The Distinguishability Barrier: Why Physical Theories Require a Record Primitive and Its Foundational Role in the VERSF Framework

Keith Taylor *VERSF Theoretical Physics Program*

Summary for the General Reader

Every scientific experiment ends the same way: something is recorded. A detector clicks, a needle deflects, a number appears on a screen. Without such records, there is nothing to compare, nothing to replicate, and nothing to call a fact. Yet most physical theories take the existence of these records for granted. They tell us what the laws of nature are, but not how the world produces the stable, readable outcomes that let us *test* those laws in the first place.

This paper asks a simple question: what is the minimum requirement for a physical fact to exist? We show that there is a sharp boundary—a **Distinguishability Barrier**—below which no stable record can form. Below this boundary, the world contains only *potentials*: possibilities that have not yet been pinned down by any irreversible process. Above it, committed, reusable distinctions become possible, and with them, empirical science.

We prove three main results. First, any theory that makes testable predictions must include, somewhere in its machinery, a rule for how records are produced—what we call a **record primitive**. Theories that lack this rule are incomplete: they describe a world of potentials but cannot connect to experiment. Second, the smallest possible record is **binary**—a single yes-or-no distinction between two stable states. This is the atom of empirical fact. Third, producing even this minimal record has an unavoidable **thermodynamic cost**: it must dissipate at least k_B T ln 2 of energy into the environment, the same bound discovered by Landauer in the context of computing. We also prove a **finite-resource no-go**: under realistic physical conditions—finite temperature, finite apparatus size, finite persistence time—the barrier cannot be made arbitrarily small. It has a strictly positive floor set by the thermodynamics of metastability.

The Void Energy-Regulated Space Framework (VERSF) takes these results as its starting point. Rather than assuming facts exist and building physics on top of them, VERSF begins at the barrier-crossing event—the first moment a potential becomes a fact—and derives space, time, mass, and gravity from the accumulation of these elementary binary commitments.

Table of Contents

Abstract

We establish the existence of a **Distinguishability Barrier**: a threshold in any physical description below which no stable record can form and, consequently, no empirical fact can be defined. We prove that any empirically meaningful theory must incorporate—either explicitly or implicitly—a primitive mechanism for stable record formation. We further demonstrate that the minimal irreducible unit of such a record is necessarily binary, corresponding to a single committed distinction between two robust macroscopic basins. These results are formalized

using information-theoretic and dynamical-systems language, yielding a Landauer-type lower bound on the entropy cost of the first factual commitment and a finite-resource no-go theorem showing that the barrier cannot be driven to zero under physically realistic conditions. We situate these results within the existing literature on the quantum measurement problem, quantum Darwinism, operational axiomatics, and device-independent approaches to foundations, and show that the Void Energy-Regulated Space Framework (VERSF) is distinguished among physical frameworks by taking this barrier-crossing mechanism as its foundational primitive, from which space, time, mass, and gravity are subsequently derived.

---

1. Introduction and Motivation

A physical theory is not merely a mathematical structure. It is a framework that maps states of the world to predicted observational outcomes—outcomes that must be certified by physically instantiated distinctions stable enough to be referenced, compared, and reused. The existence of such distinctions presupposes a mechanism for producing stable records. Without records, there are no observations; without observations, a theory possesses no empirical content.

If one descends far enough in theoretical description, one encounters a regime in which no stable record can exist. In this regime, alternatives may be mathematically well-defined, but no operational procedure is capable of certifying one alternative over another. What remains are potentials—coherent possibilities that have not been rendered factual by any irreversible commitment.

A clarification on our use of "empirical fact" is warranted, since the results of this paper depend on it. By "empirical fact" we mean the minimal operational object required for repeatable experimental comparison: a physically instantiated distinction that can be re-queried, compared across trials, and used as input to subsequent procedures. This is not an arbitrary stipulation but the weakest notion sufficient to ground empirical science. Any weaker notion—such as transient correlations that are not robust under perturbation, or mathematical distinctions that are not physically instantiated—cannot support statistical sampling, intersubjective verification, or the accumulation of experimental evidence. The definitions in Section 2 formalize this operational criterion precisely.

This paper formalizes the boundary between potential and fact. We introduce the Distinguishability Barrier: a rigorously defined threshold separating the regime of mere potentials from the regime of empirical facts. We prove that any theory with genuine empirical content must contain, either within its dynamics or within its interpretative apparatus, a primitive mechanism by which stable records form. The Void Energy-Regulated Space Framework (VERSF) distinguishes itself among physical frameworks by beginning precisely at this barrier-crossing mechanism rather than presupposing it.

1.1 Relation to Existing Work

The problem of how definite outcomes arise from quantum dynamics has generated several major research programmes. Zurek's quantum Darwinism and einselection [1,2] address how environmental decoherence selects a preferred pointer basis and proliferates information redundantly across environmental fragments. The decoherent histories programme of Griffiths, Omnès, and Gell-Mann and Hartle [3,4,5] defines consistent families of coarse-grained histories without invoking collapse. QBism (Fuchs, Mermin, Schack [6]) treats measurement outcomes as primitive personal experiences of agents, sidestepping the need for a physical mechanism of record formation.

These programmes address complementary aspects of outcome definiteness, but typically treat robust records as an emergent consequence of decoherence, consistency conditions, or agent-centered postulates. Our contribution is to formalize the structural necessity of a record variable with robustness and certification bounds in any empirically complete theory. We further show that this primitive has a unique minimal form (binary) and carries an irreducible entropic cost.

The operational or "device-independent" approach to quantum foundations—notably the informational axioms of Chiribella, D'Ariano, and Perinotti [15]—shares with our framework a commitment to deriving structural consequences from operational primitives (preparations, transformations, measurements with recorded outcomes). However, their axiomatization takes the existence of measurement outcomes as a given input, whereas the present paper asks what physical conditions must be met for those outcomes to exist in the first place. Our Lemma 3.1 and its Corollary 3.1.1 can be seen as formalizing Peres's dictum that "unperformed experiments have no results" [16]: below the Distinguishability Barrier, not only unperformed experiments but also performed-but-uncommitted interactions fail to produce results in the operationally certifiable sense.

The present paper demonstrates that the record primitive is *structurally unavoidable* in any empirically complete framework—not as an interpretative preference but as a theorem about the relationship between theory and observation. Quantum Darwinism explains *which* records are selected but presupposes that decoherence plus redundancy suffices to produce records—an assumption we show requires an additional commitment step (Section 4). Decoherent histories define families of projectors but rely on an external consistency condition to adjudicate which families are physically meaningful, implicitly importing a record criterion. QBism avoids the record problem by relocating facts to the agent, but at the cost of leaving the physical mechanism of outcome production unspecified.

The paper proceeds as follows. Section 2 introduces the formal definitions. Section 3 derives a quantitative mutual-information floor for any certified record (Lemma 3.1), establishes that the barrier is strictly positive on information-theoretic grounds, proves a finite-resource no-go showing that the barrier cannot be driven to zero under finite temperature, finite apparatus size, and nonzero persistence requirements, and illustrates these results with a concrete worked example (a symmetric double-well bit at room temperature). Section 4 establishes the necessity of a record primitive for any empirically complete theory and identifies concrete failure modes in existing frameworks that lack one. Section 5 demonstrates that the minimal irreducible record is binary and derives a Landauer-type entropic lower bound on commitment. Section 6 discusses a uniqueness conjecture for the binary commitment structure. Section 7 draws out the implications for the VERSF framework. Appendix A presents a constructive equivalence-class result showing that record formation is an unavoidable structural completion of any empirical theory.

---

2. Formal Definitions

2.1 Potential versus Fact

Definition 2.1 *(Potential State).* A potential state is a state description for which no operational procedure can certify a difference between it and any relevant alternative. Potential states are internally consistent mathematical objects, but they lack the physical anchoring required for empirical verification.

Definition 2.2 *(Factual State).* A factual state is a state for which at least one difference relative to an alternative has been certified by a stable physical record. Factual states are operationally accessible: they can be referenced, compared, and reused in subsequent procedures.

2.2 Operational Distinguishability

Definition 2.3 *(Operational Distinguishability).* Two states s_1 and s_2 are operationally distinguishable if there exists a physical interaction M such that $M(s_1)$ produces a record in R_1 and $M(s_2)$ produces a record in R_2, where R_1 and R_2 are disjoint and stable sets of record states. The key requirement is that the distinction is physically instantiated—not merely formal or mathematical—and that it persists long enough to serve as a basis for inference. This notion is made precise in Section 2.6.

2.3 Records

Definition 2.4 *(Record).* A record is a physically instantiated state whose encoded information can be re-queried and that remains robust against environmental perturbation over a nontrivial timescale. Formally, a record corresponds to a many-to-one mapping from microstates into a stable macroscopic basin of attraction—a region of phase space from which typical perturbations do not cause escape.

Remark on framework independence. Definitions 2.1–2.4 are stated in dynamical-systems language (basins of attraction, phase space) for concreteness. However, the logical content that drives the theorems of Sections 3–5 rests entirely on three operational criteria: (i) re-queryability (the record can be read more than once), (ii) robustness (the record survives perturbation over a nontrivial timescale), and (iii) bounded discrimination error (the record reliably distinguishes hypotheses). These criteria can be formulated without committing to a specific dynamical framework—they require only that one can specify a perturbation class, a timescale, and an error metric. The dynamical-systems formalization is a sufficient but not unique instantiation. This matters because the paper claims results that apply before any particular physics is specified: the definitions must not presuppose the structures the paper claims to derive. We therefore emphasize that wherever "basin of attraction" or "phase space" appears, the reader may substitute any formalization that captures the three operational criteria above.

2.4 Commitment Events

Definition 2.5 *(Commitment Event).* A commitment event is a transition in which the accessible state space of a composite system is irreversibly reduced, stabilizing one distinguishable outcome over alternatives. Commitment events are the dynamical processes by which potentials become facts. They are characterized by a loss of reversibility at the macroscopic level: once a commitment has occurred, the system cannot return to the pre-commitment superposition of alternatives without external intervention that exceeds the robustness threshold.

We note that commitment events are related to, but distinct from, Zurek's einselection [1]. Einselection identifies which pointer states survive environmental monitoring; commitment, as defined here, additionally requires that the selected outcome be stabilized into a reusable record with bounded error. The distinction matters: einselection is a *necessary* condition for commitment but, as we argue in Section 4, not a *sufficient* one.

2.5 Distinguishability Measure and Barrier

Definition 2.6 *(Distinguishability Measure).* Let D denote a quantitative measure of certified distinguishability. Representative instantiations include:

(i) the minimal mutual information between the hypothesis and the record required for reliable discrimination, and

(ii) the minimal irreducible partition number required for stable discrimination.

The **Distinguishability Barrier** D^* is defined as the infimum of D for which stable records exist.

2.6 Mathematical Formalization of Records and Certification

We now formalize the preceding operational definitions using standard information-theoretic and dynamical-systems language.

Let Ω denote the underlying microstate space of the composite world (system plus environment). Let S be the system of interest with microstate random variable X taking values in Ω_S, and let E be an environmental or apparatus subsystem with microstate random variable Y taking values in Ω_E. A measurement interaction is a physical channel (stochastic map) M that couples X into Y via dynamics

$$(X, Y) \mapsto (X', Y')$$

such that some coarse-graining of Y' acts as a robust macroscopic record.

To represent coarse-grained record states, define a surjective map $g : \Omega_E \to \mathscr{R}$, where $\mathscr{R}$ is a finite set of macroscopic record labels. The realized record label is $R := g(Y')$. Stability is captured by the existence of basins of attraction in Ω_E for which g is constant under typical perturbations.

Definition 2.7 *(Discrimination Task).* A discrimination task is specified by a binary hypothesis $H \in \{1, 2\}$ with prior probabilities (π_1, π_2) where $\pi_1, \pi_2 > 0$ (non-degenerate priors), and a pair of conditional distributions $P_1(X)$, $P_2(X)$ over system states given each hypothesis.

Definition 2.8 *(Certified Discrimination).* Given a discrimination task (H, P_1, P_2), two hypotheses are certifiably discriminable if there exists a measurement interaction M and record coarse-graining g such that a decision rule $\delta : \mathscr{R} \to \{1, 2\}$ achieves error probability

$P_e < \varepsilon$ for some fixed $\varepsilon < ½$,

where

$P_e := \pi_1 \Pr(\delta(R) = 2 \mid H = 1) + \pi_2 \Pr(\delta(R) = 1 \mid H = 2)$.

The threshold $\varepsilon < ½$ ensures that the record performs better than chance (random guessing yields $P_e = \min(\pi_1, \pi_2)$ for biased priors, or $P_e = ½$ for equiprobable hypotheses). In practice, the physically interesting regime is $\varepsilon \ll 1$ (high reliability), but the information-theoretic results of Sections 3 and 5 require only $\varepsilon < ½$.

Definition 2.9 *(Record Robustness).* A record map g is robust over timescale T if, for perturbations Π drawn from an admissible perturbation class $\mathcal{P}$ (thermal noise, small local disturbances, etc.), the record label remains invariant with high probability:

$\Pr[\, g(\Pi(Y'_t)) = g(Y'_t) \text{ for all } t \in [0, T] \,] \geq 1 - \eta$, with $\eta \ll 1$.

Definition 2.10 *(Distinguishability Measures).* We define certified distinguishability in two complementary forms.

Hypothesis-testing form: Given a discrimination task with non-degenerate priors (π_1, π_2) and a measurement-record pair (M, g) achieving certified discrimination ($P_e < \varepsilon$), define:

$D_I := I(H ; R)$,

the mutual information between the hypothesis variable H and the record label R.

Partition-based form:

$D_m := m$,

where m is the number of committed outcome classes in $\mathscr{R}$ (i.e., $|\mathscr{R}|$).

For binary irreducibility, the minimal nontrivial record alphabet has $m = 2$. VERSF identifies the first certified distinction with $D^*_m = 2$, corresponding to the existence of a stable two-basin record.

Definition 2.11 *(Task Class and Distinguishability Barrier D*).* Fix a task class $\mathcal{T}$ specified by the following parameters:

- *Prior bounds:* $\pi_i \geq \pi_{min} > 0$ for all hypotheses (non-degenerate priors bounded away from zero).
- *Certification threshold:* $\varepsilon \leq \varepsilon_{max} < \pi_{min}$ (since a trivial single-basin record achieves $P_e = \min(\pi_1, \pi_2) \geq \pi_{min}$, the condition $\varepsilon_{max} < \pi_{min}$ excludes all trivial records uniformly across the task class). The additional condition $h(\varepsilon_{max}) < H(\pi_{min})$ ensures the information-theoretic floor is nontrivial; this holds automatically for equiprobable priors and is satisfied for any ε_{max} sufficiently small relative to π_{min}.
- *Robustness parameters:* timescale $T > 0$, failure probability $\eta \ll 1$, re-query interval $\Delta t_{min} > 0$.

Let $\mathcal{M}$ denote the physically admissible set of measurement interactions and $\mathcal{P}$ the admissible perturbations. Define the set of achievable robust records as:

$\mathcal{R}_{ach} := \{ (M, g) : M \in \mathcal{M}, g \text{ robust over } T \text{ with parameters } (\eta, \Delta t_{min}) \}$.

The barrier thresholds relative to $\mathcal{T}$ are:

$D^*_I := \inf \{ I(H ; R) : (M, g) \in \mathcal{R}_{ach}, P_e < \varepsilon_{max}, \pi_i \geq \pi_{min} \}$,

$D^*_m := \min \{ |\mathcal{R}| : (M, g) \in \mathcal{R}_{ach} \text{ and } P_e < \varepsilon_{max} \}$.

The barrier statement "below D* there are only potentials" means: for $D_I < D^*_I$ (or for $m < D^*_m$), no robust discrimination procedure exists that yields a reusable record with error bounded by ε_{max} over timescale T. With these bounds locked, the prior entropy satisfies $H(H) \geq H(\pi_{min}) := -\pi_{min} \ln \pi_{min} - (1 - \pi_{min}) \ln(1 - \pi_{min}) > 0$, ensuring that the positivity results of Sections 3.2–3.3 carry explicit, non-vacuous lower bounds.

---

3. The Distinguishability Barrier

3.1 Certification Implies a Mutual-Information Floor

Lemma 3.1 *(Certification Implies a Mutual-Information Floor).* Let $H \in \{1, 2\}$ with priors satisfying $\pi_i \geq \pi_{min} > 0$. If there exists a record variable R and decision rule $\delta(R)$ with $P_e < \varepsilon_{max} < ½$, then

$I(H ; R) \geq \max\{0, H(\pi_{min}) - h(\varepsilon_{max})\}$, where $H(\pi_{min}) := -\pi_{min} \ln \pi_{min} - (1 - \pi_{min}) \ln(1 - \pi_{min})$ and $h(\cdot)$ is the binary entropy function (in nats). In particular, if $h(\varepsilon_{max}) < H(\pi_{min})$, the bound is strictly positive. For equiprobable priors ($\pi_{min} = ½$), this condition is automatically satisfied for any $\varepsilon_{max} < ½$, giving $I(H ; R) \geq \ln 2 - h(\varepsilon_{max}) > 0$.

Proof. By Fano's inequality, the conditional entropy of H given R satisfies $H(H \mid R) \leq h(P_e)$. (The standard Fano bound $H(X \mid Y) \leq h(P_e) + P_e \ln(|\mathcal{X}| - 1)$ reduces to $h(P_e)$ here because $|\{1, 2\}| = 2$ and $\ln(1) = 0$.) Since $P_e < \varepsilon_{max}$ and h is monotonically increasing on [0, ½), we have $H(H \mid R) \leq h(\varepsilon_{max})$. The prior entropy satisfies $H(H) \geq H(\pi_{min}) > 0$ for non-degenerate priors. Therefore:

$$I(H ; R) = H(H) - H(H \mid R) \geq H(\pi_{min}) - h(\varepsilon_{max}).$$

Since mutual information is always nonnegative, we may sharpen to $I(H ; R) \geq \max\{0, H(\pi_{min}) - h(\varepsilon_{max})\}$. ■

Corollary 3.1.1 *(No Facts at Vanishing Information).* If $h(\varepsilon_{max}) < H(\pi_{min})$ and $I(H ; R) = 0$ for all physically realizable robust record variables R, then no certified distinction—and hence no empirical fact (Definitions 2.1–2.2)—exists for that task.

Proof. By contrapositive: if a certified distinction exists with $h(\varepsilon_{max}) < H(\pi_{min})$, Lemma 3.1 gives $I(H ; R) \geq H(\pi_{min}) - h(\varepsilon_{max}) > 0$. ■

Remark. Lemma 3.1 and its corollary replace a purely qualitative "no facts below the barrier" statement with a quantitative bound that hooks directly into the D^*_I story: the barrier is the infimum of $I(H ; R)$, and any record that certifies a distinction must contribute at least $H(\pi_{min}) - h(\varepsilon_{max})$ nats. The substantive question is whether D^* is *physically nontrivial*—whether physical constraints force it to be not merely positive but bounded away from zero under realistic conditions. This is established next.

3.2 Nontriviality of the Barrier

Proposition 3.2 *(Strict Positivity of D*).* Under any task class $\mathcal{T}$ with $\pi_{min} > 0$, $\varepsilon_{max} < \pi_{min}$, and $h(\varepsilon_{max}) < H(\pi_{min})$, the distinguishability barrier satisfies $D^*_m \geq 2$ and $D^*_I > 0$. (The condition $\varepsilon_{max} < \pi_{min}$ ensures that $m = 1$ is excluded uniformly across the task class; $h(\varepsilon_{max}) < H(\pi_{min})$ ensures the information-theoretic floor is nontrivial. Both hold automatically when priors are equiprobable and $\varepsilon_{max} < ½$.)

Proof.

$D^*_m \geq 2$. A partition with $m = 1$ assigns all apparatus microstates to a single macroscopic basin. Any decision rule δ applied to a single-label record produces the same output regardless of the hypothesis, yielding $P_e = \min(\pi_1, \pi_2) \geq \pi_{min}$. Since $\varepsilon_{max} < \pi_{min}$ by assumption, no single-basin partition achieves $P_e < \varepsilon_{max}$ across any task in the class. Hence certified discrimination requires $m \geq 2$.

$D^*_I > 0$. By Lemma 3.1, every record achieving $P_e < \varepsilon_{max}$ within the task class satisfies $I(H ; R) \geq H(\pi_{min}) - h(\varepsilon_{max})$. Since $h(\varepsilon_{max}) < H(\pi_{min})$ by assumption, this bound is strictly positive. Since D^*_I is the infimum of $I(H ; R)$ over all achievable robust records within $\mathcal{T}$:

$$D^*_I \geq H(\pi_{min}) - h(\varepsilon_{max}) > 0.$$

This bound depends only on the task-class parameters, not on the physical details of the measurement channel. ■

Remark. Lemma 3.1 and Proposition 3.2 establish the positivity of D* on purely information-theoretic grounds: certification forces $I(H ; R)$ away from zero regardless of the physical implementation. But a referee may reasonably ask whether physical implementations can *approach* the information-theoretic floor arbitrarily closely—whether D* is positive but negligibly small. The following no-go result shows this is not the case: under finite physical resources, D* is bounded below by a constant that depends on temperature, apparatus size, and persistence requirements.

3.3 Finite-Resource No-Go: D* Cannot Be Driven to Zero

To state the physical no-go precisely, we first formalize the re-query requirement that is implicit in the notion of a "record."

Definition 3.3 *(Nonzero Re-Query Interval).* A record is re-queryable if there exists a minimum interval $\Delta t_{min} > 0$ such that the record can be read at times t and $t + \Delta t_{min}$ with bounded additional disturbance, and with the robustness condition (Definition 2.9) holding over the interval. The requirement $\Delta t_{min} > 0$ is physically necessary: any readout involves a finite physical interaction, and "re-queryable" means the record can survive at least two reads separated by a nonzero time.

Model class for the finite-resource no-go. Proposition 3.3 below applies to the following class of record implementations:

- *Two metastable macrostates:* The record consists of exactly two basins B_0, B_1 in the apparatus phase space, separated by a free-energy barrier.
- *Passive stabilisation:* The record is maintained by the metastability of the basins alone, with no active error correction, refresh, or external maintenance during the persistence interval [0, T]. (Active stabilisation protocols that periodically reset or refresh the record constitute additional resource expenditure—equivalent to increasing ΔF_{max} or apparatus size—and are excluded by the finite-resource assumption.)
- *Thermal perturbation class:* The admissible perturbations $\mathcal{P}$ include thermal noise at temperature $T_{env} > 0$, acting continuously on the apparatus.
- *Finite resource budget:* The free-energy barrier satisfies $\Delta F \leq \Delta F_{max} < \infty$, reflecting finite apparatus size and interaction strength.

This model class covers the standard physical setting of a memory element (spin, latch, potential well) at finite temperature without active maintenance—the minimal implementation of a single certified distinction.

Proposition 3.3 *(Finite-Resource No-Go for Vanishing Barrier).* Within the model class above, fix a finite temperature $T_{env} > 0$, a robustness timescale $T > 0$, a nonzero re-query interval $\Delta t_{min} > 0$, and admissible perturbations $\mathcal{P}$ that include thermal noise at T_{env}. Consider any physically admissible measurement–record implementation (M, g) whose record variable $R = g(Y')$ must be:

- *robust:* $\Pr[g(\Pi(Y'_t)) = g(Y'_t) \text{ for all } t \in [0, T]] \geq 1 - \eta$ for some $\eta \ll 1$,
- *re-queryable:* satisfying Definition 3.3 with interval $\Delta t_{min} > 0$, and
- *certifying:* achieves a discrimination error $P_e < \varepsilon$ for some fixed $\varepsilon < ½$.

Assume additionally that the apparatus has finite stabilisation resources, formalised by an upper bound on the achievable metastability barrier $\Delta F \leq \Delta F_{max} < \infty$ for its record basins.

Then there exists a constant

$$c = H(\pi_{min}) - h(p_{min}),$$

where $p_{min} = p_{min}(T_{env}, \Delta t_{min}, \Delta F_{max}, \tau_0)$ is the minimum achievable discrimination error set by physical constraints, such that for any admissible robust record implementation,

$$I(H ; R) \geq c,$$

and hence $D^*_I \geq c$. For nontrivial task classes satisfying $h(p_{min}) < H(\pi_{min})$—which holds whenever the physical noise floor is small enough relative to the prior uncertainty, and in particular for all equiprobable-prior tasks—we have $c > 0$. In particular, D^* cannot be driven to zero unless at least one of the following is allowed to diverge:

- $\Delta F_{max} \to \infty$ (infinite stabilisation energy / apparatus size),
- $T_{env} \to 0$ (zero temperature),
- $\Delta t_{min} \to 0$ (instantaneous readout—no finite re-query interval),
- or the perturbation class $\mathcal{P}$ is artificially restricted to exclude thermal noise.

This bound applies to an irreducible binary record primitive—a single certified distinction with no redundancy resources beyond ΔF_{max}. Redundancy (e.g., error-correcting codes, repeated copies) shifts the constraint into increased resource use: it effectively increases ΔF_{max} or apparatus size, which is excluded by the finite-resource assumption.

Proof. The argument proceeds in two steps that together form a sandwich: robustness forces a minimum metastability barrier, and finite metastability forces a nonzero noise floor on the record channel, which in turn forces $I(H ; R)$ bounded away from zero.

Step 1 *(Robustness forces a minimum metastability barrier).* The record must have at least two macroscopic basins B_0, B_1 (otherwise it is trivial and cannot certify). Under thermal perturbations at T_{env}, standard activated-escape results (Kramers theory [14]) give a mean escape time that scales as:

$$\tau_{escape} \sim \tau_0 \exp(\Delta F / k_B T_{env}),$$

where τ_0 is a microscopic attempt time, ΔF is the free-energy barrier separating the basins, and ~ denotes Kramers scaling (equality up to a prefactor depending on barrier curvature and friction, which we absorb into τ_0 for notational simplicity). Robustness over timescale T at failure probability η requires the escape probability over [0, T] to satisfy:

$$\Pr(\text{escape in } [0, T]) \approx 1 - \exp(-T / \tau_{escape}) \leq \eta.$$

For $\eta \ll 1$ this implies $\tau_{escape} \geq T / |\ln(1 - \eta)|$, and combining with the Kramers scaling gives a minimum barrier requirement:

$$\Delta F \geq k_B T_{env} \ln(T / (\tau_0 |\ln(1 - \eta)|)).$$

Imposing the finite-resource bound $\Delta F \leq \Delta F_{max}$, this constrains the feasible design space: for given T_{env}, T, η, the record cannot be made arbitrarily robust without a sufficiently large ΔF_{max}. With ΔF_{max} fixed and finite, robustness consumes a finite fraction of the available metastability budget.

Step 2 *(Finite metastability implies a nonzero noise floor on achievable error).* Any robust binary record is effectively a noisy binary channel whose noise is bounded below by the thermal flip rate. Let $R_t \in \{0, 1\}$ denote the basin label at time t. Even if the measurement interaction initially produces a perfect correlation between H and R_0 (i.e., $P_{e,initial} = 0$), thermal perturbations cause random flips $0 \leftrightarrow 1$ at a rate bounded below by the activated mechanism:

$$\lambda \geq \tau_0^{-1} \exp(-\Delta F_{max} / k_B T_{env}) =: \lambda_{min} > 0.$$

The key step is that re-queryability (Definition 3.3) requires the record to persist over a nonzero interval $\Delta t_{min} > 0$ during which the record is exposed to thermal noise. Over any operational re-query interval $t \geq \Delta t_{min} > 0$, the symmetric two-state Markov model gives a flip probability:

$$p_{flip}(t) = \tfrac{1}{2}(1 - \exp(-2\lambda t)) \geq \tfrac{1}{2}(1 - \exp(-2\lambda_{min} \Delta t_{min})) =: p_{min} > 0.$$

This lower bound is strictly positive because each of λ_{min}, Δt_{min}, and the exponential factor are strictly positive under the stated assumptions. The bound cannot be evaded by querying "immediately" after commitment: Definition 3.3 requires that any re-query occurs after at least Δt_{min}, during which thermal noise acts on the record.

The flip probability p_{flip} provides a lower bound on P_e at re-query time because even an initially perfect measurement ($P_{e,initial} = 0$) degrades to $P_e \geq p_{flip}$ after the record has been exposed to thermal noise over the re-query interval; imperfect initial measurements can only make P_e worse. Therefore, for any admissible record implementation under the stated resource constraints:

$$P_e \geq p_{min} > 0.$$

Combining the two results. The bound on $I(H ; R)$ now follows from two logically independent ingredients working together:

(i) *Information-theoretic (Lemma 3.1):* For any record achieving $P_e < \varepsilon_{max}$ within task class $\mathcal{T}$, the mutual-information floor gives $I(H ; R) \geq \max\{0, H(\pi_{min}) - h(\varepsilon_{max})\}$.

(ii) *Physical (Step 2 above):* Under finite resources, the achievable discrimination error satisfies $P_e \geq p_{min} > 0$, so the certification threshold ε_{max} cannot be taken below p_{min}.

Existence threshold. Under finite resources, the feasible certification regime is $\varepsilon_{max} > p_{min}$. If $\varepsilon_{max} \leq p_{min}$, no record in the model class meets the certification requirement $P_e < \varepsilon_{max}$, hence the achievable set $\mathscr{R}_{ach}$ is empty for that task class—the barrier is absolute. In the feasible regime $\varepsilon_{max} > p_{min}$, admissible records exist and Lemma 3.1 yields a nontrivial information floor for each.

In this feasible regime, any achievable record has $P_e \geq p_{min}$, and for such a record Fano gives $I(H ; R) \geq H(\pi_{min}) - h(P_e)$. Since h is monotonically increasing on $[0, ½)$, the physical content of this proposition is that it constrains which values of ε_{max} are physically achievable: no record can achieve $P_e < p_{min}$. For any feasible task class with $\varepsilon_{max} > p_{min}$ and $h(\varepsilon_{max}) < H(\pi_{min})$, every certified record satisfies $I(H ; R) \geq H(\pi_{min}) - h(\varepsilon_{max})$ by Lemma 3.1. Therefore the distinguishability barrier satisfies:

$$D^*_I \geq H(\pi_{min}) - h(\varepsilon_{max}) > 0.$$

In the limit $\varepsilon_{max} \downarrow p_{min}$ (the tightest feasible certification tolerance), this bound approaches:

$$D^*_I \geq H(\pi_{min}) - h(p_{min}) =: c.$$

For task classes where $h(p_{min}) < H(\pi_{min})$—which holds whenever the physical noise floor is small enough relative to the prior uncertainty, and in particular for all equiprobable-prior tasks—we have $c > 0$. The only routes to $c \to 0$ are exactly the unphysical limits enumerated in the proposition statement. Under finite physical resources with a nontrivial task class, the barrier is strictly positive. ■

Corollary 3.4 *(Record Floor as a Resource Inequality).* For fixed T_{env}, ΔF_{max}, and Δt_{min}, there exists a minimum achievable discrimination error $P_e \geq p_{min}(T_{env}, \Delta F_{max}, \Delta t_{min}) > 0$, hence a maximum achievable reliability for any irreducible record primitive without increasing resources. Greater certainty is not free: it is purchased by increasing ΔF (larger apparatus or stronger coupling), lowering T_{env}, or introducing redundancy (which increases total resource expenditure). This establishes a fundamental *uncertainty–resource tradeoff* for record formation: the precision of the first operational fact is bounded by the physical resources available to stabilize it.

Interpretation. Proposition 3.3 is the result that gives the Distinguishability Barrier genuine physical content. Lemma 3.1 and Proposition 3.2 show $D^* > 0$ on information-theoretic grounds; Proposition 3.3 shows *why* it cannot be made small: any attempt to reduce D^* runs into a tradeoff between robustness, thermal noise, and finite apparatus resources that enforces a strictly positive floor. Corollary 3.4 makes the tradeoff explicit: the barrier is not a mathematical artifact but a physical constraint rooted in the thermodynamics of metastability, and overcoming it has a quantifiable resource cost.

3.4 Worked Example: A Symmetric Double-Well Bit

To make Proposition 3.3 concrete, we evaluate its bounds for a minimal physical record: a symmetric double-well potential at finite temperature.

Setup. Consider a particle in a symmetric double-well potential with two basins B_0 and B_1 separated by a free-energy barrier ΔF. The apparatus is in thermal contact with a bath at T_{env} = 300 K (room temperature, $k_B T_{env} \approx 0.026$ eV). We take the microscopic attempt time $\tau_0 = 10^{-9}$ s (nanosecond, typical for mesoscale systems), the re-query interval $\Delta t_{min} = 10^{-6}$ s (microsecond readout), and require robustness over $T = 1$ s with failure probability $\eta = 0.01$. The task class has equiprobable priors ($\pi_{min} = ½$).

Step 1: Minimum barrier from robustness. The Kramers scaling requires:

$$\Delta F_{min} = k_B T_{env} \cdot \ln(T / (\tau_0 \cdot |\ln(1 - \eta)|)) \approx 0.026 \cdot \ln(10^{11}) \approx 0.026 \times 25.3 \approx 0.66 \text{ eV}.$$

Any apparatus with $\Delta F_{max} < 0.66$ eV cannot produce a record that persists for one second at room temperature with 99% reliability.

Case 1: Below the threshold (ΔF_{max} = 0.50 eV). With ΔF below ΔF_{min}, the escape time is $\tau_{escape} \approx 10^{-9} \cdot \exp(0.50/0.026) \approx 0.23$ s—far shorter than the required ~100 s for robustness. The probability that the record flips within $T = 1$ s is approximately 0.99: the record is destroyed almost immediately. No empirical fact can be certified. This is the regime below the Distinguishability Barrier.

Case 2: Above the threshold (ΔF_{max} = 1.0 eV). Now $\tau_{escape} \approx 10^{-9} \cdot \exp(1.0/0.026) \approx 5 \times 10^7$ s (~1.6 years). The minimum flip rate is $\lambda_{min} \approx 2 \times 10^{-8}\ s^{-1}$, giving a flip probability over the microsecond re-query interval:

$$p_{min} = ½(1 - \exp(-2\lambda_{min} \Delta t_{min})) \approx 2 \times 10^{-14}.$$

The mutual information floor is $c = H(\pi_{min}) - h(p_{min}) \approx \ln 2 - 10^{-12} \approx 0.6931$ nats—essentially the full bit. The record is near-perfect.

The barrier is sharp. Between $\Delta F_{max} = 0.50$ eV (no usable record) and $\Delta F_{max} = 0.70$ eV ($p_{min} \approx 10^{-9}$, $c \approx \ln 2$), the record transitions from destroyed to excellent. (For continuity: $\Delta F = 0.70$ eV $\approx 27\, k_B T_{env}$ gives $\tau_{escape} \approx 10^{-9} \cdot \exp(27) \approx 5.3 \times 10^2$ s (≈ 530 s), hence $\lambda_{min} \approx 1.9 \times 10^{-3}$ s^{-1} and $p_{min} \approx \lambda_{min} \cdot \Delta t_{min} \approx 2 \times 10^{-9}$ over a microsecond re-query interval.) This exponential sensitivity to the metastability budget is the Distinguishability Barrier made physical: there is a critical resource threshold below which no empirical fact can form, and above which facts emerge with reliability that improves exponentially with barrier height. The transition is a sharp crossover governed by the activated (Kramers) exponential, not a gradual degradation. (The near-perfect reliability in Case 2 reflects the choice of a large metastability budget, $\Delta F_{max} \approx 38\, k_B T_{env}$, which gives an escape time of ~1.6 years. This is an illustrative regime; smaller budgets closer to threshold produce intermediate reliability, as the 0.70 eV case shows.)

---

4. The Record Primitive Theorem

4.1 Empirical Content Requires Record Variables

Definition 4.1 *(Empirical Completeness: Record-Semantics Closure).* A physical framework $\mathscr{F}$ is *strongly empirically complete* relative to a task class $\mathcal{T}$ with parameters (π_{min}, ε_{max}, T, η, Δt_{min}) if, for each experimental setting θ, it supplies:

(A) A realizable record semantics consisting of:

1. *Record substrate and sample variable.* A physically realizable record substrate Y_θ generated by an admissible measurement interaction M_θ (Definition 4.1a), and an induced record variable $R_\theta \in \mathscr{R}_\theta$, where $\mathscr{R}_\theta$ is finite (or effectively finite).
2. *Decoding / coarse-graining map.* A map g_θ such that the sampled outcome is $R_\theta := g_\theta(Y_\theta(t_{read}))$, where t_{read} is the readout time. The pair (Y_θ, g_θ) is part of the framework's empirical semantics: it specifies what is being sampled in repeated trials.
3. *Robustness.* For an admissible perturbation class $\mathcal{P}$: $\Pr[g_\theta(\Pi(Y_\theta(t))) = g_\theta(Y_\theta(t))$ for all $t \in [0, T]] \geq 1 - \eta$, for all $\Pi \in \mathcal{P}$. (This ensures the record is a reusable macroscopic distinction rather than a transient correlation.)
4. *Re-queryability.* R_θ can be queried at least twice, separated by a nonzero interval $\Delta t_{min} > 0$, while preserving the robustness requirement above. This ensures that "observation" means "something that can be checked," not a one-shot microscopic event.
5. *Certification on a nontrivial admissible task family.* There exists a nontrivial admissible discrimination family $\mathcal{T}_{adm,\theta} \subseteq \mathcal{T}$, realized by varying controllable preparations/inputs within the experimental setting θ, such that for each task in $\mathcal{T}_{adm,\theta}$ (with priors $\pi_i \geq \pi_{min}$) there exists a decision rule $\delta_\theta : \mathscr{R}_\theta \to \{1, 2\}$ achieving $P_e(\delta_\theta, R_\theta) < \varepsilon_{max}$. Here "nontrivial" means $\mathcal{T}_{adm,\theta}$ contains at least one pair of hypotheses whose optimal achievable error is strictly below the constant-guess baseline, and there exists an open neighbourhood in the preparation-parameter manifold over which the certification bound $P_e < \varepsilon_{max}$ continues to hold (i.e., the certifiability is not an isolated point in parameter space but persists over a neighbourhood of experimental configurations).

(B) Empirical closure. Given (1)–(5), $\mathscr{F}$ must induce a well-defined probability model $P_{\mathscr{F}}(R_\theta \mid \theta)$ for repeated trials—that is, R_θ constitutes the experiment's sample space and the framework assigns probabilities to its outcomes. In particular, $P_{\mathscr{F}}(R_\theta \mid \theta)$ must be defined on a sample space whose elements correspond to physically instantiated, re-identifiable record states; otherwise repeated-trial frequencies are undefined.

Remark (closure vs. extra axiom). In most frameworks, (B) is not an additional postulate: once R_θ is defined as a function of the dynamics and the decoding map, the induced distribution $P_{\mathscr{F}}(R_\theta \mid \theta)$ is determined. We state (B) explicitly because frameworks often specify probabilities over microstates while leaving the induced record distribution implicit; it is precisely this gap that constitutes empirical incompleteness.

A framework that does not supply a record semantics (A) for any R_θ is *empirically incomplete* relative to $\mathcal{T}$: it cannot define the operational random variable that experiments sample and therefore cannot close (B).

Definition 4.1a *(Admissible Implementations).* An implementation at setting θ consists of a measurement interaction M_θ producing a record substrate Y_θ and a decoding map g_θ. It is admissible if it respects the task-class constraints: finite stabilization resources ($\Delta F \leq \Delta F_{max} < \infty$), finite temperature $T_{env} > 0$, and perturbations in $\mathcal{P}$. In particular, admissibility excludes implementations that rely on idealized limits ($T_{env} \to 0$, $\Delta F \to \infty$, $\Delta t_{min} \to 0$) unless explicitly permitted by $\mathcal{T}$.

Remark (Dynamics vs. semantics). A framework may be dynamically complete—it specifies evolution on a microstate space X—but empirically incomplete: it does not close the semantics $X \to R_\theta$ required to define sampled outcomes. Without such closure, the framework cannot even state what it means for two experimental runs to have produced "the same outcome," and therefore cannot underwrite reproducibility.

Remark (Why this isn't "just interpretation"). This definition is not a philosophical preference; it is a minimal closure condition required to make sense of repeated-trial statistics. Experimental science samples a random variable. If a framework does not specify the sample variable R_θ, its decoding g_θ, and the robustness/certification conditions under which R_θ is stable and re-queryable, it has not specified what probability distribution is being tested in the laboratory.

Remark (Weak vs. strong empirical completeness). Weak empirical completeness requires only items (1)–(2) and the existence of an induced $P_{\mathscr{F}}(R_\theta \mid \theta)$—no robustness or certification demands. This paper uses *strong empirical completeness* (the full (A) + (B)) because the potential/fact boundary requires reusable, error-bounded records. The distinction turns a potential objection ("robustness is too strong a demand") into a deliberate design choice: weak completeness suffices for formal probability theory; strong completeness is the condition under which the Distinguishability Barrier has content.

Under strong empirical completeness, Lemma 3.1 implies that any certified record variable satisfies $I(H ; R_\theta) \geq \max\{0, H(\pi_{min}) - h(\varepsilon_{max})\}$, making the Distinguishability Barrier an unavoidable feature of any empirically complete framework.

It is essential to distinguish this requirement from the weaker demand that a theory possess "observables" in the formal mathematical sense. A self-adjoint operator on a Hilbert space, or a measurable function on a phase space, is an abstract observable: it specifies what *could* be measured, not how a measurement outcome is physically produced, stabilized, and made available for reuse. A pointer observable is not the same as a certified, re-queryable record with robustness bounds. A framework may be rich in abstract observables or pointer variables and yet empirically incomplete if it provides no account of how any of those observables becomes a stable, re-queryable record that an experimenter can sample across trials.

4.2 The No-Go Theorem

Proposition 4.2 *(Record-Free Frameworks Fail Empirical Completeness).* Any framework that fails to supply a record semantics (Definition 4.1A) is empirically incomplete relative to $\mathcal{T}$, because it cannot define the sampled variable R_θ and therefore cannot close the empirical distribution $P_{\mathscr{F}}(R_\theta \mid \theta)$ (Definition 4.1B).

Proof. Immediate from Definition 4.1. Without a realizable record substrate, decoding map, and the robustness/re-queryability/certification conditions that make R_θ a stable, reusable sample variable, the framework does not define an empirical random variable for repeated trials. Without R_θ, the closure clause (B) cannot be satisfied: there is no operational sample space on which repeated-trial frequencies are defined. Restoring empirical content requires appending an explicit rule selecting a record substrate and specifying how robust records form—i.e., a record primitive. ■

4.3 Concrete Instances of the Failure Mode

Proposition 4.2 is structural, not interpretive. It does not assert that existing frameworks are "wrong," but that, taken at face value without supplementary postulates, they fail to supply a record semantics in the sense of Definition 4.1(A). We now indicate precisely which components are missing in several well-known approaches.

Table 1: Where standard approaches leave the record semantics implicit.

Framework (as usually stated)	Missing / underdetermined components of Definition 4.1(A)
Unitary QM + decoherence (no commitment postulate)	(1) record substrate not fixed; (2) no explicit decoding map selecting outcomes; (4) re-queryability not guaranteed; (5) certification imported via Born/readout
Decoherent histories (no realization criterion)	(2) decoding/selection of one realized history underdetermined; (4) re-queryability not enforced; (5) certification requires extra realization rule
Everett (no operational record selection rule)	(1) substrate branch-relative; (2) decoding map implicit; (5) certification/probability requires added postulate (Born-weight/typicality)

Framework (as usually stated)	Missing / underdetermined components of Definition 4.1(A)
QBism	(1) substrate relocated to agent experience; (3–4) robustness/re-queryability assumed rather than physically derived; record primitive is axiomatic

(i) Unitary quantum mechanics with decoherence but no commitment rule. Consider a framework consisting of a Hilbert space, a Hamiltonian generating unitary evolution, and environmental decoherence selecting a pointer basis. Decoherence suppresses off-diagonal elements in the reduced density matrix [1,2], yielding approximate diagonalization in a preferred basis. However:

- *Item (1) (Record substrate)* is not formally specified. The theory identifies pointer observables but does not specify which physical degrees of freedom constitute the sampled record variable R_θ.
- *Item (2) (Decoding map)* is absent. No explicit map $g_\theta : Y_\theta \to \mathscr{R}_\theta$ is defined that selects one realized outcome.
- *Item (4) (Re-queryability)* is not guaranteed. Decoherence produces suppression of interference, but it does not by itself produce a stabilized, re-identifiable macroscopic record.
- *Item (5) (Certification)* is not satisfied internally. The Born rule is applied at readout, but this step functions as an implicit commitment rule rather than something derived from the decohering dynamics.

Thus decoherence alone supplies partial structure toward robustness (Item 3, in an approximate sense) but does not close the semantics required by Definition 4.1(A). A commitment rule is imported tacitly at measurement.

(ii) Decoherent histories without a realization criterion. The consistent/decoherent histories programme [3,4,5] assigns probabilities to families of coarse-grained histories satisfying a consistency condition. However:

- *Item (2) (Decoding map)* is underdetermined. Multiple consistent families may exist; the framework does not specify which one corresponds to the realized record.
- *Item (4) (Re-queryability)* is not structurally enforced. A consistent history is a non-interfering alternative, not necessarily a physically instantiated macroscopic record.
- *Item (5) (Certification)* depends on selecting a particular history as fact, but the framework does not provide a mechanism that performs this selection.

The consistency condition constrains allowable histories but does not supply a commitment mechanism that produces a stabilized record variable R_θ. Without an explicit realization rule, Definition 4.1(A) is not satisfied.

(iii) Everettian (many-worlds) quantum mechanics without record certification. The Everett interpretation posits universal unitary evolution with branching structure. In this setting:

- *Item (1) (Record substrate)* is branch-relative. The framework identifies branch structures but does not specify, within the theory itself, which branch's record variable is the operational sample variable for a given observer.
- *Item (2) (Decoding map)* is implicit rather than explicit. The map from universal state to observer-level record is not formally specified in record-semantics terms.
- *Item (5) (Certification)* is branch-conditional. While each branch contains definite records, the theory does not, without supplementary postulates, provide a rule internal to the formalism that identifies the operational record variable R_θ whose frequencies are compared across trials.

Everettian mechanics provides a dynamical account of branching but does not, by itself, close the empirical semantics required by Definition 4.1 unless supplemented with a rule identifying operational records.

(iv) QBism. QBism treats measurement outcomes as primitive personal experiences of agents [6]. In this case:

- *Item (1) (Record substrate)* is relocated to the agent rather than specified as a physical substrate within the theory's ontology.
- *Item (3) (Robustness)* and *Item (4) (Re-queryability)* are assumed at the experiential level but not derived from the underlying physical dynamics.

QBism is empirically coherent, but the record primitive is treated as an axiomatic starting point rather than structurally derived. It therefore satisfies empirical completeness only by taking the record primitive as foundational—a choice that is, in a precise sense, the closest among existing interpretations to the approach taken by VERSF.

Structural diagnosis. In each case, the deficiency is not that the framework lacks predictive power, but that it does not internally supply a record semantics satisfying Definition 4.1(A) without importing supplementary structure. The conclusion is not that these approaches fail empirically, but that record formation is always present—either explicitly modeled or implicitly assumed. Proposition 4.2 shows that this is unavoidable: empirical completeness requires closure at the level of record variables. VERSF differs not by introducing records—every empirically viable theory must do so—but by placing the commitment mechanism at the foundational level rather than appending it at the theory–experiment interface.

4.4 The Record Primitive as a Structural Necessity

Corollary 4.3 *(Necessity of Record Formation).* Any empirically meaningful physical theory must contain, either explicitly within its dynamics or implicitly within its interpretative apparatus, a primitive specifying how stable records form.

Proof. Let a theory be specified by the quadruple $(\mathcal{X}, \mathcal{T}, \mathbb{P}, \mathscr{F})$, where $\mathcal{X}$ is a state space, $\mathcal{T}$ is an evolution rule (deterministic or stochastic), $\mathbb{P}$ is a probability assignment over outcomes, and $\mathscr{F}$ is an interpretation map connecting formal objects to empirical events.

Suppose $\mathscr{F}$ does not specify how robust records form—that is, it provides no account of how the record variable R is produced from the underlying dynamics. Then $\mathscr{F}$ is incomplete as an empirical map: it assigns probabilities to abstract "outcomes" without specifying a physically instantiable record variable. By Proposition 4.2, the framework is empirically incomplete. To restore empirical content, the theory must import an additional rule external to $(\mathcal{X}, \mathcal{T}, \mathbb{P})$: a measurement postulate, a classical cut, a coarse-graining prescription, or a decoherence-to-record selection rule. Each of these imported rules is precisely a record primitive—a specification of how the transition from potential to fact occurs.

Therefore, any empirically meaningful theory must include record formation either explicitly (within its dynamics or ontology) or implicitly (within $\mathscr{F}$ via supplementary postulates). ■

Corollary 4.4 *(Implicit Commitment in Standard Frameworks).* Any theoretical framework that does not model record formation explicitly must encode it implicitly in auxiliary interpretative structures. Standard quantum mechanics encodes it through the Born rule and the measurement postulate; decoherence programs encode it through pointer-basis selection combined with a tacit readout rule; Bohmian mechanics encodes it through the guidance equation's selection of a definite trajectory plus the quantum equilibrium hypothesis; QBism encodes it by taking the agent's experience of definite outcomes as axiomatic.

---

5. Binary Irreducibility and the Minimal Certified Distinction

Within the VERSF framework, the minimal commitment corresponds to a binary irreducible partition. Below $m = 2$, no distinction exists. At $m = 2$, the first certified distinction becomes possible. This binary threshold constitutes the first operational fact—the crossing point of the Distinguishability Barrier.

5.1 Minimal Partition Threshold

Let the record algebra be the σ-algebra generated by the partition of Ω_E induced by g. If g induces m basins of attraction $\{B_1, \ldots, B_m\}$ with $B_i \cap B_j = \emptyset$ for $i \neq j$ and $\bigcup_i B_i \subseteq \Omega_E$, then the record variable is $R \in \{1, \ldots, m\}$. The case $m = 1$ corresponds to no distinction: all microstates map to the same macroscopic label, and no information about the system is encoded.

Define the committed partition number m_c as the minimal m for which the basins are robust under admissible perturbations $\mathcal{P}$ and the discrimination error can be bounded by ε. The first certified distinction exists if and only if $m_c = 2$. This provides a mathematically precise formulation of the claim: below the barrier there is only potential; at the barrier, the first operational fact appears.

Thus D* corresponds to the minimal binary commitment event. All higher structure—particles, mass density, gravitational constraint overlap—emerges from repeated and spatially distributed commitment processes built upon this elementary binary foundation.

5.2 Binary Irreducibility Theorem

We now state the binary irreducibility claim. The goal is to isolate a minimal and defensible result: under mild axioms describing what it means for a committed record to exist, the minimal committed record is necessarily binary, and any higher-arity record implies the existence of at least one certified binary record.

Axioms (Record Existence Class):

- **A1** *(Finite Distinguishability).* In any bounded region, only finitely many mutually certifiable record classes can be realized with bounded error.
- **A2** *(Robustness).* A committed record corresponds to a basin of attraction under admissible perturbations $\mathcal{P}$, satisfying the robustness criterion of Definition 2.9.
- **A3** *(Nontriviality).* At least one certified distinction exists: there is at least one discrimination task that is certifiably discriminable with error less than ε.
- **A4** *(Irreducibility).* A committed record is irreducible if it cannot be represented as a deterministic function of a strictly smaller committed record variable without increasing the discrimination error above ε.
- **A5** *(Compositional Consistency).* Independent committed records compose via product structure: if R and R′ are independent committed records, then the joint record is (R, R′).

Remark on axiom independence. A1 constrains the record alphabet to be finite. A2 ensures physical stability. A3 excludes degenerate theories with no observational content. A4 defines the notion of a minimal record. A5 governs how records combine. The axioms are logically independent: A3 can fail (trivial theory), A1 can fail (continuous record labels), etc.

Theorem 5.1 *(Binary Irreducibility).* Within the Record Existence Class (A1–A5):

(a) No irreducible committed record has $m = 1$.

(b) Any committed record with $m \geq 3$ implies the existence of at least one certified binary record ($m = 2$).

(c) Therefore, the minimal nontrivial irreducible committed record has $m = 2$.

Proof. The argument proceeds in three steps.

Step 1 *(Excluding $m = 1$).* If $m = 1$, the record map g assigns all admissible apparatus microstates to a single macroscopic basin. For any decision rule δ, the outputs are identical under $H = 1$ and $H = 2$, yielding $P\_e \geq \min(\pi_1, \pi_2)$. This violates Axiom A3 (Nontriviality) for $\varepsilon < \min(\pi_1, \pi_2)$. Hence any nontrivial committed record has $m \geq 2$.

Step 2 *(Any $m \geq 3$ record implies a certified binary record).* Let $R \in \{1, \ldots, m\}$ with $m \geq 3$ be a committed record achieving $P_e < \varepsilon$ for some discrimination task via a decision rule $\delta : \{1, \ldots, m\} \to \{1, 2\}$. Define $S := \delta^{-1}(1) \subset \{1, \ldots, m\}$ and the binary query $Q_S := \mathbb{1}[R \in S]$. Then Q_S is a deterministic function of R that takes values in $\{0, 1\}$, and the decision rule δ can be written as $\delta(R) = 1 + Q_S(R)$ (mapping $Q_S = 1$ to hypothesis 1 and $Q_S = 0$ to hypothesis 2, or vice versa). Since δ achieves $P_e < \varepsilon$ on the original record R, the identical discrimination performance is achieved by applying the equivalent decision rule to Q_S. Hence Q_S is a certified binary record achieving $P_e < \varepsilon$.

It remains to verify that Q_S is robust. Since $Q_S = \mathbb{1}[R \in S]$ is a deterministic function of R, and R is robust under admissible perturbations (A2), Q_S inherits robustness: if the record label R remains invariant under perturbation, then so does any deterministic function of R. Formally, if $\Pr[g(\Pi(Y'_t)) = g(Y'_t)] \geq 1 - \eta$, then $\Pr[\mathbb{1}[g(\Pi(Y'_t)) \in S] = \mathbb{1}[g(Y'_t) \in S]] \geq 1 - \eta$. Hence Q_S is a robust, certified binary record.

Step 3 *(The minimal irreducible record is binary).* By Step 1, $m = 1$ is excluded. By Step 2, any committed record with $m \geq 3$ implies the existence of a certified binary record Q_S that is a coarse-graining of R. By Axiom A4, since R admits a deterministic reduction to the smaller record Q_S (with $|\{0,1\}| < m$) while preserving discrimination at error $P_e < \varepsilon$, the m-ary record R is not irreducible. Therefore no record with $m \geq 3$ is irreducible, and the minimal nontrivial irreducible committed record has $m = 2$. ■

Scope clarification: what is and is not being claimed. A natural objection to Theorem 5.1 is that the proof merely shows decision problems can be reduced to binary questions—a logical fact—rather than establishing anything about the physical world's primitives. We address this directly. The theorem does *not* claim that macroscopic reality is ontologically binary, nor that all physical records consist of a single bit. What it establishes is a claim about the structure of *certified distinction primitives*: any physical process that produces a certifiable m-ary record necessarily also produces at least one certifiable binary record, and consequently the first certified distinction that any physical system can cross is a binary one. The argument is not purely logical because it depends on the physical axioms A2 (robustness) and A3 (nontriviality): the binary query Q_S inherits physical robustness from the underlying basin structure, and the nontriviality of the discrimination task ensures that Q_S carries genuine information about the system. The conclusion is that the binary commitment is irreducible in the sense of certified physical distinction—not merely in the sense of formal decision theory. More precisely, irreducibility here is defined relative to the certification and robustness constraints (A2, A3): a record is irreducible if no proper coarse-graining of it satisfies these constraints while still carrying nontrivial information about the system.

Remark. Larger record alphabets arise by composing binary commitments. This is structurally analogous to the role of the bit in classical information theory: all finite messages decompose into binary digits, but individual messages may be arbitrarily complex.

5.2.1 Decomposition of m-ary Records into Binary Commitments

Theorem 5.1 establishes that any m-ary record *implies* the existence of at least one certified binary record. The following lemma establishes a stronger result—that m-ary records can be *implemented* as sequential binary commitments—under an explicit additional assumption about basin topology.

Definition 5.2 *(Separable Basin Topology).* A robust m-ary record with basins $\{B_1, \dots, B_m\}$ has separable basin topology if there exists a hierarchical sequence of binary partitions of the basin index set $\{1, \dots, m\}$ such that at each level of the hierarchy, the two sides of the partition correspond to unions of basins separated by a free-energy barrier ΔF_k satisfying $\Delta F_k \geq k_B T_{env} \ln(T / \tau_0)$ for all k. Informally: the basins can be recursively split into two robust groups at every stage.

Physical motivation. Separable basin topology holds generically for multi-well potentials in finite-dimensional configuration spaces where the wells are separated by saddle points of finite height. It can fail in exotic topologies (e.g., basins arranged on a ring with no natural binary partition that preserves robustness at all levels), but such configurations are non-generic in standard condensed-matter and measurement-apparatus settings.

Lemma 5.3 *(Binary Decomposition under Separable Topology).* Let $R \in \{1, \dots, m\}$ be a robust m-ary record with separable basin topology. Then R can be implemented as a sequence of $k = \lceil \log_2 m \rceil$ binary commitments $R_1, \dots, R_k$, where:

(a) Each R_i is a robust binary record corresponding to one level of the hierarchical partition.

(b) The joint record $(R_1, \dots, R_k)$ determines R via a surjective map.

(c) The total discrimination error satisfies $P_{e,total} \leq \Sigma_i P_{e,i}$, where $P_{e,i}$ is the error of the i-th binary commitment (union bound).

(d) The total entropy cost satisfies $\langle \Delta S_{env,total} \rangle \geq k \cdot (k_B \ln 2 - k_B h(\varepsilon_{max}))$, where $\varepsilon_{max} := \max_i \varepsilon_i$.

Proof sketch. The separable basin topology provides, by definition, a binary tree of depth $\lceil \log_2 m \rceil$ over the basin indices. At each node, the partition divides the remaining basins into two groups, each group being a union of robust basins separated from the other group by a free-energy barrier meeting the metastability criterion. Each such partition defines a binary commitment R_i that inherits robustness from the underlying basin structure (since each side is a union of basins and the inter-group barrier exceeds the thermal threshold). The surjectivity in (b) follows from the tree structure: the sequence of binary choices uniquely identifies a leaf (a specific basin), possibly with some leaves unpopulated when m is not a power of 2. The error bound (c) follows from Boole's inequality (the union bound): $P_{e,total} \leq \Sigma_i P_{e,i}$. This bound does not require the binary commitments to have independent errors—it holds regardless of the correlation structure among sequential commitments within a single apparatus. The bound is conservative; tighter bounds are achievable if the binary commitments are optimally coordinated or if the error correlations are favorable. The entropy bound (d) follows from applying Proposition 5.2 independently to each binary commitment. ■

Remark. Lemma 5.3 is not required for the foundational claims of this paper—Theorem 5.1 suffices to establish that the binary commitment is the irreducible primitive. The lemma shows that under physically common conditions, binary commitments are not only the irreducible building blocks but also the constructive building blocks of all higher-arity records. Whether separable basin topology is a necessary condition for robust m-ary records, or merely a sufficient one, remains an open question.

5.3 Entropic Lower Bound for Commitment (Landauer-Type Constraint)

A record is not merely a correlation; it is a stabilized distinction maintained against noise. Stabilization requires dissipation. This section derives a lower bound connecting record formation to irreversibility, adapted to the certification and robustness language developed above.

The connection to Landauer's principle [7] requires care. Bennett [8] showed that *computation* per se need not dissipate energy; only logically irreversible operations—those that erase information—carry a thermodynamic cost. A commitment event, as defined here, is logically irreversible in precisely Landauer's sense: the many-to-one mapping from pre-commitment microstates to post-commitment macroscopic basins erases the fine-grained microstate information that distinguished pre-commitment alternatives. It is this erasure that carries the entropic cost, not the act of discrimination itself.

Crucially, the claim is *not* that any measurement costs k_B T_env ln 2. A reversible measurement that merely correlates system and apparatus without stabilizing a record incurs no thermodynamic cost. The cost arises specifically from *commitment*: the requirement that the record be robust (maintained against thermal noise over timescale T) and re-queryable (readable without degradation). These requirements force the record into a metastable basin from which escape is thermally suppressed, and the resulting coarse-graining is logically irreversible. It is this combination—logical irreversibility plus physical robustness—that triggers the Landauer bound.

Assumptions for Proposition 5.2. The following conditions are assumed throughout:

- **(L1)** *Finite-temperature bath.* The apparatus is in thermal contact with an environment at temperature T_env > 0.
- **(L2)** *Finite-time stabilization.* The commitment process completes in finite time and produces a record that is stable over a specified timescale T > 0.
- **(L3)** *Nonzero robustness.* The record satisfies the robustness criterion of Definition 2.9 with $\eta \ll 1$: perturbations from the admissible class $\mathcal{P}$ do not flip the record label with high probability over [0, T].
- **(L4)** *Physical realizability of basins.* The two basins B_0 and B_1 correspond to disjoint regions of the apparatus phase space, separated by a free-energy barrier sufficient to enforce (L3).
- **(L5)** *Logically irreversible coarse-graining.* The record map g : Ω_E → {0, 1} is many-to-one: multiple distinct pre-commitment microstates map to the same record label. The fine-grained microstate cannot be reconstructed from the record label alone.

Proposition 5.2 *(Minimal Entropy Production per Binary Commitment).* Under assumptions (L1)–(L5), consider a binary commitment that maps pre-commitment microstates into two robust basins B_0 and B_1 with re-query error bounded by ε. The average entropy exported to the environment satisfies:

$\langle \Delta S\_env \rangle \geq k\_B \ln 2 - k\_B \cdot h(\varepsilon)$,

where $h(\varepsilon) := -\varepsilon \ln \varepsilon - (1 - \varepsilon) \ln(1 - \varepsilon)$ is the binary entropy function. In the high-reliability limit $\varepsilon \to 0$, this approaches k_B ln 2. Equivalently, the minimal average heat dissipated satisfies:

$\langle Q \rangle \geq k\_B\ T\_env \ln 2$

up to ε-dependent corrections.

Proof. The argument follows the information-theoretic approach to Landauer's principle developed by Sagawa and Ueda [9], adapted to the commitment setting.

Let X denote the pre-commitment microstate and $R \in \{0, 1\}$ the post-commitment record label, with $R = g(X')$ after the commitment dynamics. By assumption (L5), g is many-to-one: the record label does not determine the microstate, so $H(X \mid R) > 0$. The commitment process reduces the observer's uncertainty about the coarse-grained record degree of freedom.

Let H_prior denote the prior entropy of the record outcome—the observer's uncertainty about which basin the system will occupy before commitment occurs. (This is the entropy of the binary random variable that the commitment will resolve, not to be confused with the post-commitment record R.) For a binary commitment, $H\_prior \leq \ln 2$ (equality when the two outcomes are equiprobable; the bound holds without assuming symmetry). The maximal entropy reduction—and hence the maximal thermodynamic cost—occurs when the two commitment alternatives are equiprobable, in which case $H\_prior = \ln 2$. For biased alternatives (one outcome much more likely than the other), $H\_prior < \ln 2$ and the thermodynamic bound is correspondingly smaller: the cost of commitment reflects the actual information gained, not a fixed universal constant. After a reliable commitment with re-query error ε, the conditional entropy of the record given a readout satisfies $H(R \mid readout) \leq h(\varepsilon)$, where $h(\cdot)$ is the binary entropy function.

The mutual information gained about the record is therefore:

$I\_gained = H\_prior - H(R \mid readout) \geq H\_prior - h(\varepsilon)$.

By the generalized Landauer principle [9], any process that irreversibly reduces the entropy of a subsystem in contact with a thermal bath at temperature T_env must export at least as much entropy to the environment:

$\langle \Delta S\_env \rangle \geq k\_B \cdot I\_gained \geq k\_B\ (H\_prior - h(\varepsilon))$.

For a binary commitment with equiprobable outcomes, $H\_prior = \ln 2$, giving:

$\langle \Delta S\_env \rangle \geq k\_B \ln 2 - k\_B\, h(\varepsilon)$.

For non-equiprobable outcomes, H_prior < ln 2, and the bound is correspondingly tighter; the Landauer floor k_B ln 2 represents the worst-case (maximum-uncertainty) scenario. In the high-reliability limit $\varepsilon \to 0$, $h(\varepsilon) \to 0$, recovering the standard Landauer bound $\langle \Delta S\_env \rangle \geq k\_B \ln 2$.

Remark on feedback protocols. In the presence of feedback (Maxwell-demon-style protocols), the generalized Landauer bound takes the form $\langle Q \rangle \geq k\_B\, T\_env\, (\ln 2 - I\_feedback)$, where I_feedback is the mutual information acquired by the feedback controller about the system prior to the erasure step [9]. Our ε-dependent expression $\langle \Delta S\_env \rangle \geq k\_B(\ln 2 - h(\varepsilon))$ is already within this family: the term h(ε) plays the role of residual uncertainty after an imperfect measurement, analogous to (ln 2 − I_feedback) when the feedback channel is noisy. The key point is that even with optimal feedback, the bound cannot be driven below zero for a commitment that actually reduces uncertainty (H_prior > h(ε)), and the minimum dissipation for a reliable binary commitment ($\varepsilon \to 0$) remains k_B T_env ln 2 in the equiprobable high-reliability case (H_prior = ln 2), and scales with H_prior otherwise.

Note on units. This proof uses natural logarithms throughout (entropy in nats, consistent with the Fano-inequality calculations in Section 3). The conventional Landauer bound k_B T_env ln 2 in energy units corresponds to $\langle Q \rangle = T\_env \cdot \langle \Delta S\_env \rangle \geq k\_B\, T\_env \ln 2$, which is the same bound expressed as heat rather than entropy. ■

VERSF Interpretation. Crossing the Distinguishability Barrier carries a minimal irreversibility cost of k_B ln 2 per binary commitment. This connects the first operational fact to entropy flow, consistent with VERSF's broader thesis that all physical structure emerges from the thermodynamics of commitment. The entropic cost is not incidental but constitutive: it is what distinguishes a genuine commitment from a reversible correlation.

---

6. Toward a Uniqueness Result

The results of Sections 3–5 establish that any empirically complete theory must contain a record primitive, that the irreducible record is binary, and that commitment carries a minimum entropic cost. A natural question is whether the binary commitment mechanism is *uniquely* determined by the axioms of empirical completeness, or whether inequivalent record primitives could satisfy the same constraints.

We conjecture that, within the axiomatic class A1–A5 supplemented by a spatial locality requirement (each commitment event involves a bounded region of the substrate), the irreducible record algebra is isomorphic to $\mathbb{Z}_2$ and the commitment structure is unique up to isomorphism. A proof would require showing that the axioms fix the algebraic structure, that locality constrains the composition rule, and that the entropic cost structure is determined up to constants by thermodynamic consistency. These steps are the subject of ongoing work and will be reported separately. We note the conjecture here because, if established, it would elevate VERSF's architectural choice—beginning with binary commitment—from a well-motivated starting point to a necessary one.

---

7. Implications for VERSF

The Distinguishability Barrier provides the foundational layer upon which the broader VERSF programme is constructed. Where conventional frameworks assume the existence of facts and build upward from that assumption, VERSF begins at the first barrier-crossing event and derives physical structure from the dynamics of commitment.

7.1 From Records to Physical Structure

The results of Sections 3–5 are framework-independent: they apply to any empirically meaningful theory. What distinguishes VERSF is not that it includes a record primitive—every theory must, by Corollary 4.3—but that it *begins* with the record primitive and derives the rest of physics from it. This is an architectural choice with substantive consequences: by placing record formation at the foundation rather than importing it as an auxiliary postulate, VERSF eliminates the explanatory gap between dynamics and observation that other frameworks must bridge with supplementary rules.

Mass as commitment density. VERSF proposes that mass corresponds to the local density of irreversible binary commitments per unit of emergent spatial structure. Each commitment carries a minimal entropy cost (Proposition 5.2) and constrains the future evolution of the committed region. The accumulation of commitments would produce inertial resistance to state change—the phenomenological hallmark of mass. If this identification is correct, the entropy cost per commitment ($k_B T_{env} \ln 2$) and the rate of commitment events (set by the local entropy gradient) together determine the effective mass density. The quantitative development of this proposal is presented in a companion work [10].

Gravity as constraint overlap. Where commitment density is high, the accessible state space is more tightly constrained. Adjacent regions with high commitment density share overlapping constraints on their admissible configurations. VERSF conjectures that this mutual constraint produces an effective attraction—a restriction on relative motion that corresponds, in the continuum limit, to the curvature of spacetime that general relativity describes geometrically. A candidate derivation proceeds via a variational principle on the total commitment density; whether this yields field equations that reduce to Einstein's equations in the appropriate limit is the subject of ongoing work [11].

The quantitative development of the mass-density relation and the gravitational field equations lies outside the scope of the present paper, which is concerned solely with establishing the foundational layer—the barrier, the record primitive, and the binary irreducibility theorem. The physical consequences sketched above are developed in companion works [10,11], which should be consulted for the detailed arguments.

Time as commitment sequence. Within VERSF's interpretive framework (not as a consequence of the framework-independent results of Sections 3–5), the Distinguishability Barrier is also the barrier between the timeless regime of potentials and the temporal regime of facts. Time, in VERSF, does not flow continuously from a background parameter; it is generated by the sequential accumulation of commitment events. The ordering of commitments—which is partial, not total, reflecting the causal structure of spacetime—defines the emergent temporal ordering. This is consistent with the entropic arrow: each commitment increases the total entropy of the environment (Proposition 5.2), ensuring that the sequence of commitments is thermodynamically irreversible.

7.2 Positioning Relative to Other Programmes

The foundational literature contains several approaches that share features with VERSF's starting point but differ in architecture:

- **Quantum Darwinism** [1,2] explains *which* records survive (those that are redundantly encoded in the environment) but does not take record formation as a foundational primitive from which spatial and dynamical structure are derived. VERSF absorbs this insight—commitment events preferentially stabilize pointer-basis records—while going further by deriving gravitational and spatial structure from the commitment dynamics.
- **It from Bit** (Wheeler [12]) anticipated the idea that physical facts are ultimately informational. VERSF makes this precise by identifying the minimal informational event (binary commitment) and deriving its physical consequences.
- **Constructor Theory** (Deutsch and Marletto [13]) also emphasizes the distinction between possible and impossible tasks. The Distinguishability Barrier theorem can be read as establishing a specific impossibility: no task can produce a stable record below D*. The two programmes are complementary.

7.3 Summary

The theory roots physical law in the minimal transition from potential to operational fact. The results of this paper—the existence and nontriviality of the barrier (Lemma 3.1, Proposition 3.2), the finite-resource no-go (Proposition 3.3, Corollary 3.4), the necessity of the record primitive (Proposition 4.2, Corollary 4.3), the binary irreducibility of the minimal certified distinction (Theorem 5.1), and the entropic cost of commitment (Proposition 5.2)—constitute the axiomatic foundation from which VERSF's specific physical predictions are subsequently derived.

Appendix A: Record Formation as an Unavoidable Completion

(Equivalence Class Result)

This appendix establishes a structural result: any framework that claims to predict observable outcomes can be expressed (up to empirical equivalence) as a microdynamical model together with an explicit record map and robustness conditions. Record primitives are therefore not optional additions but unavoidable components of empirical completeness.

A.1 Setup: The Minimal Structure of an Empirical Theory

Let a micro-theory be specified by a measurable state space $(\Omega, \mathbb{F})$, an evolution family $\{T_t\}$ (deterministic flow or stochastic kernel), and an initial-state family $\{P_\theta\}$ indexed by experimental settings θ. This specification alone yields a time-dependent distribution over microstates $X_t \in \Omega$.

To connect this mathematical apparatus to experiments, we require a family of record variables R_t that can be sampled and compared across trials. A record variable must be (i) physically realizable, (ii) robust over a nontrivial timescale, and (iii) associated with a finite (or effectively finite) outcome alphabet.

A.2 The Record Completion Theorem

Theorem A.1 *(Record Completion / Empirical Equivalence).* Suppose a framework F produces empirically testable predictions; that is, for each setting θ it induces a well-defined probability distribution over observed outcomes O in repeated trials. Then there exists an empirically equivalent representation of F of the form:

$$(\Omega, \{T_t\}, \{P_\theta\}) + (g, \mathscr{R}, \mathcal{P}, T, \varepsilon, \eta),$$

where $(\Omega, \{T_t\}, \{P_\theta\})$ is a microdynamical substrate, g is a coarse-graining map from microstates to record labels, $\mathscr{R}$ is the record alphabet, $\mathcal{P}$ is an admissible perturbation class, and (T, ε, η) specify robustness and certification tolerances, such that the induced distribution over record labels $R := g(X_t)$ matches the predicted distribution over observed outcomes O up to the chosen tolerances.

Proof (constructive). The proof proceeds in four steps.

Step 1 *(Identify the observable random variable).* By assumption, F induces a distribution $P_F(O \mid \theta)$ over observed outcomes O. Treat O as a random variable in an outcome space $\mathcal{O}$. Since O is observed and comparable across trials, it is, by definition, a record variable.

Step 2 *(Represent O as a coarse-graining of microstates).* Any physical implementation of F must realize O via some physical degrees of freedom (apparatus, environment, memory). Let X_t denote the microstate of the full physical implementation at the readout time $t = t_{read}$. Define $g : \Omega \to \mathcal{O}$ by $g(X_t) = O$. This defines an explicit record map.

Step 3 *(Adjoin robustness parameters).* Because O is observable, it must be stable under ordinary perturbations of the apparatus and environment over some timescale. Introduce an admissible perturbation class $\mathcal{P}$ and parameters (T, η) capturing the probability that g remains invariant under perturbations over [0, T]. Introduce ε to bound discrimination or readout error. These parameters are not extraneous additions but a formalization of what it means for O to be an observation.

Step 4 *(Empirical equivalence).* The induced distribution of $R := g(X_{t_{read}})$ under the microdynamics matches $P_F(O \mid \theta)$ by construction, since g was defined to agree with the framework's operational readout. Thus F is empirically equivalent to a representation that includes an explicit record primitive. ■

A.3 Consequences

Unavoidability. Any empirically meaningful theory must include record formation somewhere—either inside its ontology and dynamics (explicitly) or in an interpretation map (implicitly).

Location of novelty. What distinguishes foundational programmes is not *whether* they include records, but *whether they treat record formation as primitive, derived, or merely assumed.*

VERSF's stance. VERSF places the record primitive at the deepest layer—the first barrier-crossing from potential to operational fact—and derives higher-level structures from distributed commitment.

A.4 Roadmap to a Sharper Theorem

A referee-strength version of Theorem A.1 can be sharpened along three axes:

(i) Specifying the admissible class of physical implementations.

(ii) Formalizing robustness as a stability property of basins of attraction in Ω.

(iii) Proving that any two empirically equivalent representations must share an isomorphic record algebra at the level of irreducible committed partitions.

This tightening would connect the equivalence-class result directly to the Binary Irreducibility Theorem (Theorem 5.1) and to the Distinguishability Barrier definition D*, closing the logical loop between the structural necessity of records and the minimal form that those records must take.

References

[1] W. H. Zurek, "Quantum Darwinism," *Nature Physics* **5**, 181–188 (2009).

[2] W. H. Zurek, "Decoherence, einselection, and the quantum origins of the classical," *Reviews of Modern Physics* **75**, 715–775 (2003).

[3] R. B. Griffiths, "Consistent histories and the interpretation of quantum mechanics," *Journal of Statistical Physics* **36**, 219–272 (1984).

[4] R. Omnès, "Logical reformulation of quantum mechanics," *Journal of Statistical Physics* **53**, 893–932 (1988).

[5] M. Gell-Mann and J. B. Hartle, "Classical equations for quantum systems," *Physical Review D* **47**, 3345–3382 (1993).

[6] C. A. Fuchs, N. D. Mermin, and R. Schack, "An introduction to QBism with an application to the locality of quantum mechanics," *American Journal of Physics* **82**, 749–754 (2014).

[7] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM Journal of Research and Development* **5**, 183–191 (1961).

[8] C. H. Bennett, "The thermodynamics of computation—a review," *International Journal of Theoretical Physics* **21**, 905–940 (1982).

[9] T. Sagawa and M. Ueda, "Minimal energy cost for thermodynamic information processing: Measurement and information erasure," *Physical Review Letters* **102**, 250602 (2009).

[10] K. Taylor, "Entropy-momentum foundations of quantum mechanics in the VERSF framework," in preparation (2025). Preprint available at www.versf-eos.com.

[11] K. Taylor, "Gravitational emergence from commitment constraint overlap," in preparation (2025). Preprint available at www.versf-eos.com.

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

[13] D. Deutsch and C. Marletto, "Constructor theory of information," *Proceedings of the Royal Society A* **471**, 20140540 (2015).

[14] H. A. Kramers, "Brownian motion in a field of force and the diffusion model of chemical reactions," *Physica* **7**, 284–304 (1940).

[15] G. Chiribella, G. M. D'Ariano, and P. Perinotti, "Informational derivation of quantum theory," *Physical Review A* **84**, 012311 (2011).

[16] A. Peres, "Unperformed experiments have no results," *American Journal of Physics* **46**, 745–747 (1978).