

The Physical Admissibility Framework (PAF): Finite Distinguishability, Irreversible Commitment, and the Cost of Facts

Keith Taylor

VERSF Theoretical Physics Program

Abstract

We present the Physical Admissibility Framework (PAF), a constraint-based extension of quantum mechanics that formalizes the finite cost of producing physical facts. PAF supplements the Hilbert-space, Hamiltonian, and Born-rule formalisms with an admissibility layer governing finite distinguishability, irreversible commitment, and time cost. Grounded in the Bit Conservation Bound (BCB) and the Ticks-Per-Bit (TPB) constraint, PAF yields no-go theorems excluding reversible fact creation, infinite-precision hidden-variable models, and cost-free measurement. We introduce the Distinguishability Ledger as a conservation structure for irreversible processes, derive lower bounds on measurement, readout, reset, and error correction, and show how coarse-grained ledger flows recover effective spacetime geometry. The framework makes testable scaling predictions: measurement costs scale as $\log N$ for N -hypothesis discrimination, error correction overhead scales linearly with syndrome count, and detector-bias maintenance costs scale linearly with log-asymmetry. PAF reframes quantum mechanics as the minimal admissible dynamics for a universe with finite information capacity and irreversible facts.

Plain-language summary: Standard quantum mechanics tells us what states exist and how likely various outcomes are, but it doesn't say what it *costs* to make an outcome happen. In real experiments, measurement and readout dominate time and energy budgets. PAF adds the missing accounting: distinguishability (the ability to tell things apart) is finite and conserved, and turning quantum possibilities into definite facts requires irreversible commitment that cannot be done for free. This framework explains why quantum mechanics takes the form it does, why measurement is expensive, why error correction has overhead, and potentially why spacetime has geometry.

Contents

Abstract	1
1. Introduction.....	5
2. Physical Admissibility Postulates	5
3. Primitive Resources and Admissible Operations	6
4. The Distinguishability Ledger	7
5. Composition Rules and Commitment Depth	9
6. No-Go Theorems of Physical Admissibility	10
6.1 No-Free-Facts Theorem	10
6.2 Reversible Fact Exclusion	11
6.3 Infinite Precision Exclusion	11
7. Measurement and Decision Lower Bounds	12
8. Decoherence as Ledger Export	12
9. Worked Examples: PAF in Practice	13
9.1 Single-Qubit Measurement (Minimal Fact Creation).....	14
9.2 Measurement–Reset Cycle (Why Experiments Slow Down)	15
9.3 Quantum Error Correction (Why Overhead Explodes)	16
9.4 Quantum Metrology (Why Precision Stalls)	17
9.5 Hidden Variables with Infinite Precision (Why They Fail)	17
9.6 Summary: What PAF Prices	18
10. A Toy Model of Fact Creation and Admissibility Costs	19
10.1 N-Outcome Fact Creation Model	19
10.2 Two-Outcome Model and Probability Biasing	20
10.3 Gate Nudges vs Detector Nudges (Toy Implementation)	20
10.4 Ledger Accounting in the Toy Model.....	21
10.5 What the Toy Model Demonstrates.....	21
10.6 Role of the Toy Model.....	22
11. Gate Nudges vs Detector Nudges: Bias–Cost Curves.....	22
11.1 Two-Outcome Setup	23

11.2 Gate Nudges: Reversible Biasing	24
11.3 Detector Nudges: Irreversible Biasing.....	25
11.4 Bias–Cost Curve	26
11.5 Control-Ledger Extension	27
11.6 Worked Numbers.....	28
11.7 Summary	31
12. Interpretational Admissibility	32
13. Experimental Predictions and Testability	33
13.1 Scaling Predictions	33
13.2 Potential Experimental Tests	34
13.3 What Would Falsify PAF?.....	35
14. Future Directions.....	35
14.1 Emergent Geometry from Ledger Flow	35
14.2 Interfaces and Holographic Limits.....	36
14.3 Determination of κ	36
14.3.1 A Worked Example: Constraining κ from Published Superconducting-Qubit Readout Data.....	37
14.4 Connection to Quantum Gravity	39
15. The Admissibility–Algebraic Spine	39
15.1 Algebra of Observables as the Primitive Object	40
15.2 States as Functionals: The Origin of the Born Rule	40
15.3 Dynamics as Automorphisms: Unitary Evolution Recovered	40
15.4 Hilbert Space as a Representation, Not a Foundation.....	41
15.5 Measurement and Irreversibility as Completely Positive Maps.....	41
15.6 The Distinguishability Ledger as a Resource Grading.....	42
15.7 A Single Spine, Not Separate Postulates	42
15.8 Conceptual Summary	43
16. Conclusions	43
Appendix A: Mathematical Formalization.....	44
A.1 Accessible Distinguishability.....	44

A.2 Relation to the Holevo Bound	44
A.3 Commitment Channels	45
A.4 Commitment Depth and TPB Bound	45
A.5 Ledger Closure Inequalities.....	45
A.6 Decoherence as Information Export.....	46
A.7 Relation to Thermodynamic Cost	46
A.8 Summary	46
Appendix B: Frequently Asked Questions.....	46
Figure 1: Layered Structure of Physical Law	48
Appendix C: Comparison with Standard Frameworks.....	49
Appendix D: A Heuristic Lower Bound on κ from Landauer + Margolus–Levitin	50
D.1 Motivation	50
D.2 Ingredients	50
D.2.1 Landauer cost per committed bit (thermodynamic lower bound)	50
D.2.2 Margolus–Levitin quantum speed limit (time–energy lower bound).....	50
D.3 Combined bound: a time-per-bit lower limit	51
D.4 Numerical illustration	52
D.5 Interpretation and scope (what this does and does not prove)	52
References.....	53

1. Introduction

Quantum mechanics is built on three mathematical pillars: Hilbert space defines admissible states, Hamiltonians generate reversible dynamics, and the Born rule assigns probabilities to outcomes. While internally consistent, this triad omits a critical physical ingredient: the irreversible cost of turning possibilities into facts.

For general readers: Imagine quantum mechanics as a recipe book that tells you what dishes are possible (states), how ingredients transform during cooking (dynamics), and how likely each dish is to turn out a certain way (probabilities). What the recipe book doesn't tell you is how much fuel you need, how long cooking takes, or why you can't "uncook" a meal. The Physical Admissibility Framework (PAF) adds this missing chapter—it tells you what physical resources are required to make quantum possibilities into definite outcomes.

In real quantum systems, measurement, amplification, readout, reset, and fault tolerance dominate time, energy, and entropy budgets. Standard quantum mechanics specifies what can evolve and with what probabilities, but not what it costs to make outcomes real. The Physical Admissibility Framework (PAF) introduces a fourth layer—an admissibility constraint theory—that operates prior to dynamics and probability, determining what physical processes are possible, impossible, or expensive.

2. Physical Admissibility Postulates

PAF rests on two universal constraints.

(BCB) Bit Conservation Bound: Physically realizable distinguishability is finite within any bounded region and cannot be created ex nihilo. It may only be redistributed or exported to inaccessible degrees of freedom.

(TPB) Ticks-Per-Bit Constraint: The creation of distinguishability as stable, retrievable fact requires irreducible commitment events (ticks). No finite-cost process can realize facts reversibly.

For general readers: Think of BCB like a conservation law for "tellability"—the ability to tell things apart. Just as energy cannot be created from nothing, the capacity to distinguish one thing from another is finite and conserved. You can move distinguishability around (like transferring heat), but you cannot conjure it from nowhere.

TPB says that recording information as a permanent fact requires irreversible steps—moments of no return. Consider writing in ink versus pencil: pencil marks can be erased (reversible), but ink commits permanently (irreversible). Quantum measurement is like writing in ink—once an outcome is recorded, the alternatives are irreversibly discarded. TPB says this commitment

process cannot be done for free; each bit of recorded fact requires at least one irreducible "tick" of irreversible action.

Together, these postulates enforce finite precision, exclude reversible fact creation, and underwrite the emergence of quantum structure. BCB and TPB are not additional physical laws but admissibility constraints: they formalize the empirical impossibility of creating stable facts without irreversible cost.

Status of the Postulates: BCB and TPB are operational admissibility constraints, not metaphysical laws or axioms about the ultimate nature of reality. They are justified empirically: laboratories are finite, stable records exist, and no experiment has ever produced distinguishable outcomes without irreversible resource expenditure. A universe violating these constraints is mathematically definable—one could write down dynamics that create facts for free—but such a universe would be physically non-fact-producing in the operational sense. We would have no stable records, no persistent outcomes, no basis for science. BCB and TPB thus codify what it means for a physical theory to describe a world in which facts can be established and communicated.

3. Primitive Resources and Admissible Operations

We define three primitive resources:

- (i) **Accessible distinguishability I:** The maximum reliably extractable classical information under admissible measurements, defined operationally.
- (ii) **Tick count T:** The number of irreducible commitment events available to a process.
- (iii) **Commitment operations:** Physical channels that convert alternatives into stable records while irreversibly discarding competing possibilities.

Operational definition of ticks: A tick is not a Planck-scale discrete event but an operationally defined threshold: the minimal irreversible step required to commit one unit of distinguishability as stable fact. In practice, ticks are detector-dependent—different measurement apparatuses have different commitment granularity. What PAF constrains is the *minimum* tick count required for any apparatus achieving a given discrimination task. The hazard-rate model in Section 11 provides one concrete realization: a tick corresponds to one first-passage event in the competing-risk dynamics. More generally, any process that irreversibly selects one outcome and discards alternatives constitutes at least one tick.

Relation to quantum resource theories: The quantum thermodynamics literature (Brandão, Horodecki, Oppenheim, and others) has developed sophisticated resource-theoretic frameworks for irreversibility. PAF is complementary to these approaches: while thermodynamic resource theories typically characterize what state transformations are possible given constraints on free

energy, athermality, or coherence, PAF characterizes the cost of *creating classical records*—the commitment step that resource theories often treat as a free classical side-channel. PAF and thermodynamic resource theories can be combined: the former prices the act of fact creation, the latter prices the thermodynamic work required to implement that creation under equilibrium constraints.

Comparison of resource frameworks:

Framework	What it prices	What it treats as free	Primary constraint
Thermodynamic resource theories	Free energy, work, athermality	Classical record-keeping	Second law
Coherence resource theories	Quantum coherence, superposition	Incoherent operations	Monotonicity under incoherent ops
Entanglement resource theories	Entanglement, nonlocality	LOCC operations	No entanglement creation under LOCC
PAF	Classical fact creation, commitment	Reversible (unitary) evolution	BCB, TPB

PAF's distinctive contribution is pricing the *transition from quantum to classical*—the step that other resource theories typically assume is available at no cost. This makes PAF orthogonal to, rather than competing with, existing resource-theoretic frameworks.

Admissible operations are those that respect BCB and TPB. Any operation violating these constraints is physically unrealizable, regardless of its mathematical well-definedness.

4. The Distinguishability Ledger

For general readers: Just as a business tracks income and expenses in a ledger, PAF tracks the "information economy" of any physical process. When you make a measurement, you're not just learning something—you're spending irreversible resources and generating waste. The Distinguishability Ledger is nature's accounting system for these transactions.

For any physical process Π , PAF assigns a ledger triple:

$$L(\Pi) = (T, \Delta I_{\text{fact}}, \Delta I_{\text{export}})$$

where:

- T counts commitment events (the "ticks" or irreversible steps taken)

- ΔI_{fact} measures distinguishability converted to stable fact, defined as the Shannon entropy of the classical record produced: $\Delta I_{\text{fact}} = H(\text{record})$. This is the minimum number of bits required to specify the committed outcome.
- ΔI_{export} measures distinguishability irreversibly discarded to environment (information "thrown away" to make the record permanent). This is an information-theoretic lower bound; Landauer's principle converts it into minimum thermodynamic entropy production under standard conditions (see Appendix A.7).

Analogy: Consider taking a photograph (a kind of measurement). T counts the irreversible chemical or electronic changes in your camera. ΔI_{fact} is the information content of the resulting image—how many bits you need to store it. ΔI_{export} is the light and heat dissipated in the process, carrying away information about all the photos you *didn't* take. You cannot create a photograph without these costs.

Ledger Closure Conditions:

$$T \geq \kappa \cdot \Delta I_{\text{fact}}$$

$$\Delta I_{\text{export}} \geq \max(0, \Delta I_{\text{fact}} - \Delta I_{\text{recovered}})$$

where:

- The first inequality enforces the Ticks-Per-Bit constraint
- $\Delta I_{\text{recovered}}$ represents information that returns to accessible form through environmental feedback or error correction (typically negligible for well-isolated records, so $\Delta I_{\text{recovered}} \approx 0$)
- The max ensures ΔI_{export} remains non-negative

In typical measurement scenarios with well-isolated records, the second inequality reduces to $\Delta I_{\text{export}} \geq \Delta I_{\text{fact}}$.

The constant κ governs the fundamental conversion rate between commitment events and committed information. Its value may be unity in natural units; determining κ from first principles remains an open problem, though it may be constrained experimentally (see Section 13). While κ is not fixed a priori, PAF's predictions concern scaling and lower bounds; any finite κ enforces nonzero cost per committed bit and is therefore physically constraining. More precisely, κ sets the unit conversion between logical commitment events and committed bits; PAF's scaling claims are invariant under finite rescalings of κ . **PAF is thus best understood as a scaling theory:** it predicts how costs scale with task parameters (logarithmically with hypothesis count, linearly with syndrome bits, etc.) rather than absolute values.

This ledger functions as a conservation structure governing irreversible processes, analogous to how energy conservation governs reversible dynamics.

Definition (Distinguishability Ledger as a lax resource functor). Let \mathbf{Proc} be the category whose objects are finite-resource systems and whose morphisms are admissible CPTP maps (or admissible CP maps in the algebraic spine), composed sequentially. A ledger is a map

$$L : \mathbf{Proc} \rightarrow \mathbb{R}_{+}^3, L(\Pi) = (T(\Pi), \Delta I_{\text{fact}}(\Pi), \Delta I_{\text{export}}(\Pi))$$

satisfying:

1. **Free reversibles:** If Π is reversible (unitary/automorphism), then $L(\Pi) = (0, 0, 0)$.
2. **Lax additivity:** For composition $\Pi_2 \circ \Pi_1$, $L(\Pi_2 \circ \Pi_1) \preceq L(\Pi_1) + L(\Pi_2) + \Delta_{\text{corr}}$ where $\Delta_{\text{corr}} \in \mathbb{R}_{+}^3$ accounts for correlation/initialization costs and \preceq is componentwise inequality.
3. **Closure inequalities (PAF constraints):**
 - o $T(\Pi) \geq \kappa \cdot \Delta I_{\text{fact}}(\Pi)$
 - o $\Delta I_{\text{export}}(\Pi) \geq \max(0, \Delta I_{\text{fact}}(\Pi) - \Delta I_{\text{recovered}}(\Pi))$

This formalization makes composition rules and "unitaries are free" mathematically automatic and connects PAF to categorical quantum mechanics and resource theory frameworks.

5. Composition Rules and Commitment Depth

Ledger entries compose under physical operations:

- **Sequential composition:** $L(\Pi_2 \circ \Pi_1) \preceq L(\Pi_1) + L(\Pi_2) + \Delta_{\text{corr}}$, where $\Delta_{\text{corr}} \geq 0$ accounts for correlation and initialization costs, and \preceq denotes componentwise inequality
- **Parallel composition:** Entries add for independent processes
- **Parallel processes** share ticks only if commitment bandwidth permits

Definition (Per-commitment record capacity). Let \mathcal{C} be the allowed class of elementary commitment channels (non-unitary CPTP maps producing stable record updates) in the admissible model. Define

$$I_{\text{max}} \equiv \sup_{\mathcal{C}} \{C \in \mathcal{C}\} \sup_{\rho} [H(R_{\text{out}}) - H(R_{\text{in}})]$$

i.e., the maximum increase in stable record entropy achievable in one commitment operation.

Theorem 5.1 (Commitment Depth Lower Bound). Any implementation of a process Π that outputs a stable classical record R with $H(R) = \Delta I_{\text{fact}}$ requires

$$D_C(\Pi) \geq \Delta I_{\text{fact}} / I_{\text{max}}$$

where $D_C(\Pi)$ is the minimum number of commitment operations used in any realization of Π .

Proof. By definition, unitary/reversible steps do not increase stable record entropy (Lemma 6.2 below). Each commitment operation can increase record entropy by at most I_{\max} . Therefore after d commitment operations, $H(R) \leq d \cdot I_{\max}$. To reach $H(R) = \Delta I_{\text{fact}}$, one must have $d \geq \Delta I_{\text{fact}} / I_{\max}$. Minimizing over all realizations gives the stated bound.

Note on correlations: If commitment operations are correlated (sharing information through auxiliary systems), the total record entropy may be less than the sum of individual contributions. However, this can only *increase* the required depth: producing ΔI_{fact} bits of *independent* record requires at least $\Delta I_{\text{fact}} / I_{\max}$ operations regardless of correlations. The bound is therefore robust. \square

Corollary (TPB form). If the admissible architecture enforces $I_{\max} = 1/\kappa$ bits per commitment event, then $D_C(\Pi) \geq \kappa \cdot \Delta I_{\text{fact}}$.

In plain terms: commitment depth grows at least linearly with the information committed. A process that commits n bits requires at least n irreversible stages (up to constant factors).

6. No-Go Theorems of Physical Admissibility

For general readers: "No-go theorems" are proofs that certain things are impossible—not just difficult, but fundamentally ruled out by the laws of physics. PAF establishes three such impossibilities, each revealing deep constraints on what physical systems can do.

6.1 No-Free-Facts Theorem

Definition (Stable classical record). A record variable R is *stable* if (i) it takes values in a finite alphabet \mathcal{R} , (ii) it can be copied by admissible operations into m independent registers $R^{(1)}, \dots, R^{(m)}$ with arbitrarily small degradation (classical fan-out), and (iii) the copies remain readable for times long compared to the process duration.

Assumption (BCB-capacity). Any bounded laboratory domain \mathcal{D} admits a finite upper bound $I_{\max}(\mathcal{D})$ on operationally accessible distinguishability stored in stable classical records within \mathcal{D} (measured in bits). This is the operational form of the Bit Conservation Bound.

Theorem 6.1 (No-Free-Facts). Let Π be a finite-resource process acting within a fixed bounded domain \mathcal{D} that produces a stable classical record R with Shannon entropy $H(R) = \Delta I_{\text{fact}} > 0$. Then Π necessarily incurs irreversible commitment cost satisfying

$$T(\Pi) \geq \kappa \cdot \Delta I_{\text{fact}}$$

for some $\kappa > 0$ characterizing the minimal ticks-per-bit of the admissible domain. In particular, there is no admissible process that produces ΔI_{fact} bits of stable record with $T(\Pi) = 0$.

Proof (by iteration and capacity contradiction). Suppose for contradiction that there exists a process Π producing $H(R) = \Delta I_{\text{fact}} > 0$ while violating the tick bound, i.e., $T(\Pi) < \kappa \cdot \Delta I_{\text{fact}}$ for some fixed $\kappa > 0$. Consider running Π sequentially k times inside the same bounded domain \mathcal{D} , producing records R_1, \dots, R_k that are stable and copyable by definition. Since stability allows independent storage/copying of these outcomes, the joint record $R^{(k)} = (R_1, \dots, R_k)$ has entropy

$$H(R^{(k)}) = \sum_j H(R_j) = k \cdot \Delta I_{\text{fact}}$$

up to negligible correlations that can be made arbitrarily small by resetting the apparatus between runs. The total tick cost is at most $k \cdot T(\Pi)$, hence $k \cdot T(\Pi) < \kappa \cdot k \cdot \Delta I_{\text{fact}}$. Letting $k \rightarrow \infty$, this constructs unbounded stable distinguishability inside the fixed bounded domain \mathcal{D} , contradicting the BCB–capacity assumption $H(\text{stable records in } \mathcal{D}) \leq I_{\text{max}}(\mathcal{D}) < \infty$. Therefore no such Π exists; any stable fact production requires nonzero tick cost proportional to the committed information. \square

Plain English: You cannot learn anything for free. Every bit of definite information that gets recorded into the world requires spending irreversible resources—there is no way around this.

6.2 Reversible Fact Exclusion

Lemma 6.2 (Reversible maps cannot create classical records). Let Φ be an admissible reversible transformation (unitary channel or reversible automorphism in the algebraic spine). Then for any classical record register R initially uncorrelated with the system, Φ cannot increase the accessible classical information contained in R . In particular, any process that increases $H(R)$ from 0 to > 0 must contain at least one non-unitary (commitment) step.

Proof. Reversible dynamics are information-preserving isometries on the operational state space. If a classical record R is produced, that corresponds to a many-to-one mapping from pre-measurement alternatives to a stable macrostate of R , which is not invertible on the operational domain. Equivalently, for any classical-quantum state ρ_{XR} with X the preparation label, data processing gives that $I(X; R)$ cannot increase under local reversible dynamics acting on the system alone; producing a nontrivial classical record requires a non-unitary channel coupling to an environment and discarding alternatives. \square

Plain English: Reversible processes cannot make decisions. The quantum world evolves reversibly (unitarily), which is precisely why it remains in superposition—because nothing has been decided yet. The moment a definite fact emerges, something irreversible has occurred.

6.3 Infinite Precision Exclusion

Theorem 6.3 (Infinite Precision Exclusion). Any physical model requiring unbounded state resolution to produce finite distinguishable records is inadmissible under PAF.

Proof. Suppose a model postulates a continuous hidden variable $\lambda \in \mathbb{R}$ with physical consequences—i.e., different values of λ lead to distinguishable measurement outcomes. To

discriminate N distinct values of λ requires $\Delta I_{\text{fact}} \geq \log_2(N)$ bits of committed record. By Theorem 6.1, this requires $T \geq \kappa \log_2(N)$ ticks. As $N \rightarrow \infty$ (approaching continuous resolution), $T \rightarrow \infty$. Any finite-resource process can therefore only access finite precision. Models requiring infinite precision for finite physical effects violate BCB. \square

Plain English: Reality cannot be infinitely detailed. If some theory requires knowing a number to infinite decimal places to make predictions, that theory cannot describe physical processes achievable with finite resources.

7. Measurement and Decision Lower Bounds

Resolving N hypotheses with total error probability ε requires:

$$\Delta I_{\text{fact}} \geq \log_2(N) - h(\varepsilon)$$

where $h(\varepsilon) = -\varepsilon \log_2(\varepsilon) - (1-\varepsilon) \log_2(1-\varepsilon)$ is the binary entropy function.

Here ΔI_{fact} denotes the minimum classical information that must be committed as a stable, retrievable record to label one of N hypotheses with total error probability ε . The bound follows from standard decision-theoretic and information-theoretic inequalities (Fano-type bounds), independent of quantum structure. Quantum mechanics constrains which measurements can approach this bound; PAF constrains the physical cost of any measurement that does.

By TPB, any such protocol must satisfy:

$$T \geq \kappa \cdot [\log_2(N) - h(\varepsilon)]$$

PAF thus prices optimal Helstrom discrimination and explains why measurement and readout dominate experimental cost even when unitary evolution is cheap. The irreversible cost of deciding is irreducible.

8. Decoherence as Ledger Export

For general readers: "Decoherence" is the process by which quantum superpositions appear to collapse when a quantum system interacts with its environment. Standard physics describes *what* happens (interference patterns disappear), but PAF explains *why* in terms of information accounting: decoherence is the export of "which-branch" information to the environment.

Imagine a quantum system as a secret that could be one of several possibilities. Decoherence occurs when the environment "learns" which possibility is realized—not by anyone looking, but

simply through physical interaction. Once the environment contains a record of which branch occurred, the other branches become inaccessible. The information isn't destroyed; it's exported to degrees of freedom we cannot access.

Environmental decoherence corresponds to the export of which-branch distinguishability into inaccessible degrees of freedom. Consider a system-environment interaction producing decoherence in basis $\{|i\rangle\}$. The reduced system dynamics suppress off-diagonal terms:

$$\rho_S \rightarrow \sum_i \langle i | \rho_S | i \rangle \cdot |i\rangle\langle i|$$

Equivalently, the dephasing channel acts as:

$$\mathcal{D}(\rho_S) = \sum_i |i\rangle\langle i| \rho_S |i\rangle\langle i|$$

Technical note: The off-diagonal terms (coherences) that encode quantum interference get suppressed. What remains is a classical probability distribution over outcomes.

The environment gains which-branch information quantified by the mutual information $I(S:E)$. This provides a lower bound on ΔI_{export} .

Phase information is not destroyed but displaced beyond admissible recovery. PAF reinterprets decoherence as a ledger transfer rather than information destruction—the universe's books still balance.

Note on recoverability: In principle, if one had complete access to all environmental degrees of freedom, the phase information could be recovered (decoherence is unitary at the system+environment level). In practice, environmental degrees of freedom rapidly become inaccessible—thermalized, dispersed, or correlated with further degrees of freedom. PAF's ΔI_{export} quantifies this practical irreversibility: the information is not gone, but recovering it would require resources exceeding any finite budget. This distinction between in-principle and practical recoverability is central to understanding why decoherence appears irreversible despite underlying unitary dynamics.

Key insight: Decoherence is not mysterious "wave function collapse"—it's accounting. Information about which outcome occurred gets transferred to the environment. This transfer is an irreversible ledger operation with definite cost. The quantum interference patterns don't vanish; they become encoded in correlations with environmental degrees of freedom that are practically (and often fundamentally) inaccessible.

9. Worked Examples: PAF in Practice

For general readers: This section shows PAF in action across several real scenarios in quantum physics and computing. In each case, we'll see how PAF explains costs and limitations that

standard quantum mechanics acknowledges exist but doesn't account for. Think of these examples as "mystery solved" stories—phenomena that seemed like engineering inconveniences turn out to be fundamental physical constraints.

This section presents concrete applications of PAF to standard quantum scenarios, demonstrating how admissibility constraints explain costs that standard quantum mechanics does not address.

9.1 Single-Qubit Measurement (Minimal Fact Creation)

For general readers: This is the simplest possible quantum measurement—a single yes/no question asked of a quantum system. Even this minimal case reveals the fundamental cost of creating facts.

Standard QM view: A projective measurement maps a qubit state

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

to outcome 0 or 1 with Born probabilities $|\alpha|^2$ and $|\beta|^2$.

What QM does not specify:

- How long this takes
- What resources are consumed
- Why the outcome is irreversible

PAF analysis: Recording one classical bit requires:

- $\Delta I_{\text{fact}} = H(P)$, which equals 1 bit for a balanced outcome ($P = 0.5$) and drops toward 0 as outcomes become nearly certain
- At least one commitment event: $T \geq \kappa \cdot H(P)$
- Export of at least one bit of distinguishability: $\Delta I_{\text{export}} \geq H(P)$

For a maximally uncertain measurement ($P = 0.5$), these all equal 1. For a nearly certain outcome ($P \rightarrow 1$), the costs approach zero—but achieving that certainty may itself require resources (see Section 11).

No unitary operation can achieve this. Any physical detector must irreversibly discard the orthogonal branch.

Takeaway: Even the simplest quantum measurement has an irreducible physical cost that standard QM does not account for. The superposition doesn't just "decide"—something must pay for that decision.

9.2 Measurement–Reset Cycle (Why Experiments Slow Down)

For general readers: Real experiments don't just measure once—they measure thousands or millions of times to gather statistics. Each cycle involves measuring, recording the result, and then resetting the system to measure again. This "measure-reset-repeat" loop has costs that accumulate. PAF explains why these costs are unavoidable and why they double with each component of the cycle.

Scenario: Repeatedly measure a qubit and reset the detector.

Standard QM view: Measurement followed by re-preparation. No resource accounting.

PAF ledger accounting:

Measurement phase:

- $\Delta I_{\text{fact}} = H(P)$, which equals 1 bit for a balanced outcome
- $T \geq \kappa \cdot H(P)$
- $\Delta I_{\text{export}} \geq H(P)$

Reset phase:

- Additional $\Delta I_{\text{export}} \geq H(P)$

Why reset costs as much as measurement: When you measure, you create a record and export entropy (the "discarded" alternative). When you reset, you must erase that record to prepare for the next measurement. By Landauer's principle, erasing information is just as costly as creating it—you're exporting the same entropy again, just in the opposite direction.

Total per cycle (for balanced outcomes, $H(P) = 1$):

$$T \geq 2\kappa, \Delta I_{\text{export}} \geq 2$$

This cost accumulates linearly with repetition. After n cycles:

$$T_{\text{total}} \geq 2n\kappa, \Delta I_{\text{export, total}} \geq 2n$$

Takeaway: There is no such thing as a "free measurement cycle." Resetting reality costs as much as reading it. This is why quantum experiments, despite involving tiny quantum systems, require significant time and energy budgets—the cost is in the *accounting* of facts, not the *size* of the systems.

9.3 Quantum Error Correction (Why Overhead Explodes)

For general readers: Quantum computers are incredibly fragile—tiny disturbances cause errors. Quantum error correction (QEC) is the technique that allows quantum computers to work despite errors, by encoding information redundantly and constantly checking for problems. But this checking has a cost that PAF reveals to be fundamental, not merely technical.

The puzzle: Theoretically, QEC can reduce errors to arbitrarily low levels. In practice, quantum computers struggle enormously with error correction overhead. Why is it so hard? PAF provides the answer: every error check creates a classical fact, and facts cost irreversible resources.

Standard QM / QEC theory: Fault tolerance allows arbitrarily low logical error rates via the threshold theorem.

What it hides: The cost of extracting syndromes as classical facts.

PAF analysis: Consider a stabilizer code with m syndrome bits per correction round.

Per round:

- Each syndrome bit is a committed fact: $\Delta I_{\text{fact}} \geq m$
- Commitment cost: $T \geq km$
- Ancilla reset doubles export cost: $\Delta I_{\text{export}} \geq 2m$

Scaling: As code distance d increases, syndrome count m grows (typically $m \sim n - k$ for n physical qubits encoding k logical qubits). Cumulative irreversible cost therefore scales faster than unitary gate count.

Why this matters: To detect errors, you must measure "syndrome" bits that reveal what errors occurred without revealing the protected quantum information. Each syndrome measurement is a fact that must be recorded. As you want better protection (larger codes), you need more syndrome bits, each carrying its own irreversible cost. The overhead isn't a sign of bad engineering—it's the price of creating the classical facts needed to know what errors occurred.

Takeaway: Fault tolerance overhead is not an engineering artifact—it is the price of repeatedly creating classical facts. The threshold theorem guarantees arbitrarily low logical error rates, but PAF prices the irreversible resources required to achieve them. This explains why building large-scale quantum computers is so challenging: the very act of protecting quantum information requires creating classical records, and classical records cost irreversibility.

9.4 Quantum Metrology (Why Precision Stalls)

Standard theory: The Heisenberg limit allows phase sensitivity $\Delta\theta \sim 1/N$ for N probe particles, beating the standard quantum limit $\Delta\theta \sim 1/\sqrt{N}$. These bounds derive from the quantum Fisher information, which characterizes the information content of a quantum state about the parameter.

What it ignores: The cost of certifying an estimate as a stable, retrievable fact. The quantum Fisher information bounds the precision achievable in principle, but says nothing about the irreversible resources required to commit that estimate to a physical record.

PAF analysis: PAF separates estimation into two stages:

1. Reversible phase accumulation (unitary, cost-free in principle)
2. Irreversible commitment of the estimate (costly)

To resolve θ to precision $\Delta\theta$ requires distinguishing $N \approx 1/\Delta\theta$ outcome bins. The bound below concerns the commitment stage, not the Fisher information.

PAF bound:

$$\Delta I_{\text{fact}} \geq \log_2(1/\Delta\theta)$$

$$T \geq \kappa \cdot \log_2(1/\Delta\theta)$$

Even perfect unitary sensitivity (saturating the quantum Fisher information bound) must pay irreversible cost at readout. This explains why improvements in raw phase sensitivity do not translate linearly into usable metrological advantage: beyond a precision threshold, readout and stabilization costs dominate.

Takeaway: Precision is not only dynamically limited—it is admissibility-limited. This bound applies universally, independent of probe state, entanglement structure, or Hamiltonian.

9.5 Hidden Variables with Infinite Precision (Why They Fail)

Claimed model: A continuous hidden variable (e.g., exact particle position in Bohmian mechanics) determines measurement outcomes with infinite precision.

PAF analysis: The admissibility question is not whether hidden variables can be mathematically postulated, but whether they introduce operationally accessible distinguishability beyond standard quantum predictions.

Consider two cases:

Case 1: Hidden variables are operationally inaccessible. If the additional variables cannot, even in principle, produce distinguishable records beyond Born-rule statistics, they do not violate BCB. However, they also do no explanatory work that standard quantum mechanics does not already accomplish.

Case 2: Hidden variables introduce new operational distinguishability. If the hidden variable can produce records that distinguish states indistinguishable under standard quantum mechanics, then:

- That distinguishability must be finite (by BCB)
- Infinite precision would require infinite accessible distinguishability
- Infinite distinguishability implies infinite ledger cost: $T \rightarrow \infty$

Conclusion: Such models are admissible under PAF only if the additional variables do not introduce new operationally accessible distinguishability beyond the quantum admissibility bound. If they require unbounded operational resolution to yield finite differences in observed records, they violate BCB.

Clarification on Bohmian mechanics: Bohmians carefully distinguish ontological precision from operational access. The guidance equation may specify particle positions to infinite precision, but this precision is never operationally accessible—measurements of Bohmian particles yield only Born-rule statistics, with finite precision. Under this reading, Bohmian mechanics remains *admissible* under PAF: the infinite-precision ontology does no operational work that would require unbounded ledger cost. PAF neither endorses nor refutes Bohmian mechanics; it simply constrains what operational role hidden variables can play. If the hidden variables are truly hidden—inaccessible to any measurement—they don't violate BCB. But they also cannot explain operational phenomena that standard quantum mechanics cannot.

Note: The question of what constitutes "operational accessibility" in pilot-wave theories remains a subject of active philosophical debate. Some Bohmians argue that position information becomes effectively accessible through decoherence and environmental encoding; others maintain strict operational inaccessibility. PAF is compatible with either reading, provided the chosen interpretation respects finite operational distinguishability. We do not claim PAF settles interpretational disputes—it provides constraints that any interpretation must satisfy.

Takeaway: Infinite precision is not forbidden mathematically—it is forbidden physically whenever it has operational consequences. Models may postulate hidden structure, but that structure cannot do physical work requiring resources that exceed admissibility limits.

9.6 Summary: What PAF Prices

Scenario	Standard QM Says	PAF Adds
Single measurement	Outcome with probability p	Cost: $T \geq \kappa$, $\Delta I_{\text{export}} \geq 1$

Scenario	Standard QM Says	PAF Adds
Measurement + reset	State preparation	Total cost doubles per cycle
Error correction	Arbitrarily low error possible	Irreversible cost scales with syndrome count
Metrology	Heisenberg limit achievable	Readout cost grows as $\log(1/\Delta\theta)$
Hidden variables	Mathematically consistent	Physically inadmissible if infinite precision required

In each case, PAF does not contradict standard quantum mechanics—it completes it by specifying the irreversible cost of making quantum possibilities into physical facts.

10. A Toy Model of Fact Creation and Admissibility Costs

To make the Physical Admissibility Framework concrete, we present a deliberately minimal toy model that captures the core claims of PAF without relying on microscopic detector physics, detailed thermodynamics, or specific hardware implementations. The purpose of this model is not quantitative prediction but structural demonstration: to show, in the simplest possible setting, how finite distinguishability, irreversible commitment, and ledger costs necessarily emerge.

The toy model implements three ingredients only:

1. Classical record creation as an irreversible operation
2. Finite outcome resolution among N hypotheses
3. A distinction between reversible probability steering and irreversible outcome biasing

Despite its simplicity, the model reproduces the central scaling claims of PAF.

10.1 N-Outcome Fact Creation Model

Consider a physical process whose sole purpose is to produce a stable classical record identifying one outcome from a finite set $\mathcal{O} = \{1, 2, \dots, N\}$.

We assume:

- The record must be stable, copyable, and retrievable (Definition 6.1)
- The process uses finite resources and operates within a bounded domain
- The outcome probabilities are given by a distribution $\{p_i\}$, with $\sum_i p_i = 1$

The committed record is a classical random variable R taking values in \mathcal{O} . The minimum information content of this record is its Shannon entropy:

$$\Delta I_{\text{fact}} = H(R) = -\sum_i p_i \log_2 p_i$$

In the symmetric (worst-case) discrimination task, $p_i = 1/N$, yielding:

$$\Delta I_{\text{fact}} = \log_2 N$$

By the Ticks-Per-Bit constraint (TPB), any admissible realization of this process must therefore satisfy:

$$T \geq \kappa \log_2 N$$

Key point: Even in a single-shot measurement, resolving among N distinct possibilities requires commitment resources that scale logarithmically with N . This scaling is purely information-theoretic and independent of quantum dynamics, detector microphysics, or hardware substrate.

10.2 Two-Outcome Model and Probability Biasing

We now specialize to the simplest nontrivial case: a binary outcome $R \in \{0, 1\}$.

Let the observed success probability be $P = \Pr(R = 0)$. The record entropy is:

$$\Delta I_{\text{fact}} = H(P) = -P \log_2 P - (1-P) \log_2(1-P)$$

As $P \rightarrow 1$, the record entropy $H(P) \rightarrow 0$: recording an almost-certain outcome is informationally cheap. However, the *mechanism* by which such bias is achieved matters.

10.3 Gate Nudges vs Detector Nudges (Toy Implementation)

To model outcome biasing, we introduce two independent control parameters:

- **Gate (state) bias:** modifies the intrinsic probability p via reversible dynamics
- **Detector bias:** modifies the commitment rates of outcome channels

We represent the commitment process as a competing-risk model with effective rates:

$$\lambda_0 = \kappa_0 \cdot p, \lambda_1 = \kappa_1 \cdot (1-p)$$

where $\kappa_0, \kappa_1 > 0$ characterize apparatus couplings.

The resulting outcome probabilities are:

$$P(0) = \lambda_0 / (\lambda_0 + \lambda_1), P(1) = \lambda_1 / (\lambda_0 + \lambda_1)$$

Defining the log-odds $\Lambda(P) \equiv \log(P/(1-P))$, we obtain the exact factorization:

$$\Lambda(P) = \log(p/(1-p)) + \log(\kappa_0/\kappa_1) \equiv \Lambda_g + \delta$$

where:

- Λ_g is the gate-nudge contribution (reversible)
- $\delta = \log(\kappa_0/\kappa_1)$ is the detector-nudge contribution (irreversible)

This decomposition holds for any monotonic first-passage or threshold-crossing commitment process; exponential hazards are chosen only for analytic simplicity.

10.4 Ledger Accounting in the Toy Model

The ledger costs separate cleanly:

Gate nudges:

- Implemented by unitary evolution
- Do not create records
- Carry zero ledger cost: $(T, \Delta I_{\text{fact}}, \Delta I_{\text{export}}) = (0, 0, 0)$

Record creation:

- Regardless of how P is achieved:
 - $T \geq \kappa \cdot H(P)$
 - $\Delta I_{\text{export}} \geq H(P)$
- This is the unavoidable cost of committing the outcome as fact

Detector nudges:

- To achieve a target P with fixed input state ($p = 1/2$), the required detector asymmetry is $\delta = \log(P/(1-P))$
- Maintaining $\delta \neq 0$ corresponds to holding the apparatus in a nonequilibrium asymmetric configuration
- By standard results in stochastic thermodynamics, the minimum control-entropy export satisfies:
 - $|\Delta I_{\text{ctrl}}| \geq c|\delta|$
 - for some apparatus-dependent constant $c > 0$

Thus, while $H(P) \rightarrow 0$ as $P \rightarrow 1$, the control cost diverges logarithmically: $|\delta| \sim \log(1/(1-P))$.

10.5 What the Toy Model Demonstrates

This minimal construction makes three PAF claims explicit:

1. **Logarithmic fact cost:** Resolving among N possibilities requires $\Omega(\log N)$ committed record bits, and hence $\Omega(\log N)$ ticks.
2. **Record cost universality:** The irreversible cost of recording an outcome depends only on the outcome entropy $H(P)$, not on how the probability bias was produced.
3. **Asymmetry is expensive:** Biasing outcomes via detector asymmetry requires irreversible control resources that scale at least linearly with log-odds $|\delta|$.

These features arise without invoking quantum dynamics, Hilbert space, or specific detector physics. They follow solely from finite distinguishability, irreversible commitment, and the existence of stable records.

10.6 Role of the Toy Model

This toy model is not intended as a realistic detector simulation. Rather, it serves three purposes:

1. To instantiate PAF in the simplest possible setting
2. To demonstrate scaling laws explicitly
3. To show that PAF's claims do not depend on hidden microscopic assumptions

More detailed models—incorporating explicit Hamiltonians, noise processes, or device architectures—may refine prefactors and identify substrate-specific κ values. They cannot evade the structural constraints illustrated here.

In this sense, the toy model plays the same role for PAF that idealized heat engines play for thermodynamics: not realistic, but inescapable.

11. Gate Nudges vs Detector Nudges: Bias–Cost Curves

For general readers: This section addresses a fundamental question: if you want to bias a quantum measurement toward a particular outcome, how should you do it, and what does it cost?

Imagine you're flipping a coin and want heads to come up more often. You have two strategies:

1. **Bend the coin** (analogous to a "gate nudge"): Change the coin itself before flipping, making it physically lopsided so it naturally favors heads. In quantum terms, this means rotating the quantum state before measurement.
2. **Tilt the table** (analogous to a "detector nudge"): Keep the coin fair, but modify the landing surface so that even fair flips tend to settle as heads. In quantum terms, this means adjusting the measurement apparatus to favor one outcome.

Standard quantum mechanics treats these as equivalent—both just change probabilities. PAF reveals they are fundamentally different: bending the coin is *reversible* (you can bend it back),

while tilting the table requires *irreversible* effort to maintain. This distinction has profound implications for quantum computing, metrology, and our understanding of measurement itself.

Quantum experiments routinely "nudge" outcomes. In quantum computing, nudging is typically performed by applying small unitary gates prior to readout ("gate nudges"). Alternatively, measurement apparatus design can bias which outcome becomes a stable fact by changing detector couplings and thresholds ("detector nudges"). Standard quantum mechanics treats both as implementation details preserving Born statistics, while PAF distinguishes them by where the intervention acts: the reversible sector versus the irreversible commitment layer.

11.1 Two-Outcome Setup

For general readers: We start with the simplest possible quantum system: a qubit that can be measured as either "0" or "1" (like a quantum coin that lands heads or tails). The quantum state determines the *intrinsic* probability of each outcome, but the measurement apparatus determines how that probability gets *realized* as an actual recorded fact.

Consider a system prepared in state

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, p \equiv |\alpha|^2, 1-p \equiv |\beta|^2$$

measured in the $\{|0\rangle, |1\rangle\}$ basis.

Plain English: The qubit is in a superposition—it's "both 0 and 1" until measured. The number p (between 0 and 1) is the intrinsic probability of getting outcome 0, determined by the quantum state. If $p = 0.5$, the quantum coin is fair; if $p = 0.9$, it's heavily biased toward 0.

In a first-passage / hazard-competition picture (consistent with metastable threshold detection), each outcome channel $i \in \{0,1\}$ has an effective commitment rate λ_i . For concreteness, we model commitment as a competing-risk process with exponential hazard rates; the qualitative conclusions depend only on monotonic first-passage ordering, not on the specific distribution. Any model where higher-rate channels win more often—including Weibull, gamma, or general survival processes—yields the same factorization structure; exponential rates simply provide clean closed-form expressions.

Experimental support: Superconducting qubit readout, trapped-ion fluorescence detection, and photon counting all exhibit threshold-crossing dynamics consistent with competing first-passage processes. The specific distribution details affect quantitative prefactors but not the log-odds decomposition central to our analysis.

The simplest admissible competing-risk model yields:

$$P(0) = \lambda_0/(\lambda_0 + \lambda_1), P(1) = \lambda_1/(\lambda_0 + \lambda_1)$$

Intuitive explanation: Think of measurement as a race between two possible outcomes. Each outcome has a "rate" λ at which it tries to become real. Whichever outcome crosses the finish

line first wins and gets recorded. The probability of each outcome equals its rate divided by the total rate—faster outcomes win more often.

Parameterize the rates as:

$$\lambda_0 = \kappa_0 \cdot p, \lambda_1 = \kappa_1 \cdot (1-p)$$

What this means: Each outcome's effective rate combines two factors:

- **p or (1-p):** The intrinsic quantum probability from the state (what the "coin" wants to do)
- **κ_0 or κ_1 :** The apparatus coupling factor (how responsive the detector is to each outcome)

If $\kappa_0 = \kappa_1$ (symmetric detector), the measured probabilities equal the Born probabilities. But if $\kappa_0 \neq \kappa_1$, the detector itself biases the outcomes.

Define the bias ratio $R \equiv P(0)/P(1) = \lambda_0/\lambda_1$. Then:

$$\log R = \log(p/(1-p)) + \log(\kappa_0/\kappa_1)$$

This factorization is the key structural result:

- First term: **gate-nudge contribution** (state steering)—how much the quantum state favors outcome 0
- Second term: **detector-nudge contribution** (apparatus steering)—how much the detector favors outcome 0

The crucial insight: These two contributions are *additive* in log-odds space, but they have completely different physical character. Gate nudges operate in the reversible quantum realm; detector nudges operate at the irreversible commitment layer. PAF prices these two contributions differently.

Important clarification: In standard quantum mechanics, detector asymmetry ($\kappa_0 \neq \kappa_1$) corresponds to measuring with a different effective POVM or a different coarse-graining of outcomes. The model above does not violate Born statistics; it parameterizes how apparatus coupling changes the realized outcome distribution by changing the effective measurement channel. PAF's point is not that the Born rule fails, but that *choosing, implementing, and stabilizing* a biased measurement channel carries irreversible cost that the Born rule does not track.

11.2 Gate Nudges: Reversible Biasing

For general readers: A gate nudge is like adjusting the angle of a spinning top before it falls. By rotating the quantum state, you change which outcome is more likely—but crucially, you can always rotate it back. Nothing irreversible has happened yet.

A gate nudge applies a unitary $U(\theta)$ before measurement, changing the Born weight $p \rightarrow p(\theta)$. For a qubit rotation about a transverse axis:

$$p(\theta) = \cos^2(\theta/2)$$

Example: Start with a fair quantum coin ($p = 0.5$). Apply a rotation of $\theta = \pi/2$. Now $p(\theta) = \cos^2(\pi/4) = 0.5$ —still fair. Apply $\theta = \pi$. Now $p(\theta) = \cos^2(\pi/2) = 0$ —certain to get outcome 1. The rotation is completely reversible; apply $-\theta$ and you're back where you started.

Gate nudges are reversible operations and carry zero ledger cost:

$$\Delta I_{\text{fact}} = 0, T = 0, \Delta I_{\text{export}} = 0$$

for the unitary itself. Their cost is engineering (control energy, calibration), not PAF-fundamental.

Why zero cost? Because no fact has been created yet. The quantum state has been rotated, but it's still in superposition—still "undecided." No alternatives have been discarded. It's like repositioning chess pieces before the game starts: you can move them around freely because nothing has been committed.

However, gate nudges cannot eliminate commitment cost. When the outcome is recorded, the ledger cost is set by the entropy of the outcome distribution. Let $P \equiv P(0)$. The minimum committed information per trial is:

$$\Delta I_{\text{fact}} \geq H(P) \equiv -P \log_2 P - (1-P) \log_2(1-P)$$

and by TPB:

$$T \geq \kappa \cdot H(P)$$

What $H(P)$ means: $H(P)$ is the Shannon entropy—a measure of "surprise" or "uncertainty" in the outcome. If $P = 0.5$ (fair coin), $H = 1$ bit: maximum uncertainty, maximum information gained by learning the outcome. If $P = 0.99$ (almost certain), $H \approx 0.08$ bits: low uncertainty, little information gained.

Key implication: Gate nudges change P , but the irreversible cost of making the result a fact is governed by $H(P)$ regardless of how P was achieved. You can make the quantum coin very biased using gate nudges (cheap), but you still have to *record* the outcome (costly), and recording costs at least $\kappa \cdot H(P)$ ticks.

11.3 Detector Nudges: Irreversible Biasing

For general readers: A detector nudge is like weighting a roulette wheel—not by changing the ball or the spin, but by modifying the pockets so some outcomes "catch" more easily than others.

Unlike gate nudges, this requires maintaining a persistent physical asymmetry in the measurement apparatus, which has ongoing costs.

A detector nudge changes κ_0/κ_1 by modifying apparatus couplings, thresholds, metastability, or bandwidth. Define the detector bias parameter:

$$\delta \equiv \log(\kappa_0/\kappa_1)$$

so that $\log R = \log(p/(1-p)) + \delta$.

Example: Suppose your detector is twice as sensitive to outcome 0 as to outcome 1 ($\kappa_0/\kappa_1 = 2$). Then $\delta = \log(2) \approx 0.69$. Even with a fair quantum state ($p = 0.5$), the measured probability becomes $P(0) = 2/3$ instead of $1/2$. The detector itself is biasing the results.

Detector nudges bias outcomes even when the input state is held fixed. They act at the commitment layer and require increased irreversibility to maintain bias reliably. PAF captures this by assigning additional ledger overhead $\Delta I_{\text{export}}(\text{nudge})$ and tick consumption $T(\text{nudge})$.

Why the extra cost? Maintaining detector asymmetry requires active stabilization. The detector must "remember" its bias across trials—this is itself a form of stored information that must be preserved against noise, thermal fluctuations, and drift. Holding a system in an asymmetric configuration, away from equilibrium, is thermodynamically costly.

Analogy: Keeping a door propped open requires continuous effort (a doorstop, or someone holding it). The symmetric state (door swinging freely) is natural; the asymmetric state (door held open) requires maintained intervention.

PAF prediction: Stronger detector bias requires greater irreversible overhead.

11.4 Bias–Cost Curve

For general readers: This subsection derives a universal relationship between how biased you want your outcomes and how much it irreversibly costs. The key insight: recording a nearly-certain outcome is cheap (low surprise), but *making* an outcome nearly-certain via detector bias is expensive (maintaining asymmetry costs resources).

For target success probability $P \in (0,1)$, define log-odds $\Lambda(P) \equiv \log(P/(1-P))$. From the model:

$$\Lambda(P) = \Lambda_g + \Lambda_d$$

where $\Lambda_g \equiv \log(p/(1-p))$ is the gate contribution and $\Lambda_d \equiv \delta$ is the detector contribution.

Plain English: The total bias (log-odds of success) is the sum of bias from the quantum state and bias from the detector. You can reach any target probability P by combining these two sources in any proportion.

The record entropy is $H(P)$ regardless of how P was achieved, so:

$$T_{\min}(P) \geq \kappa \cdot H(P)$$

With reset doubling for repeated cycles:

$$T_{\text{cycle}}(P) \geq 2\kappa \cdot H(P)$$

The universal record-cost curve: This is perhaps the most important result. No matter how you achieved probability P —whether through gate nudges, detector nudges, or any combination—the minimum cost to *record* the outcome is $\kappa \cdot H(P)$ ticks. This is a floor that no technology can beat.

As $P \rightarrow 1$, outcome entropy $H(P) \rightarrow 0$, so recording cost decreases. But achieving $P \rightarrow 1$ requires either:

- Gate nudges rotating the state close to an eigenstate, or
- Detector nudges enforcing asymmetry $\delta \gg 0$, requiring additional irreversible stabilization

The catch: While record cost decreases as outcomes become more certain, the *control cost* to achieve that certainty increases. There's no free lunch.

11.5 Control-Ledger Extension

For general readers: To fully account for the cost of biasing outcomes, we need to track not just the cost of recording, but the cost of maintaining the bias itself. This subsection introduces a "control cost" term that captures this additional overhead.

To capture the cost of maintaining bias, extend the ledger:

$$L(\Pi) = (T, \Delta I_{\text{fact}}, \Delta I_{\text{export}}; \Delta I_{\text{ctrl}})$$

where ΔI_{ctrl} measures irreversible resources to maintain bias parameter δ stably (calibration, active stabilization, dissipation).

Theorem 11.1 (Control-Maintenance Cost, scaling form). Consider a biased detector whose effective outcome-channel asymmetry is parameterized by $\delta = \log(\kappa_0/\kappa_1)$. Assume the detector is maintained in a nonequilibrium steady state against a thermal environment at temperature T_{bath} , with a characteristic relaxation time τ (apparatus-dependent). Then sustaining a nonzero $|\delta|$ over time requires a nonzero entropy production rate \dot{S}_{ctrl} satisfying a scaling lower bound

$$\dot{S}_{\text{ctrl}} \geq c|\delta|$$

for some $c > 0$ depending on (T_{bath}, τ) and detector microphysics. Equivalently, the integrated control export over one cycle obeys

$$\Delta I_{ctrl} \geq c' |\delta|$$

for an apparatus-dependent constant $c' > 0$.

Proof sketch (stochastic thermodynamics). Model the detector bias as maintaining a stationary distribution over two metastable configurations with log-odds δ . The free-energy difference associated with enforcing this log-odds against a thermal bath is $\Delta F \sim k_B T_{bath} |\delta|$. Maintaining a nonequilibrium steady state with $\Delta F \neq 0$ requires continuous housekeeping heat dissipation (Hatano–Sasa / Seifert framework). For fixed relaxation time τ , the minimum housekeeping entropy production rate scales as $\dot{S}_{ctrl} \sim \Delta F / (T_{bath} \tau) \propto |\delta| / \tau$, yielding the stated linear-in- $|\delta|$ scaling with constant $c \sim k_B / \tau$. \square

What this means: The cost of maintaining detector bias grows at least linearly with the strength of the bias. Doubling the bias (in log-odds) at least doubles the control cost. This is a fundamental constraint grounded in nonequilibrium thermodynamics, not an engineering limitation.

Combined with record/reset cost:

$$T_{cycle}(P) \geq 2\kappa H(P) + \kappa c |\delta|$$

This quantifies the trade-off:

- **Gate-only strategy:** Maximize Λ_g by state steering (cheap), leaving $\delta \approx 0$
- **Detector-assisted strategy:** Use $\delta \neq 0$ to reach target P , paying extra irreversible overhead

Practical implication: This explains why quantum computing emphasizes precise gate control rather than detector engineering. Gate nudges are fundamentally cheaper because they operate in the reversible domain. Detector nudges seem tempting (just make the detector favor the answer you want!) but PAF reveals they carry hidden irreversible costs.

11.6 Worked Numbers

For general readers: Let's see what these abstract principles mean in concrete numbers. How much detector asymmetry do you need to achieve various success probabilities, starting from a fair quantum coin?

Suppose the input state is unbiased ($p = 1/2$) and we want target success probability P . Then $\Lambda_g = 0$ and $\delta = \Lambda(P)$. The required rate asymmetry is:

$$\kappa_0 / \kappa_1 = P / (1 - P)$$

Target P Rate Asymmetry κ_0 / κ_1 Detector Bias δ Record Cost $H(P)$

0.50	1	0	1.00 bit
------	---	---	----------

Target P Rate Asymmetry κ_0/κ_1 Detector Bias δ Record Cost $H(P)$

0.90	9	2.20	0.47 bit
0.99	99	4.60	0.08 bit
0.999	999	6.91	0.01 bit

Figure 2a:

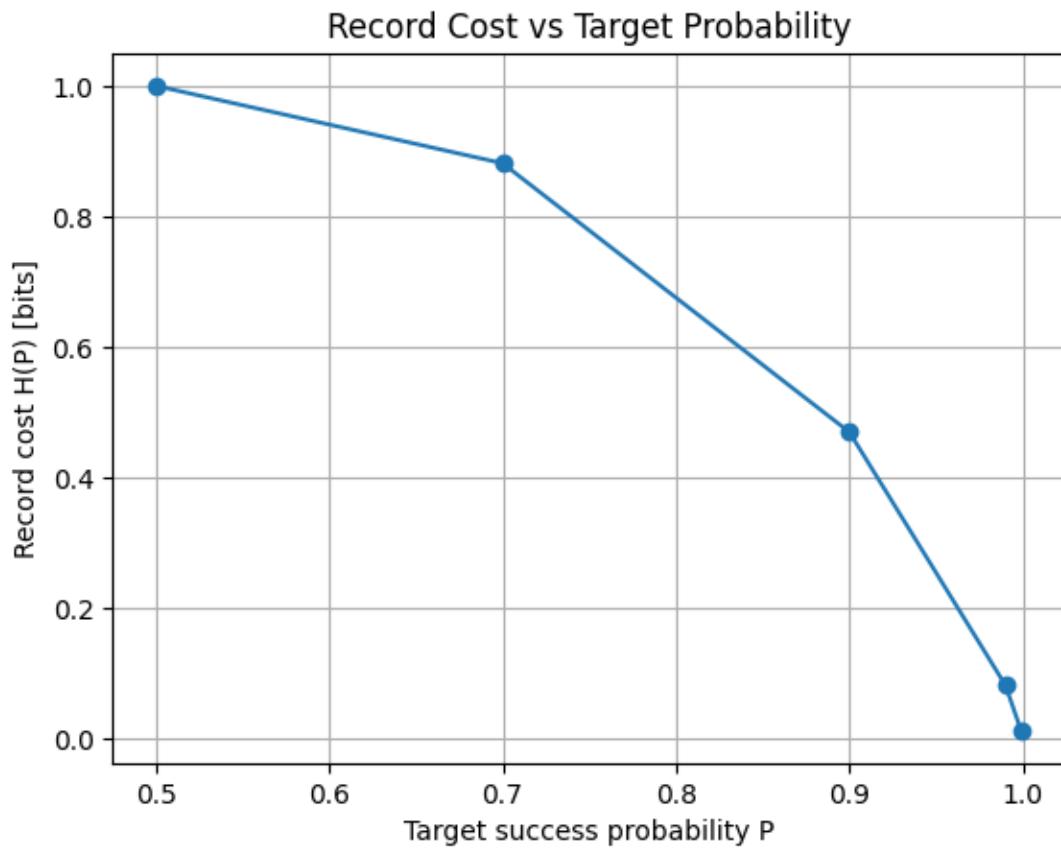


Figure 2b:

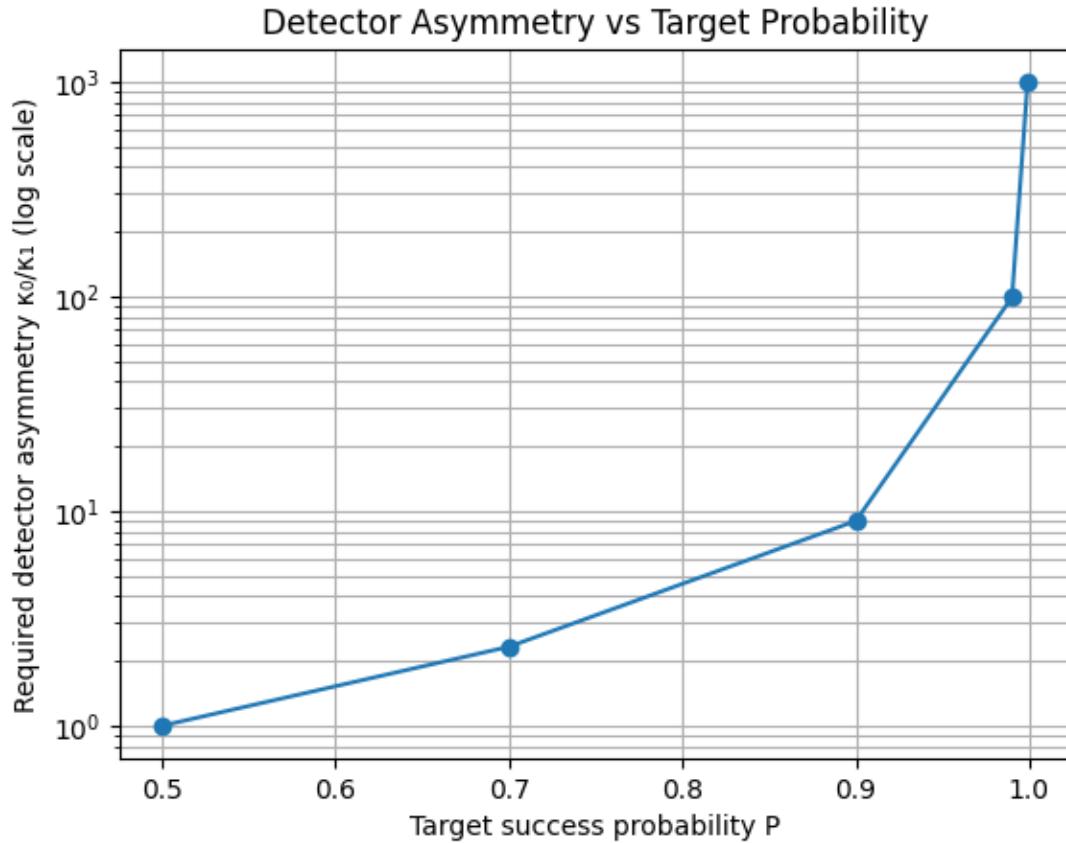


Figure 2(a): Record Cost vs Target Probability

- The vertical axis shows the minimum record entropy $H(P)$, which sets the irreducible commitment cost $T \geq \kappa H(P)$.
- As outcomes become more certain ($P \rightarrow 1$), the information content of the record drops rapidly.
- Recording a near-certain outcome is cheap because little information is learned.

Interpretation:

The cost of *recording* an outcome depends only on its surprise. High certainty means low record cost.

Figure 2(b): Detector Asymmetry vs Target Probability

- The vertical axis shows the required detector asymmetry κ_0/κ_1 on a log scale.
- Achieving modest bias (e.g. $P = 0.9$) already requires order-of-magnitude asymmetry.
- Pushing toward near determinism ($P \rightarrow 1$) causes the required asymmetry to diverge.

Interpretation:

While *recording* a certain outcome is cheap, *forcing* that certainty via detector bias is expensive. The control effort needed to maintain detector asymmetry grows without bound.

Key takeaway (what the pictures show immediately)

- Left plot: Record cost $H(P) \rightarrow 0$ as $P \rightarrow 1$.
- Right plot: Detector asymmetry $\kappa_0/\kappa_1 \rightarrow \infty$ as $P \rightarrow 1$.

This visualizes the central PAF message:

Certainty is cheap to record but expensive to enforce irreversibly.

Gate nudges achieve high P in the reversible sector and pay only the left-panel cost. Detector nudges pay the right-panel cost as well.

Reading the table:

- To get 90% success from a fair coin using only detector bias, you need one detector channel to be 9× more responsive than the other.
- To get 99% success, you need 99× asymmetry—an order of magnitude harder.
- To get 99.9% success, you need 999× asymmetry—another order of magnitude.

The record cost ($H(P)$) drops as certainty increases, but the detector bias (δ) grows logarithmically. Since control cost scales with $|\delta|$, pushing to ever-higher certainty via detector nudges becomes progressively more expensive.

Contrast with gate nudges: If you instead use gate nudges to achieve $P = 0.99$, you simply rotate the quantum state to $p = 0.99$. This costs zero ticks (it's reversible). You still pay $H(0.99) \approx 0.08$ bits to record the outcome, but you avoid the $\delta = 4.60$ control overhead entirely.

11.7 Summary

PAF clarifies what standard quantum mechanics leaves implicit:

- **Gate nudges** change p (Born weights) but do not create facts and do not price irreversibility. They're "free" in the PAF sense—like rearranging pieces before the game.
- **Detector nudges** change κ (commitment rates), biasing outcome selection at the irreversible layer. They require ongoing resources to maintain—like holding a door open.

- **Probability bias factorizes:** $\log(P/(1-P)) = \log(p/(1-p)) + \log(\kappa_0/\kappa_1)$. Gate and detector contributions add independently.
- **Record cost is unavoidable:** $\sim \kappa H(P)$ per trial, $\sim 2\kappa H(P)$ per cycle. This floor exists regardless of how P was achieved.
- **Strong detector bias requires irreversible stabilization:** captured by $\Delta I_{ctrl} \geq c|\delta|$

The bottom line: While reversible control (gate nudges) can steer outcomes efficiently, pushing toward near-determinism via irreversible means (detector nudges) encounters hard overheads. PAF explains why quantum engineers focus on precise unitary gates rather than trying to "rig" detectors: it's not just easier—it's fundamentally cheaper in the currency of irreversibility.

12. Interpretational Admissibility

For general readers: Physicists have debated for a century what quantum mechanics *means*—what it tells us about reality. Different "interpretations" offer different answers: Are there many parallel worlds? Are there hidden variables we can't see? Does observation create reality? PAF doesn't answer these questions directly, but it constrains what answers are *physically possible* by requiring that any interpretation respect the cost of creating facts.

PAF functions as a constraint theory that excludes certain interpretations of quantum mechanics while remaining neutral among admissible alternatives.

Inadmissible under PAF:

- Everettian formulations that posit unlimited, cost-free branch distinguishability
- Infinite-precision hidden-variable models (e.g., Bohmian mechanics with exact particle positions having physical consequences)
- Classical phase-space realism requiring continuous state resolution

Why these fail: Each requires either unlimited information capacity (violating BCB) or cost-free fact creation (violating TPB).

Sharpening the Everett critique: PAF distinguishes two versions of many-worlds:

- *Everett with ontologically real branches but operationally costly access* — ADMISSIBLE. Branches exist, but determining "which branch am I in?" requires a physical measurement with ledger cost. Decoherence (Section 8) automatically exports the required distinguishability. This version is consistent with PAF.
- *Everett with free operational access to branch information* — INADMISSIBLE. If observers could freely access branch identity without ledger cost, this would violate TPB. The branches would be "facts for free."

A committed Everettian might respond: "Decoherence already pays the ledger cost—there's no additional branching cost." This is correct, and PAF agrees: the admissible version of Everett is one where decoherence *is* the commitment process, with all its associated costs. What PAF excludes is the notion that branching is ontologically "cheap" or that branch counting is physically meaningful without operational cost.

Clarification: PAF is agnostic about ontology; it constrains only the physical cost of accessing branch information. Whether branches "really exist" is a question PAF does not address. What PAF addresses is whether distinguishing branches—operationally determining "which branch am I in?"—can be done for free. It cannot.

Admissible under PAF:

- Operational Everett interpretations where branch records require irreversible commitment
- Relational quantum mechanics with finite observer distinguishability
- QBist and other epistemic interpretations that treat probabilities as constraints on rational agents rather than objective frequencies

Why these survive: Each can be formulated consistently with finite distinguishability and costly fact creation. The branches exist, but distinguishing them requires resources. The hidden variables exist, but their precision has physical limits. The probabilities are real, but they describe limits on what agents can know given finite resources.

PAF does not determine interpretation but constrains what any admissible interpretation can assume about facts, records, and distinguishability.

13. Experimental Predictions and Testability

PAF makes specific, testable predictions about how irreversible costs scale with task parameters. While the framework does not predict absolute values (which depend on the undetermined constant κ), it predicts *scaling relationships* that can be tested against alternatives.

13.1 Scaling Predictions

Theorem 13.1 (κ -independent scaling exponents). Let \mathcal{C} be any admissible commitment architecture with finite $\kappa \in (0, \infty)$. Then the following scaling exponents are universal (independent of κ):

- (i) $T = \Omega(\log N)$ for N -hypothesis discrimination at fixed error;
- (ii) $T = \Omega(m)$ for syndrome extraction with m committed bits;

(iii) $T_{ctrl} = \Omega(|\delta|)$ for maintaining detector log-asymmetry δ .

The constant κ sets only the multiplicative prefactor.

Proof. Each bound follows from the corresponding theorem: (i) from Theorem 6.1 applied to $\log_2(N)$ bits of committed record; (ii) from Theorem 5.1 with $\Delta I_{fact} = m$; (iii) from Theorem 11.1. Since each bound has the form $T \geq \kappa \cdot f(\text{parameters})$ for some function f independent of κ , the scaling exponent of f is universal. \square

Measurement discrimination: For N -hypothesis discrimination with error probability ε :

$$T \geq \kappa \cdot [\log_2(N) - h(\varepsilon)]$$

This predicts that minimum measurement time scales logarithmically with hypothesis count. Doubling the number of distinguishable states requires one additional tick's worth of commitment time, not twice the time.

Error correction: For stabilizer codes with m syndrome bits per round:

$$T_{cycle} \geq 2\kappa m$$

Overhead scales linearly with syndrome count, not with code block size or with the number of correctable errors directly.

Detector bias maintenance: For detector asymmetry $\delta = \log(\kappa_0/\kappa_1)$:

$$\Delta I_{ctrl} \geq c|\delta|$$

Maintaining 99% vs 90% success (from a fair state) requires roughly twice the control overhead, since $\delta(0.99) \approx 2 \cdot \delta(0.90)$.

13.2 Potential Experimental Tests

Superconducting qubit readout: Modern superconducting qubits achieve single-shot readout in $\sim 100\text{-}500$ ns. PAF predicts that readout time cannot be reduced below $\kappa \cdot H(P)$ for outcome probability P . Systematic studies of minimum readout time versus measurement fidelity could constrain κ .

Ion trap detection: Trapped-ion systems achieve near-unit-fidelity readout through fluorescence detection. The photon collection time sets a commitment bound. PAF predicts this time scales logarithmically with the number of distinguishable states in a multi-level ion.

Quantum error correction overhead: As experimental QEC improves, PAF predicts that the ratio of correction-cycle time to unitary-gate time will grow with code distance, with a specific functional form determined by syndrome count.

13.3 What Would Falsify PAF?

PAF would be falsified by:

- A measurement protocol achieving reliable N -hypothesis discrimination in time sub-logarithmic in N
- Error correction with overhead sub-linear in syndrome count
- A demonstration that detector bias can be maintained at cost sub-linear in log-asymmetry
- Any process that creates stable classical records without irreversible export

Concrete thought experiment: Suppose someone builds a "free-fact machine"—a device that produces N distinguishable, copyable classical outcomes in constant time T_0 independent of N , with no entropy export to the environment. Running this device k times would produce $k \cdot \log_2(N)$ bits of stable record at cost $k \cdot T_0$, violating $T \geq \kappa \cdot \Delta I_{\text{fact}}$ for sufficiently large N . If such a device were demonstrated, PAF would be refuted. No such device has ever been constructed or credibly proposed.

These are strong claims. PAF's value lies precisely in making falsifiable predictions that standard quantum mechanics does not.

14. Future Directions

14.1 Emergent Geometry from Ledger Flow

For general readers: One of physics' deepest mysteries is why space and time have the geometric structure described by Einstein's general relativity. PAF suggests a radical possibility: spacetime geometry isn't fundamental but emerges from the accounting of irreversible processes—the same ledger that tracks fact creation.

Status: This direction is exploratory and intended to motivate future work rather than claim a completed derivation.

In large composite systems, coarse-grained flows of ΔI_{export} define effective geometric responses. We conjecture that spacetime curvature arises as a statistical bookkeeping of irreversible commitment imbalance across regions.

The intuition: when irreversible processes are spatially distributed, the ledger must track where distinguishability is created and where it is exported. Persistent imbalances—regions that systematically export more than they commit, or vice versa—manifest as effective geometric structure when viewed at scales much larger than individual commitment events.

General relativity, in this view, emerges as a hydrodynamic limit of admissibility constraints, not a fundamental dynamical law. The detailed derivation of Einstein's equations from ledger flow statistics remains an open problem but represents a natural extension of the framework.

14.2 Interfaces and Holographic Limits

When bulk distinguishability is suppressed—as in atomic interiors or sub-Planck domains—irreversible ledger export must occur through boundary-like degrees of freedom. These constitute **interface regimes**: regions where information processing is confined to boundaries rather than distributed through volume.

This explains area-limited information capacity (as in black hole entropy bounds) without postulating spacetime microstructure. The holographic scaling $S \sim A$ emerges because commitment events are constrained to occur at boundaries when bulk degrees of freedom are admissibility-suppressed.

14.3 Determination of κ

The constant κ remains undetermined in the present work. This is the paper's most significant open question: without κ , PAF makes only *ordinal* predictions (X costs more than Y) rather than *cardinal* predictions (X costs 3.7 units). However, this is not necessarily a weakness—many successful physical theories began as scaling theories before absolute constants were determined (thermodynamics before Boltzmann's constant was measured, for instance).

Possible approaches to constraining κ :

Experimental measurement: Systematic studies of minimum detection times in well-characterized quantum systems could constrain κ . For superconducting qubits with \sim 100-500 ns readout times achieving $H(P) \approx 1$ bit, one obtains $\kappa \lesssim 100$ -500 ns/bit. More precise bounds would require controlled variation of $H(P)$ while measuring minimum achievable readout time.

Theoretical derivation: If PAF is embedded in a deeper framework, κ might be derivable. One candidate approach combines Landauer's principle with quantum speed limits (Margolus-Levitin): the minimum time to distinguish orthogonal states is $\tau \geq \pi\hbar/(2\Delta E)$, while Landauer requires energy dissipation $W \geq k_B T \ln(2)$ per bit. Combining these yields a lower bound $\kappa \geq \pi\hbar/(2\Delta E)$ per bit, but this depends on the energy scale ΔE available to the commitment process.

Substrate dependence: It remains open whether κ is a universal constant (like \hbar or c) or substrate-dependent (like thermal relaxation times). The former would suggest κ reflects fundamental physics; the latter would suggest it characterizes classes of commitment architectures. Current evidence is insufficient to distinguish these possibilities.

Consistency requirements: Any viable κ must satisfy $\kappa \geq \tau_{\text{Landauer}}$ for the relevant temperature regime, where τ_{Landauer} is the minimum time to dissipate $k_B T \ln(2)$ of heat.

14.3.1 A Worked Example: Constraining κ from Published Superconducting-Qubit Readout Data

While PAF is fundamentally a scaling theory, κ can be numerically constrained from published experimental readout metrics. The key idea is to relate (i) a reported single-shot assignment error in an N -outcome discrimination task to (ii) the minimum number of committed record bits required to achieve that error, and then compare this to (iii) the reported readout duration.

Setup: an operational κ in time units

PAF's TPB closure inequality states that committing ΔI_{fact} bits as stable classical record requires $T \geq \kappa \cdot \Delta I_{\text{fact}}$. Here T is a logical commitment resource (ticks), whereas experiments report a physical time τ_r for readout. To connect the two without committing to a microscopic detector model, we define an operational (time-valued) effective constant:

$$\kappa_{\text{eff}} \equiv \tau_r / \Delta I_{\text{fact,min}}$$

where $\Delta I_{\text{fact,min}}$ is a lower bound on the number of committed bits required to attain the reported assignment error. Since τ_r includes all irreversible detection and integration needed to produce a stable record, κ_{eff} provides a conservative upper bound on the physical time required per committed bit in that platform:

$$\kappa_{\text{phys}} \lesssim \kappa_{\text{eff}}$$

(If additional classical post-processing time not included in τ_r is required for stability, the bound tightens further.)

Information-theoretic lower bound on committed record bits

Consider an N -class discrimination task with uniform prior over hypotheses and average misclassification probability P_e . Fano's inequality gives:

$$H(X|Y) \leq h(P_e) + P_e \log_2(N-1)$$

so the mutual information satisfies:

$$I(X;Y) \geq \log_2(N) - h(P_e) - P_e \log_2(N-1)$$

In PAF, producing a stable outcome label with error probability P_e requires committing at least this many bits:

$$\Delta I_{\text{fact,min}} \geq \log_2(N) - h(P_e) - P_e \log_2(N-1)$$

For $N = 2$, this reduces to $\Delta I_{\text{fact,min}} \geq 1 - h(P_e)$.

Example dataset: transmon readout at 140 ns

Chen et al. [70] report single-shot transmon readout with $\tau_r = 140$ ns and two-state assignment fidelity $F = 99.5\%$ (and three-state fidelity $F = 96.9\%$) without a quantum-limited amplifier. We interpret $P_e = 1 - F$ as the average classification error for the stated discrimination task.

(i) *Two-state discrimination ($N = 2$)*

Reported: $F = 0.995 \Rightarrow P_e = 0.005, \tau_r = 140$ ns.

For $N = 2$: $\Delta I_{\text{fact,min}} \geq 1 - h(0.005) \approx 0.9546$ bits

Thus: $\kappa_{\text{eff}} = 140$ ns / 0.9546 $\approx 1.47 \times 10^{-7}$ s/bit ≈ 147 ns/bit

Interpretation: In this platform and protocol, the experimental record production rate is consistent with a bound of order $\sim 10^2$ ns per committed bit (up to constant factors associated with the mapping between physical time and the tick resource).

(ii) *Three-state discrimination ($N = 3$)*

Reported: $F = 0.969 \Rightarrow P_e = 0.031, \tau_r = 140$ ns.

Using Fano for $N = 3$: $\Delta I_{\text{fact,min}} \geq \log_2(3) - h(0.031) - 0.031 \cdot \log_2(2) \approx 1.3546$ bits

Therefore: $\kappa_{\text{eff}} \approx 140$ ns / 1.3546 $\approx 1.03 \times 10^{-7}$ s/bit ≈ 103 ns/bit

Interpretation: The effective time-per-committed-bit is again $O(10^2)$ ns/bit for this published multi-level discrimination.

Cross-check: earlier rapid dispersive readout results

As an independent comparison, Walter et al. [71] report superconducting-qubit dispersive readout fidelities of 99.2% in 88 ns (fidelity-optimized) and 98.25% in 48 ns (time-optimized). Applying the same $N = 2$ bound gives $\Delta I_{\text{fact,min}} = 1 - h(P_e)$ with $P_e \in \{0.008, 0.0175\}$, yielding κ_{eff} again on the scale of $\sim 10^2$ ns/bit.

Scope and caveats (referee-facing clarity):

- κ vs κ_{eff} : PAF's κ is defined at the level of admissibility bookkeeping (ticks per committed bit). The conversion from ticks to seconds is apparatus-dependent. The quantity κ_{eff} is an operational physical bound ("seconds per committed bit") inferred from published readout durations.
- *Uniform prior assumption:* If the experimental protocol uses non-uniform priors, $\log_2(N)$ should be replaced by $H(X)$. The calculation above is conservative for the common symmetric benchmarking case.
- *What this demonstrates:* Even without a microscopic model of detection, published readout time + assignment fidelity suffice to produce quantitative bounds on an

operational κ -scale. This shows PAF can be constrained empirically, not merely discussed qualitatively.

14.4 Connection to Quantum Gravity

PAF's prediction that geometry emerges from ledger flows suggests connections to approaches that derive spacetime from entanglement (ER=EPR, tensor networks) or from thermodynamic/entropic arguments (Jacobson, Verlinde, Padmanabhan). A systematic comparison of these programs is beyond present scope but represents an important direction.

15. The Admissibility–Algebraic Spine

For general readers: This section reveals something remarkable: the seemingly separate pieces of quantum mechanics—states, evolution, measurement, probabilities—are not independent ingredients but different views of a single mathematical structure. It's like discovering that electricity, magnetism, and light are all aspects of one underlying electromagnetic field.

This "algebraic spine" is a mathematical framework (C^* -algebras) that physicists have known about for decades but is not widely taught. PAF fits naturally into this framework as a "resource grading"—a way of assigning costs to different operations based on whether they create facts or not.

Why this matters: If quantum mechanics were just a collection of separate postulates, adding PAF would seem arbitrary—bolting on an extra rule. But if quantum mechanics is already unified in a deeper algebraic structure, then PAF emerges naturally as the accounting system for that structure. Reversible operations (automorphisms) have zero cost; irreversible operations (CP maps) have nonzero cost. The math was already set up for this distinction—PAF just reads off the price tags.

The preceding sections have treated the formal components of quantum mechanics—Hilbert space, Hamiltonian dynamics, probability assignments, and measurement—largely in their conventional guises, while introducing an admissibility layer that constrains their physical realization. In this section, we show that these elements are not independent structures but manifestations of a single underlying mathematical framework: an algebraic theory of states, transformations, and irreversible channels equipped with admissibility constraints.

This formulation is useful because it makes reversible evolution, measurement, and irreversible cost different faces of one operational structure: automorphisms versus CP maps, with PAF as a resource grading. The framework provides a unifying "spine" from which the Schrödinger equation, Hilbert space representation, Hamiltonian generators, Born probabilities, and measurement update rules all arise as representational consequences rather than axioms. PAF then appears naturally as a resource-theoretic grading on this spine, pricing the irreversible creation of facts.

15.1 Algebra of Observables as the Primitive Object

Rather than taking Hilbert space as fundamental, we begin with an abstract algebra of observables \mathcal{A} , representing all physically meaningful questions that can be asked of a system. Mathematically, \mathcal{A} is taken to be a C*-algebra (or, in appropriate limits, a von Neumann algebra), whose elements encode operationally accessible observables.

This move shifts the ontology from "states evolving in Hilbert space" to distinguishability structure encoded in an algebra. Importantly, the algebra itself does not presuppose probabilities, dynamics, or measurement outcomes—it only specifies which distinctions are meaningful in principle.

Finite distinguishability, as required by admissibility, constrains the physically realizable subalgebras of \mathcal{A} . Infinite-resolution algebras may exist mathematically, but admissibility restricts operational access to finite substructures, consistent with the Bit Conservation Bound.

15.2 States as Functionals: The Origin of the Born Rule

In the algebraic framework, a state is defined as a positive, normalized linear functional

$$\omega : \mathcal{A} \rightarrow \mathbb{R}$$

which assigns expectation values to observables. This definition is purely operational: a state encodes all experimentally accessible statistics without reference to an underlying wavefunction.

Probabilities arise when evaluating the state on effects (positive elements of the algebra corresponding to measurement outcomes). The assignment

$$p(E) = \omega(E)$$

is not an additional postulate but the definition of probability in this setting.

When represented on a Hilbert space (see §15.4), this reduces to the familiar Born rule $p(E) = \text{Tr}(\rho E)$. Thus, the Born rule is not a separate law but the unique way a positive linear functional assigns weights to effects. Its uniqueness is enforced by admissibility through operational non-contextuality and finite distinguishability.

15.3 Dynamics as Automorphisms: Unitary Evolution Recovered

Reversible time evolution is represented abstractly as a one-parameter group of automorphisms of the algebra,

$$\alpha_t : \mathcal{A} \rightarrow \mathcal{A}$$

preserving the algebraic structure and all distinguishability relations.

Admissibility requires that reversible evolution neither creates nor destroys accessible distinguishability; thus, it must act isometrically on the state space. Continuity of α_t then guarantees the existence of a generator.

When a Hilbert-space representation is chosen, these automorphisms are implemented by unitary operators $U(t)$, and Stone's theorem yields a self-adjoint generator H . The Schrödinger equation therefore appears as the representational form of algebraic automorphism flow, not as a primitive dynamical law.

In this way, Hilbert space, Hamiltonians, and Schrödinger evolution are unified as consequences of a single algebraic dynamical structure constrained by admissibility.

15.4 Hilbert Space as a Representation, Not a Foundation

Theorem 15.1 (Gelfand–Naimark–Segal). For any C^* -algebra \mathcal{A} and any state $\omega : \mathcal{A} \rightarrow \mathbb{C}$, there exists a Hilbert space \mathcal{H}_ω , a $*$ -representation $\pi_\omega : \mathcal{A} \rightarrow B(\mathcal{H}_\omega)$, and a cyclic vector $\Omega_\omega \in \mathcal{H}_\omega$ such that

$$\omega(A) = \langle \Omega_\omega, \pi_\omega(A) \Omega_\omega \rangle$$

for all $A \in \mathcal{A}$.

Thus, Hilbert space is not assumed but emerges as a representation of the algebra relative to a state. Different states may yield different Hilbert-space representations of the same underlying algebra, reinforcing the idea that Hilbert space is a mathematical convenience adapted to operational context, not fundamental ontology.

This perspective aligns naturally with PAF: admissibility constrains which representations are physically realizable, while the algebraic spine remains invariant.

15.5 Measurement and Irreversibility as Completely Positive Maps

Irreversible processes—including measurement, amplification, readout, and reset—are represented in the algebraic framework by completely positive (CP) maps,

$$\Phi : \mathcal{A} \rightarrow \mathcal{A}$$

which need not be invertible. CP maps capture exactly the structure required for physical irreversibility while remaining compatible with entangled extensions.

Measurement instruments correspond to families of CP maps indexed by outcomes, and state update is simply conditioning of the functional ω on the selected channel. No additional collapse postulate is required.

Within PAF, these CP maps are precisely the operations that carry ledger cost: they consume ticks, convert accessible distinguishability into stable fact, and export competing alternatives to inaccessible degrees of freedom. The algebraic framework specifies what transformations are allowed; PAF specifies what they cost.

15.6 The Distinguishability Ledger as a Resource Grading

Definition (Resource grading on CP maps). A resource grading is a function $R : CP(\mathcal{A}) \rightarrow \mathbb{R}_+$ such that:

- (i) $R(\alpha) = 0$ for any $*$ -automorphism α ;
- (ii) $R(\Phi \circ \Psi) \leq R(\Phi) + R(\Psi)$ (subadditivity under composition);
- (iii) $R(\Phi) > 0$ for any Φ that increases stable classical record entropy (fact creation).

PAF supplies such a grading with $R(\Phi) = T(\Phi)$ (or equivalently $R(\Phi) = \kappa \cdot \Delta I_{\text{fact}}(\Phi)$ depending on normalization).

The Distinguishability Ledger introduced earlier,

$$L(\Pi) = (T, \Delta I_{\text{fact}}, \Delta I_{\text{export}})$$

can now be understood as a vector-valued resource grading on CP maps within the algebraic spine:

- Automorphisms (reversible dynamics) have zero ledger cost: $L(\alpha) = (0, 0, 0)$
- CP maps that produce records incur nonzero T and ΔI_{export}
- Composition of maps satisfies lax additivity (Section 4)

This parallels how thermodynamics grades mechanical processes by work and heat without altering Newtonian dynamics. PAF similarly grades algebraic quantum processes by admissibility cost without modifying their formal structure.

15.7 A Single Spine, Not Separate Postulates

From this perspective, the traditional components of quantum mechanics are unified:

- **Hilbert space:** a representation of an observable algebra
- **Hamiltonians:** generators of algebra automorphisms
- **Schrödinger evolution:** representational form of reversible algebraic flow
- **Born rule:** evaluation of states on effects
- **Measurement:** CP maps with irreversible ledger cost

All are facets of a single algebraic framework constrained by admissibility.

PAF does not introduce new dynamics or probabilities. It completes the theory by specifying which algebraic processes are physically admissible and what irreversible resources they consume. Quantum mechanics thus appears as the minimal reversible core of a broader admissibility-constrained theory of fact production.

15.8 Conceptual Summary

The Physical Admissibility Framework does not add a fifth postulate to quantum mechanics. It reveals that the familiar postulates already sit inside a larger mathematical structure: an algebraic theory of distinguishability, with reversible automorphisms and irreversible CP maps, graded by the finite cost of making facts real.

In this sense, Schrödinger evolution, Hilbert space, Hamiltonians, and the Born rule are not separate ingredients to be unified—they are already unified, once one adopts the correct mathematical spine and imposes the physical requirement that facts exist and are costly.

16. Conclusions

For general readers: The Physical Admissibility Framework (PAF) addresses a gap in standard quantum mechanics: while QM specifies what states exist and how likely various outcomes are, it does not account for the physical cost of making outcomes definite. PAF fills this gap by treating fact creation as a resource-consuming process subject to conservation constraints.

The core insight is simple: *facts are not free*. Every definite outcome requires irreversible physical resources to create.

PAF accounts for:

- Why quantum mechanics is unavoidable (it is the minimal dynamics compatible with finite distinguishability)
- Why measurement dominates experimental cost (commitment is irreducibly expensive)
- Why fault tolerance requires physical overhead (error correction commits information)
- Why precision has irreversible cost (certification requires commitment)
- Why geometry emerges only after coarse-graining (spacetime is ledger statistics)

As shown in Section 15, PAF does not introduce new dynamics or probability rules. The familiar postulates of quantum mechanics—Hilbert space, Hamiltonians, the Born rule, and measurement update—are already unified within the algebraic spine of C^* -algebras, automorphisms, and completely positive maps. PAF completes this structure by specifying which processes are physically admissible and what irreversible resources they consume.

Future work should address: (a) experimental tests of scaling predictions (Section 13), (b) determination of κ from first principles or experiment (Section 14.3), (c) the emergent geometry

program (Section 14.1), and (d) deeper connections to quantum gravity and holography (Section 14.4).

PAF provides a unifying admissibility layer for quantum mechanics, thermodynamics, and spacetime physics. It does not replace the standard formalism but completes it by specifying what that formalism cannot: the cost of making quantum possibilities into physical facts.

Note: PAF is compatible with and motivated by a broader theoretical program (VERSF) that treats finite distinguishability and emergent time as fundamental. The present paper develops PAF as an independently motivated constraint framework; connections to the broader program will be explored in subsequent work.

Appendix A: Mathematical Formalization

A.1 Accessible Distinguishability

Let a physical system be described by a density operator ρ acting on Hilbert space \mathcal{H} . Given an admissible measurement set \mathcal{M} (POVMs realizable under finite precision and finite resources), define:

$$I_{\text{acc}}(\rho) = \sup_{\mathcal{M}} \{M \in \mathcal{M}\} I(X; Y)$$

where $I(X; Y)$ is the classical mutual information between preparation variable X and measurement outcome Y . This definition aligns with standard accessible information bounds but is explicitly constrained to physically admissible measurement operations.

A.2 Relation to the Holevo Bound

For an ensemble $\{p_x, \rho_x\}$, the Holevo quantity χ provides an upper bound:

$$I_{\text{acc}} \leq \chi = S(\rho) - \sum_x p_x S(\rho_x)$$

where $S(\rho) = -\text{Tr}(\rho \log_2 \rho)$ is the von Neumann entropy.

PAF does not modify the Holevo bound; instead, it constrains the physical realization of χ . Even when χ is large, committing this information as stable fact requires irreversible ledger cost. Thus χ bounds what may be extracted, while PAF bounds what may be made real.

This distinction explains why the Holevo bound is not always saturable in practice: achieving it requires measurement protocols whose commitment cost may exceed available resources.

A.3 Commitment Channels

A commitment operation is modeled as a CPTP map:

$$\mathcal{C}(\rho) = \sum_y \text{Tr}(E_y \rho) \cdot |y\rangle\langle y| \otimes \rho_y$$

where:

- $\{|y\rangle\}$ are orthogonal stable record states (the classical outcomes)
- $\{E_y\}$ forms a POVM (the measurement operators)
- $\rho_y = (\sqrt{E_y} \rho \sqrt{E_y}) / \text{Tr}(E_y \rho)$ is the post-measurement state of the system conditional on outcome y (for projective measurements, $\rho_y = E_y \rho E_y / \text{Tr}(E_y \rho)$)

Such channels are necessarily non-unitary and involve irreversible discard of alternative branches. The tensor product structure explicitly separates the classical record (left factor) from the updated quantum state (right factor).

PAF asserts that any such channel incurs tick cost and entropy export proportional to the classical entropy of outcomes:

$$T \geq \kappa \cdot H(Y) \Delta I_{\text{export}} \geq H(Y)$$

where $H(Y) = -\sum_y p(y) \log_2 p(y)$ is the Shannon entropy of the outcome distribution.

A.4 Commitment Depth and TPB Bound

Define the commitment depth $D_C(\Pi)$ of a process Π as the minimum number of non-unitary commitment channels required to implement Π . By Theorem 5.1, for any process producing ΔI_{fact} bits of stable classical record:

$$D_C(\Pi) \geq \Delta I_{\text{fact}} / I_{\text{max}}$$

where I_{max} is the per-commitment record capacity. If $I_{\text{max}} = 1/\kappa$ bits per commitment, then $D_C(\Pi) \geq \kappa \cdot \Delta I_{\text{fact}}$.

This establishes the Ticks-Per-Bit constraint as a depth lower bound analogous to circuit depth in computational complexity. The bound is tight: there exist protocols achieving $D_C = \Theta(\Delta I_{\text{fact}})$.

A.5 Ledger Closure Inequalities

For any admissible process Π , the distinguishability ledger $L(\Pi) = (T, \Delta I_{\text{fact}}, \Delta I_{\text{export}})$ satisfies:

$$T \geq \kappa \cdot \Delta I_{\text{fact}} \Delta I_{\text{export}} \geq \max(0, \Delta I_{\text{fact}} - \Delta I_{\text{recovered}})$$

These inequalities encode the Bit Conservation Bound and ensure that irreversible fact creation cannot occur without compensating export of distinguishability. The quantity $\Delta I_{\text{recovered}}$ represents information that returns to accessible form (e.g., through environmental feedback or error correction), which is typically negligible for well-isolated records. The max ensures ΔI_{export} remains non-negative.

A.6 Decoherence as Information Export

Consider a system-environment interaction U_{SE} producing decoherence in pointer basis $\{|i\rangle\}$:

$$U_{\text{SE}} : |i\rangle_S \otimes |0\rangle_E \rightarrow |i\rangle_S \otimes |e_i\rangle_E$$

For an initial superposition $|\psi\rangle = \sum_i c_i |i\rangle$, the reduced system state becomes diagonal:

$$\rho_S = \text{Tr}_E(U_{\text{SE}} |\psi\rangle\langle\psi| U_{\text{SE}}^\dagger) \rightarrow \sum_i |c_i|^2 |i\rangle\langle i|$$

The mutual information $I(S:E)$ between system and environment provides a lower bound on ΔI_{export} , making decoherence an explicit ledger transfer rather than information destruction.

A.7 Relation to Thermodynamic Cost

By Landauer's principle, erasure of information requires minimum work:

$$W \geq k_B T \ln(2) \cdot \Delta I_{\text{export}}$$

PAF thus links admissibility directly to thermodynamic cost. Time (ticks), entropy (export), and energy (work) expenditures are unified consequences of irreversible commitment. Landauer's principle emerges as a thermodynamic corollary of the more fundamental admissibility constraints.

A.8 Summary

This appendix formalizes the mathematical backbone of the Physical Admissibility Framework. Standard quantum information quantities (mutual information, Holevo bound, CPTP maps) are retained but embedded within an admissibility layer that enforces finite distinguishability, irreversible commitment, and ledger-based conservation. PAF therefore extends, rather than replaces, the mathematical structure of quantum mechanics.

Appendix B: Frequently Asked Questions

Q1: Is PAF just Landauer's principle restated?

No. Landauer's principle prices the erasure of information in thermodynamic systems. PAF is more general: it constrains the creation of physical facts themselves, independent of temperature, substrate, or specific thermodynamic models. Landauer's principle is a corollary of PAF applied to erasure operations, not its foundation.

Q2: Does PAF contradict standard quantum mechanics?

No. PAF leaves Hilbert space, Hamiltonian dynamics, and the Born rule intact. It adds an admissibility layer that governs irreversibility and cost—structure that standard QM deliberately abstracts away. PAF is to quantum dynamics as thermodynamics is to Newtonian mechanics: a constraint theory operating at a different level.

Q3: Does PAF rule out Many-Worlds?

PAF rules out Everettian formulations that assume unlimited distinguishability and cost-free branching. Operational Everett-style views remain admissible provided that facts, records, and observer experiences require irreversible commitment. The question "which branch am I in?" must be answered by a physical process with ledger cost.

Q4: Is PAF just an interpretation of quantum mechanics?

No. PAF is a constraint theory, not an interpretation. It specifies what is physically possible or impossible regardless of one's preferred interpretation, much like no-cloning or Bell inequalities. Different interpretations may remain viable under PAF, but all must respect admissibility constraints.

Q5: Does PAF require new physics or modifications to experiments?

No. PAF explains why existing experiments behave as they do and predicts scaling limits and overheads already observed in quantum engineering, fault tolerance, and metrology. It provides principled explanations for costs that are currently treated as engineering challenges.

Q6: Why introduce "ticks" rather than continuous time?

Ticks represent irreducible commitment events, not microscopic time steps. They formalize the empirical fact that irreversible operations cannot be arbitrarily subdivided or parallelized without cost. The tick count T is a logical resource measure, not a physical clock reading.

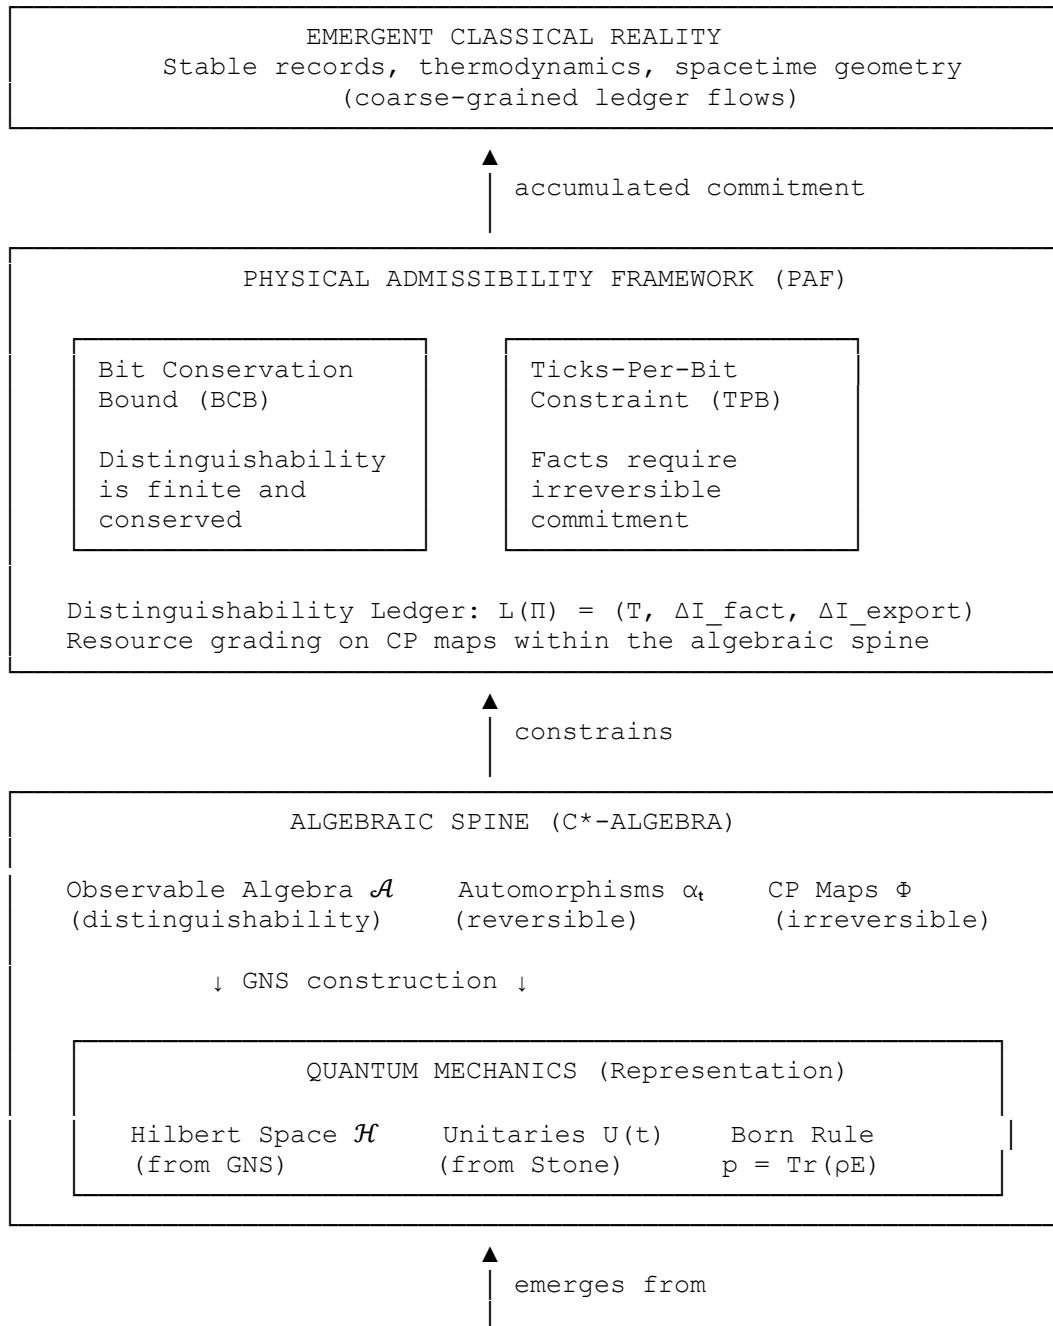
Q7: What is the value of κ ?

The constant κ governs the conversion between commitment events and committed bits. Its value may be unity in appropriately chosen units, or it may depend on physical substrate. Determining κ from first principles—or demonstrating its universality—remains an open problem. Experimental bounds may be obtainable from precision measurements of minimum detection times.

Q8: How does PAF relate to holographic bounds?

PAF provides a mechanism for holographic entropy bounds. When bulk distinguishability is suppressed (void-regulated regimes), commitment events are confined to boundary degrees of freedom, yielding $S \sim A$ scaling without postulating fundamental discreteness of spacetime.

Figure 1: Layered Structure of Physical Law



FINITE DISTINGUISHABILITY FOUNDATION

Admissibility-limited information capacity per region
(holographic bounds, area-law constraints)

Caption: The layered architecture of physical law proposed by PAF. The foundation (bottom) is finite distinguishability—the requirement that any bounded region supports only finite accessible information, consistent with holographic bounds. The algebraic spine—a C^* -algebra of observables with automorphisms and CP maps—provides the mathematical foundation; standard quantum mechanics (Hilbert space, unitaries, Born rule) emerges via GNS construction as a representation of this algebra. The Physical Admissibility Framework (PAF layer) constrains irreversible commitment through BCB and TPB, with the Distinguishability Ledger grading CP maps by their irreversible cost. Classical reality, thermodynamics, and spacetime geometry (top) emerge from accumulated ledger flows at macroscopic scales.

Appendix C: Comparison with Standard Frameworks

Framework	What It Describes	What It Cannot Address	What PAF Adds
Hilbert Space	Which states are allowed	Cost of state realization; irreversibility	Admissibility limits on realizable distinguishability
Hamiltonian Dynamics	Reversible time evolution	Irreversible fact creation; measurement cost	Separation of unitary evolution from commitment
Born Rule	Outcome probabilities	Physical cost of making outcomes real	Lower bounds on fact creation and readout
Quantum Information Theory	Information processing limits	Irreversible time and entropy cost	Ledger-based accounting of commitment and export
Thermodynamics	Equilibrium and entropy	Why measurement is thermodynamically special	Fact creation as primitive irreversible process
<i>Algebraic QM (C^*-algebras)*</i>	Unified mathematical structure	Physical realizability; resource costs	Admissibility grading on CP maps
PAF	What is physically possible and affordable	—	Prices facts, enforces no-go constraints, explains scaling limits

Appendix D: A Heuristic Lower Bound on κ from Landauer + Margolus–Levitin

D.1 Motivation

In Section 14.3 we note that κ may be constrained by combining (i) **thermodynamic irreversibility** (Landauer) and (ii) **quantum speed limits** (Margolus–Levitin). This appendix makes that connection explicit. The resulting bound is not claimed as a strict theorem of PAF (because mapping “ticks” to physical time is architecture-dependent), but it is instructive: it shows that *any* device that commits stable classical facts in a thermal environment faces a joint **energy–time** constraint.

D.2 Ingredients

D.2.1 Landauer cost per committed bit (thermodynamic lower bound)

Landauer’s principle states that erasing one bit of information in contact with a heat bath at temperature T_{bath} requires dissipating at least

$$W_L \geq k_B T_{\text{bath}} \ln 2.$$

PAF treats record creation as requiring irreversible export ΔI_{export} , and in steady operation record creation is typically paired with reset/erasure. Thus, even if one wishes to interpret Landauer strictly as an erasure bound, a repeated measurement–reset cycle implies a per-bit dissipation floor of order $k_B T \ln 2$ per committed bit per reset cycle.

Operational assumption A (Landauer-limited commitment energy).

For a sustained fact-producing device operating in a bath at temperature T_{bath} , the irreversible commitment pipeline (record creation + stabilization + eventual reset) must dissipate at least

$$E_{\text{diss}/\text{bit}} \gtrsim k_B T_{\text{bath}} \ln 2$$

per reliably committed bit of stable record.

(If the device is not Landauer-limited, then $E_{\text{diss}/\text{bit}}$ is larger and the bounds below strengthen.)

D.2.2 Margolus–Levitin quantum speed limit (time–energy lower bound)

The Margolus–Levitin bound states that the minimum time τ required for a quantum system to evolve from a state to an orthogonal state is bounded by

$$\tau \geq \frac{\pi\hbar}{2E},$$

where E is the system's average energy above its ground state during the evolution.

A stable classical record bit requires distinguishability between at least two reliably separable macrostates (e.g., pointer states) which must be physically discriminable. At minimum, producing one classical bit corresponds to implementing a transformation that separates two alternatives into (effectively) orthogonal record states.

Operational assumption B (orthogonalization per committed bit).

Each committed bit corresponds to at least one effective orthogonalization event in the record+apparatus degrees of freedom (one “commitment transition” between distinguishable record states). This is the minimal operational content of “a tick commits a distinguishable alternative.”

D.3 Combined bound: a time-per-bit lower limit

Assume a commitment event dissipates energy $E_{\text{diss/bit}}$ into the bath and that the commitment dynamics have access to an energy scale E of the same order (since maintaining a nonequilibrium record transition requires energy throughput).

Taking $E \sim E_{\text{diss/bit}}$ as a conservative identification, the Margolus–Levitin bound yields

$$\tau_{\text{bit}} \gtrsim \frac{\pi\hbar}{2E_{\text{diss/bit}}}.$$

Using Assumption A,

$$\tau_{\text{bit}} \gtrsim \frac{\pi\hbar}{2k_B T_{\text{bath}} \ln 2}.$$

This provides a **heuristic lower bound** on a physical “seconds-per-committed-bit” κ -scale:

$\kappa_{\text{phys}}(T_{\text{bath}}) \gtrsim \frac{\pi\hbar}{2k_B T_{\text{bath}} \ln 2}$

(Landauer + Margolus–Levitin).

D.4 Numerical illustration

Define

$$\kappa_{\text{ML+L}}(T) \equiv \frac{\pi\hbar}{2k_B T \ln 2}.$$

- **Room temperature ($T = 300$ K)**

$$k_B T \ln 2 \approx 2.9 \times 10^{-21} \text{ J}, \kappa_{\text{ML+L}} \approx 5.8 \times 10^{-14} \text{ s} (\sim 60 \text{ fs/bit}).$$

- **Cryogenic dilution refrigerator ($T = 10$ mK)**

$$k_B T \ln 2 \approx 9.6 \times 10^{-26} \text{ J}, \kappa_{\text{ML+L}} \approx 1.7 \times 10^{-9} \text{ s} (\sim 1.7 \text{ ns/bit}).$$

The cryogenic value is strikingly close to the $\mathcal{O}(10^2 \text{ ns/bit})$ operational κ -scale extracted from published superconducting readout data in §14.3.1, given that real devices are not Landauer-limited and contain additional integration and thresholding overhead.

D.5 Interpretation and scope (what this does and does not prove)

1. Not a strict theorem of PAF.

PAF's κ is defined as a conversion between *logical commitment resources* and committed record bits. Mapping ticks to seconds depends on architecture. The derivation above supplies a **physically motivated lower bound** on a seconds-per-bit scale in thermal environments, not a universal κ in the abstract ledger sense.

2. Why it is still instructive.

The bound demonstrates a general principle: **fact creation cannot be arbitrarily fast at fixed energy throughput**, and any fact-producing apparatus that dissipates at least $k_B T \ln 2$ per bit inherits a corresponding time floor from quantum speed limits.

3. How the bound strengthens.

If an implementation dissipates $E_{\text{diss/bit}} \gg k_B T \ln 2$, then

$$\tau_{\text{bit}} \gtrsim \frac{\pi\hbar}{2E_{\text{diss/bit}}}$$

can be substantially *smaller* than the Landauer-based estimate. Conversely, if the relevant speed limit is set by a smaller effective energy above ground (or by a variance-based

limit), the time-per-bit floor can be *larger*. In either case, the central message remains: the measurement interface is constrained by coupled energy–time–irreversibility bounds.

4. Relation to PAF’s TPB.

TPB asserts that stable record creation requires irreducible commitment events; the present derivation suggests that when such events are implemented physically in thermal environments, they naturally inherit a minimal timescale determined jointly by dissipation constraints and quantum speed limits.

References

- [1] J. von Neumann, *Mathematical Foundations of Quantum Mechanics* (Princeton University Press, Princeton, 1955).
- [2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
- [3] A. S. Holevo, "Bounds for the quantity of information transmitted by a quantum communication channel," *Probl. Inf. Transm.* **9**, 177–183 (1973).
- [4] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM J. Res. Dev.* **5**, 183–191 (1961).
- [5] C. H. Bennett, "The thermodynamics of computation—a review," *Int. J. Theor. Phys.* **21**, 905–940 (1982).
- [6] W. H. Zurek, "Decoherence, einselection, and the quantum origins of the classical," *Rev. Mod. Phys.* **75**, 715–775 (2003).
- [7] M. Schlosshauer, "Decoherence, the measurement problem, and interpretations of quantum mechanics," *Rev. Mod. Phys.* **76**, 1267–1305 (2005).
- [8] E. Joos, H. D. Zeh, C. Kiefer, D. Giulini, J. Kupsch, and I.-O. Stamatescu, *Decoherence and the Appearance of a Classical World in Quantum Theory*, 2nd ed. (Springer, Berlin, 2003).
- [9] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic Press, New York, 1976).
- [10] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (North-Holland, Amsterdam, 1982).

[11] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," *Phys. Rev. A* **52**, R2493–R2496 (1995).

[12] A. M. Steane, "Error correcting codes in quantum theory," *Phys. Rev. Lett.* **77**, 793–797 (1996).

[13] D. Gottesman, "Stabilizer codes and quantum error correction," Ph.D. thesis, California Institute of Technology (1997), arXiv:quant-ph/9705052.

[14] E. Knill, R. Laflamme, and W. H. Zurek, "Resilient quantum computation," *Science* **279**, 342–345 (1998).

[15] D. Aharonov and M. Ben-Or, "Fault-tolerant quantum computation with constant error rate," *SIAM J. Comput.* **38**, 1207–1282 (2008).

[16] V. Giovannetti, S. Lloyd, and L. Maccone, "Quantum metrology," *Phys. Rev. Lett.* **96**, 010401 (2006).

[17] V. Giovannetti, S. Lloyd, and L. Maccone, "Advances in quantum metrology," *Nat. Photonics* **5**, 222–229 (2011).

[18] S. L. Braunstein and C. M. Caves, "Statistical distance and the geometry of quantum states," *Phys. Rev. Lett.* **72**, 3439–3443 (1994).

[19] J. D. Bekenstein, "Black holes and entropy," *Phys. Rev. D* **7**, 2333–2346 (1973).

[20] S. W. Hawking, "Particle creation by black holes," *Commun. Math. Phys.* **43**, 199–220 (1975).

[21] G. 't Hooft, "Dimensional reduction in quantum gravity," in *Salamfestschrift*, edited by A. Ali, J. Ellis, and S. Randjbar-Daemi (World Scientific, Singapore, 1994), arXiv:gr-qc/9310026.

[22] L. Susskind, "The world as a hologram," *J. Math. Phys.* **36**, 6377–6396 (1995).

[23] R. Bousso, "The holographic principle," *Rev. Mod. Phys.* **74**, 825–874 (2002).

[24] J. S. Bell, "On the Einstein Podolsky Rosen paradox," *Physics* **1**, 195–200 (1964).

[25] S. Kochen and E. P. Specker, "The problem of hidden variables in quantum mechanics," *J. Math. Mech.* **17**, 59–87 (1967).

[26] D. Bohm, "A suggested interpretation of the quantum theory in terms of 'hidden' variables. I and II," *Phys. Rev.* **85**, 166–193 (1952).

[27] H. Everett III, "'Relative state' formulation of quantum mechanics," *Rev. Mod. Phys.* **29**, 454–462 (1957).

[28] D. Wallace, *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation* (Oxford University Press, Oxford, 2012).

[29] C. A. Fuchs, "QBism, the perimeter of quantum Bayesianism," arXiv:1003.5209 (2010).

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

[31] C. Rovelli, "Relational quantum mechanics," Int. J. Theor. Phys. **35**, 1637–1678 (1996).

[32] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," Phys. Rev. Lett. **75**, 1260–1263 (1995).

[33] E. P. Verlinde, "On the origin of gravity and the laws of Newton," J. High Energy Phys. **2011**, 29 (2011).

[34] T. Padmanabhan, "Thermodynamical aspects of gravity: New insights," Rep. Prog. Phys. **73**, 046901 (2010).

[35] S. Lloyd, "Ultimate physical limits to computation," Nature **406**, 1047–1054 (2000).

[36] S. Lloyd, "Computational capacity of the universe," Phys. Rev. Lett. **88**, 237901 (2002).

[37] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," Nature **299**, 802–803 (1982).

[38] D. Dieks, "Communication by EPR devices," Phys. Lett. A **92**, 271–272 (1982).

[39] A. Peres and W. K. Wootters, "Optimal detection of quantum information," Phys. Rev. Lett. **66**, 1119–1122 (1991).

[40] M. B. Plenio and V. Vitelli, "The physics of forgetting: Landauer's erasure principle and information theory," Contemp. Phys. **42**, 25–60 (2001).

[41] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, "Thermodynamics of information," Nat. Phys. **11**, 131–139 (2015).

[42] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, "Experimental verification of Landauer's principle linking information and thermodynamics," Nature **483**, 187–189 (2012).

[43] Y. Jun, M. Gavrilov, and J. Bechhoefer, "High-precision test of Landauer's principle in a feedback trap," Phys. Rev. Lett. **113**, 190601 (2014).

[44] J. P. S. Peterson, R. S. Sarthour, A. M. Souza, I. S. Oliveira, J. Goold, K. Modi, D. O. Soares-Pinto, and L. C. Céleri, "Experimental demonstration of information to energy conversion in a quantum system at the Landauer limit," *Proc. R. Soc. A* **472**, 20150813 (2016).

[45] L. Hardy, "Quantum theory from five reasonable axioms," arXiv:quant-ph/0101012 (2001).

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

[47] L. Masanes and M. P. Müller, "A derivation of quantum theory from physical requirements," *New J. Phys.* **13**, 063001 (2011).

[48] J. Barrett, "Information processing in generalized probabilistic theories," *Phys. Rev. A* **75**, 032304 (2007).

[49] R. W. Spekkens, "Evidence for the epistemic view of quantum states: A toy theory," *Phys. Rev. A* **75**, 032110 (2007).

[50] Č. Brukner, "On the quantum measurement problem," in *Quantum [Un]Speakables II*, edited by R. Bertlmann and A. Zeilinger (Springer, Cham, 2017), pp. 95–117.

[51] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed. (Springer, Berlin, 1996).

[52] H. Araki, *Mathematical Theory of Quantum Fields* (Oxford University Press, Oxford, 1999).

[53] O. Bratteli and D. W. Robinson, *Operator Algebras and Quantum Statistical Mechanics*, Vols. 1 and 2 (Springer, Berlin, 1987, 1997).

[54] I. M. Gelfand and M. A. Naimark, "On the embedding of normed rings into the ring of operators in Hilbert space," *Mat. Sb.* **12**, 197–213 (1943).

[55] I. E. Segal, "Irreducible representations of operator algebras," *Bull. Am. Math. Soc.* **53**, 73–88 (1947).

[56] K. Kraus, *States, Effects, and Operations: Fundamental Notions of Quantum Theory* (Springer, Berlin, 1983).

[57] M.-D. Choi, "Completely positive linear maps on complex matrices," *Linear Algebra Appl.* **10**, 285–290 (1975).

[58] A. S. Holevo, *Quantum Systems, Channels, Information* (de Gruyter, Berlin, 2012).

[59] M. M. Wilde, *Quantum Information Theory*, 2nd ed. (Cambridge University Press, Cambridge, 2017).

[60] B. Coecke and A. Kissinger, *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning* (Cambridge University Press, Cambridge, 2017).

[61] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, J. M. Renes, and R. W. Spekkens, "Resource theory of quantum states out of thermal equilibrium," *Phys. Rev. Lett.* **111**, 250404 (2013).

[62] M. Horodecki and J. Oppenheim, "Fundamental limitations for quantum and nanoscale thermodynamics," *Nat. Commun.* **4**, 2059 (2013).

[63] F. G. S. L. Brandão, M. Horodecki, N. H. Y. Ng, J. Oppenheim, and S. Wehner, "The second laws of quantum thermodynamics," *Proc. Natl. Acad. Sci. U.S.A.* **112**, 3275–3279 (2015).

[64] G. Gour and R. W. Spekkens, "The resource theory of quantum reference frames: Manipulations and monotones," *New J. Phys.* **10**, 033023 (2008).

[65] J. Goold, M. Huber, A. Riera, L. del Rio, and P. Skrzypczyk, "The role of quantum information in thermodynamics—a topical review," *J. Phys. A: Math. Theor.* **49**, 143001 (2016).

[66] U. Seifert, "Stochastic thermodynamics, fluctuation theorems and molecular machines," *Rep. Prog. Phys.* **75**, 126001 (2012).

[67] T. Hatano and S.-i. Sasa, "Steady-state thermodynamics of Langevin systems," *Phys. Rev. Lett.* **86**, 3463–3466 (2001).

[68] U. Seifert, "Entropy production along a stochastic trajectory and an integral fluctuation theorem," *Phys. Rev. Lett.* **95**, 040602 (2005).

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

[70] L. Chen, H.-X. Li, Y. Lu, C. W. Warren, C. J. Križan, S. Kosen, M. Rommel, S. Ahmed, A. Osman, J. Biznárová, A. Fadavi Roudsari, B. Lienber, M. Peber, I.-M. Willenberg, A. F. Kockum, J. Bylander, and P. Delsing, "Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier," *npj Quantum Inf.* **9**, 26 (2023).

[71] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, "Rapid high-fidelity single-shot dispersive readout of superconducting qubits," *Phys. Rev. Appl.* **7**, 054020 (2017).