, not to an independent geometric parameter.

This perspective supports event-based frameworks in which time is fundamentally discrete, emerging from commitment density rather than presupposed as a continuum.

10. Excluded Physical Frameworks

The no-go theorem isn't merely abstract philosophy—it has teeth. This section identifies which approaches to fundamental physics are ruled out as complete descriptions of reality, and which remain viable.

10.1 What Infinite Distinguishability Would Entail

To appreciate what the theorem excludes, consider what a universe with infinite distinguishability would actually be like:

A universe with infinite distinguishability would lack:

- Genuine facts (all states recoverable in principle)
- Meaningful entropy increase (information never lost)
- Final measurement outcomes (always refinable)
- Persistent records (always reversible)
- Temporal direction (no asymmetry in commitment)

Such a universe could not contain observers, since observation requires irreversible recording.

10.2 Classification of Frameworks

The following table classifies physical frameworks by their compatibility with irreversible commitment:

Framework	Infinite Distinguishability?	Admits IC?	Status
Classical mechanics (exact)	Yes	No	Excluded as complete, fact-producing description

Framework	Infinite Distinguishability?	Admits IC?	Status
Unitary QM (no collapse)	Effectively yes	No	Incomplete
QM + projection postulate	No (finite outcomes)	Yes	Admissible
Many-worlds (all branches real)	Yes (globally)	No	Relocates commitment to branch-relative records
Block universe (infinite precision)	Yes	No	Excluded
Block universe (finite info)	No	Yes	Admissible
Discrete spacetime models	No	Yes	Admissible
Holographic theories	No (area-bounded)	Yes	Admissible

Reading the table: "Excluded" means the framework cannot be a complete description of a universe containing facts. "Incomplete" means additional structure is needed. "Admissible" means the framework is compatible with irreversible commitment. Note that exclusion doesn't mean useless—classical mechanics remains extraordinarily powerful as an effective theory.

10.3 Interpretation of Exclusions

"Excluded" does not mean mathematically inconsistent or empirically falsified. It means *incompatible with being a complete description of fact-producing reality*.

Excluded frameworks may remain extraordinarily useful as:

- Effective theories (accurate within domains)
- Calculational tools (mathematically tractable)
- Limiting cases (idealized descriptions)

But they cannot be literal descriptions of a universe containing irreversible facts.

10.4 Constraint-Based Physics

The exclusions identified here are not philosophical preferences. They follow from the empirical existence of irreversible facts. Finite distinguishability functions as an admissibility constraint, analogous to:

- Energy conservation
- Causal locality
- Lorentz invariance

Such constraints identify necessary structural features of physical reality without specifying particular dynamics.

11. Empirical Content and Constraints

11.1 Is This Empirically Meaningful?

One might ask whether an admissibility argument has empirical content. It does, in three senses:

Diagnostic: The argument explains *why* certain physical structures exist. Discreteness scales, entropy bounds, and measurement finality are not contingent features but necessary conditions for facts.

Constraining: Any proposed fundamental theory must respect finite distinguishability. Theories postulating infinite precision, exact reversibility, or unbounded information density are ruled out as complete descriptions.

Empirical Constraints: The framework implies:

1. A fundamental minimum scale for distinguishability (independent of technology)
2. Minimum entropy production for any fact-establishing process
3. Discreteness in any complete spacetime theory
4. Area-scaling (not volume-scaling) of maximum information content

These implications align with quantum uncertainty, Landauer's bound, holographic entropy bounds, and approaches to quantum gravity—suggesting the framework captures genuine physical structure.

11.2 Relation to Existing Physics

The finite distinguishability requirement does not conflict with established physics. Rather, it provides conceptual foundations for features often introduced heuristically:

Physical Feature	Standard Status	Present Framework
Quantum uncertainty	Fundamental postulate	Consequence of FD
Landauer's bound	Empirical/thermodynamic	Requires FD
Bekenstein bound	Quantum gravity result	Admissibility requirement
Measurement finality	Interpretive addition	Enabled by FD
Second Law	Statistical/axiomatic	Requires FD

12. Conclusion

The argument of this paper can be summarized simply: facts require finitude. If the universe could be carved infinitely finely, nothing would ever be settled—every apparent outcome could be undone by looking closer. The existence of genuine facts, memories, records, and the arrow of time all depend on there being a limit to physical distinguishability. This isn't a discovery about how physics happens to work; it's a requirement for any universe that contains facts at all.

12.1 Summary of Results

This paper has established:

1. **No-Go Theorem:** Irreversible commitment is impossible in systems with infinite distinguishability.
2. **Necessity Result:** Finite distinguishability is required for the existence of facts, entropy increase, measurement outcomes, records, causation, and temporal direction.
3. **Framework Exclusions:** Physical pictures relying on infinite precision, exact reversibility, or unbounded information density are excluded as complete descriptions of reality.
4. **Conceptual Foundations:** Information-theoretic bounds (Bekenstein, holographic) and thermodynamic limits (Landauer) are revealed as admissibility requirements, not merely empirical regularities.

12.2 Broader Implications

The result has implications for:

Quantum Foundations: Any solution to the measurement problem must provide a mechanism for finite-outcome commitment that cannot be refined away.

Quantum Gravity: Discrete spacetime structures are not optional modeling choices but admissibility requirements.

Philosophy of Time: Temporal direction is constituted by commitment accumulation, not imposed by boundary conditions.

Computation: Physical computation is bounded by finite distinguishability, not merely by energy or time constraints.

12.3 Outlook

The framework developed here is intentionally minimal. It identifies necessary conditions without proposing specific dynamics. This opens a programmatic path:

1. Derive specific physical structures as minimal realizations of finite distinguishability
2. Determine the relationship between distinguishability bounds and known physics (Planck scale, holographic bounds)
3. Develop quantitative models of commitment-based temporal emergence

The present result serves as a foundation: it establishes *why* such structures must exist, leaving *how* they are realized to further investigation.

Appendix A: Formal Statement and Proof

This appendix provides two formal treatments of the no-go theorem: a metric-space formulation (§A.1–A.4) and a measure-theoretic formulation using σ -algebras (§A.5–A.7). The latter is included for mathematical rigor and to connect with standard probability theory.

A.1 Metric Space Formulation: Definitions

Let (S, d) be a metric space of physical states.

Definition A1. S has ε -*distinguishability* if for all $s_1, s_2 \in S$ with $d(s_1, s_2) > \varepsilon$, there exists a physically realizable measurement distinguishing s_1 from s_2 .

Definition A2. S has *infinite distinguishability* if it has ε -distinguishability for all $\varepsilon > 0$.

Definition A3. A map $P: S \rightarrow S$ exhibits *irreversible commitment* if:

- $\exists s_1 \neq s_2 : P(s_1) = P(s_2)$
- \nexists physically realizable $Q : Q \circ P = \text{id}$ on $\{s_1, s_2\}$

A.2 Theorem (Metric Formulation)

Theorem. If S has infinite distinguishability, then no map $P: S \rightarrow S$ exhibits irreversible commitment.

A.3 Proof

Let S have infinite distinguishability. Suppose $P: S \rightarrow S$ with $P(s_1) = P(s_2) = s^*$ for some $s_1 \neq s_2$.

Let $\varepsilon = d(s_1, s_2)/2 > 0$. By infinite distinguishability, \exists measurement M with:

- $M(s_1) = m_1$
- $M(s_2) = m_2$

- $m_1 \neq m_2$

Define the extended state space $S' = S \times M$ where M is the outcome space of M .

Define $P': S' \rightarrow S'$ by $P'(s, m) = (P(s), M(s))$.

Then:

- $P'(s_1, m_1) = (s^*, m_1)$
- $P'(s_2, m_2) = (s^*, m_2)$
- $(s^*, m_1) \neq (s^*, m_2)$ since $m_1 \neq m_2$

Hence P' is injective on $\{(s_1, m_1), (s_2, m_2)\}$.

Define $Q': S' \rightarrow S'$ by $Q'(s^*, m_i) = (s_i, m_i)$. Then $Q' \circ P' = \text{id}$.

The map Q' is physically realizable since M is (by hypothesis) physically realizable and Q' merely conditions on M 's outcome.

Thus P does not exhibit irreversible commitment; the apparent information loss is recoverable via the finer description (S', P') . ■

A.4 Corollary

Corollary. Irreversible commitment implies finite distinguishability.

Proof: Contrapositive of theorem. ■

A.5 Measure-Theoretic Formulation: σ -Algebras

For mathematical precision, we reformulate the result using σ -algebras, which provide the standard framework for probability and measurability.

Definition A4 (Physically Decidable Events). Let S be a state space. A σ -algebra of physically decidable events is a σ -algebra $\mathcal{A} \subseteq 2^S$ such that for each $A \in \mathcal{A}$, there exists a physically realizable procedure that determines whether $s \in A$ for any state s .

Definition A5 (Distinguishability via σ -algebras). Two states s_1, s_2 are \mathcal{A} -distinguishable if there exists $A \in \mathcal{A}$ with $s_1 \in A$ and $s_2 \notin A$.

Definition A6 (Refinement of σ -algebras). A σ -algebra \mathcal{A}' refines \mathcal{A} (written $\mathcal{A} \subseteq \mathcal{A}'$) if every set in \mathcal{A} is also in \mathcal{A}' .

Definition A7 (Infinite Distinguishability, σ -algebra version). A system has *infinite distinguishability* if for every physically decidable σ -algebra \mathcal{A} , there exists a finer physically decidable σ -algebra $\mathcal{A}' \supset \mathcal{A}$ that separates at least one pair of states not separated by \mathcal{A} .

Equivalently: the supremum of physically decidable σ -algebras is the discrete σ -algebra 2^S .

Definition A8 (Measurable Commitment). A measurable map $P: (S, \mathcal{A}) \rightarrow (S, \mathcal{A})$ exhibits \mathcal{A} -irreversible commitment if:

- $\exists A \in \mathcal{A}$ with $|P^{-1}(A)| > 1$ containing \mathcal{A} -distinguishable states
- \nexists measurable $Q: (S, \mathcal{A}) \rightarrow (S, \mathcal{A})$ with $Q \circ P|_{\{P^{-1}(A)\}} = \text{id}$

A.6 Theorem (σ -Algebra Formulation)

Theorem. If a system has infinite distinguishability (Definition A7), then no measurable map exhibits \mathcal{A} -irreversible commitment for any physically decidable \mathcal{A} .

A.7 Proof

Suppose $P: (S, \mathcal{A}) \rightarrow (S, \mathcal{A})$ appears to exhibit irreversible commitment: $P(s_1) = P(s_2) = s^*$ for \mathcal{A} -distinguishable s_1, s_2 .

By infinite distinguishability, there exists a finer σ -algebra $\mathcal{A}' \supset \mathcal{A}$ and a set $B \in \mathcal{A}'$ with $s_1 \in B$ and $s_2 \notin B$.

Extend P to $P': (S, \mathcal{A}') \rightarrow (S, \mathcal{A}')$ by preserving the indicator of B :

$$P'(s) = (P(s), 1_B(s))$$

where we adjoin the Boolean value $1_B(s)$ to the output.

Then $P'(s_1) \neq P'(s_2)$ since $1_B(s_1) = 1 \neq 0 = 1_B(s_2)$.

The inverse Q' defined by $Q'(s^*, b) = s_i$ where $1_B(s_i) = b$ is \mathcal{A}' -measurable.

By induction over refinements, any σ -algebra \mathcal{A} admits an extension under which apparent commitment becomes recoverable.

Conclusion: Irreversible commitment requires a *terminal* σ -algebra—one that cannot be further refined by physically decidable procedures. This is precisely finite distinguishability. ■

Remark (Codomain extension). The construction formally extends the codomain from S to $S \times \{0,1\}$. This is legitimate because we are asking whether *any* physically accessible description renders the process invertible. If the finer σ -algebra \mathcal{A}' includes the indicator 1_B , and 1_B is physically decidable (which it is, by the definition of \mathcal{A}'), then the extended output $(P(s), 1_B(s))$

is physically accessible. The "output space" of a physical process includes all accessible information about the post-process state, not merely a pre-designated mathematical codomain. The construction makes explicit the information that infinite distinguishability guarantees is available.

A.8 Connection to Information Theory

The σ -algebra formulation connects directly to information theory:

- A σ -algebra \mathcal{A} on S induces a partition $\mathcal{P}_{\mathcal{A}}$ (the atoms of \mathcal{A})
- $H(X|\mathcal{P}_{\mathcal{A}})$ measures uncertainty about X given \mathcal{A} -resolution
- Infinite distinguishability: for all \mathcal{A} , $\exists \mathcal{A}'$ with $H(X|\mathcal{P}_{\mathcal{A}'}) < H(X|\mathcal{P}_{\mathcal{A}})$
- Finite distinguishability: \exists terminal \mathcal{A}^* with no proper refinement

The entropic formulation of §4 and the σ -algebra formulation are thus formally equivalent.

Appendix B: Classical Illustration—Bat and Ball

To illustrate how irreversible commitment operates even in classical contexts, consider a bat striking a ball.

B.1 Standard Causal Narrative

Conventionally: the bat swings, contacts the ball, transfers momentum, and the ball departs along a trajectory. The bat "causes" the ball's motion. Time orders these events passively.

B.2 Pre-Impact: Admissible Histories

Prior to contact, the system does not occupy a unique microstate. Microscopic variations exist in:

- Bat angle and speed
- Contact point geometry
- Ball position and spin
- Material deformation profiles

These variations define a *space of admissible interaction histories*, not a single predetermined trajectory.

Before impact, "which point on the bat will strike which point on the ball" has no factual answer—only a range of admissible possibilities.

B.3 Impact as Commitment

The moment of contact constitutes irreversible commitment:

- Momentum is redistributed
- Microscopic deformations occur
- Sound and heat are generated
- Information about alternative geometries is lost
- Entropy increases

One interaction history becomes factual; alternatives are eliminated.

B.4 Retrospective Fixation

Only *after* commitment can a definite causal story be told:

"A specific point on the bat struck a specific point on the ball with a particular force profile, producing this trajectory."

These details were not pre-existing facts revealed by impact. They are *fixed by* the committed outcome. The event selects which prior configuration was actualized.

B.5 No Backward Causation

This account involves no backward-propagating signals. The bat does not receive information from the future. Rather, the temporal ordering itself is *generated by* commitment.

Asking "does cause precede effect?" presupposes a temporal structure that commitment creates.

B.6 Classical Intuition Explained

At macroscopic scales, commitment occurs rapidly with large entropy production. Alternative histories are eliminated almost instantaneously, creating the appearance of continuous deterministic causation.

Classical cause–effect reasoning emerges as effective bookkeeping over dense commitment sequences. It is not fundamental but derivative—a coarse-grained description of underlying commitment dynamics.

Appendix C: Glossary of Key Terms

Admissibility: The condition of being physically realizable using finite resources; distinct from mathematical possibility.

Commitment: The irreversible transition from multiple admissible possibilities to a single factual outcome.

Distinguishability: The capacity to reliably differentiate states using finite physical procedures.

Finite Distinguishability (FD): The condition that bounded regions contain only finitely many distinguishable states.

Infinite Distinguishability: The condition that arbitrarily fine distinctions remain physically accessible.

Irreversible Commitment (IC): A process producing outcomes from which prior alternatives cannot be recovered by any physically realizable operation.

Physical Realizability: Implementability using finite energy, time, spatial extent, and without presupposing inaccessible information.

Appendix D: Operational Lift, Admissibility Tightening, and Empirical Contours

D.1 Motivation and Scope

This appendix strengthens two aspects of the main text that are likely to attract the greatest scrutiny from skeptical readers:

- (1) The construction in §3.3 (Step 3), where apparent irreversibility is resolved by lifting the description to include finer-grained accessible degrees of freedom.
- (2) The connection between finite distinguishability and concrete physical bounds (e.g., Planck-scale or holographic limits), which in the main text is intentionally gestural in order to preserve generality.

The purpose of this appendix is not to modify the core no-go theorem, but to make explicit why the lifted description is operationally forced rather than mathematically optional, and to clarify

the empirical contours implied by finite distinguishability without over-claiming specific numerical scales.

D.2 The Operational Lift Principle

In §3.3 the proof appeals to an extended or “lifted” description of the process output in order to track distinctions that remain physically accessible under infinite distinguishability. This step may appear, at first glance, to introduce auxiliary structure by hand. The present subsection makes explicit that no such addition is being made.

Definition (Operational Transcript).

For any physically realizable process P acting on an initial state s within a bounded region, define the operational transcript $O(P, s)$ as the totality of degrees of freedom within the admissible domain that are physically accessible after the process has occurred. This includes, but is not limited to:

- pointer variables of measuring devices,
- emitted radiation or fields remaining within the domain,
- thermodynamic changes (heat, work, dissipation),
- any other degrees of freedom that can, in principle, be interrogated by admissible procedures.

The operational transcript is not a mathematical construct but a physical one: it is simply “everything that is left available to be known” after the process completes.

Operational Lift Principle (OLP).

The physically relevant output of a process is not an arbitrarily chosen coarse-grained variable, but the full operational transcript $O(P, s)$. Any claim of information loss or irreversible commitment must therefore be evaluated with respect to this transcript.

D.3 Reframing Step 3 of §3.3

With the Operational Lift Principle in place, Step 3 of §3.3 can be reformulated without introducing an explicit product space $S \times \Phi$.

Let P be a process such that, at a coarse description level, $P(s_1) = P(s_2)$ for distinct input states $s_1 \neq s_2$.

Define the operationally lifted map:

$$\tilde{P} : S \rightarrow \Omega, \quad \tilde{P}(s) := O(P, s),$$

where Ω is the space of possible operational transcripts.

If s_1 and s_2 are physically distinguishable under infinite distinguishability, then by definition there exists an admissible procedure capable of distinguishing them. If the operational transcripts were identical, $O(P, s_1) = O(P, s_2)$, then no admissible post-process procedure could distinguish whether s_1 or s_2 occurred. In that case, the distinction between s_1 and s_2 would have been eliminated from the admissible domain.

But elimination of a physically accessible distinction is precisely irreversible commitment. Therefore, under the hypothesis that no irreversible commitment occurs at this refined level, it must be the case that:

$$O(P, s_1) \neq O(P, s_2).$$

This establishes injectivity of the operationally lifted map \tilde{P} on distinguishable inputs, without appeal to hidden registers or artificial extensions.

D.4 Lemma: No Erased-but-Accessible Distinctions

Lemma.

If a distinction between two states is physically accessible prior to a process and no irreversible commitment occurs during that process, then the distinction must remain recoverable from the operational transcript.

Proof.

Suppose a distinction is physically accessible but not recoverable from the operational transcript. Then, by definition, no admissible post-process procedure can retrieve it. The distinction has therefore been eliminated from the admissible domain, which constitutes irreversible commitment. Contradiction. ■

This lemma makes explicit the dichotomy at the heart of the argument: either distinctions persist in accessible transcripts, or commitment has occurred.

D.5 Environment and Nonlocal Correlations

A common response invokes environmental or nonlocal correlations as repositories of “lost” information. The Operational Lift Principle clarifies that such correlations only matter insofar as they are physically accessible within the admissible domain.

If environmental degrees of freedom encoding the distinction are accessible, then the operational transcript differs and reversibility is restored. If they are not accessible, then the distinction is operationally destroyed and finite distinguishability holds in the relevant domain.

Thus, appeals to environment or global unitarity do not evade the theorem; they merely relocate which domain exhibits finite distinguishability.

D.6 Empirical Contours Without Scale Fixing

The no-go theorem is structural rather than dynamical. It does not determine the numerical value of the minimum distinguishability scale. However, it does imply empirical contours that any fundamental theory must respect.

These include:

- A finite upper bound on the number of stably distinguishable states in any bounded region.
- A maximum mutual information that a region can retain about its past.
- A non-zero minimum entropy production associated with establishing a new irreversible record.
- A fundamental coarse-graining scale that cannot be surpassed by improved technology.

Existing physics already contains candidate realizations of these contours, including Bekenstein-type entropy bounds, holographic area scaling, and Landauer's principle. The present result explains why some bound of this general kind must exist, independent of its precise numerical value.

D.7 Relation to Planck and Holographic Scales

The theorem does not derive the Planck length, Planck area, or holographic entropy bounds. Instead, it provides a conceptual explanation for why any successful quantum-gravitational theory is expected to implement a finite information bound per bounded region.

Determining whether the realized bound coincides with Planck-scale discreteness, holographic area laws, or another regulator is a separate empirical and theoretical task. The admissibility result constrains the space of viable theories but does not fix their detailed implementation.

D.8 Falsifiability Criterion

The framework developed in this paper would be falsified if one could demonstrate a physically realizable process that simultaneously:

- (1) Produces stable, reproducible records (irreversible commitment), and
- (2) Permits unbounded refinement of physically distinguishable states within the same bounded admissible domain.

No such process is currently known.

References

Reversibility and Irreversibility

- Loschmidt, J. (1876). Über den Zustand des Wärmegleichgewichtes eines Systems von Körpern mit Rücksicht auf die Schwerkraft. *Sitzungsberichte der Akademie der Wissenschaften, Wien*, 73, 128–142. [The reversibility objection to Boltzmann's H-theorem]
- Poincaré, H. (1890). Sur le problème des trois corps et les équations de la dynamique. *Acta Mathematica*, 13, 1–270. [Recurrence theorem: any bounded Hamiltonian system returns arbitrarily close to its initial state]

Information and Thermodynamics

- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191. [Minimum energy cost of bit erasure: $kT \ln 2$]
- Bennett, C. H. (1982). The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12), 905–940. [Logical reversibility and physical reversibility]
- Bennett, C. H. (2003). Notes on Landauer's principle, reversible computation, and Maxwell's Demon. *Studies in History and Philosophy of Modern Physics*, 34(3), 501–510.

Information Bounds and Holography

- Bekenstein, J. D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23(2), 287–298. [Entropy bound: $S \leq 2\pi kR E/\hbar c$]
- 't Hooft, G. (1993). Dimensional reduction in quantum gravity. In *Salamfestschrift* (pp. 284–296). World Scientific. [Holographic principle origin]
- Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics*, 36(11), 6377–6396. [Holographic bound: ~1 bit per Planck area]
- Bousso, R. (2002). The holographic principle. *Reviews of Modern Physics*, 74(3), 825–874. [Comprehensive review]

Quantum Measurement and Decoherence

- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3), 715–775. [Decoherence and pointer states]
- Zurek, W. H. (2009). Quantum Darwinism. *Nature Physics*, 5(3), 181–188. [Redundant recording and apparent classicality]
- Schlosshauer, M. (2005). Decoherence, the measurement problem, and interpretations of quantum mechanics. *Reviews of Modern Physics*, 76(4), 1267–1305.

Measurement Theory

- Busch, P., Lahti, P., & Mittelstaedt, P. (1996). *The Quantum Theory of Measurement* (2nd ed.). Springer. [POVM formalism]
- Nielsen, M. A., & Chuang, I. L. (2010). *Quantum Computation and Quantum Information* (10th anniversary ed.). Cambridge University Press. [Standard reference for quantum information]
- Holevo, A. S. (1982). *Probabilistic and Statistical Aspects of Quantum Theory*. North-Holland. [Quantum information limits]

Foundations and Interpretation

- Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press. [Objective collapse and gravitational decoherence]
- Zeh, H. D. (2007). *The Physical Basis of the Direction of Time* (5th ed.). Springer. [Arrow of time and decoherence]
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press. [Modern Everettian interpretation]

Hidden Variables and Pilot-Wave Theory

- Dürr, D., & Teufel, S. (2009). *Bohmian Mechanics: The Physics and Mathematics of Quantum Theory*. Springer. [Comprehensive treatment of Bohmian mechanics]
- Goldstein, S. (2017). Bohmian Mechanics. *Stanford Encyclopedia of Philosophy*. [Accessible overview of pilot-wave theory]
- Valentini, A. (2010). Inflationary cosmology as a probe of primordial quantum mechanics. *Physical Review D*, 82(6), 063513. [Non-equilibrium Bohmian mechanics]