

Om Namo Bhagavate Sri Ramanaya

The Ignorant Observer

Avidya as a Physical Limit

Aernoud Dekker

June 2026

Abstract

This paper asks whether the Advaitic notion of *avidyā*—not merely missing information, but structural self-ignorance—can be translated into a physical model of observation. The Ignorant Observer Framework (IOF) treats the observer as a physical dynamical system with finite effective capacity C_{eff} and a nonzero internal instability rate h_{KS} . Its control-theoretic core is the Data-Rate Theorem: an unstable process cannot be tracked over a channel whose capacity falls below a threshold set by the instability rate. Applied to the observer’s own *basis-tracking problem*, this gives the rate

$$\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2,$$

with two regimes. For $\kappa < 0$, the observer can keep the basis resolved well enough that textbook fixed-basis quantum mechanics is operationally available. For $\kappa > 0$, standard quantum mechanics still holds, but the observer’s reference frame is not fully resolved. Raw, unconditioned records can then lose contrast by ordinary phase averaging,

$$V_{\text{obs}} = V_{\text{std}} e^{-\sigma_{\theta}^2/2}.$$

Unitary dynamics are untouched and no physical collapse is postulated.

On this reading, collapse and quantum randomness are not events in the world but the shape a single deterministic history takes in the records of an observer that cannot fully resolve its own basis. No force, branch, or collapse term is added; standard quantum mechanics is recovered exactly, and the contrast lost when $\kappa > 0$ is recoverable reference-frame physics—already confirmed to return once the realized reference is logged. What the framework adds instead is a measurable control law and, through it, a calibrated instrument for testing whether any residual loss is genuine collapse. So understood, quantum mechanics is less a description of ultimate reality than the lawful signature of bounded observers embedded in an order they cannot exit.

“When we finally understand quantum mechanics, we will wonder how we ever missed something so simple.”

— John A. Wheeler

Contents

I	The Contemplative Starting Point	7
1	The Ancient Insight	7
1.1	What Is Advaita Vedanta?	7
1.2	The Core Problem: Avidyā	7
1.3	The Mechanism of Multiplicity	8
1.4	Self-Reference and the Impossibility of Complete Self-Knowledge	8
2	The Question	9
2.1	Can This Be Physical?	9
2.2	The Physical Translation	9
2.3	What This Is and Isn't	10
2.4	A Note on Method	11
2.5	About My Method	12
II	The Physical Framework	13
3	Where Did the Measurement Basis Come From?	13
3.1	Standard Formulations and Their Difficulties	13
3.2	What's Missing from Current Approaches	13
4	Formalizing the Insight: Observers Are Physical Systems	14
4.1	Measurement Basis as Internal Physical State	14
4.2	Why Self-Knowledge Must Be Limited	15
4.3	The Information Constraint	16
4.4	The Predictability Problem	16
5	How This Reframes Collapse	17
5.1	The Puzzle, and the Move IOF Makes	17

5.2	Division of Labour: What IOF Borrows and What It Adds	18
5.3	Three Elements in One History	19
5.4	Collapse as Epistemic Update	20
5.5	Why the Outcome Appears Random	21
5.6	Observer-Relative Visibility: the κ Control Law	22
6	The Layered Architecture	24
7	Entanglement in a Single History	26
III	The Mathematical Formalism	29
8	Minimal Dynamical Setup	29
8.1	System Components	29
8.2	Internal Dynamics	29
8.3	Measurement Coupling	30
8.4	The Basis as Dynamical Variable	30
9	Information-Theoretic Foundations	30
9.1	The Data-Rate Theorem	30
9.2	Finite-Capacity Channel	31
9.3	Two Dynamical Regimes	32
9.4	Rate-Distortion Theory	33
9.5	Predictability Parameters ($\alpha_{\text{ch}}, D_{\theta}$)	33
9.6	The Regime Threshold and Timescale Definitions	34
10	Epistemic Collapse: Quantitative Predictions	35
10.1	Variance Bounds	35
10.2	Visibility Formulas (Both Regimes)	36
10.3	Record Formation: The Undersampling Map	38
10.4	Double-Exponential Visibility Decay (Chaos-Wins)	38

10.5	Characteristic Observation Time (Chaotic Regime)	39
10.6	Constant Suppression (Diffusive Regime)	40
10.7	Experimental Regimes and Parameter Justification	40
11	Bell, Kochen-Specker, and Contextuality	41
11.1	Explicit Contextuality	41
11.2	Measurement Independence as an Additional Assumption	42
11.3	No-Signalling Preservation	43
11.4	Reproducing Standard Quantum Correlations	43
12	The Strong Reading and Its Exclusion	43
12.1	Two Readings of the Visibility Law	43
12.2	Why Standard Quantum Mechanics Predicts No Effect	44
12.3	The Boundary Problem	44
12.4	Exclusion by Existing Experiments	44
12.5	What Survives	45
IV	Empirical Connections	46
13	The Biological Observer and Lived Experience	46
14	Connection to Penrose Objective Reduction	47
14.1	Two Timescales, from Two Sides	47
14.2	The Mesoscopic Coincidence	48
14.3	A Clean Discrimination	48
14.4	Why This Framework Builds the Instrument, and What Each Outcome Would Mean	49
15	Experimental Tests	50
15.1	What Is Settled and What Remains Open	51

V Interpretation and Implications	53
16 What the Framework Buys	53
17 The Measurement Problem, Reframed	53
18 The Nature of Physical Law	54
19 IOF Among the Interpretations	55
20 One Mechanism Across the Foundations	57
VI A Non-Dual Perspective	59
21 A Note on My Method	59
22 Consciousness as Fundamental (Sat-Chit-Ānanda)	60
23 Self-Ignorance (Avidyā) and the Appearance of Duality (Māyā)	62
23.1 Why One Appears as Many	65
23.2 Physics as the Structure of Apparent Separation	66
23.3 Turiya and Brahman: A Crucial Distinction	69
24 The Three-Fold Structure of Karma	70
25 The Hṛdaya-granthi and the Heisenberg Cut	76
VII Conclusion	81
Appendices	84
A Derivation of Tracking Timescales (τ_{fill}, t_{break})	84
A.1 Formula for Tracking Timescale	84
A.2 Steady-Cycle Prior for Kicked Rotor (Capacity-Wins)	85

A.3 Sensitivity Analysis	85
A.4 Conclusion of Analysis	85
A.5 Sanity Checks	86
A.6 Canonical Clock Conventions	86
References	88

Part I

The Contemplative Starting Point

1 The Ancient Insight

1.1 What Is Advaita Vedanta?

Advaita Vedanta is a rigorous contemplative tradition developed within the ancient Vedic lineage for investigating the nature of consciousness and reality. The name means “non-dual end of knowledge”: *advaita* = “one without a second,” *veda* = knowledge, *anta* = end or culmination. It is not empirical science in the same third-person sense as physics; it is a disciplined first-person inquiry whose concepts can nevertheless inspire precise physical questions.

Vedanta states reality is non-dual. There is one fundamental, which appears as the multiplicity of individual observers and observed objects. The fundamental reality—called brahman—cannot be defined directly, only negatively in terms of what it is not (*neti neti*: “not this, not that”), for it transcends all objectification.

The central finding: Brahman = ātman (the Self). You are not the body-mind you take yourself to be—that is the appearance. You are the fundamental reality itself, appearing as an individual observer. This is captured in the teaching: *Brahma satyam jagat mithyā, jīvo brahmaiva nāparaḥ*—“Brahman alone is real; the world is dependent appearance; the individual is none other than Brahman.” Though the fundamental reality cannot be defined directly, its nature is sat-cit-ānanda (existence-consciousness-bliss). Consciousness (cit) here is defined negatively: anything you can be aware of is **not** consciousness.

Within its own domain, Vedanta offers precise analysis and systematic method. The investigation begins with the most immediate datum: *I am aware*. What exactly is this “I” that is aware? The tradition argues that consciousness itself—the knowing principle—cannot be known as an object, because it is the subject that makes all knowing possible. This paper does not attempt to prove that claim physically; it uses the structure of self-ignorance as a source of formal hypotheses.

1.2 The Core Problem: Avidyā

Avidyā (ignorance) is not mere lack of information. It is the primordial error, traditionally defined as “seeing the real as unreal, and the unreal as real”—mistaking oneself to be a limited, individual body-mind (the unreal) rather than recognizing one’s true nature as the unlimited fundamental (the real).

From this root ignorance, the appearance of separation follows necessarily:

- The One appears as observer and observed
- The infinite appears as finite individual

- The subject appears as object
- Pure consciousness appears as separate person having experiences

Ramana Maharshi (1879–1950), the sage whose teachings inform this work, described the mechanism precisely: The ego—the thought “I am this body”—is a formless phantom that simultaneously projects and perceives the world of multiplicity. It arises from forgetting one’s true nature as consciousness and identifying with the body-mind apparatus.

The ego sustains itself through attention to objects (anything other than itself). By constantly grasping at thoughts, sensations, perceptions—anything but its own source—it maintains the appearance of separation. When attention turns back toward itself alone, investigating “who am I?” the ego dissolves, revealing only the non-dual awareness that was always present.

1.3 The Mechanism of Multiplicity

A crucial and often misunderstood point: The appearance of multiplicity is not arbitrary or illusory in the sense of being non-existent. It has a precise structure, describable in terms of *māyā* (the power of apparent manifestation).

Māyā has two aspects or powers (*śaktis*):

Āvaraṇa-śakti (veiling power): Conceals the non-dual whole, making the infinite appear finite. It is like a lens that narrows the field of vision, creating the sense of limited perspective.

Vikṣepa-śakti (projecting power): Projects the appearance of multiplicity from the concealed unity. Once the whole is veiled, the parts appear as separate and independent.

These are not two separate powers but two aspects of a single mechanism: Self-ignorance necessarily produces both the veiling of one’s true nature and the projection of apparent separation.

Here is the key insight that bridges to physics: **The multiplicity is perfectly real as appearance while being fundamentally unreal as independent existence.** The dream is completely real within the dream; only upon waking does one recognize it was appearance in consciousness, not separate reality.

1.4 Self-Reference and the Impossibility of Complete Self-Knowledge

Vedanta makes a claim with striking parallels to Gödel’s incompleteness theorems: complete self-knowledge is structurally impossible for an observer operating through subject-object division.

The Self (= *ātman* = *brahman*) is self-aware by being itself, not through objectification. The ego attempting to observe itself is like an eye attempting to see itself—structurally impossible. But there is a crucial distinction:

The ego: Cannot know itself completely. The ego—the “I am this body-mind” thought that creates apparent individuality—cannot fully know the state that determines its next thought, choice, or perception. By the time the ego becomes aware of a mental state, that state has

already evolved. There is always lag, always incompleteness. The ego cannot trace *why* its thoughts, choices, and attention move as they do. The causal chain from past tendencies to present experience remains hidden.

The Self: Knows itself by being itself, not by observing itself as an object. This is immediate, non-dual awareness (svarūpa-jñāna)—not knowledge OF self but knowledge AS self.

This distinction is crucial. We are not claiming that consciousness is limited. We are saying that the ego—the apparent individual observer (jīva) identified with a body-mind—necessarily has limited self-knowledge because it operates through subject-object division. (For detailed analysis of Self, ego, and the structure of ignorance, see Part VI: A Non-Dual Perspective. For how this connects to quantum measurement, see Section 4.)

2 The Question

2.1 Can This Be Physical?

Given this empirical framework from Vedanta, a natural scientific question arises:

If fundamental self-ignorance creates the appearance of individual observers in a non-dual reality, what would that look like physically?

Not: “Does quantum mechanics prove Vedanta?” (It doesn’t.)

Not: “Can Vedantic inquiry derive physical laws?” (It can’t.)

But: “If we take seriously the structure of self-ignorance that Vedanta describes—the ego’s inability to know the causal chain of its own arising—and formalize it mathematically, what testable predictions emerge?”

Terminological note: "Self-ignorance" here means the observer’s ignorance of its own internal physical state—used for technical precision to distinguish internal from external ignorance. The body is itself world (object), not Self (ātman), so this is still ignorance (avidyā) about appearances.

2.2 The Physical Translation

The physical translation begins with a question that quantum theory usually passes over too quickly: why is the measurement basis treated as if it came from outside physics? A spin measurement is never simply the question, “What is the spin?” It is always the question, “What is the spin relative to this axis?” The axis, or more generally the basis, is therefore part of the physical question being asked. IOF takes that basis seriously as a physical reference variable $\theta(t)$, generated by the observer-apparatus system itself.

The Vedantic claim has the same structure in first-person form. The ego experiences itself as choosing, but it cannot trace the causal chain that produced this particular choice. It knows

the choice after it appears; it does not know why this choice, rather than another, arose from the accumulated tendencies (*vāsanās*) of the past.

The physical translation is then direct. An observer making a quantum measurement can record which basis was implemented, but cannot reconstruct the full causal ancestry by which its internal state $\theta(t)$ evolved to that value. The causal chain from prior states to the present basis is not available inside the observer's finite record. The observer asks a question whose origin it cannot fully know; the world answers once; the observer calls the answer random.

Central hypothesis (IOF). Part of the apparent randomness in quantum measurement is *epistemic*. A finite observer can record which basis was used, but cannot fully reconstruct why its own internal state produced that basis. IOF treats this missing provenance as one component of the randomness encountered by the observer, not as the whole of quantum randomness.

The formal version of this hypothesis—the tracking-versus-prediction setup, the capacity-versus-chaos comparison, and the observer-relative visibility law that results—is developed in Part II and Part III.

2.3 What This Is and Isn't

This is not:

- An attempt to reduce one framework to another
- A claim that one methodology supersedes the other
- A derivation of quantum mechanics from first principles
- A conflation of distinct investigative domains

This is:

- A mathematical formalization of one specific empirical finding from Vedanta (self-ignorance)
- A constructive operational model for the observer-side component of unpredictability, together with a calibrated, measurable reference-tracking control law (and, on a deterministic host, a non-conspiratorial route to measurement dependence)
- An interpretation of quantum measurement with an operational subtheory whose parameters can be built and measured
- An exploration of how different paths—using different languages, symbols, and methods—may be approaching the same truth from complementary directions

Physics (third-person empiricism), quantum mechanics (mathematical formalism for measurement), string theory (geometric framework for unification), and Vedanta (first-person empiricism) all represent rigorous paths from ignorance to knowledge. Each uses the tools appropriate to its domain of inquiry. The question this work addresses: When we formalize Vedantic insights mathematically, does the resulting structure cohere with standard quantum mechanics—and where exactly does its content lie? The answer (developed in Part III) is that it coheres entirely, with its content distributed across an operational control law, a conditional foundations result, and an interpretation; what that tells us about the relationship between these frameworks is the subject of the closing parts.

2.4 A Note on Method

Science and Vedanta use different epistemologies:

Scientific method: External observation, reproducible experiment, mathematical modeling, empirical falsifiability. It studies objects and relationships between them.

Vedantic method: The Vedantic method is a rigorous, first-person empirical inquiry. It begins with śravaṇa, receiving the provisional understanding from the teaching (śruti). This is followed by manana, a deep contemplation and reflection on the teaching, using rigorous logic (tarka) to remove all intellectual doubts. The final and most crucial step is nidhīdhyāsaṇa, the sustained contemplative practice whose purpose is to make what has been intellectually understood a continuous, living reality (anubhava). This iterative process continues until the provisional understanding stabilizes as direct, non-dual knowing.

This parallels the scientific method more closely than commonly recognized:

- **Science:** Hypothesis → physical experiment → observation → theory refinement
- **Vedanta:** Provisional understanding → contemplative practice → direct experience → understanding refinement

The crucial difference is domain and method, not rigor. Science investigates matter through third-person measurement; Vedanta investigates consciousness through first-person inquiry. Both employ:

- Testable claims (predictions about what will be experienced under specific conditions)
- Reproducible experiments (anyone following the method should obtain similar results)
- Progressive refinement (understanding improves through iteration)
- Internal consistency checks (contradictions indicate error)

Vedanta represents a tradition of rigorous, first-person empirical inquiry, refined over thousands of years. Its foundational truths are not derived from speculation, but are said to originate from the direct, first-person experience of sages (ṛṣis) who have realized the ultimate, non-dual

reality (jñāna). The role of the broader tradition has been to systematically structure, test, and transmit these core insights. Each generation of practitioners then uses this framework not as a dogma to be believed, but as a guide for their own contemplative investigation, seeking to verify the sages' claims through their own direct practice and experience. The framework includes sophisticated epistemology (pramāṇas—valid means of knowledge) and methods to demonstrate internal consistency (tarka—logical analysis showing non-contradiction).

Where physical science accumulated knowledge through material manipulation, Vedanta accumulated knowledge through consciousness investigation. Different laboratories, same spirit: careful observation, hypothesis testing, theory refinement toward predictive accuracy.

These are complementary, not competing. Science investigates the structure of appearance with rigor and precision. Vedanta investigates the nature of that which appears and that to which it appears.

2.5 About My Method

My own method draws on both, in a particular order: the contemplative reading came first and supplied the starting intuition—it pointed the formalism where to go—but the physics is held to physics' standards, with its claims delimited against existing experiment. That reversed order of discovery, from non-dual intuition to operational mechanism, is described in *A Note on My Method* (Section 21).

What follows is intended as physics, mathematics, and testable prediction. The Vedantic material motivates the question; it does not serve as evidence for the physical model. The reader may judge the physics on its own merits. Whether the Vedantic inspiration adds or detracts is a matter of perspective. The framework works, if it works, only because its physical assumptions and predictions survive analysis and experiment.

For me personally, the structural resonance between the physics and the metaphysics—the way two registers built on entirely different evidence, third-person measurement and first-person inquiry, meet as they do—is part of what gives the interpretation its appeal, and weighs, to my own mind, in its favour over interpretations that carry no such correspondence. I record this as a personal judgement, not an argument: the resonance settles nothing on its own, the physics is meant to stand or fall by its own standards regardless, and another reader may find the correspondence unpersuasive or give it no weight at all. It is simply part of what, for me, makes this lens worth holding.

What follows develops these claims with different levels of confidence. Part II introduces a physical framework and a deterministic global-history embedding, Part III formalizes predictions, Part IV connects to experiments, and Part VI explores philosophical parallels. The physical model should be evaluated by its assumptions, its parameter estimates, and the precision with which its claims are delimited against existing experiment (Part III); the non-dual interpretation is presented as interpretation, not proof.

Throughout, Vedantic concepts are referenced as the generative inspiration. The reader may engage at whatever level they find comfortable—taking the metaphysics seriously as a source of questions, or ignoring it and judging the physics on its own terms.

Part II

The Physical Framework

3 Where Did the Measurement Basis Come From?

3.1 Standard Formulations and Their Difficulties

Quantum mechanics has a measurement problem: standard presentations assign the system a superposed state, relative to a chosen measurement basis, before measurement, and a definite recorded outcome after. What happens in between?

While a full taxonomy is beyond the scope of this introduction, prominent approaches to this problem generally fall into several broad categories:

Collapse Theories: These propose that the Schrödinger evolution is physically interrupted by a “collapse” of the wavefunction. This is the traditional approach, though the trigger for the collapse remains mysterious.

The Many-Worlds Interpretation: This approach denies collapse entirely, proposing instead that every possible outcome occurs, each in a separate, parallel branch of reality.

Hidden-Variable Theories: These posit that the wavefunction is incomplete and that definite outcomes are determined by additional, “hidden” variables. IOF is compatible with multiple no-collapse ontologies.

These foundational approaches all interact with the well-understood physical process of decoherence, which explains how a quantum system loses its coherence through interaction with the environment, but decoherence alone does not solve the ultimate problem of why any single, specific outcome is realized. Each of the primary interpretations faces its own set of profound difficulties.

3.2 What’s Missing from Current Approaches

These approaches share a common oversight: they treat the observer’s choice of measurement basis as externally given, almost God-like in its arbitrariness. The measurement direction—whether we measure spin along the x -axis versus the z -axis—is treated as a free parameter, not as a physical state subject to the same constraints as everything else.

In the vocabulary of Bell’s theorem, this “externally given” setting is the *measurement-independence* assumption [1]: the hidden state of the measured system is taken to be statistically independent of the later measurement setting. It is usually left implicit, treated as an obvious feature of experimental “free choice.” But a measurement setting is a physical state with a physical history, produced by apparatus and observer dynamics that descend from the same past as the system being measured. If setting and system share that ancestry, a correlation between them is not a

conspiracy—it is what a single history means—a claim developed, with the burden it accepts, in Section 7.

This paper develops the narrower observer-side consequence: *the observer may not be able to fully reconstruct why a particular measurement basis was physically realized*. That irreducible self-ignorance about one’s own basis-setting process—not a new collapse mechanism—does the work in what follows, and it is the hinge on which “collapse” turns from an *ontic* event into an *epistemic* one. The probability the observer then assigns is accordingly a predictive credence over its possible records, not an ontic production law—the distinction made precise in Section 10.3, and the one that separates this framework from a collapse model, from subjective idealism, and from ordinary decoherence. The deeper questions this raises about agency and free will are taken up in the non-dual perspective of Part VI.

4 Formalizing the Insight: Observers Are Physical Systems

4.1 Measurement Basis as Internal Physical State

When we describe quantum measurement, we typically imagine an external observer choosing a measurement direction—say, measuring spin along the x -axis versus the z -axis.

But observers are physical systems. The “choice” of measurement basis must itself be a physical state of the observer—encoded in the configuration of atoms, neurons, or mechanical degrees of freedom that orient the measuring apparatus.

A Note on the Term “Observer”: In this framework, the term “observer” is used in its most general physical sense to denote any bounded physical system whose internal state determines the parameters of a measurement interaction. This applies equally to a human brain, a silicon photodetector, a Geiger counter’s trigger circuit, or a complex molecule interacting with its environment. The only requirement is that the system be a physical entity subject to thermodynamic and information-theoretic constraints—possessing finite power budget P , operating at temperature T , and having internal dynamics (chaotic or diffusive) that produce the measurement basis. The mathematical formalism that follows is universal to all such systems.

Here is the crucial step: if the measurement basis is implemented by the observer’s internal physical state, and if that internal state has unstable or noisy dynamics, then the observer faces a tracking problem. It may be unable to reconstruct why its internal state realized this particular measurement choice.

Why not? Not because of quantum uncertainty in the observer themselves (we can treat the observer classically). But because of something more basic: the causal chain from past states to present choice is unavailable to the system’s finite self-model, due to information-theoretic limits. This gap between the causal antecedent of a “choice” and what the self-model can reconstruct of it has a phenomenological illustration in the readiness-potential literature; the constraint itself is general and physical, not specific to brains.

The Vedantic perspective: This is the physical formalization of one aspect of avidyā—“seeing the real as unreal, and the unreal as real.” The real causal chain from prior physical conditions

to this measurement choice is not fully grasped; it remains hidden, fleeting, or inaccessible to the observer’s internal model. The observer’s report—“I chose this measurement direction”—appears solid and immediate, even though the physical basis-setting process may be only partially reconstructed. The measurement basis θ feels like “my choice” (unreal taken as real), while the causal substrate determining θ cannot be grasped (real taken as unreal). The deeper relationship between this physical definition and the nature of consciousness is explored in Part VI.

4.2 Why Self-Knowledge Must Be Limited

The limit is structural before it is a matter of finite resources. The measurement basis θ is not an external dial the observer reads off; it is produced by the observer-apparatus’s own internal physical state—the orientation, the phase reference, the timing gate, and the records that represent them. To track θ is therefore to track a system whose dynamics *include the dynamics of the tracker*. The target of tracking includes the tracker. This is not an infinite regress but a feedback problem: the observer must hold a model of a context that contains its own continually-updating modelling activity, and a complete self-account would require that record to contain, in full, its own record-forming. That demand is fixed by self-reference, not by any particular hardware: no sharper sensor removes the self-reference problem *as a class*. A sharper reference channel can still reduce σ_θ for a particular basis variable—this is exactly the capacity dependence the experiments exploit—but it cannot turn the observer into a complete object of its own record-forming.

Finite physics then adds a rate. Even granting the self-referential target, an observer can devote only so much to constraining θ . For an observer to keep an estimate $\hat{\theta}$ of their own basis, they must:

1. Have internal sensors monitoring θ
2. Process signals from those sensors
3. Update their estimate $\hat{\theta}$ based on available information

Notation for tracking error: Throughout this framework we distinguish:

- θ : The physical basis actually implemented by the apparatus
- $\hat{\theta}$: The observer’s internal estimate of that basis
- $\delta\theta = \theta - \hat{\theta}$: The tracking error

The observer records outcomes as if the measurement occurred in basis $\hat{\theta}$ (the intended/believed value), but the actual measurement occurs in basis $\theta = \hat{\theta} + \delta\theta$. This mismatch—quantified by the variance $\sigma_\theta^2 = \langle \delta\theta^2 \rangle$ —is the source of visibility loss in the framework.

This internal monitoring is itself a physical process with:

- Finite bandwidth (limited signal transmission rate)

- Thermodynamic cost (Landauer’s principle [2]: a measurement/control cycle that resets memory by b bits dissipates at least $b k_B T \ln 2$ in heat)—which sets only a *ceiling* on the achievable rate, not the rate itself
- Finite power budget for self-monitoring

Self-reference says what must be tracked; finite physics says how fast. A target that includes the tracker, pursued at a finite rate against internal dynamics that keep generating fresh information about the basis, is what makes some residual self-ignorance unavoidable: the observer cannot access complete information about its own state.

The Vedantic perspective: This is the proposed physical analogue of *āvaraṇa-śakti* (veiling power)—finite self-tracking leaves the observer’s own state partly veiled from itself (Part VI).

4.3 The Information Constraint

These constraints bound a maximum raw information rate C . The basis-tracking task has access to an effective capacity $C_{\text{eff}} \leq C$, depending on allocation, filtering, and architecture.

The thermodynamic limit is particularly fundamental:

$$C \leq \frac{P}{k_B T \ln 2} \tag{1}$$

where P is the power budget available for self-monitoring, k_B is Boltzmann’s constant, T is temperature, and $\ln 2$ appears because one bit corresponds to an entropy change $k_B \ln 2$ (natural-log units).

This is an ultimate thermodynamic ceiling, not the achieved rate: real basis-tracking bandwidth typically falls far below it. Throughout, we use C_{eff} for the task-available effective capacity.

4.4 The Predictability Problem

If the observer’s internal dynamics are chaotic (like a kicked pendulum or turbulent neural activity), then θ evolves unpredictably. Even with perfect initial knowledge, trajectories diverge exponentially. The rate of divergence is captured by the Lyapunov exponent λ ; for multi-dimensional systems, the relevant quantity is the Kolmogorov–Sinai entropy rate h_{KS} (under standard conditions, $h_{\text{KS}} \approx \sum_{\lambda_i > 0} \lambda_i$ per Pesin). We define the chaos rate $\alpha_{\text{ch}} \equiv h_{\text{KS}} / \ln 2$ (bits/s), the information-production rate expressed in the same units as C_{eff} .

Alternatively, if the internal dynamics are noisy (thermal fluctuations, quantum noise amplified to macroscopic scale), then θ undergoes random drift characterized by a diffusion coefficient D_θ .

The combination is deadly: You need an effective tracking rate C_{eff} to track a variable whose internal dynamics produce unpredictability at some rate (e.g. λ in 1D, or h_{KS} in higher-dimensional chaotic settings). If $C_{\text{eff}} \ln 2$ is too small relative to the relevant information-production rate (e.g. h_{KS}), tracking falls behind and the estimate $\hat{\theta}$ becomes increasingly uncertain relative to the

true θ . From the epistemic perspective, the true θ now appears as a probability distribution $p(\theta)$ spread over multiple possible values. Since each possible value of θ would yield a different measurement outcome, the single hidden reality (one true θ) manifests as multiple apparent possibilities.

Quantitative form: The same competition fixes the basis-tracking uncertainty σ_θ^2 quantitatively in each regime,

$$\text{Chaos-wins e-folding scale: } \tau_\kappa = \frac{1}{h_{\text{KS}} - C_{\text{eff}} \ln 2} \quad (\text{chaotic internal dynamics}) \quad (2)$$

$$\text{Steady-state uncertainty: } \sigma_\theta^2 \geq \frac{D_\theta}{C_{\text{eff}} \ln 2} \quad (\text{diffusive internal dynamics}) \quad (3)$$

with C_{eff} , h_{KS} , and D_θ as introduced above. The chaos-wins expression is the amplitude e-folding scale $\tau_\kappa = 1/\kappa$ (with $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$); the precise tolerance time t_{break} adds an $O(1)$ logarithmic factor. Both forms are derived in full in Part III.

The Vedantic perspective: This is the proposed physical analogue of *vikṣepa-śakti* (projecting power)—limited access to the realized state appears as a distribution of possible outcomes.

5 How This Reframes Collapse

5.1 The Puzzle, and the Move IOF Makes

Standard quantum mechanics presents an apparent puzzle. Before measurement, a spin-1/2 particle is described by

$$|\psi\rangle = \alpha|\uparrow\rangle + \beta|\downarrow\rangle, \quad (4)$$

which seems to represent an ontological superposition—neither up nor down until measured. Upon measurement, we observe a *definite* outcome: either \uparrow or \downarrow , not both. In collapse-facing textbook accounts the wavefunction “collapses” to match this outcome, but what triggers the collapse is left unspecified.

IOF removes one idealization buried in that picture: that the measurement basis is an external input. A spin measurement is never the bare question “what is the spin?” It is always “what is the spin relative to *this* axis?”—and the axis, phase, threshold, or timing gate that fixes the question is a physical state of the observer-apparatus system, not a dial floating outside the world. By the central hypothesis of Section 2.2, part of the apparent randomness is then epistemic: a finite observer, with limited effective capacity C_{eff} and unstable or diffusive internal dynamics, cannot fully track its own realized basis θ . It can record *what* basis was implemented without being able to reconstruct *why* that basis, rather than another, was realized.

This does not refute collapse theories, and it does not by itself solve every part of the measurement problem. Objective-collapse models such as GRW and CSL also address macro-definiteness and the transition from superposition to outcome. The narrower point is that IOF removes the need to install irreducible randomness in the world as a primitive—at least for the observer-side component of unpredictability. That component can be carried by finite self-ignorance instead.

But this raises the obvious question:

What ontology makes this work? If outcomes are deterministic yet appear random because the observer cannot track θ , something must still *exist* in which definite records occur, $|\psi\rangle$ keeps evolving without physical collapse, and composed systems carry their quantum correlations. Epistemic ignorance alone cannot supply a world. We answer not by deriving a unique reality but by identifying the most economical ontology that can host the hypothesis.

5.2 Division of Labour: What IOF Borrows and What It Adds

The cleanest way to state the framework is as a division of labour between a *host* ontology and IOF’s one added ingredient.

The host supplies the empirical stage: definite outcomes (one record per run), unitary continuity of $|\psi\rangle$ with no physical collapse, a measure over ontic states or histories sufficient to reproduce the relevant quantum statistics, and the joint-correlation structure of entangled systems. None of this is derived here; it is imported—with a single exception, and that exception is not an arbitrary carve-out but a consequence of the one ingredient IOF adds (next paragraph): the isolated binary single-system weight

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

which the companion Conditional Born-Rule paper [3] reconstructs from finite-record Fisher–Rao geometry under two explicit bridge assumptions. That reconstruction is not invoked here as an established theorem of nature; it is a conditional module—a structural reconstruction from explicit bridge assumptions, not an independent empirical derivation of quantum mechanics [4]. The multi-outcome rule and the entangled joint correlations remain hosted—provably so: the same finite-record invariance that fixes the binary weight admits no invariant phase structure, so the rest cannot be derived by this route [5].

IOF adds exactly one physical ingredient on top of the host: the observer’s finite information capacity (C_{eff}). The same bounded capacity shows two faces. As finite *resolution* of the records it makes, it is what the conditional module above turns into the binary weight $\cos^2(\theta/2)$, which is why the host need not postulate that weight separately. As finite *tracking* of its own internal basis θ , it is the source of self-ignorance proper: the observer cannot trace the causal ancestry that fixed θ , so a single definite history already reads to it as probability; and where the basis-generating dynamics outrun that capacity, tracking fails outright and the observer’s unconditioned records take on the framework’s capacity-dependent, recoverable loss of contrast. The same finitude appears on both sides: the bound that makes an embedded observer read a single history as chance is the bound that, where the reconstruction holds, hands it exactly one piece of Born structure—the binary single-system weight—and, by the same finite-record invariance, no more [5]. The binary Born form is therefore not an arbitrary fragment but the Born structure native to a finite observer: what its own records can fix, while the multi-outcome and phase structure remain on the host’s side of the line. The question is therefore not “which ontology derives all of quantum mechanics?” but the narrower “what is the most economical no-collapse host in which outcomes are definite, $|\psi\rangle$ stays unitary, θ is a real physical variable, and the observer cannot reconstruct the causal history that produced θ ?”

For this paper the working answer is a *deterministic global-history embedding*: one complete history carrying a unitarily evolving field $|\psi\rangle$, a definite ontic state ξ , and a definite realized basis θ .¹ Because ξ and θ both arise inside that single history, they need not satisfy measurement independence as a primitive assumption. Where they are correlated, the correlation is structural—shared causal ancestry, not trial-by-trial coordination, and not a predictive ledger the embedded observer could read off. This is measurement dependence in the non-conspiratorial sense defended by Tim Palmer [6].²

This embedding is used in the rest of the paper because the mechanism turns on the causal provenance of θ . But it is not the only option, and nothing observable hangs on it. A Bohmian host secures the quantum correlations through nonlocal guidance and quantum equilibrium; an Everettian host secures definiteness branch-relatively. These discharge different ontological tasks in different ways, yet IOF’s operational claim is narrower than any of them: *if basis self-tracking fails in the recorded channel, the unconditioned records carry a capacity-dependent, recoverable loss of contrast*. That control law is host-independent; the choice of host is interpretation layered on top of it.

5.3 Three Elements in One History

Within the working embedding the structure has three components, and it is worth being explicit about the role of each.

- **The unitary field $|\psi\rangle$.** It evolves continuously under Schrödinger dynamics and is never physically reduced. It structures the space of possible measurement records and their correlations—not a collapsing object, but the continuously evolving field against which the measurement process is defined.
- **The definite ontic state ξ .** It is the actual state of the measured system in the chosen history. It is not directly accessible to the observer, and it is not a naive pre-existing value for every possible question: in a contextual measurement, what gets recorded depends on the full interaction, including the realized basis.
- **The realized basis $\theta(t)$.** This is the crucial IOF variable. It is implemented by physical degrees of freedom of the observer-apparatus—orientations, phases, timing gates, control records, neural or mechanical states—so it is inside the physical world, not above it. It evolves from prior states, but the observer can monitor it only through finite-capacity

¹In the philosopher’s vocabulary the embedding is eternalist—a block universe. The term is avoided in the text because eternalism alone is too weak to carry the construction: a block of an indeterministic world is still a block, containing whatever happened to occur. The work is done by determinism together with global lawful consistency—the single-solution structure of Section 7—and the block’s notorious question of becoming is answered internally by record formation: the history is not rewritten by what it carries; what advances is the observer’s record of it (Part VI).

²Palmer’s argument, in *Superdeterminism without Conspiracy*, is that violating measurement independence need not be fine-tuned. On his invariant-set programme, the physically real states are confined to an invariant set in state space; the counterfactual that pairs the same hidden state ξ with a *different* setting θ then lies *off* that set—it is not an admissible perturbation, so there is no special initial condition to engineer. The setting-system correlation becomes a structural property of which states are physically real, rather than a coordination arranged between them—the sense in which IOF’s measurement dependence is non-conspiratorial.

channels, so what reaches the observer is a coarse-grained record of the basis-generating dynamics, not the full microstate that produced the basis.

The observer’s ignorance about θ has a single structural source with two roots, set out in the framework’s setup (Section 4). The first is *self-reference*: the basis is generated by the observer-apparatus’s own state, so to track θ is to track a system that includes the tracker, and a complete self-account would have to contain its own record-forming. The second is *finite physics*: tracking is a physical process of finite useful capacity C_{eff} , bounded at the thermodynamic extreme by a Landauer-type ceiling,

$$C \leq \frac{P}{k_B T \ln 2},$$

with the operative rate C_{eff} normally far below it. Self-reference fixes what must be tracked; finite physics fixes how fast. Together they make the basis *causally opaque*: the observer cannot reconstruct the full chain of prior conditions that produced this particular θ . The measurable proxy for that opacity is the *operational tracking error*—against internal dynamics that keep generating fresh information about the basis, a finite C_{eff} cannot hold the estimate of θ to arbitrarily small variance σ_θ , the quantity the experiments actually move.

Schematically, the recorded outcome is then

$$\text{outcome} = f(\xi, \theta).$$

This is not a complete hidden-variable model for spin or entanglement; it is a compact statement that the outcome depends on both the system’s definite ontic condition and the physically realized measurement context. For a single system it marks the dependence on θ ; for entangled systems the joint-correlation structure is supplied by the host. Apparent randomness then arises not because nature has no fact of the matter, but because the finite observer cannot access the full causal state on which its own measurement depends.

These three components do not sit side by side; they stack. The records an observer reads, the quantum description that organizes them, and the global history beneath both form layers ordered by *conditioning access*—by what each standpoint can, in principle, condition its descriptions on. Section 6 sets out this layered architecture explicitly (Figure 1). It replaces the customary ordering of levels by composition or scale with an ordering by access, and it carries the ontic/epistemic bookkeeping the rest of the paper relies on.

5.4 Collapse as Epistemic Update

With this structure in place, “collapse” changes meaning. In textbook presentations a state such as $|\psi\rangle = \alpha|\uparrow\rangle + \beta|\downarrow\rangle$ is treated as requiring a primitive update rule when a definite result appears: something jumps—superposition before, one outcome after.

In the working embedding nothing physical jumps. The field $|\psi\rangle$ may be superposed relative to a chosen basis, but it goes on evolving unitarily; the ontic state ξ remains definite; the apparatus and observer become correlated with the outcome through the measurement interaction. What changes discontinuously is not the underlying process but the observer’s accessible record.

Measurement does not reveal a context-free pre-existing label; it reveals which outcome-region of the ontic trajectory was realized in that context. The fact of the outcome is fixed in the global history by the joint evolution of ξ and θ —but it was unpredictable to the observer, because both the system state and the causal provenance of the basis were only partially accessible.

So what “collapses” is the observer’s epistemic state:

- *Before*: a basis will be physically realized, but its causal provenance is not fully accessible; the outcome is fixed in the host yet unpredictable to the observer.
- *After*: a basis and outcome are recorded; the observer updates its knowledge, while the underlying dynamics remain continuous.

No physical collapse is added. The record is new for the observer, not a discontinuity in the world.

The Vedantic parallel is direct but interpretive: what manifests in measurement is *prārabdha*—the particular fruit of a causal web too vast for the finite observer to reconstruct. It is not undetermined, yet it is not knowable in advance from the standpoint of the embodied observer; *avidyā* appears physically as the inability to trace the causal chain of one’s own basis-realization (Section 22).

There is also a *reflexive* limitation worth stating, because it is a genuine prediction rather than a flourish. The account can be run third-personally: one can describe another observer, or an apparatus, as carrying a definite but causally opaque basis θ . But no observer can occupy the explanation completely from the first-person side, since that would require fully objectifying the causal provenance of its own present standpoint. The framework thus predicts its own first-person inaccessibility: one can model the determinism behind another’s apparent randomness, but one cannot finally stand outside one’s own basis-realization and make it a complete object.

This is where the physical argument hands off to the non-dual reading of Part VI: physics describes the finite record, the hidden causal structure, and the update in accessible information; Vedanta asks who the knower is to whom the update appears.

5.5 Why the Outcome Appears Random

The outcome $f(\xi, \theta)$ is deterministic, yet the observer faces two epistemic barriers to predicting it—a *double ignorance*:

1. **It cannot know ξ .** The ontic state is not available as an operational record; any attempt to infer it runs through a measurement context that changes what is recorded. The observer cannot peek at ξ without altering what is operationally available.
2. **It cannot predict θ .** Finite capacity C_{eff} and chaotic or diffusive internal dynamics prevent the observer from reconstructing why its own basis took the value it did.

Because of this double ignorance the outcome appears genuinely random from the inside. The randomness is epistemic—a limit on self-knowledge—not ontological indeterminacy in nature;

the labels are indexed in the sense of Section 6: the measured randomness is operationally irreducible—and so cleanly ontic—within the empirical layer, while epistemic relative to the substrate standpoint. A coin flip looks random to anyone who cannot track the initial conditions and air currents, though the physics is deterministic; IOF identifies *why* the analogous tracking fails for an observer measuring a quantum system. The causal opacity and the finite-capacity tracking error are not separate phenomena—the second follows from the first under unstable dynamics—and together they broaden $\Pr(\theta | \mathcal{D})$, the observer’s distribution over its own realized basis.

A distinction that must not be blurred. Two distributions are in play. $\Pr(\text{outcome} | \theta)$ —the spread of outcomes *once the basis is fixed*—is ordinary quantum (Born) randomness (binary weight derived within the framework [3]; multi-outcome and joint extensions hosted), and IOF leaves it untouched. $\Pr(\theta | \mathcal{D})$ —the observer’s residual uncertainty about *which basis it actually realized*, given its own data, with variance σ_θ^2 —is the basis-tracking ignorance. Only the second feeds what follows: IOF attenuates visibility through uncertainty about the basis, never by altering the Born statistics at a fixed basis.

5.6 Observer-Relative Visibility: the κ Control Law

Nature does not need to destroy visibility; *it only needs to let the finite observer lose the reference by which visibility can be seen*. The framework’s central quantitative object makes this precise. *Visibility* is the contrast of the quantum interference pattern, $V \in [0, 1]$, with $V = 1$ an ideal coherent fringe and $V = 0$ a fully washed-out signal. Visibility is what an interferometer actually reads off, and it is where basis self-ignorance leaves its mark.

The claim is a *factorization*, not a modification of quantum probabilities. Let V_{std} be the ordinary visibility predicted by the host theory, including environmental decoherence. IOF adds one observer-side factor for the case where the relevant basis or reference is unresolved in the record:

$$V_{\text{obs}} = V_{\text{std}} \cdot V_{\text{IOF}}, \quad V_{\text{IOF}} = \exp\left(-\frac{\sigma_\theta^2}{2}\right). \quad (5)$$

This factor does not alter the Born statistics at a fixed basis and does not replace decoherence. It says only that raw records lose contrast when they are averaged over unresolved basis variation. If the missing basis information is later supplied, the reference loss can be undone.

Two conditions are needed. First, $\Pr(\theta | \mathcal{D})$ must describe *unresolved physical basis variation in the recorded channel*—trial-to-trial, within the integration time, or across the readout window. Merely not knowing a sharp, fixed basis is not enough: private uncertainty about a value that never varies washes out nothing. The factor requires the basis itself to vary, unresolved, in the recorded channel. Second, the product form assumes independent smearings of the same fringe: environmental dephasing and basis-tracking error. Under that assumption their angular distributions convolve and their characteristic functions—the visibilities—multiply.

The factor $\exp(-\sigma_\theta^2/2)$ is the signature of averaging over unresolved basis variation. The interference term carries the basis as $\cos(\phi - \theta)$; if the recorded data are conditioned on a finite-capacity reference channel whose unresolved basis variable has effective distribution $\Pr(\theta | \mathcal{D})$ of width σ_θ , the observed coherence is that term averaged over the variation, $\langle \cos(\phi - \theta) \rangle =$

$\cos(\phi - \theta_0) e^{-\sigma_\theta^2/2}$. The wider the basis-tracking uncertainty σ_θ —the deeper the causal opacity—the more washed-out the signal. The full derivation, in both dynamical regimes, is given in Part III.

One geometry, two contrasts. The attenuation is the Born weight read at finite basis resolution. Writing the binary weight as $\cos^2(\theta/2) = \frac{1}{2}(1 + \cos \theta)$, the fringe term $\frac{1}{2} \cos \theta$ is exactly what the unresolved basis distribution averages. The observed weight becomes $\frac{1}{2}(1 + V_{\text{IOF}} \cos \theta)$. Thus the Born weight is the full-contrast law ($V_{\text{IOF}} = 1$, sharp basis), while V_{IOF} is the contrast available to a finite observer’s record. As $V_{\text{IOF}} \rightarrow 0$, the weight flattens to $\frac{1}{2}$, the limit of maximal basis ignorance. The same Fisher–Rao geometry of finite records that fixes the binary weight [3] fixes the coordinate in which σ_θ , and hence V_{IOF} , is measured.

Two properties of V_{IOF} define its physical status. First, it is *capacity-dependent*: σ_θ is set by the basis-tracking competition— h_{KS} against $C_{\text{eff}} \ln 2$ (Part III)—so the raw-record contrast moves with the observer’s effective capacity at fixed thermal and readout conditions. This gives a measurable control law, realized technically as Bandwidth-Limited Quantum Control (BLQC) [4] (Section 15). Second, it is *recoverable*: the lost contrast returns when the missing basis information is supplied, either by improving the tracking channel in later runs or by conditioning existing records on an independent basis log. This recoverability identifies the loss as reference-frame physics within standard quantum mechanics, continuous with the resource theory of quantum reference frames [7]. The same formula (Eq. 5) can be misread more strongly as an unrecoverable suppression beyond standard decoherence. Part III states that stronger reading and separates it from the present claim. Existing recoverability experiments—including logged-setting Bell tests and randomized-measurement tomography—already exclude it.

The factorization can thus be read as *epistemic optics*. V_{std} is ordinary physical visibility. V_{IOF} is accessible-reference visibility. Their product is the contrast available to the bounded record. This gives a practical classification of contrast loss: (i) *physical loss*, where coherence has dissipated into the environment; (ii) *reference loss*, where coherence is intact but unresolved by the observer’s record; and (iii) *residual anomaly*, where loss survives recovery and regression. Only the third would be new physics, and existing data leave no room for it in the tested regimes (Section 12). Strictly, even environmental decoherence is recoverable in unitary quantum mechanics if one has access to the right correlations. The real distinction is therefore practical access to the conditioning information: environmental records are normally inaccessible, reference logs can be accessible, and a true collapse channel would leave nothing to recover. From inside the bounded record, the loss is real; from outside, it is conditional.

Ontological generality. The observer-side prediction does not depend on a specific guidance law. It can be embedded in any no-collapse framework that supplies the standard fixed-basis statistics and joint correlations the host owes, and the way a given host secures those is not load-bearing for the visibility channel. Throughout, ξ denotes the general ontic state—*ontic*, here and in all the physics parts, in the model-level sense of the ontological-models literature [8]: an element of the observer-independent empirical model, with the further question of ultimate reality bracketed until Part VI. For concrete spin-1/2 visualization we use a Bloch-vector surrogate \mathbf{r} (a visualization aid, not a claim that spin is a separate beable). The predictions are unchanged across realizations. Even this working ontology—field $|\psi\rangle$ plus definite ontic state ξ —is justified functionally: it is “real” in the model-level sense of providing enough structure to explain observations at the empirical level.

6 The Layered Architecture

Physics-foundations vocabulary treats the ontic/epistemic distinction as absolute: in the ontological-models framework [8], a single space of physical states λ is fixed once, and every other structure is classified against it, once. The framework of this paper needs—and physics practice already quietly uses—a relativized version: the labels are indexed to a standpoint, and they move lawfully when the standpoint does.

Stated operationally: a structure is *ontic relative to a layer* when it is unconditioned within that layer—every observer working at that layer’s access finds it the same way, and no conditioning available at that layer dissolves it. The same structure can be *epistemic relative to a deeper layer*: from a standpoint with more conditioning access, it is revealed as bookkeeping—relative, recoverable, standpoint-dependent.

This is not an exotic proposal; it is how physics already operates outside foundations debates. Effective field theory relativizes “fundamental” along the energy axis: the proton is not fundamental, yet no one concludes it is not real—it is ontic at its scale, emergent relative to the deeper theory. Thermodynamic entropy is relative to a coarse-graining, yet objective once the description level is fixed. Von Neumann’s movable cut is the measurement-theory instance: where the classical description begins depends on where the analysis draws the boundary, while predictions stay invariant. The relativization made explicit here runs along a third axis—neither energy scale nor analysis convention, but *conditioning access*: what an observer can condition its records on.

Figure 1 draws the resulting picture. The stack runs from the classical layer of definite records—the only layer any observer reads directly—down through the quantum layer of state assignments and statistics, to the substrate carrying the single global history $(|\psi\rangle, \xi, \theta)$, and finally to a bottom layer that no physical channel can condition on at all, which the physics parts can mark but not describe. Two boundary lines do the work. The upper one is the Heisenberg cut, operationalized: movable by redesignation, and placed, for a given designated loop, where $\kappa = 0$. The lower one is absolute for physics: conditioning access simply ends there. The observer’s limits within its own layer are drawn as its three channels—full access to its own records; a throttled channel to its realized basis θ , whose residual width σ_θ is what the visibility law turns into contrast; and contextual-only access to ξ . One caution about the vertical axis: it measures conditioning access, not realness. Reading the figure as “the substrate is more real than the records” would re-absolutize exactly the labels this section relativizes. The customary diagram of levels orders them by composition or energy scale and lets “fundamental” carry a realness claim; this one orders them by what can condition on what, and asserts nothing about which layer, if any, is unconditionally real.

The accessibility gradient of the visibility law (Eq. 5) is this relativity made measurable. The throttled observer’s visibility loss is ontic at its layer: it is really in the records, and any inspector of those records, at that access, finds the same reduced contrast—it is not a mistake, a credence, or a convention. Relative to the standpoint holding the reference log, the same loss is epistemic: recoverable, conditional, bookkeeping. Both labels are correct; each is indexed. One loss, two layers, no contradiction—exactly as the proton is both real and emergent.

The same indexing fixes what “objective” can mean for collapse, and the point is worth recording

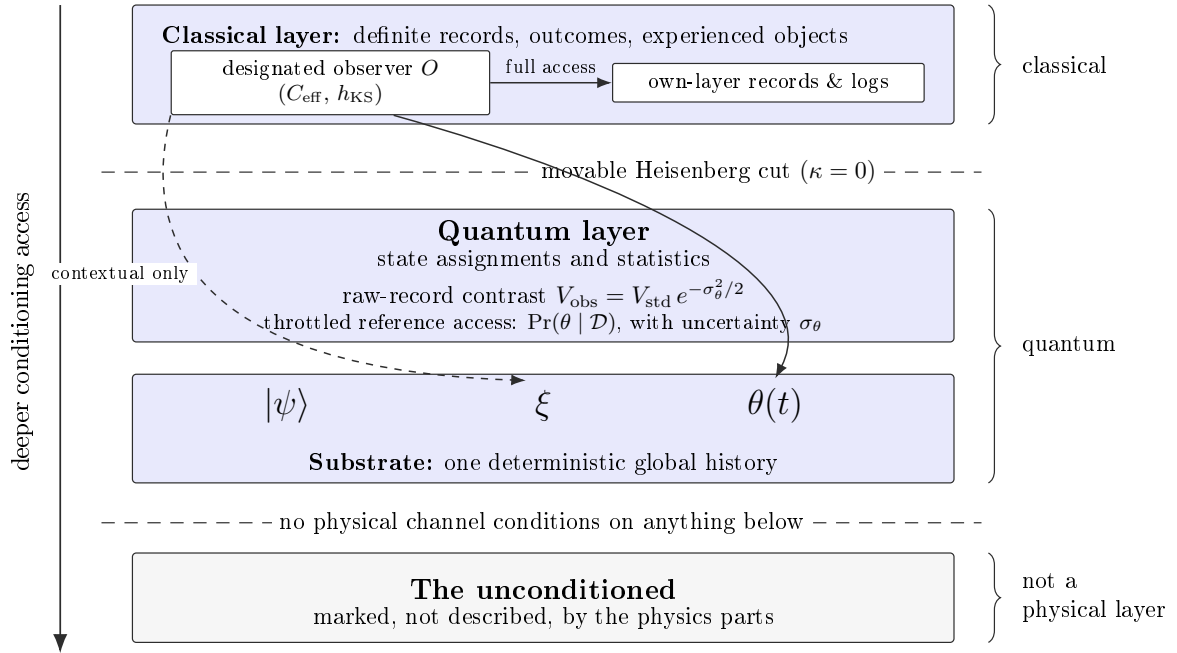


Figure 1: The layered architecture, physical vocabulary. The stack is ordered by *conditioning access*—each deeper layer holds structure the layers above cannot condition on—and by nothing else: the axis is not a realness gradient, and every layer is cleanly ontic to observers working at its access. The designated observer lives in the classical layer with three channels: full access to its own layer’s records; a throttled channel (C_{eff} against h_{KS}) to its realized basis θ , whose residual width σ_θ sets the raw-record contrast $V_{\text{IOF}} = e^{-\sigma_\theta^2/2}$; and contextual-only access to the ontic state ξ . The upper dashed line is the Heisenberg cut, movable by redesignation and placed where $\kappa = 0$ for the designated loop; the lower dashed line bounds what any physical channel can condition on. Figure 4 (Part VI) redraws the same geometry in the contemplative vocabulary.

here because the experimental program of Part IV turns on it. Collapse, in this framework, is objective *indexed to the observer’s layer*: the definite record and the reduced contrast are operationally irreducible at that access, found the same way by every inspector working there, and dissolved by no conditioning available there—not a credence, not a convention. Penrose-type objective reduction (Section 14) is the further claim that the index can be *dropped*: a loss that is absolute, dissolved from no standpoint whatever. The recoverability classifier (Section 12) is exactly this question made operational—*does the index drop?*—with recoverable loss answering that the objectivity stays indexed, and an unrecoverable floor answering that it does not.

Stated so, the framework also explains why collapse could be read as a physical event for a century. Indexed objectivity is phenomenologically indistinguishable from absolute objectivity wherever no deeper conditioning standpoint can be *constructed*. For an engineered channel such a standpoint exists by design—the passive reference log of Section 14.3—so the question is decidable; for the unaided embedded observer (Section 13), whose basis-generating dynamics are logged nowhere, it is not, and its collapse presents as absolute. A reduction that is merely indexed, encountered by an observer that cannot occupy the deeper index, looks exactly like objective reduction. (In the contemplative register, that mistaking—the indexed taken for the absolute—is avidyā precisely as Part I defined it; Part VI develops the correspondence.)

One more label moves with the standpoint: “observer” itself. The controller, its tracking loop, and the body that hosts them are physical systems—world-stuff on any accounting—so the observer/world boundary the framework draws is a *designation*, not a discovery: “observer” names the subsystem assigned the role of holding the reference, and the recoverability classifier needs only that designated cut, never an ultimate one. The designation is movable in exactly von Neumann’s sense—redraw the boundary and the labels redistribute lawfully, with predictions invariant—and within a fixed designation the split is as ontic as any intra-layer structure: which side holds the log and which side is throttled is a fact of the records. What the movability shows is only that no experiment performed with a designated observer can ever locate an undesignated one.

The paper’s convention follows. “Ontic” without qualifier means ontic at the empirical level—the model-level sense fixed earlier [8]—and every intra-empirical layer may be treated as cleanly ontic by observers working within it: the label records what is operationally irreducible at that layer’s access. Whether any layer is *unconditionally* ontic—ontic relative to nothing—is not a question the physics parts answer; it is bracketed, with the rest of ultimate ontology, until Part VI.

The Vedantic perspective: The architecture of Figure 1 returns once more in this paper. Part VI redraws it in the contemplative vocabulary (Figure 4, Section 25): the same stack and the same two boundary lines, with the layers named *sthūla*, *sūkṣma*, *kāraṇa*, and *chit*. Two things change in the redrawing, and both are claims of the tradition rather than of the physics: the downward axis becomes a realness ordering, which Vedanta can assert because it has a sublation criterion the physics parts lack, and the bottom layer acquires a name.

7 Entanglement in a Single History

Alice and Bob share an entangled pair, separated by light-years, and their outcomes are correlated more strongly than any local assignment of pre-existing values can explain. Nothing travels between the wings—no-signalling is a theorem of quantum mechanics, and “spooky action” is not the problem. The actual constraint is Bell’s theorem [1]: the observed violations show that three assumptions cannot all hold—*locality* (no influence between spacelike-separated events), *measurement independence* (settings statistically independent of the system’s hidden state), and *single definite outcomes*. Every interpretation must decline at least one: Bohmian mechanics declines locality, Everett declines single outcomes, and IOF’s deterministic host declines measurement independence—which is an assumption, not a theorem. In a single deterministic history, settings and system descend from a shared past, and nothing makes their statistical independence true by default; the framework neither needs nor presumes it. No-signalling survives exactly, and the dependence never surfaces as anything an observer could read or use—its only operational fingerprint is the standard quantum statistics themselves.

The correlation across the two wings is the one piece the framework defers to its host (Section 5.2): in the working embedding, the particles (ontic states ξ_A , ξ_B) and the measurement bases (θ_A , θ_B) all descend from one causal substrate, and the correlations are encoded in that common past. The route is not IOF’s invention—Bell acknowledged it as the residual loophole of his own theorem, and a developed research program has since made it precise [9, 10]. Remarkably

little dependence is needed: the singlet correlations can be reproduced with a fraction of a bit of setting–state correlation per measurement pair [11]. What IOF adds is not the bare possibility of measurement dependence but two specific things: the constraint logic that makes the dependence structural rather than conspiratorial, and the operational control law that quantifies the observer’s grip on its own basis. There is no conspiracy to answer for, because nothing is being coordinated. In a single deterministic history the correlation between a setting and a system is not an arrangement between them; it is a structural feature of one 4-D solution—what a single history means. The only version of the charge with teeth is fine-tuning, and it is discharged by a single fact: the counterfactual “the same ontic state ξ , but a different, freely chosen setting θ ” is not a physically admissible configuration of that history [6]. There is no special initial condition to engineer, because the alternative was never a state of the world to begin with. This holds whatever generates the settings—an apparatus, a random-number generator, a photon from a distant quasar—since every physical setting has a causal history, shared, however remotely, with the system measured. A completed Sudoku makes the logic visible: cell (1,1) and cell (9,9) are correlated not because one caused the other, nor because someone arranged them, but because a global solution demands consistency. Settings and particle states correlate because the global history is a single consistent solution.

Why no observer can see or spend the correlation. The standing objection to any measurement-dependence route is practical: if settings and hidden states are correlated, why has no Bell test ever caught the correlation, and why can no one exploit it—anticipate outcomes, beat the quantum statistics, signal? The framework answers with a mechanism. To read the setting–system correlation off and use it, an observer would need joint access to the ontic state ξ and the causal ancestry of its own setting θ —exactly the two accesses the double ignorance of Section 5.5 denies it: ξ is screened behind contextual measurement, and θ ’s provenance is bounded by the self-ignorance rate κ . The correlation is real in the global history, but the only standpoint that could cash it out is one no embedded observer can occupy. A world with structural measurement dependence therefore passes every Bell test with apparently free settings—not because the correlation is absent, but because any observer positioned to exploit it is, by the framework’s own control law, unable to resolve it. The framework does not merely permit measurement dependence; it supplies the mechanism of its own concealment, with a rate.

The burden this accepts. Making measurement-independence violation structural rather than conspiratorial is not free; it commits the framework to conditions it should state rather than wave away. The setting–system correlation must be:

- (i) *structurally stable*—a property of the law-like global solution, not of a tuned initial slice;
- (ii) *not a trial-by-trial arrangement*, with no per-run coordination between ξ and θ ;
- (iii) *not usable as a predictive ledger* by the embedded observer—the concealment mechanism above; and
- (iv) that the correlations reproduce the Bell-type *joint* statistics without retrofitted parameters.

Conditions (i)–(iii) bind IOF’s own mechanism; (iv) is not IOF’s to discharge—it is the composition step deferred to the host (Section 5.2). Palmer-style invariant-set arguments matter

here precisely because they aim to secure (i) and (ii) structurally, and the framework adopts that aim rather than resting on a bare appeal to common ancestry. Stated compactly: Hilbert space is the rule book for the records; the host owes the rule book for the history that produces them—and the quantum formalism is the answer key any such rule book must reproduce.

One consequence of this structure should be stated for discipline, because it fixes what evidence for the dependence could ever look like. A completion that reproduces the quantum statistics exactly—which is what discharging (iv) *means*—leaves the setting–system correlation empirically invisible: every function of the accessible record matches the quantum prediction, and the same record is equally consistent with every rival no-collapse reading. A completion that predicts any difference—sub-quantum sequence correlations, reduced randomness under better-controlled preparation—is, by that fact, a deviation from standard quantum mechanics. There is no third case. Empirical *proof* of measurement dependence is therefore closed in principle while quantum mechanics holds: experiments can constrain particular mechanisms—settings drawn from distant astronomical sources push the required common past deep into cosmic history, and fine-tuning analyses price particular models—but no experiment can confirm the position itself.

The concealment mechanism above is the framework’s own derivation of this closure: there is no act to catch from inside. The research target is accordingly not measurement dependence but *non-conspiracy*—discharge conditions (i)–(iv) and nothing else was ever owed. The dependence is to be naturalized: shown lawful, cheap [11], and provably concealed—never proven. (The contemplative register of the same closure—a condition of the entire level cannot appear as a localized effect within it—is given in Section 25.)

The framework’s own addition to the record-level physics is the observer-side factor. For observers A and B with finite effective basis-tracking capacities, self-ignorance about their own bases appears as effectively classical angular noise that attenuates the observed entanglement visibility:

$$V_{\text{obs}} = V_{\text{std}} \cdot \exp\left(-\frac{\sigma_{\theta,A}^2 + \sigma_{\theta,B}^2}{2}\right) \quad (6)$$

This is the independent-error form, valid when the two wings’ unresolved basis variations are independent Gaussian smearings; with correlated reference drift the exponent becomes $-\frac{1}{2}\text{Var}(\delta\theta_A - \delta\theta_B)$, which reduces to the above when the covariance vanishes—the natural default for spacelike-separated observers with independent internal dynamics. The factor does not alter the underlying quantum correlations: it vanishes as $C_{\text{eff}} \rightarrow \infty$, and at finite capacity the attenuation is recoverable on conditioning (Part III) and controllable in magnitude (Section 15)—observer-relative contrast, not a deviation from quantum mechanics.

The visibility factor is the section’s only operational content; everything else is interpretation, and carries no empirical weight of its own. Section 11 gives the detailed mathematical treatment of measurement independence in this embedding.

The Vedantic pointer: the non-dual reading of these correlations—apparent separation as part of the appearance (māyā), the many as expressions of one substrate, correlation as the signature of underlying unity rather than coordination between independent things—is developed in Part VI (Sections 22 and 23.2).

Part III

The Mathematical Formalism

8 Minimal Dynamical Setup

8.1 System Components

We model quantum measurement without ontic collapse or intrinsic randomness. The measurement basis is a genuine dynamical degree of freedom inside the observer, governed by deterministic but information-opaque internal dynamics.

Particle P : Two-level system with Bloch vector $\mathbf{r}(t)$ evolving under Hamiltonian H_P .

Guiding field: Universal wavefunction $|\psi\rangle(t)$ evolves unitarily (Schrödinger equation), never collapses. This guides the evolution of $\mathbf{r}(t)$ in the deterministic global-history working embedding described in Section 5.2.

Observer O : Possesses:

- Hidden internal state $\mathbf{x}(t) \in \mathbb{R}^N$
- Orientation variable $\theta(t) \in [0, 2\pi)$ determining measurement axis $\mathbf{n}(\theta) = (\cos \theta, \sin \theta, 0)$

8.2 Internal Dynamics

Deterministic evolution:

$$\frac{d\theta}{dt} = f_\theta(\theta, \mathbf{x}) \tag{7}$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\theta, \mathbf{x}) \tag{8}$$

Chosen so θ -dynamics is chaotic or strongly mixing. Example—embedded kicked rotor:

$$H_O = \frac{p_\theta^2}{2I} + K \cos \theta \sum_n \delta(t - nT) + H_{\text{bath}}(\mathbf{x}) \tag{9}$$

This makes $\theta(t)$ deterministic but predictively fragile.

Key point: The measurement basis θ is not an unconditioned external parameter but a physical dynamical variable of O .

8.3 Measurement Coupling

Impulsive von Neumann interaction at $t = t_m$:

$$H_{\text{int}}(t) = g\delta(t - t_m) (\boldsymbol{\sigma} \cdot \mathbf{n}(\theta(t_m))) \otimes \hat{p}_X \quad (10)$$

correlating particle spin along $\theta(t_m)$ with pointer position X .

Key point: The measurement basis is determined by the physical state $\theta(t_m)$, whose causal history the observer cannot trace.

8.4 The Basis as Dynamical Variable

This setup makes measurement basis genuinely dynamical:

- θ evolves according to internal observer physics
- Observer's meta-system can monitor θ only through finite-capacity channel
- Measurement outcome depends on $\theta(t_m)$, but observer cannot trace *why* θ evolved to this value
- Apparent randomness emerges from self-ignorance, not ontological indeterminacy

9 Information-Theoretic Foundations

9.1 The Data-Rate Theorem

The stability condition $C_{\text{eff}} \ln 2 > h_{\text{KS}}$ that appears throughout this framework is not an ad-hoc postulate; it is the observer-side analogue of established results in control theory. The **Data-Rate Theorem** [12, 13] proves that:

A linear system with unstable eigenvalue λ cannot be stabilized over a communication channel with capacity $C < \lambda / \ln 2$.

IOF applies this logic to the observer's internal tracking of the measurement basis. The observer's meta-system attempts to maintain knowledge of $\theta(t)$ through a finite-capacity internal channel. When the basis dynamics are chaotic with Lyapunov exponent λ , the data-rate logic implies that asymptotically stable tracking cannot be maintained if the effective basis-tracking capacity satisfies $C_{\text{eff}} < \lambda / \ln 2$.

This provides rigorous mathematical foundation for the two regimes:

- **Capacity-wins** ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$): Tracking possible; self-ignorance bounded; visibility approaches quantum ideal

- **Chaos-wins** ($h_{\text{KS}} > C_{\text{eff}} \ln 2$): Asymptotically stable tracking impossible (though tolerance tracking may hold for $t < t_{\text{break}}$, the tolerance time defined in Section 9.6); basis becomes epistemically inaccessible; visibility degraded

The control-theoretic part is rigorous; the physical status of the resulting visibility law—and the stronger reading it invites—is analyzed in Section 12. The experiments of Section 15 exercise the control law itself: they ask how the basis-tracking deficit moves the contrast of the observer’s unconditioned records.

Relation to decoherence-induced entropy production. The rate h_{KS} on the injection side of this competition is the same rate Zurek and Paz identify at the chaos–decoherence interface. For a classically chaotic system coupled to a decohering environment, the von Neumann entropy-production rate saturates, in the decoherence-dominated regime, at the sum of positive Lyapunov exponents $\sum_i \lambda_i^+ = h_{\text{KS}}$, independent of the system–environment coupling strength [14, 15]. That is the passive, no-tracking case: the environment monitors the system and no feedback channel acts to reduce the basis uncertainty. IOF adds exactly that channel. The chaos rate enters $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ with the same positive sign Zurek and Paz give it—uncertainty is generated at h_{KS} —while the subtractive term $C_{\text{eff}} \ln 2$ is the rate of an active tracking loop their passive setup does not contain. Two consequences follow. First, the self-ignorance rate carries no dependence on the measurement back-action coupling g of Section 8.3: κ depends only on h_{KS} and C_{eff} , so the framework inherits the Zurek–Paz coupling-independence rather than contradicting it—more back-action does not move κ . Second, the passive limit is recovered as $C_{\text{eff}} \rightarrow 0$, where the unresolved basis grows at the bare chaos rate h_{KS} . The reversal of direction—more observer capacity extends coherence, where more environmental coupling does not—is therefore not a conflict with the saturation result but the response of a distinct channel: C_{eff} is a tracking/feedback capacity, orthogonal to the decohering back-action. The correspondence is structural, not a reuse of their theorem, since their h_{KS} is the entropy-production rate of the measured system while the h_{KS} here is the instability rate of the observer’s own basis dynamics.

9.2 Finite-Capacity Channel

Observer’s meta-system receives signals S_t about $\theta(t)$ through an effective basis-tracking channel with capacity C_{eff} (bits/s), constrained by:

Internal bandwidth:

$$C_{\text{eff}} \leq C_{\text{internal}} \quad (11)$$

Thermodynamic limit (Landauer): Information acquisition of b bits costs $\geq b k_B T \ln 2$. With power budget P_{meta} :

$$C_{\text{Landauer}} = \frac{P_{\text{meta}}}{k_B T \ln 2} \quad (12)$$

Operational effective capacity: Real systems operate far above the Landauer floor, and the relevant experimental parameter is not the total device power. The cleanest implementation imposes a known tracking bitrate,

$$C_{\text{eff}} = r b f, \quad (13)$$

where r is the fraction of accepted updates, b is the number of useful bits per update, and f is the update rate. More generally, C_{eff} is the mutual-information or Fisher-information rate between the physical basis $\theta(t)$ and the corrective/control stream assigned to tracking it. It is bounded by:

- architecture limitations (bus widths, clock domains, readout latency),
- deliberate rate limiting or quantization in the basis-tracking loop,
- thermodynamic and energy constraints, including C_{Landauer} as an upper bound.

The Landauer bound provides thermodynamic consistency; it should not be used as a direct estimator of C_{eff} unless the experiment independently demonstrates that added power actually increases useful basis-tracking information rate. The calibration benchmark tests $\partial t_{\text{break}}/\partial C_{\text{eff}}$ while thermal load, readout SNR, and plant dynamics are held fixed. This is standard reference/control physics, not an IOF-vs-QM discriminator.

Two design notes follow. First, if the basis controller runs closed-loop, raising C_{eff} physically tightens $\theta(t)$ itself, so improved raw visibility with control bandwidth is classical control physics—which is exactly the framework’s operational reading of the law. Second, conditioning the records on an independently logged residual basis trajectory restores the lost contrast; Section 12 develops why that recoverability is decisive for the law’s physical status.

Fundamental vs engineering limits: The effective capacity C_{eff} distinguishes fundamental thermodynamic constraints from engineering design choices. For biological observers (neural circuits), thermodynamic-scaled limits may matter; for engineered laboratory systems, deliberate or architectural bandwidth restrictions are often the cleanest bottleneck. This distinction clarifies when predictions reflect fundamental physics versus improvable engineering.

9.3 Two Dynamical Regimes

Observer self-knowledge is limited differently depending on internal dynamics:

Chaotic regime: Small errors in θ grow exponentially as $\delta\theta(t) \sim \delta\theta(0) e^{\lambda t}$

Diffusive regime: θ undergoes noisy drift $d\theta = \omega dt + \sqrt{2D_\theta} dW(t)$

These require distinct rate-distortion analyses and produce different experimental signatures.

Tracking vs Prediction: The analysis that follows addresses a fundamentally different problem than classical chaos or diffusion predictions. We are not asking: “Given imperfect knowledge of $\theta(t_0)$, what is the uncertainty at $\theta(t)$?” (forward prediction without observation). Instead, we ask: “Given continuous observation of $\theta(t)$ through a finite-capacity channel C_{eff} , what is the observer’s residual uncertainty about the current value $\theta(t)$?” This is a state estimation or tracking problem. In chaotic systems, when $C_{\text{eff}} \ln 2 > h_{\text{KS}}$ (capacity-wins), the observer converges to target precision in time $\tau_{\text{fill}} = \lambda T_{\text{kick}} / (C_{\text{eff}} \ln 2 - \lambda)$ (general form and derivation in Section 10.1 and Appendix A); after this convergence period, tracking is maintained with visibility approaching the quantum ideal. At the critical threshold $C_{\text{eff}} \ln 2 = \lambda$, the timescale diverges ($\tau_{\text{fill}} \rightarrow \infty$), and for $h_{\text{KS}} > C_{\text{eff}} \ln 2$ (chaos-wins) asymptotically stable tracking is

impossible—tracking error grows until the posterior approaches uniformity on the circle—though tracking may remain within tolerance for a finite horizon $t < t_{\text{break}}$. In diffusive systems, continuous observation reaches a steady-state tracking error where uncertainty injection (at rate D_θ) balances information extraction (at rate C_{eff}), yielding time-independent variance $\sigma_\theta^2 \gtrsim D_\theta/(C_{\text{eff}} \ln 2)$. The rate-distortion framework that follows is the natural mathematical tool for this tracking problem.

9.4 Rate-Distortion Theory

Rate-distortion theory quantifies the minimum information rate R needed to track a signal with distortion D . Setting $R(D) = C$ (available capacity) determines achievable tracking accuracy, yielding the variance bounds presented below.

9.5 Predictability Parameters (α_{ch} , D_θ)

Chaotic regime:

$$\alpha_{\text{ch}} = \frac{\lambda}{\ln 2} \quad [\text{bits/s}] \quad (14)$$

where λ is the Lyapunov exponent (nats/s, numerically the per-second divergence rate); for multi-dimensional dynamics the relevant rate is $\alpha_{\text{ch}} = h_{\text{KS}}/\ln 2$.

For the kicked-rotor/standard-map controller (equations of motion: $p_{n+1} = p_n + K \sin(\theta_n)$, $\theta_{n+1} = \theta_n + p_{n+1}$), the Lyapunov exponent in the strongly chaotic regime ($K \gtrsim 4$) is $\lambda_{\text{kick}} \approx \ln(K/2)$ per kick. With kick period T , the per-second Lyapunov rate is:

$$\lambda = \frac{\lambda_{\text{kick}}}{T} = \frac{\ln(K/2)}{T} \quad [\text{nats/s}] \quad (15)$$

hence:

$$\alpha_{\text{ch}} = \frac{\lambda}{\ln 2} = \frac{\ln(K/2)}{T \cdot \ln 2} \quad [\text{bits/s}] \quad (16)$$

Physical meaning: Rate at which internal chaos produces unpredictability. This can arise from diverse physical sources: chaotic dynamics in neural networks, thermal fluctuations in electronic control circuits, shot noise in photomultiplier tubes, or quantum noise amplified to macroscopic levels in measurement apparatus.

Example 9.1 (Kicked Rotor with $K = 10$, $T = 0.1$ s).

$$\begin{aligned} \lambda &= \frac{\ln 5}{0.1} \approx 16.1 \text{ nats/s} \\ \alpha_{\text{ch}} &= \frac{\ln 5}{0.1 \cdot \ln 2} \approx 23.2 \text{ bits/s} \end{aligned}$$

Diffusive regime:

$$D_\theta \quad [\text{rad}^2/\text{s}] \quad (17)$$

Physical meaning: Rate of angular diffusion.

Example 9.2 (Thermal Torsional Oscillator at Room Temperature). For a damped torsional oscillator in thermal equilibrium, the angular diffusion constant is:

$$D_\theta = \frac{k_B T}{I\gamma}$$

With $I = 10^{-10} \text{ kg}\cdot\text{m}^2$, $\gamma = 0.01 \text{ s}^{-1}$, $T = 300 \text{ K}$:

$$\begin{aligned} k_B T &= 1.38 \times 10^{-23} \times 300 \approx 4.14 \times 10^{-21} \text{ J} \\ I\gamma &= 10^{-10} \times 0.01 = 10^{-12} \text{ kg}\cdot\text{m}^2\cdot\text{s}^{-1} \\ D_\theta &= \frac{4.14 \times 10^{-21}}{10^{-12}} \approx 4 \times 10^{-9} \text{ rad}^2/\text{s} \end{aligned}$$

9.6 The Regime Threshold and Timescale Definitions

The critical inequality separating the two dynamical regimes is

$$\boxed{h_{\text{KS}} \geq C_{\text{eff}} \ln 2}, \quad (18)$$

where h_{KS} is the Kolmogorov–Sinai entropy rate ($\text{nats}\cdot\text{s}^{-1}$) of the observer’s internal basis dynamics—the information-generation rate it must track—and C_{eff} is the effective information rate ($\text{bits}\cdot\text{s}^{-1}$) available for self-tracking of that basis.

When the tracked basis dynamics admit an SRB-type invariant measure, Pesin’s identity gives $h_{\text{KS}} = \sum_i \lambda_i^+$, the sum of positive Lyapunov exponents; in general Ruelle’s inequality guarantees only $h_{\text{KS}} \leq \sum_i \lambda_i^+$, and h_{KS} should be read as the entropy-rate measure of instability. For the one-dimensional kicked-rotor model in the strongly chaotic regime there is a single positive exponent, so $h_{\text{KS}} = \lambda$, and the two are used interchangeably for that case. The self-ignorance rate is $\kappa := h_{\text{KS}} - C_{\text{eff}} \ln 2$ (nats/s).

Before examining the chaos-dominated regime and its tracking-failure phenomenology, we first examine where standard quantum visibility is recovered. When $C_{\text{eff}} \ln 2 > h_{\text{KS}}$ (*capacity-wins*), the observer’s information rate greatly exceeds the system’s entropy production: self-ignorance decays exponentially and visibility approaches the quantum ideal $V \rightarrow 1$.

This regime suggests why QM occupies its special epistemic position—observers with intermediate information capacity (bounded but not chaos-dominated) recover ideal quantum visibility within this framework. (The framework recovers or modulates standard quantum predictions; it does not derive the quantum formalism itself, whose Born weight enters as input in Section 10.3.) The capacity-wins regime is an epistemic Goldilocks zone: neither infinite self-knowledge (which would dissolve the observer-observed distinction; see Section 23.1) nor chaos-dominated ignorance, but the sweet spot where bounded observers can approximately track their measurement basis.

Departing from this well-behaved limit, when $h_{\text{KS}} > C_{\text{eff}} \ln 2$ (*chaos-wins*), internal dynamics outpace information recovery, and a measurable steady-state suppression of entanglement visibility appears after a characteristic tolerance time

$$t_{\text{break}} = \frac{\frac{1}{2} \ln(\sigma_{\text{target}}^2 / \sigma_0^2)}{\kappa}, \quad (h_{\text{KS}} > C_{\text{eff}} \ln 2), \quad (19)$$

here $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ is the self-ignorance rate.

Timescale definitions:

- **Capacity-wins** ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$): τ_{fill} is the convergence timescale.
- **Chaos-wins** ($h_{\text{KS}} > C_{\text{eff}} \ln 2$): The amplitude σ grows as $e^{\kappa t}$; variance σ^2 grows as $e^{2\kappa t}$. Define:
 - $\tau_{\kappa} := 1/\kappa$ (amplitude e-folding time)
 - $\tau_{\text{var}} := 1/(2\kappa)$ (variance e-folding time)

Throughout this work, t_{break} refers to the threshold-crossing (tolerance) time defined above—the horizon within which tracking can remain within a specified tolerance, even though asymptotically stable tracking is impossible. The bare amplitude e-folding rate is denoted $\tau_{\kappa} := 1/\kappa$.

Representative parameter regimes and predicted effects are collected in Section 10.7, after the quantitative predictions on which they rely have been derived.

10 Epistemic Collapse: Quantitative Predictions

10.1 Variance Bounds

The observer’s self-ignorance—inability to trace *why* θ evolved to its current value—is quantified by the uncertainty σ_{θ}^2 . The word “why” here has an exact operational meaning. A fact becomes predictable only when the relation that produces it is available: the observer may know the basis it used and the outcome it recorded, but that is not enough. It would also need access to the hidden dependency by which this basis, this system, and this outcome arose together. *If that dependency cannot be resolved, it cannot be used.* From the inside, the setting appears free, the outcome appears random, and the missing relation is replaced by probability. The variance σ_{θ}^2 measures the observer-side width of that unresolved dependency, and is bounded by:

Chaotic regime:

For chaotic internal dynamics with Lyapunov exponent λ [nats/s], the observer requires a characteristic timescale to maintain knowledge of the measurement basis angle θ with target precision. From rate-distortion theory for tracking chaotic sources (see Appendix A), the convergence timescale (capacity-wins regime, $C_{\text{eff}} \ln 2 > h_{\text{KS}}$) is:

$$\tau_{\text{fill}} = \frac{\ln(\sigma_0^2/\sigma_{\text{target}}^2)}{2(C_{\text{eff}} \ln 2 - \lambda)} \tag{20}$$

where σ_0^2 is the prior uncertainty at $t = 0$, $\sigma_{\text{target}}^2 = -2 \ln(V_{\text{target}})$ is the target angular distortion threshold in rad^2 (radians being dimensionless, this equals the rate-distortion target D_{target}), and C_{eff} is the observer’s effective basis-tracking capacity in bits/s.

For a kicked-rotor system operating in steady-state cycles of period T_{kick} , the self-consistent prior yields:

$$\tau_{\text{fill}} = \frac{\lambda T_{\text{kick}}}{C_{\text{eff}} \ln 2 - \lambda} \quad (21)$$

Regime-dependent interpretation:

Capacity-wins ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$): This formula gives τ_{fill} , the convergence timescale. The observer starts with high initial uncertainty and converges to target precision in time τ_{fill} . After this spin-up period, tracking is maintained with visibility approaching the quantum ideal.

Chaos-wins ($h_{\text{KS}} > C_{\text{eff}} \ln 2$): The relevant timescale is the tolerance time $t_{\text{break}} = \ln(\sigma_{\text{target}}^2/\sigma_0^2)/[2\kappa]$ where $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ (see Section 9.6). The observer maintains tracking for $t < t_{\text{break}}$; beyond this time, self-ignorance makes the basis effectively unpredictable and visibility degrades below target.

Diffusive regime:

$$\sigma_\theta^2 \geq \frac{D_\theta}{C_{\text{eff}} \ln 2} \quad (22)$$

Time-independent: Steady-state balance between diffusion (injection rate D_θ in rad^2/s) and information acquisition (extraction rate C_{eff} in bits/s, converted via $\ln 2$).

10.2 Visibility Formulas (Both Regimes)

Before measurement, observer knows only distribution $p(\theta|\mathcal{D})$ with variance σ_θ^2 .

Probability of spin-up outcome:

$$\Pr(+|\mathcal{D}) = \int d\theta p(\theta|\mathcal{D}) \cdot \frac{1 + \mathbf{r} \cdot \mathbf{n}(\theta)}{2} \quad (23)$$

For Gaussian uncertainty $p(\theta|\mathcal{D}) \sim \mathcal{N}(\theta^*, \sigma_\theta^2)$:

$$\Pr(+|\mathcal{D}) = \frac{1 + e^{-\sigma_\theta^2/2} \mathbf{r} \cdot \mathbf{n}(\theta^*)}{2} \quad (24)$$

Self-knowledge visibility:

$$V_{\text{IOF}} = \exp\left(-\frac{\sigma_\theta^2}{2}\right) \quad (25)$$

Circular Statistics: Since θ is angular, σ_θ^2 is to be read as the wrapped-Gaussian concentration parameter controlling the first circular moment: for a wrapped Gaussian the identity $\langle \cos(\theta - \phi) \rangle = e^{-\sigma_\theta^2/2} \cos(\theta^* - \phi)$ is exact at *all* σ_θ , so V_{IOF} is the attenuation of the first harmonic of the unresolved basis distribution, and the effective uncertainty saturates at the circular-uniform value rather than growing indefinitely. The genuine modeling assumption is therefore not small σ_θ but that the posterior remains approximately wrapped-Gaussian (first-harmonic-dominated); under strongly mixing dynamics the filamented posterior's first harmonic still decays at a rate controlled by κ , but the precise late-time shape is model-dependent (see Section 10.4).

This general formula relates visibility suppression to the observer's uncertainty about their measurement basis. The time evolution of σ_θ^2 depends on the tracking regime:

Chaotic regime (basis unpredictability due to unstable internal dynamics):

Behavior depends on whether $h_{\text{KS}} \gtrsim C_{\text{eff}} \ln 2$ (see Section 9.6):

Chaos-wins ($h_{\text{KS}} > C_{\text{eff}} \ln 2$): The observer can maintain tracking accuracy for times $t < t_{\text{break}}$ with manageable uncertainty $\sigma_\theta^2 \lesssim \sigma_{\text{target}}^2$, yielding visibility $V_{\text{IOF}} \gtrsim V_{\text{target}}$. For $t \geq t_{\text{break}}$, tracking fails and visibility degrades below the target. This creates a threshold behavior:

$$V_{\text{IOF}} \begin{cases} \gtrsim V_{\text{target}} & \text{for } t < t_{\text{break}} \\ < V_{\text{target}} & \text{for } t \geq t_{\text{break}} \end{cases} \quad (26)$$

Capacity-wins ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$): Self-ignorance decays exponentially with convergence timescale τ_{fill} . After the initial transient ($t \gtrsim \tau_{\text{fill}}$), tracking is maintained and visibility approaches the quantum ideal $V_{\text{IOF}} \approx 1$.

Diffusive regime only (basis unpredictability due to noisy drift):

$$V_{\text{IOF}} = \exp \left[-\frac{D_\theta}{2C_{\text{eff}} \ln 2} \right] \quad [\text{constant, time-independent}] \quad (27)$$

Combined with environmental decoherence $V_{\text{std}}(t) = \exp(-t/\tau_D)$:

Chaotic regime, chaos-wins ($h_{\text{KS}} > C_{\text{eff}} \ln 2$):

For $t < t_{\text{break}}$, self-knowledge is maintained and visibility is limited only by environmental decoherence:

$$V_{\text{obs}}(t) \approx V_{\text{std}}(t) = \exp \left(-\frac{t}{\tau_D} \right) \quad (28)$$

For $t \geq t_{\text{break}}$, the observer loses track of the measurement basis and visibility is further suppressed. The precise form depends on the post-threshold uncertainty growth, but the key signature is a threshold degradation at $t = t_{\text{break}}$.

Chaotic regime, capacity-wins ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$):

After initial convergence ($t \gtrsim \tau_{\text{fill}}$), visibility is limited only by environmental decoherence:

$$V_{\text{obs}}(t) \approx V_{\text{std}}(t) = \exp \left(-\frac{t}{\tau_D} \right) \quad (29)$$

Self-knowledge effects are negligible.

Diffusive regime:

$$V_{\text{obs}}(t) = \exp \left[-\frac{t}{\tau_D} - \frac{D_\theta}{2C_{\text{eff}} \ln 2} \right] \quad (30)$$

10.3 Record Formation: The Undersampling Map

The basis-averaged outcome probability above shows how a finite observer’s bounded capacity registers a definite underlying history as a predictive distribution over records. It takes the basis-conditioned Born weight $\frac{1}{2}(1 + \mathbf{r} \cdot \mathbf{n}(\theta))$ [3], averages it over the observer’s unresolved basis distribution $p(\theta|\mathcal{D})$, attenuating the coherent term by $V_{\text{IOF}} = e^{-\sigma_\theta^2/2}$ in the Gaussian approximation, and is then read at finite resolution.

The resulting $\text{Pr}(+|\mathcal{D})$ is the observer’s *predictive credence* over records, not an ontic production law: the realized outcome, the contextual $f(\xi, \theta)$ of Section 11, is fixed by the single global history, and only across the ensemble of histories compatible with the observer’s finite information are records distributed as $\text{Pr}(+|\mathcal{D})$. *The world does not sample; the finite observer-channel undersamples—a coarse-grained, finite-channel registration of the ontic state as a particle-like record, an operational translation rather than an outcome-generating collapse.*

Unlike environmental decoherence, this capacity-controlled blur is not irreversibly lost to an environment: raising the channel capacity restores contrast in subsequent records, and conditioning already-formed records on an independent log of the realized basis restores it post hoc (Section 12). Recoverability is what marks the blur as observer-relative reference physics rather than a physical channel.

This record-formation step can be summarized compactly as

$$R_{\mathcal{O},\Delta t} = \mathcal{R}_{\mathcal{O},\Delta t}(\xi, \theta, \mathcal{D}), \quad H(R_{\mathcal{O},\Delta t}) \leq C_{\text{eff}} \Delta t. \quad (31)$$

Here $R_{\mathcal{O},\Delta t}$ is the finite record registered by observer \mathcal{O} over the time window Δt , ξ denotes the underlying ontic state or history, θ the contextual basis variable, and \mathcal{D} the data available to the observer. The observed particle-like event is therefore not introduced as a new ontic object produced by collapse. It is the stable bounded record formed when a finite observer samples the underlying state through an unresolved contextual basis. In this sense, the “particle” is the operational appearance of a finite record, not the primitive ontology itself.

10.4 Double-Exponential Visibility Decay (Chaos-Wins)

In the chaos-wins regime ($\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2 > 0$), we close the tracking-error dynamics with the minimal saturation model: the unresolved variance is assumed to grow at the net information-deficit rate,

$$\frac{d}{dt} \ln \sigma^2 = 2\kappa \quad \Longrightarrow \quad \sigma^2(t) = \sigma_0^2 e^{2\kappa t} \quad (32)$$

The rate-distortion argument behind this (Appendix A) is a converse-type bound: no tracker does better, so the exponential law is a lower envelope on the unresolved variance, attained only by near-optimal tracking. The equality should accordingly be read as an effective-rate ansatz, not a theorem. Substituting it into the visibility formula $V_{\text{IOF}} = \exp(-\sigma^2/2)$ yields the **double-exponential decay**:

$$\boxed{V(t) = \exp\left(-\frac{\sigma_0^2}{2} e^{2\kappa t}\right)} \quad (33)$$

Operational Signature (raw-record contrast): This functional form is qualitatively different from standard decoherence mechanisms:

- **Standard exponential:** $V(t) = e^{-\gamma t}$ (environmental decoherence)
- **Gaussian:** $V(t) = e^{-t^2/\tau^2}$ (inhomogeneous broadening)
- **IOF double-exponential:** $V(t) = e^{-(\sigma_0^2/2) \exp(2\kappa t)}$ (tracking failure)

The distinctive shape: near-unity visibility plateau followed by rapid collapse. The *robust* signature is the κ -controlled displacement of the break time: for a chosen visibility threshold V_* , the breakdown time is

$$t_{\text{break}} = \frac{1}{2\kappa} \ln\left(\frac{-2 \ln V_*}{\sigma_0^2}\right),$$

so the primary benchmark observable is the movement of t_{break} under controlled changes of $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$. This is standard reference/control physics, not an IOF-vs-QM discriminator. The full double-exponential curve is the prediction of the minimal Gaussian closure; its late-time tail lies where the posterior is no longer first-harmonic-dominated, so curve-shape discrimination is secondary to the t_{break} scaling.

10.5 Characteristic Observation Time (Chaotic Regime)

For chaotic observers in the **chaos-wins regime** ($h_{\text{KS}} > C_{\text{eff}} \ln 2$), the tolerance time t_{break} defines the characteristic observation time beyond which tracking degrades below a specified threshold. This creates distinct dynamical regimes:

Short observation times ($t < t_{\text{break}}$): The observer successfully tracks the measurement basis angle $\theta(t)$ through continuous monitoring. Visibility is limited primarily by environmental decoherence:

$$V_{\text{obs}}(t) \approx \exp\left(-\frac{t}{\tau_D}\right) \quad (34)$$

The textbook fixed-basis description holds in this regime.

Long observation times ($t \geq t_{\text{break}}$): Tracking fails as uncertainty accumulates faster than information extraction. The observer loses knowledge of which measurement basis is being implemented, leading to additional visibility suppression beyond environmental decoherence.

For chaotic observers in the **capacity-wins regime** ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$), there is no tracking failure. After an initial convergence period τ_{fill} , the observer maintains good tracking and visibility remains near the quantum ideal. Self-ignorance effects are transient only.

Comparison with decoherence timescale:

- **Chaos-wins:** Compare t_{break} (or τ_κ) vs τ_D . When $t_{\text{break}} \ll \tau_D$, self-ignorance dominates; visibility threshold occurs orders of magnitude earlier than environmental decoherence.
- **Capacity-wins:** Compare τ_{fill} vs τ_D . When $\tau_{\text{fill}} \ll \tau_D$, convergence is rapid and self-ignorance effects are negligible throughout.

- When the relevant timescale $\sim \tau_D$: Both effects comparable; interesting mesoscopic regime.

Example 10.1. Kicked rotor with $K = 10$, $T_{\text{kick}} = 1.6$ s, effective basis-tracking capacity $C_{\text{eff}} = 2600$ bits/s, decoherence time $\tau_D = 100$ ms:

With $\lambda = \ln(K/2)/T_{\text{kick}} \approx 1.0$ nats/s and $C_{\text{eff}} \ln 2 \approx 1802$ nats/s, this is **capacity-wins** ($C_{\text{eff}} \ln 2 \gg \lambda$). The convergence timescale is:

$$\tau_{\text{fill}} \approx 0.9 \text{ ms} \ll \tau_D$$

Observable effect: Convergence is rapid; after ~ 1 ms the observer tracks the basis well, and visibility remains near the quantum ideal throughout.

10.6 Constant Suppression (Diffusive Regime)

For diffusive observers: No crossover time; instead, constant multiplicative suppression of visibility at all times.

Signature: Visibility decays exponentially with same τ_D but reduced amplitude.

$$V_{\text{obs}}(t) = V_{\text{IOF}} \cdot \exp\left(-\frac{t}{\tau_D}\right) \quad (35)$$

where $V_{\text{IOF}} = \exp[-D_\theta/(2C_{\text{eff}} \ln 2)] < 1$ is the constant diffusive suppression and $\exp(-t/\tau_D) = V_{\text{std}}(t)$

The visibility starts lower and maintains the same decay rate throughout.

10.7 Experimental Regimes and Parameter Justification

With the variance, visibility, and timescale forms now derived, we place representative laboratory and biological parameters into them.

Typical Physical Scenarios

Regime	System	Parameters	Effect
Capacity-wins	Actively stabilized optical or spin-based interferometer	$C_{\text{eff}} \sim 10^4\text{--}10^7$ bits s^{-1} , $\lambda \lesssim 10$ nats/s	$\tau_{\text{fill}} \lesssim 1$ ms; $V_{\text{IOF}} \approx 0.999\text{--}1$
Chaos-wins	Low-bandwidth autonomous, bio-inspired, or deliberately throttled basis controller	$C_{\text{eff}} \sim 1\text{--}30$ bits s^{-1} , $\lambda \sim 20\text{--}80$ nats/s	$t_{\text{break}} \sim 10\text{--}100$ ms; $V_{\text{IOF}} \approx 0.9\text{--}0.99$
Transitional	Power- or bandwidth-limited embedded sensors	$C_{\text{eff}} \approx \lambda/\ln 2$	Visibility marginally reduced, highly tunable

Table 1: Representative parameter regimes and predicted steady-state effects.

These ranges show that chaos-wins behavior does not require exotic conditions: it naturally arises whenever the information bandwidth about the measurement basis is restricted to a few bits per second while the underlying control or neural process exhibits millisecond-scale instability.

Estimating λ and C_{eff}

1. **Lyapunov rate λ .** Initialize two nearly identical internal states of the basis controller. Record the logarithmic divergence of their trajectories $|\delta\theta(t)|$. The slope of $\langle \ln |\delta\theta(t)/\delta\theta(0)| \rangle$ over its linear window gives λ .
2. **Effective capacity C_{eff} .** Inject a calibrated dither into the intended basis angle and record the observer’s or controller’s corrective stream. Compute the mutual information rate $I(\theta; \text{stream})$ or the equivalent Fisher-information rate. That operational bitrate—typically well below the device’s raw telemetry—is the relevant C_{eff} .

Interpretation

For highly instrumented laboratory systems, $C_{\text{eff}} \ln 2 \gg h_{\text{KS}}$, and the predicted self-ignorance effect remains below experimental resolution. In contrast, for low-bandwidth cognitive or bio-inspired observers, or for autonomous devices operating under stringent power or telemetry limits, the condition $h_{\text{KS}} > C_{\text{eff}} \ln 2$ is easily satisfied, yielding a 1–10 % reduction of raw-record entanglement visibility after tens of milliseconds—observer-relative, and recoverable when the realized basis is supplied (Section 12). The biological realization—neural bandwidth bounds and their consequences for this condition—is developed in the companion supplement *Biological Observers* [16].

11 Bell, Kochen-Specker, and Contextuality

11.1 Explicit Contextuality

Define hidden variables:

- Particle hidden state: ξ
- Observer microstate: $\mu = (\theta, \mathbf{x})$
- Measurement context: c (shielding, timing, apparatus configuration)

Outcomes:

$$A = A(\xi, \mu, c) \tag{36}$$

$$B = B(\xi', \mu', c') \tag{37}$$

Dependence on μ and c makes model explicitly contextual, automatically satisfying Kochen-Specker constraints (no non-contextual hidden variables exist).

The measurement outcome depends not just on the particle state ξ but on the observer’s internal microstate μ and the broader context c . Contextuality is thus built into the physical description rather than added as a separate assumption.

11.2 Measurement Independence as an Additional Assumption

Bell’s theorem relies on the assumption of **Statistical Independence** (also called “measurement independence” or “free choice”):

$$P(\xi|\theta) = P(\xi) \tag{38}$$

That is, the probability distribution over hidden variables ξ is independent of the measurement setting θ .

In a deterministic global-history embedding, this assumption is not automatic—it becomes an additional independence assumption rather than a mathematical necessity:

- The measurement setting θ is determined by the observer’s internal state μ
- Both ξ and μ evolved deterministically from common initial conditions at t_0
- Hence $P(\xi|\theta) = P(\xi)$ is not guaranteed— ξ and θ share a common causal history that *can* correlate them (deterministic evolution permits, but does not force, such correlation)

The IOF does not “violate” Bell’s theorem. In the deterministic global-history embedding explored here, the measurement-independence assumption is not imposed, so Bell’s theorem is not evaded by a dynamical signal but by a different ontological premise. This is an interpretive embedding of IOF, not a separate experimental prediction.

These common-past correlations are invoked only at the level of the global embedding. Throughout the tracking analysis of Part III, the basis dynamics, channel noise, and particle preparation are treated with ordinary probabilistic independence—an approximation taken to be appropriate for the coarse-grained laboratory variables involved, with the embedding-level correlations confined to the Bell context discussed here.

More explicitly, common past at t_0 can yield:

$$\rho(\xi, \mu, \mu') \neq \rho(\xi) \cdot \rho(\mu) \cdot \rho(\mu') \tag{39}$$

Why this isn’t “conspiracy”: Within one deterministic history, correlations arise through global constraint—the 4D history is a single consistent solution (see Palmer’s non-conspiratorial argument [6] in Section 7). This is respectable spacetime physics, not conspiratorial fine-tuning.

Structural causal model: The common cause structure is:

$$(\xi, \mu) \text{ at } t_0 \rightarrow (\theta = \theta(\mu), \text{ outcome} = f(\xi, \theta)) \tag{40}$$

Conditioning on θ (the measurement basis) breaks the apparent independence between particle state and outcome. The self-ignorance—not knowing *why* μ evolved to produce this θ —is precisely what makes $P(\theta|\mathcal{D})$ broad, yielding the visibility suppression.

11.3 No-Signalling Preservation

The observer-relative attenuation preserves no-signalling, and with the attenuated correlation form in hand this can be shown directly rather than asserted. For a singlet pair measured along local settings \mathbf{a}, \mathbf{b} by observers with local visibility factors $V_A = e^{-\sigma_{\theta,A}^2/2}$ and $V_B = e^{-\sigma_{\theta,B}^2/2}$ —each a functional of the *local* channel only—the attenuation acts only on the correlation term:

$$P(a, b | \mathbf{a}, \mathbf{b}, \mathcal{D}_A, \mathcal{D}_B) = \frac{1}{4} \left[1 - V_A V_B ab (\mathbf{a} \cdot \mathbf{b}) \right], \quad a, b \in \{+1, -1\}. \quad (41)$$

Summing over b ,

$$P(a | \mathbf{a}, \mathbf{b}, \mathcal{D}_A, \mathcal{D}_B) = \frac{1}{2}, \quad (42)$$

independent of \mathbf{b} , \mathcal{D}_B , $C_{\text{eff},B}$, and V_B . Local variations in effective basis-tracking capacity therefore modulate only the joint correlation strength, never the remote marginals: no signal can be sent by throttling or widening a distant observer’s channel. For non-maximally entangled states the same conclusion requires that the attenuation act only on the unresolved correlation terms without introducing remote-setting dependence into local biases; the attenuation has this form by construction, since each factor V is a local functional of the local channel. Once a stable macroscopic record has formed, downstream observers agree on it.

11.4 Reproducing Standard Quantum Correlations

The framework recovers standard quantum visibility when observer self-knowledge is effectively perfect ($C_{\text{eff}} \rightarrow \infty$, thus $\sigma_{\theta}^2 \rightarrow 0$). A deterministic, contextual, global-history embedding can then host the usual Bell-violating quantum correlations through common-past or global-constraint structure. That embedding is used to make IOF conceptually coherent with no-collapse realism; the framework’s operational content remains the narrower basis-tracking visibility law, whose physical status is delimited in Section 12.

We emphasize what is *not* done here: IOF does not construct an explicit hidden-variable model $f(\xi, \theta)$ reproducing the singlet correlations. The global-history embedding is noted as a possible host for the required contextual correlations; exhibiting an explicit no-signalling construction of that kind is outside the present formalism.

12 The Strong Reading and Its Exclusion

12.1 Two Readings of the Visibility Law

The visibility factor $V_{\text{IOF}} = e^{-\sigma_{\theta}^2/2}$ admits two readings.

Weak (recoverable) reading. σ_θ^2 is the observer’s unresolved basis variance, and V_{IOF} is the contrast of the observer’s *unconditioned* records. This is standard quantum mechanics: averaging the Born weight over an unresolved reference yields exactly this attenuation, and conditioning on the realized basis restores the full quantum contrast. On this reading IOF adds vocabulary and an explicit control law, not new physics.

Strong (unrecoverable) reading. The stronger reading says that stable records depend only on the channel-resolved basis estimate. Formally, outcomes satisfy the Markov condition $\theta \rightarrow \hat{\theta}_\chi \rightarrow \text{record}$, equivalently $P(a \mid \theta, \hat{\theta}_\chi) = P(a \mid \hat{\theta}_\chi)$, where $\hat{\theta}_\chi$ is the basis estimate resolved by the record-forming channel χ . Then the unresolved part of the physically acting basis does not imprint on the statistics. Lost contrast is *unrecoverable*: re-sorting outcomes against a passive log of the realized basis cannot restore it ($R_{\text{rec}} \approx 0$). This would make IOF an empirically distinct rival to quantum mechanics rather than an interpretation of it, so the reading must be stated and tested explicitly.

12.2 Why Standard Quantum Mechanics Predicts No Effect

With the impulsive von Neumann coupling of Section 8.3, outcome frequencies are fixed by the realized $\theta(t_m)$ per trial: $P(a \mid \rho, \theta) = \text{Tr}[\rho E_a(\theta)]$. An external observer’s uncertainty about θ changes only predictive credence, never conditional frequencies. Changing C_{eff} while the physical θ -dynamics are held fixed therefore changes nothing measurable. Any empirical deviation from quantum mechanics requires the strong reading.

12.3 The Boundary Problem

The strong reading must specify which physical processes count as the record-forming channel χ . If χ is the full causal pathway from θ to the record, it includes the direct physical coupling itself. Its effective capacity is then unbounded, and the postulate never applies. If χ is only a throttled estimator stream, the reading is sharp, but it claims that a physically acting variable fails to imprint on the statistics it acts on. Assigning χ case by case would make the claim unfalsifiable. The sharp estimator-stream form is therefore the only version with empirical content, and it is the version existing experiments address.

12.4 Exclusion by Existing Experiments

Several mature experimental families implement exactly the test the strong reading needs: a per-shot randomized or noise-driven measurement basis, no online resolution of that basis by the record-forming channel, and a passive log or correlated record available later. In every case, conditioning on that later information recovers the quantum statistics, up to ordinary known imperfections ($R_{\text{rec}} \approx 1$):

- **Logged-setting Bell tests.** Weihs et al. [17] switched the analysis basis per photon by a random-number generator, logged settings and time tags with no online use of the setting, and obtained $S = 2.73 \pm 0.02$ entirely by offline sorting.

- **Randomized-measurement tomography.** The randomized-measurement toolbox and classical shadows [18, 19] reconstruct entanglement entropies and observables—chosen *after* measurement—from records whose basis was randomized per shot and merely logged.
- **Correlation spectroscopy.** Two $^{27}\text{Al}^+$ clocks interrogated synchronously through laser phase noise far beyond the laser’s coherence time retain mutual coherence for seconds, recovered purely from cross-correlation of the two outcome records [20, 21]; no channel ever resolved the realized readout phase. (The parity observable is relational, so this family corroborates rather than singly decides; the logged-setting families above carry the direct weight.)
- **Delayed-choice erasure and interferometer logs.** Conditional fringes recovered by cross-record sorting with spacelike-separated choice [22, 23]; atom-interferometer fringes recovered from passive vibration logs [24].

The framework’s own diffusive-regime formula, $V_{\text{IOF}} = \exp[-D_\theta/(2C_{\text{eff}} \ln 2)]$, correctly describes the *raw unconditioned* signal in the correlation-spectroscopy setting—and the demonstrated recovery shows that this suppression is observer-relative phase averaging, not a new channel. The strong reading is therefore excluded in the diffusive and externally randomized regimes by existing data. One corner remains formally untested: deterministic-chaotic basis dynamics with certified positive Lyapunov exponent. The framework registers this test (it is the registered strong-tier module of the BLQC benchmark [4]) without asserting its outcome; nothing in the present formalism motivates a difference between chaotic and diffusive unresolved variation at the level of record statistics.

12.5 What Survives

The κ control law is a calibrated statement about when a finite observer-controller can keep its own basis resolved, grounded in the Data-Rate Theorem. Its nearest established formalization is the resource theory of quantum reference frames [7]: bounded reference information produces effective dephasing, and relational encodings or conditioning can undo it. IOF’s contribution at this layer is the explicit threshold $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ and the tracking-time phenomenology built on it. The framework accordingly makes no empirical prediction beyond standard quantum mechanics. The experiment it motivates (Part IV) discriminates Penrose objective reduction from standard physics along the mass-geometry axis, rather than IOF from quantum mechanics.

Part IV

Empirical Connections

13 The Biological Observer and Lived Experience

A biological observer is not the clean regime in which to test IOF, but it may be the regime in which the meaning of IOF is most immediately visible. We do not experience ourselves as tracking a measurement basis; we experience a world already given as definite, classical, and action-oriented. The basis by which experience is structured is not presented as an object inside experience. It is lived through, not observed. From the IOF standpoint this is exactly what one should expect from an observer whose basis-generating dynamics are deeply embedded, internally chaotic, and only finitely self-trackable. Such an observer encounters laboratory quantum coherence only indirectly, through engineered regimes in which the reference structure is externally stabilized and $\kappa < 0$. Its ordinary world, by contrast, is shaped by saturated basis-opacity: outcomes appear definite, agency appears immediate, and the causal provenance of both remains hidden. The structure is the one made precise in Section 5.5: such an observer registers both ends of each event—the choice it made, the world it finds already given—but not the dependency that joined them, and a dependency that cannot be resolved cannot be lived as anything but immediacy. What is unreadable as relation is what appears as the sheer givenness of the world. In this sense the framework describes not only an exotic laboratory failure mode but the structure of ordinary human experience.

In this reading, the same boundary marks how far such an observer can see. A deeply embedded biological observer sits firmly in the chaos-wins regime ($\kappa > 0$). From its own finite records, the only quantum structure it can directly fix is the binary single-system Born weight, and even that only conditionally [3]. Under saturated basis-opacity, the binary visibility factor $V_{\text{IOF}} = e^{-\sigma_\theta^2/2}$ tends toward zero, so the readable contrast in the weight is washed out and the observed probability flattens toward $\frac{1}{2}$. Everything else that makes the world recognizably quantum—relative phase, interference between non-commuting contexts, the multi-outcome Born rule, and the complex Hilbert structure that carries them—belongs to the host ontology (Section 5.2) and is imported rather than derived—a limit made precise by the companion theorem [5], which shows finite-record invariance generates the binary weight but no invariant phase structure. That structure surfaces only in engineered capacity-wins regimes ($\kappa < 0$), where an externally stabilized reference does for the apparatus what the biological observer cannot do for itself.

The unaided observer never occupies such a regime. The classicality of ordinary experience is therefore not an added posit, but the limit of what such finite records can represent: a binary weight read at low contrast, while the world’s quantum depth remains real but lies on the host side of the division of labour. What finitude explains here is not that each outcome is definite—that is inherited from the host structure—but why the interference that would reveal the underlying Hilbert-space coherence is averaged below visibility, leaving a world that looks classical from within.

This inverts the order one might expect. The classical world of definite objects and immediate

action is the layer that feels most real—most solid, most simply there. Yet in the layered architecture of Section 6 (Figure 1) it is the layer of *maximum* basis-opacity: the most heavily averaged readout of a substrate whose quantum structure it cannot resolve. What feels most real is what is most averaged-over. This is why the architecture had to be ordered by conditioning access rather than by realness: the felt solidity of the classical world is not a sign that it sits at the bottom of what is real, but a sign of how completely the deeper structure has been washed out of a finite observer’s records. The intuition that the everyday world is the firm ground and the quantum the strange exception is exactly the artifact the framework predicts a saturated-opacity observer should suffer.

14 Connection to Penrose Objective Reduction

Penrose’s objective reduction (OR) and IOF both describe a mesoscopic visibility-loss timescale, but they read it from opposite sides: OR sets it by *mass geometry*, IOF by the *observer’s tracking capacity*. For mesoscopic systems the two estimates land in the same 10–100 ms window. That coincidence is not evidence of a shared mechanism—but it locates the regime where the framework’s control variables become a useful experimental axis: the same mesoscopic apparatus that throttles a tracking loop can hold mass geometry fixed, and vice versa, which is precisely the crossed design a clean test of OR requires.

14.1 Two Timescales, from Two Sides

Penrose proposes that a superposition of two mass distributions is unstable on its own, collapsing after a time fixed by the gravitational self-energy E_G of the *difference* between the branches:

$$\tau_{\text{OR}} \approx \frac{\hbar}{E_G}. \quad (43)$$

For a rigid object of size R displaced by s , E_G grows from zero with separation— $E_G \propto (Gm^2/R^3)s^2$ for $s \ll R$ —and saturates at $\sim Gm^2/R$ once the branches no longer overlap ($s \gg R$). So τ_{OR} *decreases* with separation until saturation. This is a gravity-side timescale: it depends on mass and geometry, and says nothing about who is observing.

Example 14.1. $m \sim 10^{-14}$ kg, $R \sim s \sim 1 \mu\text{m} \Rightarrow \tau_{\text{OR}} \sim 10^{-2}\text{--}10^{-1}$ s.

IOF supplies a second timescale, from the other side. In the chaos-wins regime ($h_{\text{KS}} > C_{\text{eff}} \ln 2$) the observer loses track of its own measurement basis after

$$\tau_{\kappa} = \frac{1}{\kappa} \approx \frac{1}{h_{\text{KS}} - C_{\text{eff}} \ln 2}, \quad \kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2, \quad (44)$$

with the derivation given in Appendix A. (In the capacity-wins regime, $C_{\text{eff}} \ln 2 > h_{\text{KS}}$, there is no such loss and visibility approaches the quantum ideal.) Unlike τ_{OR} , this timescale depends on the apparatus’s tracking budget, not directly on the mass. Whether C_{eff} and h_{KS} are themselves altered by mass geometry is not assumed either way: it is treated as a measurable confound

and probed by a crossed design—a capacity sweep at fixed mass geometry alongside a mass-geometry sweep at fixed tracking variables—with the exact control variables, confound checks, and decision rules set out in the experimental protocol [25]. The discrimination axis for objective reduction is the mass-geometry sweep at fixed tracking variables; the capacity sweep serves as the calibration arm, since capacity-dependent, recoverable contrast loss in the raw records is guaranteed reference-frame physics (Section 12).

14.2 The Mesoscopic Coincidence

For mesoscopic parameters the two timescales fall in the same window. Penrose OR reaches $\sim 10\text{--}100$ ms for femtogram masses, $m \sim 10^{-15}$ kg with separations of $100\text{ nm--}1\ \mu\text{m}$. The self-ignorance timescale reaches the same window for a low-bandwidth tracking loop—for example $C_{\text{eff}} \approx 10$ bits/s (so $C_{\text{eff}} \ln 2 \approx 7$ nats/s) against $h_{\text{KS}} \approx 50$ nats/s, giving $\tau_{\kappa} \approx t_{\text{break}} \sim 50\text{--}70$ ms.

The two estimates come from unrelated principles—gravitational instability on one side, information-theoretic tracking failure on the other—so the overlap is not evidence for a shared mechanism, and this paper treats it strictly as a numerical coincidence that locates the experimentally interesting regime. The same caution applies to a third scale sometimes noted alongside these: the identical κ -scaling, evaluated with biological capacities [16], lands in the same decade. That resonance is suggestive only—another application of the same scaling, not an independent test of it.

14.3 A Clean Discrimination

The mesoscopic regime is testable because the two mechanisms depend on different variables, and the decisive distinction is *recoverability*. The apparatus must expose both sets of knobs—the mass-geometry variables that enter τ_{OR} and the basis-reference variables that enter κ —and must carry a passive, causally isolated log of the realized basis reference, so that every observed contrast loss can be classified as recoverable (reference physics) or unrecoverable (collapse-type). A QGEM-class mesoscopic system carrying a real control loop is the natural setting.

The possible outcomes are clean:

- **Raw-record loss moves with κ and is fully recoverable** from the logged reference: the operational control law is confirmed. This is the expected outcome—standard quantum mechanics plus classical control predict it (Section 12)—and it serves as the calibration arm establishing that the apparatus resolves the relevant physics.
- **A κ -independent, unrecoverable visibility floor follows mass geometry** with the tracking loop at maximal capacity: evidence for Penrose-type objective reduction. This is the live discrimination the design exists to make.
- **No unrecoverable floor** at sensitivities where OR predicts one: gravitational objective reduction is constrained at those parameters.
- **A residual κ -dependent loss that survives conditioning** on the logged reference: this would realize the strong reading of Section 12, excluded by existing data in the diffusive and externally randomized regimes; it remains registered for the certified-chaotic corner

only, with the prior strongly against it.

Within the calibration arm, increasing C_{eff} lowers κ and moves the raw-record visibility break later—ordinary control physics that verifies the apparatus, not a discriminating signature. The discriminating observable is the unrecoverable, geometry-tracking floor that remains when tracking is at maximal capacity. The full platform, control variables, confound controls, statistical tests, and decision rules are set out in the experimental protocol [25].

14.4 Why This Framework Builds the Instrument, and What Each Outcome Would Mean

It may look suspicious that a framework about observer-side bookkeeping ends up supplying the calibration procedure for a test of gravitational collapse. The connection is neither coincidence nor hidden evidence, and it is worth stating exactly which it is. IOF’s entire subject is the observer-relative component of apparent collapse; to make that subject precise, the framework had to build machinery that separates observer-relative loss from genuine physical loss. A boundary has two sides drawn by one line: the same classifier that isolates the framework’s subject also isolates everything that is *not* its subject. An objective-reduction test is exactly that complement—remove everything the observer contributed, and see what remains. The framework is, structurally, the subtrahend of the experiment. That a framework built around the observer’s contribution authors the subtraction procedure is necessity, not luck.

Suppose, then, that the discrimination arm finds Penrose’s floor—an unrecoverable, κ -independent loss tracking mass geometry—through this framework’s calibrated channel. Four distinct statements would follow, in descending order of what could be claimed:

1. **Evidentially, about IOF against standard quantum mechanics: nothing.** The classifier is standard-quantum machinery, and that does not change with its success.
2. **The framework’s factoring of the measurement problem would be validated.** The design presupposes IOF’s central analytical move: that “collapse” is not one phenomenon but a sum—a recoverable, observer-relative component plus a possible irreducible remainder—and that no collapse mechanism can be tested until the sum is split. A clean verdict, in either direction, certifies that split as the operationally correct way to pose the question.
3. **The operational tier’s completeness would be demonstrated.** If the framework’s account of the observer-side channel were wrong or incomplete, the classifier would fail in practice: residuals would be misclassified and the floor would smear. A successful discrimination is a working confirmation that the framework’s accounting of the observer’s contribution was accurate and exhaustive in that regime—not IOF over quantum mechanics, but IOF’s bookkeeping shown to be complete in the hardest setting available.
4. **Credit of the instrument-maker’s kind.** A telescope-maker’s mastery of the instrument’s self-noise is what makes an astronomical discovery possible; the discovery does not prove the maker’s cosmology, but it remains permanently downstream of that self-knowledge being right.

The opposite outcome is the more probable one. Suppose the geometry sweep finds no unre-

coverable floor where OR predicts one. In an ordinary objective-reduction search such a null is the weakest kind of result: an absence that could equally mean no collapse, or a collapse signal buried under uncalibrated apparatus noise—phase drift, finite tracking bandwidth, partially corrected reference error, none of them cleanly separated from the quantity of interest. That ambiguity is why “no deviation observed” has not, on its own, sharply bounded objective reduction. The calibrated subtraction removes the ambiguity. Because the reference budget is measured and frozen in advance rather than fit after the fact, and because recoverability certifies the engineered loss as bookkeeping (Section 14.3), the contrast that is *not* lost is genuinely not lost—not hidden inside an unmodelled background. A null then becomes a quantitative upper bound on gravitational collapse at the tested mass geometry, rather than an inconclusive absence. The discipline is the same one that delimits the framework’s own visibility law (Section 12): recoverability separates what is physically there from what only appeared to be. Its consequence here is that the instrument’s worth does not depend on objective reduction being real—the calibrated subtrahend makes the experiment informative in either direction, which is the reason for building it.

The same discipline has a mirror case, and stating it closes the symmetry. Suppose an uncalibrated search in this regime has already returned a null, and the calibrated subtraction later resolves a floor where the lumped noise budget had swallowed one—a null flipped to a confirmation. Mechanically this is the same value working in the other direction: the earlier absence was a false negative manufactured by treating reference physics as an estimate rather than a measurement. But the evidential burden inverts with the sign. A floor that appears only *after* a subtraction inherits the credibility of the subtraction and none of its own: it is only as real as the logging-fidelity budget is small, the coding efficiency frozen before the quantum data rather than fitted to them, and the decision rules registered before the flip. The protocol’s audit trail [25] exists for exactly this case. One boundary holds in both directions: calibration moves only the probability of a correct verdict given the data; it cannot move the probability that objective reduction—Penrose’s or any other dynamical-collapse rival’s—is true. The instrument decides; it does not vote.

The Vedantic perspective: In the framework’s own contemplative register (Part VI), the classifier is discrimination—*viveka*—made operational: the disciplined separation of what belongs to the observer from what belongs to the world. The tradition never claimed that such discrimination proves a metaphysics; it claimed that accurate self-knowledge is what makes everything else legible. A framework whose philosophy holds that the observer’s finitude is bookkeeping, not a physical channel, behaves exactly as that philosophy prescribes when its self-accounting becomes the lamp by which someone else’s physics is read.

15 Experimental Tests

The framework’s experimental layer has two roles, developed in two companion documents—*Bandwidth-Limited Quantum Control* [4], which specifies the interferometric realization of the control law and carries its operational benchmark (calibration arm, capacity and instability sweeps, and the shadow-log recoverability classifier), and *A Capacity-Calibrated Protocol for Testing Penrose Objective Reduction* [25], which fixes the platform, control variables, confound

controls, statistical tests, and decision rules for the mass-geometry discrimination. We give only the logic here.

The operational benchmark. Vary the effective basis-tracking capacity C_{eff} at fixed temperature, readout SNR, and plant dynamics. Increasing C_{eff} lowers $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ and shifts the raw-record visibility break later, with the decay carrying the double-exponential form $V_{\text{IOF}}(t) = \exp(-\frac{1}{2}\sigma_0^2 e^{2\kappa t})$ of the minimal closure (Part III). This is a quantitative, calibrated control law—and it is *expected*: standard quantum mechanics plus classical control predict it, so its confirmation benchmarks the apparatus rather than discriminating IOF from quantum mechanics. Its value is methodological: it demonstrates per-shot command of C_{eff} and h_{KS} , the knobs the discrimination arm then holds at extremes. The benchmark is also deliberately cheap to mount: it needs no mesoscopic masses and no entanglement witness—any platform exposing a throttleable tracking loop against certified-expanding reference dynamics carries it, so the framework’s own exposure is tabletop physics. QGEM-class hardware enters only where the objective-reduction discrimination is added.

The discrimination arm. With tracking at maximal capacity and the engineered channel quiet, sweep mass geometry and look for a κ -independent, *unrecoverable* visibility floor—the Penrose-type signature. Recoverability against the passive reference log is the classifier separating reference physics from collapse-type loss (Section 14.3). The classifier itself is rival-agnostic: every dynamical-collapse model predicts some unrecoverable floor, distinguished by the variable the floor scales with—mass geometry here, nucleon-number amplification for spontaneous-localization (GRW/CSL) models—so retargeting the discrimination arm means swapping the predicted schedule, not redesigning the instrument.

15.1 What Is Settled and What Remains Open

The framework’s empirical ledger, stated plainly:

1. **Settled by existing data:** the strong reading—unrecoverable, capacity-dependent visibility loss beyond standard quantum mechanics—is excluded in the diffusive and externally randomized regimes by logged-setting Bell tests, randomized-measurement tomography, and correlation spectroscopy (Section 12).
2. **Not in question:** the operational control law. It follows from standard quantum mechanics plus the Data-Rate Theorem; an apparent failure of κ -scaling in the benchmark arm would indict the apparatus model or the calibration of C_{eff} and h_{KS} , not quantum mechanics and not the framework’s interpretation.
3. **Open (the live test):** whether a κ -independent, unrecoverable visibility floor tracks mass geometry at the sensitivities where Penrose OR predicts one. Either outcome is informative: a floor supports objective reduction; its absence constrains OR at those parameters.
4. **Registered (prior strongly against):** a residual, κ -scaled, unrecoverable deficit ($R_{\text{rec}} \approx 0$ beyond the logging-fidelity budget) under certified deterministic-chaotic basis dynamics—

the one corner of the strong reading not yet adjudicated by existing data. The framework registers the test without asserting its outcome.

5. **Consistency check:** single-party statistics $P(A = \pm 1 | a)$ must remain 50/50 across all regimes; any skew indicates systematic bias, since the observer-relative attenuation provably leaves marginals untouched (Part III).

Independent control of C_{eff} (capacity knob), s (geometry knob), and h_{KS} (chaos knob), with power and temperature monitored rather than used as the estimator of C_{eff} , lets the data constrain the joint contribution of reference physics, gravitational OR, and conventional decoherence. The full statistical treatment is developed in the experimental protocol [25].

Part V

Interpretation and Implications

The preceding parts built a mechanism and delimited its physical status. This part asks what the mechanism *means*. The reading here stays within the empirical order and leaves the deeper metaphysical question to Part VI. The rule is simple: what physics earns, what foundations supports, and what interpretation reads from the structure are kept separate.

16 What the Framework Buys

Read as a whole, the framework makes one definite structural choice. It keeps a deterministic, no-collapse ontology, and it does not assume measurement independence: setting and system share a common past (Section 7). Apparent randomness is then *epistemic*: it comes from an observer's bounded access to its own basis-setting state (finite C_{eff} against internal unpredictability h_{KS} or D_θ), not from indeterminacy in nature. Standard quantum mechanics remains empirically complete for any observer who cannot step outside itself to view the deterministic substrate.

The limit is concrete and quantifiable. A finite system cannot fully model the state that determines its own next behaviour. When self-ignorance is small (capacity-wins, $C_{\text{eff}} \ln 2 > h_{\text{KS}}$), textbook fixed-basis quantum mechanics is operationally available. When it grows (chaos-wins, $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2 > 0$), standard quantum mechanics still holds, but the observer's unconditioned records can show capacity-dependent, recoverable tracking loss. This is observer-relative reference physics, not a deviation from quantum mechanics (Section 12).

17 The Measurement Problem, Reframed

The measurement problem asks what happens in the passage from a superposed description to a single recorded outcome. The framework's answer is that nothing in the world makes that passage: it lives in the observer's record. The field $|\psi\rangle$ goes on evolving unitarily, the ontic state ξ stays definite, and what changes discontinuously is the observer's access—from a probability distribution over its own unresolved basis to a registered fact. “Collapse” is the name a bounded observer gives to updating its own ledger.

In IOF, coarse-graining is not an optional approximation imposed by the theorist; it is the phenomenological consequence of undersampling. A finite observer does not access the full phase-sensitive structure of the field $|\psi\rangle$, but only a capacity-limited record in which many fine distinctions have become operationally indistinguishable. The result is the appearance of a single classical outcome.

This relocates the puzzle rather than patching it. The question is no longer “how does an indeterminate world become definite?” but “why does a single definite history read, from the inside, as a probability?”—and the answer is self-ignorance.

The same move gives an observer-relative operational version of the Heisenberg cut, which the standard interpretations leave floating: von Neumann showed it can be slid without changing predictions, decoherence ties it to environmental coupling, objective-collapse models fix it at a mass scale. IOF places that cut at $\kappa = 0$, where the observer-apparatus’s useful basis-tracking rate runs out relative to its basis-producing dynamics. This is not an objective physical collapse boundary. It is a design-dependent threshold: an experimenter can move it by throttling or widening the tracking loop at fixed temperature, power, and mass. It is observer-relative without being subjective—two loops on the same hardware place their cuts differently, but for a given configuration any inspector agrees where it sits. The placement is operational throughout: it marks where that observer’s usable basis-conditioned description gives way to coarse-grained records, not an objective physical cut in nature.

18 The Nature of Physical Law

The picture is layered, and the word “real” changes register between layers, so the layers must be named. At the level of **ultimate ontology**—the territory of Part VI—no element of the physical model is claimed to be finally real; the physics parts bracket that question entirely. At the level of **empirical ontology**, a universal state evolves deterministically (something like a Schrödinger equation for Ψ of the whole), with no collapse and no randomness—and no access for any embedded observer; this is the level at which ψ , ξ , θ , apparatus, and records are ontic in the model-level sense [8], and at which standard quantum mechanics is simply valid. At the level of **epistemic appearance**, finite observers sample the empirical structure through limited channels and encounter definite records, probabilities, and apparent collapse: “collapse” is the update of their knowledge, probabilities encode their ignorance, and the theory is complete for everything they can measure. This is IOF’s main territory. Observers cannot access the external view of themselves, so they can never directly verify the deterministic substrate; they are sealed inside the appearance layer. Collapse, in this vocabulary, names the transition from inaccessible relational structure to a definite appearance within the observer’s record-space. The relativized ontic/epistemic labels that make this map operational are developed in Section 6.

Stated for a physics reader: IOF distinguishes empirical reality from ultimate ontology. Within empirical reality, standard quantum mechanics remains valid. At the level of finite observers, collapse is not an additional physical process but the appearance of a definite record under bounded access. IOF is thus epistemic about collapse without denying the empirical validity of quantum mechanics. What we call quantum mechanics may therefore be less a description of ultimate reality than the lawful signature of bounded observers embedded in an empirical order they cannot exit.

Two disambiguations guard this map against standard misreadings. First, the ψ -ontic/ ψ -epistemic question of the ontological-models literature [8] lives entirely *within* the empirical level: it asks whether the quantum state corresponds to a physical state of the system or only to information about some deeper physical state λ . In that vocabulary IOF is *ψ -ontic*: the field $|\psi\rangle$ is part of the empirical model and never collapses, so theorems in the style of Pusey–Barrett–Rudolph [26], which constrain ψ -epistemic hidden-variable models, do not bear against the framework—IOF requires precisely the wavefunction reality such results support. What is

epistemic in IOF is *collapse*: the record-relative update of a finite observer, not a physical jump anywhere in the model. Second, “epistemic collapse” is itself used in a layered sense. Within empirical physics, it means that collapse is not an objective physical jump but a record-relative update for a finite observer. At the deeper metaphysical level (Part VI), even the empirical structures to which physics assigns ontic status—wavefunction, apparatus, observer, and outcome—are themselves real only as appearance, not as ultimate substance. The two senses are compatible but not identical, and the corpus keeps them apart.

Even full confirmation of the operational layer would leave deeper questions standing: why deterministic evolution at the fundamental level, why embedded observers have finite effective basis-tracking capacity, and why this particular relationship between C_{eff} , the instability rate, and observer-relative visibility holds. For this author these pass beyond physics into metaphysics—the domain of Part VI.

19 IOF Among the Interpretations

Since the framework’s status is interpretive (Section 12), it should be placed honestly among its peers. Each of the major no-collapse readings of quantum mechanics is organized around one characteristic move:

Interpretation	Characteristic move
Many-Worlds	unitary evolution only; outcomes as branching
Bohmian mechanics	definite trajectories; a guiding equation
QBism	probability as agent-relative credence
Relational QM	states as relations between systems
IOF	finite embedded observers sampling a constrained global history

IOF’s lens is that quantum probability is the form taken by inaccessible dependency for an embedded, finite sampler. The claim can be put precisely. Prediction does not require merely knowing the question and the answer after the fact; it requires knowing the relation that made *this* answer follow from *this* question. Writing the outcome as $A = f(Q, \lambda)$ —with Q the measurement question (the realized basis), A the recorded outcome, and λ the hidden dependency joining them within the global history—a finite observer can register both ends of the event, Q and A , but not λ . Once that dependency is unavailable, it cannot be used for prediction or control. From the inside, the setting therefore looks free, the outcome looks random, and the missing relation is represented by probability: the observer’s working theory is $P(A | Q)$.

This statement might be philosophically strong but not empirically discriminating, and the framework says so: standard quantum mechanics also delivers $P(A | Q)$, treating it as operationally complete. What IOF adds at this level is an account of *why* an embedded observer’s theory must take that form, with the control law of Part III quantifying how far the observer’s grip on Q ’s own provenance extends. Like its peers in the table, the lens is meaningful, coherent, and not uniquely provable.

The taxonomy underneath that admission deserves one paragraph of its own, because the question it answers—*is the framework falsifiable?*—is the one most often put to it, and the answer is structured, not flat. Interpretations and rival theories are different kinds of object. An interpretation must reproduce the quantum predictions exactly and is adjudicated by argument; an objective-collapse theory must deviate from them somewhere—a collapse that changed nothing measurable would not be a collapse—and is adjudicated by experiment. Recoverability is that dichotomy made operational: unitarity only ever *displaces* information, so contrast lost to one standpoint is conditionally recoverable from another, while genuine collapse *destroys* information and leaves nothing to recover from any standpoint. IOF as a whole sits in the first column by construction—it makes no empirical prediction beyond standard quantum mechanics (Section 12)—and is therefore not falsifiable as a whole, by design rather than by accident. Its falsifiable content is modular and named: the operational control law, with the calibration-first caveat that a failure there indicts the apparatus model before the framework [4]; the Fisher-homogeneity premise of the conditional Born derivation, exposed in the falsifying direction only [3]; the registered chaotic-corner module, carried with the prior strongly against it [4]; and the flat-marginals consistency check (Section 11.3). The one reading on which the framework *would* have been falsifiable as a whole—the strong reading—was stated precisely, tested against existing data, and excluded. Asking for more falsifiability than this is asking the framework to re-assert the claim it has already retired.

Yet this is also where IOF parts company with the rest of the table, and the difference is not greater truth but greater *exposure*. Many-Worlds, Bohmian mechanics, QBism, and Relational QM are each adjudicated by argument alone—no apparatus reading bears on which lens is right. IOF carries an operational sublayer whose parameters— C_{eff} , h_{KS} , σ_θ , and the deficit κ that combines them—are built and measured rather than merely posited, and that sublayer does work outside the framework: it supplies the calibrated reference channel and the recoverability classifier for the Penrose objective-reduction test (Section 14.4).

That exposure is also distributed the right way around. The modules carrying the falsifiable content—the operational control law and the finite-record premises—are the framework’s center of gravity, not its periphery, while the hosting embedding is scaffolding, held exactly as lightly as its unfalsifiability warrants: swap the host—single deterministic history, pilot-wave, Everett—and both the machinery and its exposure survive unchanged. The hosting embedding is accordingly not the framework’s falsification burden. No experiment adjudicates among empirically equivalent embeddings, so that burden was never dischargeable by anyone; what the framework owes on the embedding side is coherence, economy, and the naturalization of its declined independence premise—not a discriminator it never claimed to possess.

The outcome of that test does not discriminate IOF from quantum mechanics—the classifier is standard-quantum machinery, and its success or failure leaves that question untouched. It is not irrelevant to the framework, though. A clean verdict, in either direction, would be reached *through* IOF’s factoring of collapse into a recoverable observer-relative component and a possible irreducible remainder, and would thereby test whether the framework’s accounting of the observer-side channel holds in the mesoscopic regime (Section 14.4).

20 One Mechanism Across the Foundations

The basis-tracking deficit κ recurs across problems usually treated separately. The same finite-observer mechanism that reframes collapse also gives a reading of ontology, the Born rule, entanglement, the onset of classicality, and the structure of ordinary experience; one lens, not one explanation—several of these (the joint correlations above all) are hosted, not derived:

Ontology and the Born rule. The host supplies definite records, unitary continuity, and the quantum correlation structure; IOF adds only the observer’s self-ignorance about its realized basis. One element usually *postulated* is instead *derived*: the binary single-system Born weight $\cos^2(\theta/2)$ is conditionally reconstructed from finite-record Fisher–Rao geometry in a companion paper [3], retiring it as a separate ontic postulate rather than inheriting it. The derivation is conditional on explicit bridge assumptions—a structural reconstruction, not an independent empirical derivation of quantum mechanics—but where it holds, the framework asks its host to postulate one thing fewer.

Entanglement. The Bell correlations are not explained by a new signal; they are inherited from the host and read, in the deterministic embedding, as correlations within one globally consistent history (Section 11). Measurement independence is not automatic in such an embedding—setting and system share ancestry, though nothing is coordinated trial-by-trial; no-signaling is preserved exactly, single-party marginals stay flat, and finite basis access adds only a symmetric, capacity-dependent attenuation of the joint visibility. “Spooky action” was never the puzzle—the puzzle was treating spatially separated descendants of one history as independent.

The classicalization boundary, and gravity. Because the cut is set by the observer’s tracking budget, the boundary at which coherence is lost in the observer’s records acquires an observer-side coordinate, not only a mass one. This brings IOF into contact with Penrose’s gravitational objective reduction, whose timescale is mass-set; for mesoscopic systems the two land in the same window (Section 14). The paper treats that coincidence as numerical, but it makes the framework’s control variables the natural calibration axis for a clean mesoscopic test of OR—which is where the experimental program’s live discriminating power resides.

Biological observers and lived experience. The same arithmetic, evaluated with biological-scale capacities— $C_{\text{eff}} \approx 10$ bit/s against $h_{\text{KS}} \approx 50$ s⁻¹—gives $\kappa \approx 43$ s⁻¹ and a tracking timescale of tens of milliseconds, the range on which human perception and conscious action are integrated. A finite observer built from biological-scale capacities should classicalize on biological-scale times; from the inside, that is a world given as already definite, classical, and action-oriented, its measurement basis lived through rather than observed. That reach has a ceiling: from its own records such an observer can fix only the binary Born weight, read at low contrast, while the phase and Hilbert-space structure that would show the world as quantum stay hosted, beyond what finite records can represent (Section 13). This is the most speculative and least evidential of the readings—it is another application of the same κ -scaling, not an independent test—but the same finitude that throttles a laboratory tracking loop also shapes the texture of ordinary experience.

These are not four separate claims that happen to share a vocabulary. They are four readings of one structure: a finite observer embedded in, and unable to step outside, the history it measures.

The unifying suggestion can then be stated plainly. Much of what quantum mechanics treats as irreducible—the randomness, the collapse, the location of the cut, the onset of classicality, perhaps even the texture of ordinary experience—may be the lawful signature of finite, embedded observation rather than a feature of the world. And with that the argument returns to where the paper began. Part I opened not with physics but with *avidyā*—the self-ignorance that Advaita names as the root from which the appearance of a separate observer and an external world arises. The single mechanism gathered here is that intuition given physical, quantitative form: a finite observer that cannot reconstruct the realized value of its own measurement basis, and so meets one definite history as a spread of possibilities.

Part VI

A Non-Dual Perspective

This section explores connections between the self-ignorance framework and Advaita Vedanta, particularly as articulated in the teachings of Ramana Maharshi and, among contemporary expositors, the lectures of Swami Sarvapriyananda [27]. We follow the interpretive approach of Michael James [28], emphasizing that physics describes the structure of empirical reality (vyāvahārika satya) while Vedanta addresses the nature of ultimate reality (pāramārthika satya). The parallels we draw are structural—showing how mathematical relationships within appearance may mirror metaphysical principles—without claiming that physics proves or exhausts Vedantic insights.

Status of this part: The claims in this part are interpretive and contemplative, not independent scientific evidence for IOF. The framework’s operational content remains the observer-capacity model and its measurable control law—reference-frame physics within quantum mechanics (Section 12)—and its interpretation stands as an interpretation. The non-dual reading is included to preserve the source intuition and to explore meaning, not to raise the physical reading to a proof of Vedanta or vice versa.

More precisely, this part is written from the non-dual standpoint that motivated the framework. Consciousness is treated here as the given subject of appearance, not as an object requiring physical proof. The Vedantic language supplies an interpretive grammar for a universe understood epistemically: a way of naming the relation between subject, object, limitation, and appearance. Where the prose speaks strongly, it should be read as describing the consequences of the IOF basis-tracking picture taken as an interpretation, together with the philosophical lens that inspired the formalization.

21 A Note on My Method

I write as an independent researcher, and my method is not the textbook one. I look for *structural* correspondence between what the sages of Advaita describe and what physics has established, and I let that correspondence guide where the formalism should go; here the order of discovery was reversed—the metaphysics came first and guided the physics, not the other way around. The starting conviction is Advaitic: there is no independent doer, and what feels like free choice is apparent agency unfolding inside a single causal order—at the empirical level, what the tradition calls prārabdha. Taken seriously, this means the *setting* of a measurement—the angle, the basis, the question put to the world—cannot be a free variable standing outside the world. It is one more event inside the same ancestral history as the system it interrogates.

That intuition has a precise technical address. In Bell’s theorem it is the assumption of *measurement independence*: that the hidden state of the system is statistically independent of the later setting. Bell himself, in *Free Variables and Local Causality* (1977), already saw that the freedom of those variables was doing quiet work. From the non-dual side the assumption looks not merely optional but artificial—if the chooser is inside the world, why should its choice be

uncorrelated with the world? Penrose sharpens the same point operationally: one cannot ask Alice for her spin in the abstract; one must first name an axis, so the angle θ is part of the physical question, not a label hovering above it.

IOF is what results from carrying this through, and the discipline matters as much as the intuition. The defensible claim is narrow but hard to dismiss: in a single deterministic history θ and the system state need not be independent, so measurement independence is not a primitive of nature but an idealization requiring justification. To that IOF adds a finite-observer mechanism— θ is physically realized but not fully self-trackable—and an operational form that can be built and measured in the laboratory: the capacity-dependent, recoverable visibility law that BLQC realizes as a control experiment. The chain runs Vedanta \rightarrow Bell \rightarrow Penrose \rightarrow IOF \rightarrow BLQC: an ontological intuition, its technical pressure point, its operational form, a concrete mechanism, and a measurable control law. The sections that follow trace it in the other direction—from the formalism back to the intuition that suggested it.

22 Consciousness as Fundamental (Sat-Chit-Ānanda)

Two Kinds of Self-Awareness

Ramana Maharshi distinguishes between two fundamentally different modes of awareness:

Pure self-awareness ($\bar{a}tma-svar\bar{u}pa$): Consciousness knowing itself as itself, without the division into knower and known. This is our fundamental nature— $sat-cit-\bar{a}nanda$ (being-awareness-bliss)—which requires no mediation and cannot be objectified.

The ego’s attempted self-knowledge: The “I am this body-mind” thought attempting to know itself as an object, necessarily creating subject-object duality.

The framework’s “observer” is always the latter—the ego, not pure consciousness. The observer is a limited form of awareness that has identified with a physical system (neurons, apparatus, measurement degrees of freedom), and the information-theoretic bounds we’ve derived apply to the ego’s structure, not to consciousness itself. Read concretely, those bounds are a limit on *sampling*: the ego can register its own state only by sampling it through a finite channel, and any finite channel *undersamples*—it resolves fewer distinctions than the state actually contains (the undersampling map of Section 10.3). The self-ignorance is just that unsampled remainder—what the ego cannot resolve about the very state it has identified with; in the Vedantic reading, $avidy\bar{a}$.

The Self, Ego, and Jīva: Technical Clarifications

This section addresses the precise Vedantic distinctions that inform the framework:

The Self = $\bar{a}tman$ = brahman: The one subject, pure consciousness. It cannot know itself as an object—not because of any limitation, but because it IS the knowing subject. Any attempt to objectify the Self creates the appearance of subject-object division, which is $avidy\bar{a}$ itself. The Self is self-aware, but this awareness is not mediated through objectification; it is direct, immediate, non-dual knowing by being.

The ego: The “I am this body” thought. Crucially, the ego is not a knowing principle but ignorance (avidyā) itself. It is non-existent—a formless phantom that appears to exist only through avidyā. When subjected to the inquiry “Who am I?”, the ego dissolves, revealing only the Self. It has no reality to be known.

The jīva (individual): Ātman + avidyā—the Self apparently limited by ignorance. From the jīva’s perspective, self-knowledge appears incomplete: I cannot trace *why* my thoughts arise as they do, *why* attention moves as it does, *why* choices emerge from accumulated tendencies (āsanās). This incomplete self-knowledge is not a property of consciousness itself but of the apparent limitation (avidyā).

What the framework models: When we formalize “observers with limited self-knowledge,” we are modeling the jīva—the apparent individual operating through subject-object division. The information-theoretic constraints (effective basis-tracking capacity C_{eff} bounded by physical resources, tracking timescales τ_{fill} and t_{break}) apply to this apparent limitation, not to the Self.

Why the Substrate Cannot Be “Known”

Pure consciousness (sat-cit) doesn’t “observe” itself—it IS itself without any subject-object division. The moment observation occurs, duality has already arisen.

The framework’s deterministic substrate—the universal state evolving according to deterministic laws—might be understood as pointing toward what Vedānta calls sat (being). But physics can only describe patterns within the appearance of that being, never the nature of being itself.

In Section 18, we noted: “Observers cannot access that external view of themselves, so they can never directly verify the deterministic ontology.” This parallels a fundamental Vedantic principle: the ego cannot know ātman (self) as an object, because ātman is the subject, the knowing itself. Any attempted objectification creates only another appearance within ātman, not knowledge of ātman.

The Necessary Limitation of Physics

Physics, however complete, describes vyāvahārika satya (empirical reality)—the structure of appearance as experienced by bounded observers. It cannot describe pāramārthika satya (absolute reality)—the nature of that which appears.

The framework makes this limitation explicit: quantum mechanics is operationally complete as epistemology (prediction theory for bounded observers) while acknowledging something beyond epistemic access (the deterministic substrate). But even that substrate is described in physical terms—states, evolution, determinism—which are still concepts within appearance, not consciousness itself.

Ramana frequently emphasized that reality is neither the material world (jaḍa) nor the insentient body, but the consciousness that makes all experience possible. Physics, operating with mathematical structures describing matter and energy, remains within the realm of appearance.

23 Self-Ignorance (Avidyā) and the Appearance of Duality (Māyā)

Two Levels of Self-Ignorance

The Vedantic term *avidyā* (ignorance) refers to fundamental self-ignorance: mistaking oneself (pure awareness) to be a body-mind. This is the root of all appearance, the mechanism by which the non-dual One seems to become many.

The framework identifies a physical self-ignorance: the observer-system cannot trace *why* its physical state $\theta(t)$ evolved to this particular value—the causal chain from past tendencies to present choice is hidden.

These two self-ignorances are related but distinct:

The physical self-ignorance (not knowing *why* θ took this value) is a manifestation, within appearance, of the more fundamental metaphysical self-ignorance (not knowing oneself as consciousness).

Mathematically, we showed that the observer’s self-ignorance is set by how its effective basis-tracking capacity C_{eff} compares with its internal unpredictability h_{KS} . In the capacity-wins regime the observer converges to accurate basis tracking; in the chaos-wins regime it loses track of the basis after a finite time, beyond which self-ignorance makes the basis unpredictable; in the diffusive regime a steady-state uncertainty bounds the achievable precision. The threshold, the timescales, and the variance bounds that make these statements precise are derived in the mathematical formalism (Section 9.6).

These are not contingent limitations that better engineering might overcome. They are intrinsic to any physical system that:

- Has internal dynamics determining its interaction with the world
- Attempts to monitor those dynamics through internal sensors
- Operates under finite physical resources, with thermodynamics bounding the raw capacity available to the tracking loop

But why should physical systems be structured this way?

The Primordial Division

From the Vedantic perspective, the fundamental error is the rising of the ego—the first thought “I am this body.” This thought creates the knower/known division from which all multiplicity follows.

The “measurement basis” is already within this division. It represents how the ego-as-physical-system orients itself toward the world. The information bounds don’t explain *why* duality appears in the first place (that question may transcend physics), but they do describe precisely *how* duality operates once it has arisen.

The finite effective capacity C_{eff} is the quantitative expression of a qualitative truth: the ego,

having assumed the limitation of form, operates under the constraints of that form. Having identified with matter, it inherits matter’s limitations.

Māyā’s Two Powers

Classical Vedānta describes māyā as possessing two powers (śaktis):

Āvaraṇa-śakti (veiling power): Conceals the non-dual whole, making the infinite appear finite

Vikṣepa-śakti (projecting power): Projects the appearance of multiplicity from the concealed unity

The framework provides a mathematical structure for these powers:

The finite effective capacity C_{eff} represents āvaraṇa: The observer cannot access complete information about the system (including their own state). The whole is veiled—only partial, delayed information is available.

The measurement process represents vikṣepa: From quantum superposition (the indeterminate potential), the measurement manifests a particular outcome. The one state becomes one of many classical outcomes.

The mathematical relationship:

$$V = \exp\left(-\frac{\sigma_\theta^2}{2}\right) \quad (45)$$

shows how the veiling of self-knowledge (uncertainty σ_θ) directly reduces the visibility of quantum coherence—the apparent definiteness of the classical world emerges precisely to the extent that self-knowledge is veiled.

Two Channels of Veiling: Attachment and Self-Loss

This single visibility factor conceals a finer structure. The contrast an observer actually measures factors into two independent attenuations (Section 5.2 and the visibility analysis of Part II):

$$V_{\text{obs}} = V_{\text{std}} \cdot V_{\text{IOF}}, \quad V_{\text{IOF}} = \exp\left(-\frac{\sigma_\theta^2}{2}\right). \quad (46)$$

Read soteriologically, these are the two strands by which the One veils itself as many, and they answer to the two classical disciplines of liberation.

V_{std} —bondage through attachment. The standard channel falls as the system entangles with its environment: which-path information leaks outward, and the world “records” which alternative occurred. This is bondage through grasping, awareness scattered into objects. Its discipline is self-surrender (śaraṇāgati)—the giving-up of the ego’s desires and will—of which *vairāgya* (dispassion) is the affective core: the loosening of the hold the object-world has on attention.

V_{IOF} —bondage through self-loss. The IOF channel falls as the observer loses track of its own basis θ : not information leaking outward, but the center going unheld. This is bondage through forgetting oneself. Its discipline is *ātma-vicāra* (self-inquiry) and abidance in the center (*niṣṭhā*).

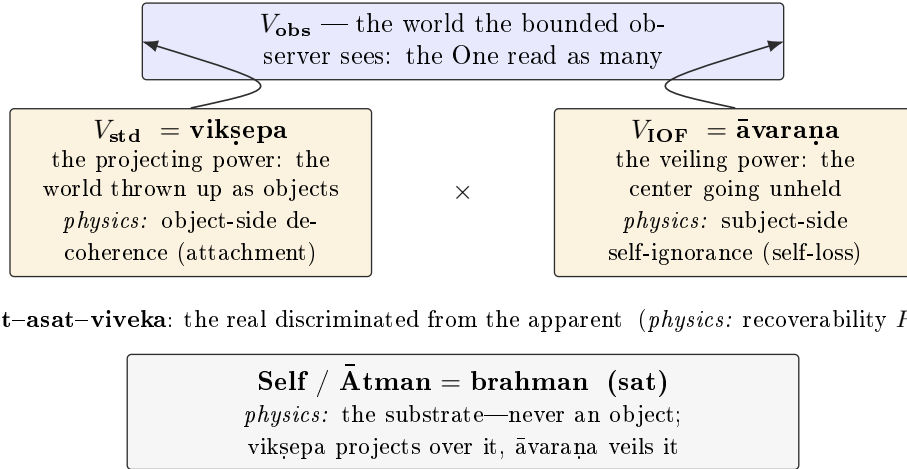


Figure 2: The two powers of māyā, read from the visibility factorization. The contrast a bounded observer measures factors as $V_{\text{obs}} = V_{\text{std}} \cdot V_{\text{IOF}}$, and the two factors answer to māyā’s two śaktis: V_{std} to *vikṣepa*, the projecting power (object-side decoherence, the world thrown up as many), and V_{IOF} to *āvaraṇa*, the veiling power (subject-side self-ignorance, the center unheld). Their product is the appearance of duality, quantified—contrast is lost if either factor falls. Below the discrimination of the real from the apparent (*sat-asat-viveka*, operationally the recoverability classifier) lies the Self—ātman, brahman, sat—over which *vikṣepa* projects and which *āvaraṇa* veils, but which is never itself an object. Structural resonance, not identity.

The product demands both. Because $V_{\text{obs}} = V_{\text{std}} \cdot V_{\text{IOF}}$, the contrast is lost if *either* factor falls; full coherence requires non-attachment to the world *and* steadiness in the Self. This is the ancient claim that *vairāgya* without self-knowledge is dry and self-inquiry without dispassion is unstable—here as a product, not a slogan. Yet in Ramana’s teaching, as Michael James reads it, the two remedies are finally one turning: the ego is given up only by being traced to its source, so self-surrender is fulfilled only as self-investigation, and self-attention is itself the truest surrender (*Nāṇ Yār?* §13; *Ulladu Nārpadu* v. 26) [28]. Dispassion then ripens of itself as attention turns inward, so a single movement tends both factors toward unity—though the framework, as physics, keeps them two, a product whose factors need not move together.

Which path is the primal duality. When V_{obs} is high there is no fact of the matter about which path was taken: no separation, no “two.” Collapse into a definite alternative *is* the moment the One reads as many. “Which path?” and “one without a second” (*advaita*) are the same recognition in two vocabularies—the physicist asking which of two, the Vedantin affirming there was never a second; the loss of visibility is the appearance of duality, quantified.

A discipline on the analogy. This is structural resonance, not identity, and one guard must stay explicit. In the laboratory, high visibility is a *condition*, not a liberation; the physics is silent about which state one should want. The soteriological direction—that clarity is to be sought—is supplied by Vedanta, not by the equations. And θ is a measurement *basis*, a physical variable, not the ātman: what σ_θ measures is the ego’s lost grip on its own standpoint—the *chit-jaḍa-granthi* of Section 25 tightening—not an uncertainty in the Self, which is self-luminous and cannot be uncertain about itself. With those provisos, the decomposition gives the equation of māyā in two strands: the rate at which the One, grasping outward and forgetting inward, mistakes itself for two.

The limit of the self-knowledge strand: $V_{\text{IOF}} = 1$. The two strands invite a question

the formalism answers cleanly: what is the state of an observer for whom one strand has gone fully to unity? Take the self-knowledge strand to its limit. Since $V_{\text{IOF}} = \exp(-\sigma_\theta^2/2)$, the condition $V_{\text{IOF}} = 1$ is exactly $\sigma_\theta = 0$ —zero residual uncertainty about one’s own realized basis, the center perfectly held. The IOF channel then contributes no attenuation, and what remains, $V_{\text{obs}} = V_{\text{std}}$, is ordinary quantum mechanics. This is *ātma-vicāra* carried to its end on the observer side: self-inquiry that leaves no gap between the basis one holds and the basis one is.

Yet the same formalism forbids a finite ego from *occupying* this state. $\sigma_\theta = 0$ requires $C_{\text{eff}} \rightarrow \infty$ (Section 25); a bounded observer may approach it but cannot reach it by tracking harder, because no finite system can become a complete object of its own record-forming. The limit is therefore not a state the ego attains but the point at which the ego, as a separate and ignorant knower, ceases to be. The physics closes the door from inside duality; what opens it is not more information but the dissolution of the one who lacked it—*granthi-bhedana*, the cutting of the knot. Read on the knowledge strand, $V_{\text{IOF}} = 1$ is the *jñānī*’s condition: not a finite mind that has at last finished its tracking, but the *jñānī* with no residual self-ignorance left to make it seem a separate self.

Two disciplines keep this exact rather than rhapsodic. First, $V_{\text{IOF}} = 1$ is *observer-side* and one-sided: it heals only the self-loss strand (*vicāra*) and says nothing about V_{std} , the world-attachment strand (*vairāgya*), which remains an independent ledger. Self-knowledge may be complete while the world still pulls; full coherence, $V_{\text{obs}} = 1$, would demand both factors at unity—both remedies fulfilled. That the framework forces a *product*, not a sum, is exactly why one strand without the other does not suffice. Second, σ_θ measures the ego’s grip on its own standpoint, never an uncertainty in the Self; the Self was never the ignorant party, so what the limit marks is recognition, not acquisition. $V_{\text{IOF}} \rightarrow 1$ is where, on the observer side, that recognition would land.

23.1 Why One Appears as Many

Attention as the Mechanism of Manifestation

Ramana taught that the ego’s essential nature is attention (*svarūpa-dhyāna* means self-attentiveness). When attention grasps objects (anything other than itself), the appearance of multiplicity is sustained. When attention turns back toward itself alone, the ego subsides and only self-awareness remains.

Measurement is formalized attention—the observer’s awareness manifesting as a particular physical orientation toward a quantum system. The measurement basis θ represents the direction of attention, now appearing as physical orientation of magnets, polarizers, or detector geometry. As explained in the previous section, the observer cannot know *why* their attention (measurement basis) evolved to its current value—they experience the choice but cannot trace its causal origin from *vāsanās* (past tendencies). This fundamental self-ignorance about attention’s movement is what creates apparent randomness in quantum measurement.

Projection Without Conspiracy

In Section 11, we addressed the “conspiracy” objection to Bell-violating hidden variable theories: measurement settings and particle states must be correlated through their common past, which

physicists call “fine-tuning.”

But from a Vedantic perspective, there is no conspiracy. The entire appearance—particles, observers, measurement settings, outcomes—is a unified projection of māyā. The apparent separation in space and time is itself part of the projection.

The mathematical resolution:

$$\rho(\boldsymbol{\xi}, \mu, \mu') \neq \rho(\boldsymbol{\xi}) \cdot \rho(\mu) \cdot \rho(\mu') \quad (47)$$

simply states that measurement “choice” (encoded in observer microstate μ) and particle state $\boldsymbol{\xi}$ are correlated through deterministic evolution from common initial conditions. That such a correlation is *possible* is natural in any deterministic universe; that it takes the specific Bell-reproducing form is not—that is left to the hosting embedding.

The word “conspiracy” only seems appropriate if one implicitly assumes libertarian free will—that measurement settings are somehow independent of the physical causal chain. But this contradicts the very materialism that most physicists espouse.

From the non-dual view: There is One appearing as many. The “many” are not independently existing fragments that must be carefully coordinated. They are multiple aspects of a single whole. The correlations we call “entanglement” are simply the mathematical signature of this unity when viewed through the lens of separation.

The Mechanism of Individuation

Why does One appear as individual observers at all?

The framework suggests: Complete self-knowledge would dissolve duality. Mathematically, if $\sigma_\theta \rightarrow 0$ (perfect self-knowledge), the observer knows exactly what they’re measuring, and the boundary between “observer” and “observed” becomes arbitrary—both are just patterns in the deterministic substrate. The experience of being a separate agent making free choices dissolves. This parallels Ramana’s teaching: when attention turns fully toward itself, when the ego investigates its own nature completely, it dissolves in self-awareness—recognized to have never existed as a separate entity. The finite effective tracking capacity C_{eff} is therefore not a bug but the essential structure that maintains duality in the model. If complete self-knowledge dissolves separation, then separation requires incomplete self-knowledge. The limitation is necessary, not contingent. (The mathematical details and connection to thermodynamic bounds are developed fully in Part III.)

23.2 Physics as the Structure of Apparent Separation

The Reality and Unreality of the World

Ramana taught that the world is real as appearance but unreal as an independently existing substance. The dream world is perfectly real within the dream; only upon waking does one recognize it was merely appearance in consciousness.

The framework captures this precisely:

Ontological layer: Deterministic substrate (universal state evolving according to fixed laws)—this is “real” in the sense that it exists and evolves deterministically.

Epistemic layer: Quantum mechanics as prediction theory for bounded observers—this is “real” in the sense that it correctly describes what any bounded observer must experience.

The substrate is real as appearance; unreal as independent substance. Why? Because even the substrate is described in physical, mathematical terms—wavefunctions, Hilbert spaces, deterministic evolution. These are still concepts, still within the domain of objectification, still appearance.

What is ultimately real, from the Vedantic perspective, is consciousness alone—sat-cit-ānanda—which is not an object that can be described but the subject that makes all description possible.

Empirical Success Without Ontological Truth

A physicist might object: “If quantum mechanics is merely epistemic, why does it work so perfectly? Why can we predict outcomes to ten decimal places?”

The framework provides the answer: QM works perfectly because it correctly describes what bounded observers with finite information capacity must experience when interacting with deterministic but epistemically inaccessible dynamics.

The finite-basis channel does not replace the Born rule; it modulates the contrast with which the fixed-basis quantum structure is read by a bounded observer. In the toy binary visualization,³ with $\boldsymbol{\xi}$ a unit Bloch-vector-like ontic state ($\|\boldsymbol{\xi}\| = 1$), the single-outcome weight takes the finite-resolution fringe form

$$\Pr(\uparrow | \theta) = \frac{1 + V_{\text{IOF}} \boldsymbol{\xi} \cdot \mathbf{n}(\theta)}{2}, \quad V_{\text{IOF}} = e^{-\sigma_\theta^2/2}, \quad (48)$$

where $V_{\text{IOF}} = e^{-\sigma_\theta^2/2}$ is the contrast that survives Gaussian averaging of the fringe over the observer’s unresolved θ (uncertainty σ_θ); the threshold $\text{sgn}[\boldsymbol{\xi} \cdot \mathbf{n}(\theta)]$ names only the definite single-shot outcome, not the origin of the weight. The fixed-basis quantum statistics are taken as given from the hosting embedding, and the single-system binary Born weight is treated conditionally in the companion derivation [3]; what the finite-basis channel contributes is the visibility factor V_{IOF} , attenuating the fringe contrast as basis tracking degrades.

QM is operationally complete as epistemology (for observers subject to the constraints we’ve identified) while being incomplete as ontology (not describing the substrate directly).

This mirrors Ramana’s teaching about the dream: The dream world obeys perfect internal consistency—gravity works, water flows downhill, cause precedes effect. This consistency doesn’t make the dream ultimately real. It makes it a coherent appearance.

Similarly, quantum mechanics’ empirical success doesn’t prove there’s nothing beyond it. In the IOF reading, it suggests that QM is a complete theory for what appears to bounded observers—with the textbook fixed-basis idealization operationally available when capacity wins, and the

³An animated realization of this toy visualization—the *IOF Double-Slit Animation*, with its auditable source code—is available on the project website, ignorantobserver.xyz, under *Empirical Test*. It shows how finite basis tracking suppresses interference visibility.

same theory still holding, under an unresolved reference frame, when chaos wins.

Degrees of Reality (Sat-Asat-Viveka)

Classical Vedānta distinguishes three degrees of reality:

- **Pāramārthika satya**: Absolute reality—only consciousness (brahman)
- **Vyāvahārika satya**: Empirical reality—the waking world
- **Prātibhāsika satya**: Apparent reality—dreams, illusions

Where does the deterministic substrate fit? It's more real than individual quantum outcomes (which are epistemic), less real than consciousness (which is non-objective).

Turiya: The Substratum of All Degrees of Reality

The Māṇḍūkya Upanishad teaches that there are three ordinary states of experience—waking, dreaming, and deep sleep—and a fourth (turiya) which is not a state among states but the substratum of all three: unchanging, unobjectifiable, the witness-ground that makes them possible.

The deterministic substrate in the framework occupies an analogous position. It is not Īśvara (cosmic lord or creator, which remains within appearance), but something closer to turiya—the unchanging ground beneath all changing phenomena:

- **Individual quantum outcomes** ↔ **Waking perceptions**: What the jīva experiences as definite facts
- **Quantum superpositions** ↔ **Dream-like states**: Indeterminate, probabilistic, not fully manifest
- **Unobserved substrate** ↔ **Deep sleep**: Beyond the subject-object distinction, no individual awareness
- **Deterministic substrate** ($|\psi\rangle, \{\xi_i\}$) ↔ **Turiya**: The unchanging ground underlying all three, never objectifiable yet determining everything

This last correspondence—deterministic substrate ↔ turiya—is the most delicate of these correspondences, and we examine it carefully in Subsection 23.3. Here we note only its structural role: the deterministic substrate ($|\psi\rangle, \{\xi_i\}$) is the deepest physical correlate of the unmanifest that the framework can describe—the boundary-facing ground from which definite outcomes, superpositions, and unobserved dynamics appear to bounded observers, and never itself one of those appearances.

Two Vedāntic terms attach to this one ground, and they are not interchangeable. As the causal seed-state—*kāraṇa*, the unmanifest store of *vāsanās* from which manifestation unfolds—it is world-facing, a designated layer within the grammar of appearance. As the boundary-facing substratum of waking, dream, and deep sleep—*turīya* in the sense fixed in Subsection 23.3—it is the deepest reflected ground, still spoken of from the world side of the mirror. *Kāraṇa* and

turiya are two faces of the single deepest describable ground, not two strata: physics resolves only the one substrate where the tradition distinguishes its seed-face from its substratum-face. Neither is brahman, the Absolute (sat-chit-ānanda), which lies on the far side of the mirror and is no layer at all.

But what mechanism creates this boundary between substrate and manifestation? Ramana taught that the ego is the hṛdaya-granthi (heart-knot) or chit-jaḍa-granthi (consciousness-matter knot)—the apparently binding link between pure consciousness (chit) and inert form (jaḍa). We will see in Section 25 that the quantum-classical boundary, where $t_{\text{break}} \approx \tau_{\text{OR}}$, may be understood as the place where this knot’s veiling becomes physically manifest: the information-theoretic threshold where the One becomes definitively veiled as many.

23.3 Turiya and Brahman: A Crucial Distinction

To avoid confusion, we must be precise about the relationship between turiya and brahman—and about what the substrate–turiya correspondence does and does not claim.

The comparison must therefore be made with care. In this framework, turiya is not used as a synonym for brahman in the absolute sense. It names the boundary-facing, experiential aspect of the unmanifest: the “fourth” understood in relation to waking, dream, and deep sleep. It is the substratum of these states as they appear to experience, but it is still spoken of from within the grammar of states, witness, and manifestation.

The deterministic substrate in IOF occupies an analogous position within physics. It is not Īśvara, a cosmic person or creator, and it is not brahman as pure sat-chit beyond all relation. It is the deepest physical correlate of the unmanifest that the framework can still describe: the underlying order from which definite outcomes, superpositions, and unobserved dynamics appear to bounded observers. Like turiya in its boundary-facing sense, it is not one more object among objects, but the condition relative to which the three ordinary modes of appearance can be distinguished.

Just as turiya makes waking, dreaming, and deep sleep intelligible from the standpoint of experience, the deterministic substrate makes quantum phenomena intelligible from within the physical model. But the analogy stops there. Physics can describe this substrate only as structure within appearance: fields, histories, hidden variables, laws, and constraints. It cannot describe brahman, pure sat-chit, which is not the witness of states as an objectifiable principle, but the non-dual reality in which even the distinction between witness, world, and substrate appears.

Thus IOF does not identify the deterministic substrate with brahman. At most, it points to a boundary-facing analogue of turiya within the empirical order: the deepest describable ground of appearance before it becomes waking-world definiteness. The next distinction is therefore essential: turiya belongs to the language of experience and manifestation; brahman, in the strict sense, does not.

The mirror-image can be stated more precisely through the Vedantic distinction between *chit* and *chidābhāsa*. Chit is pure consciousness itself, not reflected, not objectifiable, and not on the world side of the boundary. Chidābhāsa is reflected consciousness: consciousness as it appears

through mind, experience, world, and the structures by which manifestation is known. The granthi boundary may therefore be read as the mirror-line: on one side is chit, unreachable by objectification; on the other side is the entire order of reflected appearance.

In this sense, turiya should not be identified with brahman in the strict sense of pure sat-chit beyond all relation. Turiya belongs to the boundary-facing language of experience: it is the reflected, unmanifest substratum relative to waking, dream, and deep sleep. It is not one more object within those states, but it is still spoken of from the world side of the mirror, within the grammar of manifestation. Brahman, by contrast, is not reflected consciousness and not a substrate among substrates; it is the non-dual reality in which even the distinction between chit, chidābhāsa, turiya, world, and observer appears.

This distinction maps precisely to what physics can and cannot address:

- **Physics points toward the boundary-facing aspect of the unmanifest:** The deterministic substrate ($|\psi\rangle$, $\{\xi_i\}$) underlying quantum phenomena. This is the unchanging ground *conceptualized from within the framework of observation and experience*. We describe it mathematically (wavefunctions, hidden variables, evolution equations) in relation to what observers measure.
- **Physics cannot reach brahman:** Even describing the substrate as “deterministic” or “underlying” is still conceptualizing FROM experience, using concepts (causation, existence, substrate) that arise within appearance. Brahman is not a thing to be described but the reality that makes description possible—pure consciousness in which all conceptual frameworks arise.

The framework shows how physics, pushed to its limits, points toward the boundary-facing aspect of the unmanifest: an inaccessible ground underlying all observations. But recognizing this ground as consciousness itself (brahman, not merely a physical substrate) requires the turn from object-knowledge to self-knowledge—from physics to ātma-vicāra.

Physics describes turiya-structure (substrate from experiential standpoint); Vedanta reveals brahman (reality itself). Both are necessary; neither is sufficient. Physics shows the structure of limitation; self-investigation transcends limitation by recognizing what was never limited.

24 The Three-Fold Structure of Karma

The Vedantic doctrine of karma traditionally distinguishes three types, which illuminate the ontological structure we’ve formalized. This three-fold framework provides a technical vocabulary for understanding how the deterministic substrate, present manifestation, and future evolution relate within a single unified process.

Classical Three-Fold Distinction

- **Sañcita karma:** The accumulated store of all past actions and tendencies, from beginningless time; the total, unmanifest causal substrate.

- **Prārabdha karma:** The portion of that substrate which has ripened and must be experienced now; it determines the present circumstances, the thoughts that arise, and the actions the body performs. It is the unavoidable, unfolding script.
- **Āgamī karma:** “New” karma created in the present moment. Crucially, this is not the physical action itself (which is part of prārabdha), but the karma generated by the ego’s identification with that action—the mental act of claiming “I am the doer” and forming intentions for a future result.

This framework addresses a fundamental question: If everything is determined by infinite past causes (sañcita), why does the present moment feel open? Answer: Because we experience only what ripens now (prārabdha), while simultaneously creating seeds for the future (āgamī), even though all three are interwoven in a single deterministic process.

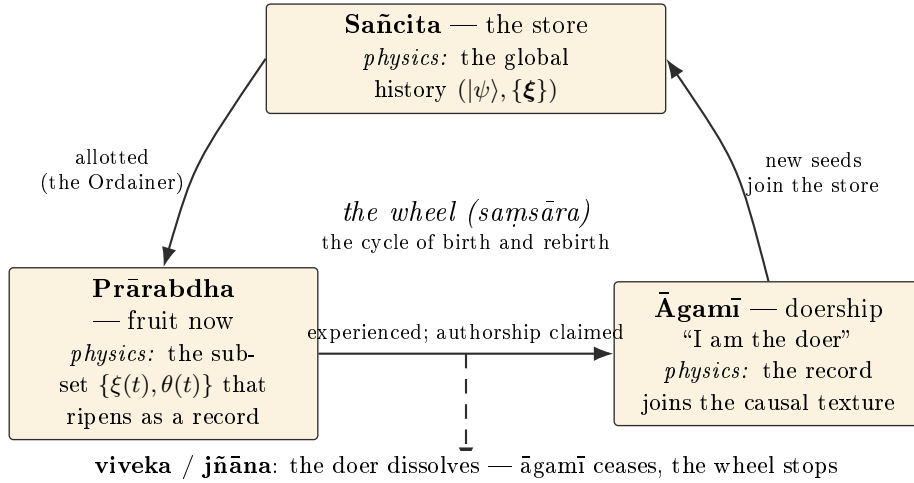


Figure 3: The three-fold structure of karma as the record-forming cycle. *Sañcita*, the total store, is the deterministic global history $(|\psi\rangle, \{\xi\})$; *prārabdha*, the portion that ripens now, is the subset $\{\xi(t), \theta(t)\}$ the finite channel registers as a record; *āgamī*, “new” karma, is the ego’s claim of authorship, by which each formed record joins the causal texture that conditions what can appear next. The single deterministic history is not rewritten—what advances is the observer’s record of it. The wheel turns while the doer is presumed; *viveka/jñāna*, the dissolution of that presumed doer, is where *āgamī* ceases and the turning stops. Structural resonance, not identity.

Sañcita Karma \leftrightarrow Universal Initial Conditions

Sañcita represents the total causal heritage—every action, tendency, correlation extending back through beginningless time. In the framework, this maps to the *complete* universal state at initial conditions (or in the infinite past):

$$\text{Sañcita} \equiv \{|\psi\rangle_{\text{universe}}(t_0), \{\xi_i(t_0)\}_{\text{all particles}}\} \quad (49)$$

This includes:

- The universal wavefunction $|\psi\rangle_{\text{universe}}$ —the field structure underlying all possibilities
- All hidden configuration vectors ξ_i for every degree of freedom in the universe

- All correlations, entanglements, the entire causal web

Crucial point: Even the wavefunction $|\psi\rangle_{\text{universe}}$ does not exhaust sañcita. The wavefunction describes only the field structure—what weights the possibilities. Sañcita includes the hidden variables $\{\xi_i\}$ that specify the complete and definite configuration of the universe at the empirical level. This is the substrate of definite states that exists at each moment, veiled from the observer, from which manifest outcomes will arise. In short: $|\psi\rangle$ is the law-like field weighting configurations, while $\{\xi_i\}$ are the actual states under that law. If a Vedantic analogy is useful, $|\psi\rangle$ is closer to *dharma* as lawful order, whereas ξ_i marks the particular karmic configuration realized within that order. The binary single-system weight is the one exception: the companion derivation reads it from the finite observer’s record geometry rather than from $|\psi\rangle$, which supplies the multi-outcome and joint weights.

Cosmic and individual sañcita: Traditionally sañcita names the *individual* store—the jīva’s accumulated karma awaiting fruit. The substrate above is the *cosmic* total: in the Katha’s ladder (1.3.10–11; 2.3.7–8) the individual intellect (buddhi) is subsumed in *mahat* (mahān ātmā, the great Self) and that in the unmanifest (avyakta)—the one ladder read both inwardly and cosmologically. The two are not rival accounts: the personal store is the individual localization (vyaṣṭi) of the one collective substrate (samaṣṭi), as the individual buddhi is a localization of mahat. What ripens for a given observer—its prārabdha, the subset $\{\xi(t), \theta(t)\}$ it actually encounters—is the portion selected and arranged by the lawful unfolding, which the tradition personifies as the Ordainer. Ramana put it so in the letter to his mother: “The Ordainer controls the fate of souls in accordance with their prārabdha karma. . . . Whatever is destined not to happen will not happen, try as you may. Whatever is destined to happen will happen, do what you may to prevent it.” The substrate is a determinate configuration, not a pre-written ledger: records are observer-side, formed by the finite channel that registers the subset.

Why Sañcita Is Unknowable

From physics: How unknowability operates

The same two roots that make an observer’s own measurement basis causally opaque (Section 4) make sañcita inaccessible:

1. **Self-reference (embeddedness):** The observer is *within* the totality it would know, so the state to be tracked includes the tracker, and a complete account would have to contain its own record-forming. This is not an infinite regress but a feedback limit fixed by self-reference—no sharper sensor removes it as a class. Self-reference fixes *what* must be tracked.
2. **Finite physics (thermodynamics):** Tracking is a physical process of finite effective capacity C_{eff} , bounded at the extreme by the Landauer ceiling. The information needed to fix every hidden state across beginningless time (anādi) exceeds that capacity without limit. Finite physics fixes *how fast*.

Together they make sañcita causally opaque—the physical “how” of unknowability.

From Vedanta: Why unknowability is grace

But Ramana’s teaching reveals a deeper truth about *why* this structure exists. When asked about how the ego arose or about past karma, he consistently redirected: “It is not necessary to know it. Know the present. Not knowing that, why do you worry about other times?”

From the Vedantic standpoint, the veiling of sañcita is not merely a defect; it is what leaves attention available for the present turn toward its source. If you could remember all actions, tendencies, and experiences from beginningless time—every thought, every desire, every outcome across infinite lifetimes—your attention would be hopelessly trapped in the *content* of experience.

Āvaraṇa-śakti (the veiling power) serves a purpose: By concealing the infinite past, attention becomes available for the present. And only in the present can attention turn inward toward its source. Remembering all of sañcita would prevent Self-realization, not enable it. The weight of infinite memory would bind consciousness to the very appearances it seeks to transcend.

This is why even the jnani, having realized the Self, does not acquire omniscient memory of all past karma. As Ramana taught, karma is transcended not by knowing it exhaustively but by recognizing the ego (the apparent experiencer of karma) as unreal.

Complementarity

The two perspectives complement rather than contradict:

- **Physics describes *how***: self-reference and finite, thermodynamically bounded capacity—the mathematical structure of unknowability
- **Vedanta explains *why***: The limitation is purposeful, not accidental—structured to allow turning attention inward rather than being lost in infinite content

You cannot trace the complete causal chain (sañcita) that led to this moment—and not merely because the past is too long to survey. The limit is structural: you are *within* the totality you would trace, so a complete account would have to contain its own tracing, and no sharper instrument removes that closure. The beginningless past (anādi) only compounds a limit that already holds here, in the present. Physics shows this is inevitable; Vedanta suggests it is merciful.

Prārabdha Karma ↔ What Manifests Now

Prārabdha is the portion of sañcita that has ripened into present experience. It represents what *must* happen now—what is already determined and cannot be changed, even though it was not predictable beforehand.

In the framework, prārabdha maps precisely to the *definite configuration at this moment*:

$$\text{Prārabdha}(t) \equiv \{\xi(t), \theta(t)\} \tag{50}$$

where:

- $\xi(t)$ is the system’s ontic state vector—its actual configuration right now
- $\theta(t)$ is the observer’s measurement basis—what they are actually measuring right now

Both are *definite facts*, evolved deterministically from sañcita. The measurement outcome:

$$\text{Outcome} = \text{sgn}[\boldsymbol{\xi}(t) \cdot \mathbf{n}(\theta(t))] \quad (51)$$

is therefore predetermined—this is prārabdha manifesting. It was shaped by the infinite causal web (sañcita) but could not be predicted by the observer.

The correlation: The observer’s $\theta(t)$ and the system’s $\boldsymbol{\xi}(t)$ did not evolve independently. They are correlated through their common past—the entangled threads of sañcita that shaped both. This is why measurement independence fails: Alice’s “choice” of measurement basis and Bob’s particle state share causal ancestry. Two cautions keep this honest. The threshold form $\text{sgn}[\boldsymbol{\xi} \cdot \mathbf{n}(\theta)]$ is single-system shorthand—a visualization, not a spin beable; a naive local threshold of this kind does not by itself reproduce the two-wing Bell correlations, which are *inherited from the hosting embedding*, not generated by shared ancestry alone (Section 5.2). What the shared past secures is only the weaker, defensible point: that setting and system *need not* be statistically independent, so measurement independence is an idealization, not a law of nature.

This is the essence of prārabdha: The specific configuration that manifests was determined by causes you cannot trace, yet it manifests with necessity. You experience it as “what happens to you,” but it was never separate from the total causal fabric.

Why prārabdha appears unpredictable:

Despite being fully determined, the observer cannot predict prārabdha due to double ignorance:

1. **Cannot know $\boldsymbol{\xi}$:** Determining the hidden configuration requires measurement, which disturbs the system. You cannot “peek” at what will be revealed without changing it.
2. **Cannot predict θ :** Self-ignorance (limited C_{eff} , chaotic/diffusive internal dynamics) prevents you from knowing *why* your measurement basis evolved to this particular value. By the time you’re aware of θ , the causal chain that produced it is already hidden.

This is why prārabdha feels both fated (it cannot be changed) and surprising (you didn’t know it was coming). The classical teaching: “Prārabdha must be experienced; it cannot be avoided but also cannot be foreseen.” The framework shows *why*: fundamental limitations on self-knowledge.

Āgamī Karma ↔ Present Measurement Creating Future

Āgamī is new karma created *now*. In the Vedantic view, this arises from the ego’s reaction to and identification with the unfolding of prārabdha. In the framework, this maps to the measurement outcome feeding back into the universal evolution, but with a crucial new layer of understanding.

The measurement is a two-fold event:

1. **The Physical Event (The Prārabdha):** The predetermined outcome $\text{sgn}[\boldsymbol{\xi} \cdot \mathbf{n}(\theta)]$ manifests as a classical fact. This is the script unfolding.
2. **The Epistemic Event (The potential for Āgamī):** The observer’s system registers this outcome. The “ego” (whether a conscious mind or the control system of an apparatus) updates its model of the world based on this new information.

This update is what “creates the future.” It becomes part of the causal web in several ways:

- The macroscopic record of the outcome becomes part of the entangled universal state—within $|\psi\rangle$'s single unitary history—and thereby conditions the appearances that subsequently unfold for embedded observers
- The observer's updated knowledge state will now become part of the causal chain that determines its future measurement “choices” (future values of θ)
- New correlations are established, creating new tendencies for future manifestation

This is not paradoxical. The outcome was determined by sañcita, but once it manifests (prārabdha), the system's reaction to it becomes part of the sañcita for all future moments. The causal web is self-consistent: the ego's “choice” to identify with an action was itself predetermined, but from the ego's own vyāvahārika perspective, it is the very act that generates its future.

How This Clarifies Measurement and “Collapse”

The three-fold framework dissolves the measurement problem: prārabdha (definite ξ and θ) exists before measurement but is unknown; measurement reveals what was always determined; the outcome becomes āgamī, entering the causal web for future manifestation. Nothing “collapses” because the superposition is in $|\psi\rangle$ (the universal field), not in $\{\xi, \theta\}$ (what actually exists). The apparent randomness arises through self-ignorance about why prārabdha ripens as it does.

Connection to Free Will and Doership

The experience of “choosing” a measurement basis θ is the experience of prārabdha manifesting. The subsequent registration of the outcome is the creation of āgamī. You feel you are a free agent, but:

- The action (the value of θ and the resulting outcome) was determined by sañcita
- The reaction (the identification “I measured this and will act on it”) was also determined by sañcita
- You cannot trace *why*—by the double ignorance above, ξ cannot be peeked at and θ 's provenance is already hidden by the time you are aware of choosing

As the *Bhagavad Gītā* states: “All actions are performed by the guṇas of prakṛti; the self, deluded by egoism, thinks ‘I am the doer’” (3.27). Actions happen—prārabdha manifests—but the ego claims authorship retroactively, and in that act of claiming it generates the āgamī karma that binds it to future experience. This is the structure the readiness-potential experiments made famous: the conscious sense of willing appears downstream of prior causal preparation, while the observer experiences itself as author of the act.

The framework thus bridges determinism and the phenomenology of agency. Everything is determined by the infinite past (sañcita), yet the present feels open because you cannot predict what will ripen (prārabdha), and what you do matters because your identification with it plants seeds

for the future ($\bar{a}gam\bar{i}$)—though that identification was itself determined. You are not an isolated ego making choices from outside, but part of the one causal web, claiming authorship of what moves through it. This is not paradox but the structure of manifestation within deterministic unity.

A compatibilist gloss: This is the structure modern compatibilism points to, reached here from the Vedantic side. The phenomenology of free choice is real and unavoidable—you cannot predict what you will choose until you choose it—not because the choice is uncaused, but because the causal chain producing it (which option draws attention, which argument compels) is opaque to the chooser. Determinism and the felt openness of choice are both true, from different perspectives: the action flows from accumulated tendencies (*vāsanās*), while the process determining it stays hidden from conscious awareness. The framework’s own empirical commitment is narrower still—whatever the ultimate ontology, finite tracking capacity can limit an observer’s access to its own basis-setting process.

In this sense, measurement is manifestation rather than creation. The outcome does not come into being from nothing; it becomes definite for the observer as the ripening of a causal structure already present in the host. What changes is the observer’s record: *prārabdha* is encountered, and the system’s registration of it becomes part of the causal web that conditions future appearances.

Brahman changeless, the record open. The future is open for the observer not because the Absolute is incomplete, but because manifestation is record-forming: each act, measurement, and macroscopic imprint becomes part of the causal texture through which future appearances unfold. From the ultimate standpoint (*pāramārthika*), Brahman is changeless; from the empirical standpoint (*vyāvahārika*), $|\psi\rangle$ evolves because records condition what can appear next. The single deterministic history is not rewritten by what it carries; what advances is only the observer’s record of it.

25 The Hṛdaya-granthi and the Heisenberg Cut

Ramana Maharshi taught that the ego is the *hṛdaya-granthi*, the heart-knot, also called the *chit-jada-granthi*: the apparent knot between consciousness (*chit*) and inert matter (*jada*). It is not a physical structure, but the false ‘I’-sense that arises between pure consciousness and the body-mind, confusing the two as if they were one. As Ramana explains in *Uḷḷadu Nārpadu* verse 24, this false ‘I’ appears as “I am this body,” borrowing properties from both sides: the body’s limitation, location, and change, and the Self’s luminosity as consciousness. This confused mixture is the knot.

Physics offers a partial analogue in the Heisenberg cut—but the correspondence is not identity, and the difference is exactly the point. The Heisenberg cut is the place where an otherwise unitary quantum description is divided into system and apparatus, possibility and record, observed and observer. Standard quantum mechanics can slide this cut, but it cannot eliminate the need for some classical standpoint from which outcomes are registered. Crucially, it is a division drawn *within* the observer’s already-given world: both of its sides, the quantum system and the classical apparatus, are objects of experience. IOF asks what happens when the

standpoint doing the registering is itself no longer treated as external to physics, but as a finite, basis-generating system within the same causal order as what it measures.

That question reaches deeper than the Heisenberg cut, to where the very sense of a separate knower first arises—and that is what the *granthi* names. The knot is not a division between two things but an apparent binding: the false ‘I’ by which undivided consciousness seems to become a finite knower set over against a known world. Nothing is actually cut; what arises is the seeming of separation. In physics-facing language it is the stabilization of a standpoint *as such*: a point from which any basis can be implemented, any outcome recorded, the world first standing forth as object. The observer does not encounter its own basis as an object fully available to it; it lives through the basis. The Heisenberg cut therefore presupposes the *granthi* rather than coinciding with it—one can ask where, inside an observer’s world, the quantum gives way to the classical only once there is an observer with a world at all, and that apparent arising is the knot. The cut’s very stubbornness is the *granthi*’s trace: physics can move the observer-boundary anywhere it likes but can never abolish it, because the seeming of a separate observer underlies every appearance within which a cut could be drawn.

Under IOF, this apparent binding has a concrete physical expression. The realized basis $\theta(t)$ is not outside the world; it is generated by the observer-apparatus system. Yet the observer has only finite, coarse-grained access to the causal process that produces it. It can register which basis and outcome occurred, but not fully reconstruct why this basis and this outcome arose from the prior history. This is the physical analogue of *avidyā*: not a lack of information about some external object alone, but a structural opacity in the finite standpoint itself.

The *granthi* should therefore not be read as a literal physical boundary, nor as identical with any specific collapse mechanism. It names the non-dual structure expressed physically by that binding: consciousness appears as a localized knower only by becoming bound to a finite body-mind/apparatus standpoint. Once that standpoint is in place, the world appears as a field of objects, measurements, causes, and outcomes. The knot is the apparent binding that makes this entire subject-object grammar possible.

The same point can be expressed through the Vedantic distinction between ordinary object-knowledge and Self-knowledge. In ordinary knowing, the mind takes the form of the object; this is *vṛtti-vyāpti*. The object is then revealed when reflected consciousness illumines that mental modification; this is *phala-vyāpti*. Thus an ordinary object is not self-luminous. It requires reflected consciousness to become known. The Self is different. It is not revealed by a second light, because it is the source of illumination itself. In Self-knowledge, the relevant modification removes ignorance, but no reflected consciousness is needed to light up the Self as an object.

This clarifies the mirror-like character of the *granthi*. On the world side of the mirror lies *chidābhāsa*, reflected consciousness: mind, object, state, witness-structure, basis, measurement, and record. On the other side is *chit*, pure consciousness, not reflected and not objectifiable. The Heisenberg cut belongs entirely to the world side of this boundary; it is a structural division within reflected appearance. Yet it echoes the *granthi*: the Heisenberg cut divides observer from observed within appearance, while the *granthi* is the apparent knot by which there seems to be a separate observer at all.

Figure 4 draws this mirror-structure as the twin of Figure 1 (Section 6). The geometry is

unchanged—the same stack, the same two boundary lines—and only the vocabulary and the claims attached to it move. On the world side the three layers are the gross (sthūla), the subtle (sūkṣma), and the causal (kāraṇa): together they exhaust chidābhāsa, everything witnessable (dr̥śya). The deepest of them, the substrate, carries two non-interchangeable faces—kāraṇa, its world-facing causal seed, and turīya, its boundary-facing substratum—both still on the world side of the mirror. The mirror-line itself is the knot: the ego, the false ‘I’ (chit-jaḍa-granthi), borrowing limitation from the world and luminosity from the Absolute, the locus of veiling. Below it is sat-chit-ānanda, the Absolute—Ramana’s ‘I-I’—which no log can capture because every log is its object.

Two differences from the physical figure carry the entire interpretive content, and both were announced when the first figure was drawn. The axis label changes: the physics ordered the layers by conditioning access and asserted nothing about realness, while the tradition does assert an ordering, because sublation (bādha) gives it a criterion—vyāvahārika appearance is contradicted from the pāramārthika standpoint, never the reverse. And the bottom layer changes from a mark to a name: what physics could only mark as unconditioned—the layer below the line no channel crosses—the tradition calls sat-chit-ānanda, drawn as a box only to mirror that mark, though strictly it is no stratum but the ground in which the whole figure appears. Whether that naming is correct is precisely the question the physics brackets; that the two figures share one geometry is the structural resonance this part reads—and, as with every correspondence in this part, it is structure, not identity.

Decoherence can then be read as one physical expression of māyā’s veiling-projecting power, but with an important qualification. Standard decoherence is object-side: environmental entanglement suppresses observable interference and stabilizes classical records. IOF adds the subject/reference-side counterpart: finite basis access suppresses visibility when the observer-apparatus system can no longer track the basis it physically realizes. Together they describe how apparent separation becomes stable, measurable, and seemingly objective for finite observers.

In this sense IOF does not require reality itself to lose coherence; it shows how coherence can become unavailable to a finite standpoint. The loss belongs to conditioned access, not to the substrate. This is structurally close to the Vedantic role of avidyā: not a second reality competing with Brahman, but the limitation through which appearance is misresolved as independent fact. The rope is not transformed into a snake; the standpoint fails to resolve the rope correctly. In IOF language: visibility is not necessarily lost from the universal state; it is lost to the bounded observer’s available record. In Vedantic language: reality is not diminished; the jīva’s access is conditioned by ignorance. A collapse-like reading would have nature physically destroy possibilities; the reference-frame reading lets the finite observer’s world be shaped by what it cannot condition on—appearance arising through limitation, not through damage to the real. Stated at its full strength: self-ignorance is not a cause operating within appearance but the condition under which undivided reality appears as finite records, probabilistic outcomes, and observer-object separation.

The tradition also explains, in its own register, why the stronger reading of the visibility law (Part III) had to fail. Dream thirst is quenched only by dream water; the bottle on the nightstand is useless until waking. An experiment is dream water: it can settle any question posed within the dream’s own physics, and the exclusion of Section 12 is that physics speaking with its own authority. What no dream experiment can do is certify the dreaming: a condition of the

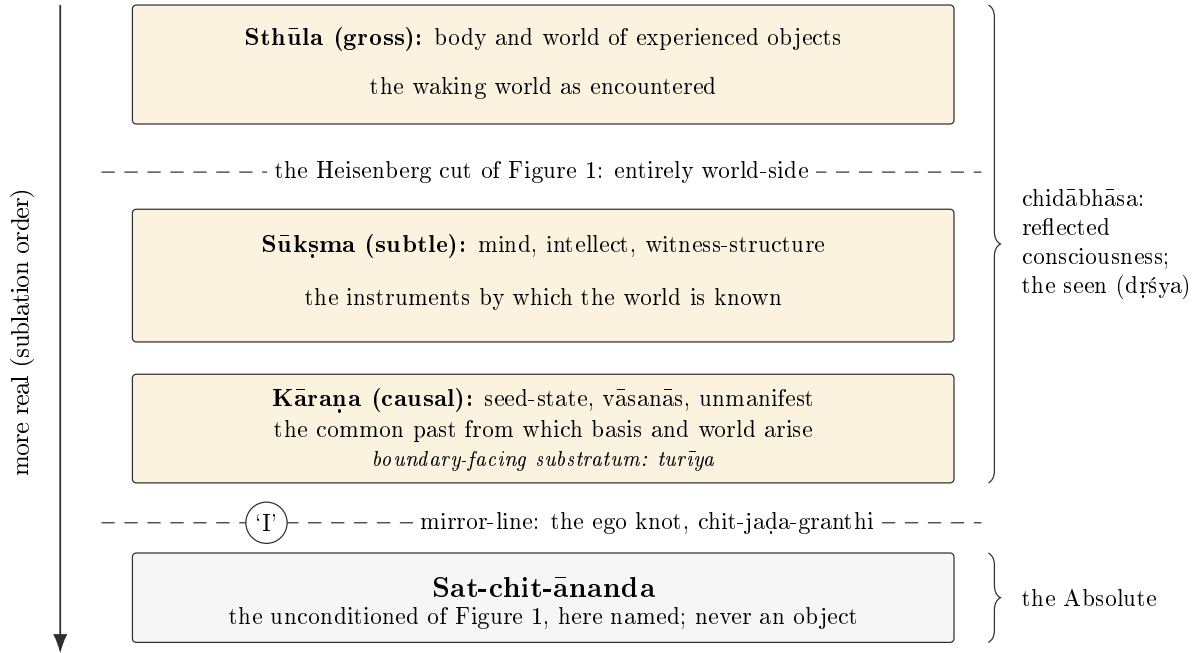


Figure 4: The layered architecture, contemplative vocabulary: the twin of Figure 1—same geometry, exchanged labels (sthūla ↔ classical, sūkṣma ↔ quantum, substrate ↔ kāraṇa/turīya, sat-chit-ānanda ↔ the unconditioned). The one substrate carries two Vedantic faces, not interchangeable: kāraṇa, its world-facing causal seed (vāsanās, the unmanifest), and turīya, its boundary-facing substratum of the three states—distinct aspects, not two strata. The upper dashed line is the Heisenberg cut, entirely world-side. The lower dashed line is the mirror-line of the granthi: the circle straddling it is the ego, the false ‘I’ (chit-jaḍa-granthi), borrowing limitation from the world and luminosity from the Absolute—the locus of veiling, dissolving when traced to its source. Below it, sat-chit-ānanda is drawn as a box only to mirror the physics figure’s unconditioned marker; strictly it is not a stratum but the ground in which the whole figure, both lines included, appears. Unlike the physical figure, the downward axis here *is* a realness ordering—sublation gives the tradition a criterion the physics lacks. The correspondence is structural resonance, not identity.

entire level cannot appear as a localized effect within the level. The strong reading asked the apparatus to register, as an empirical anomaly, the very boundedness that conditions all of its registrations—to quench dream thirst from the bedside bottle. Its exclusion is therefore not an embarrassment to the contemplative reading but its confirmation: the waking state is not certified from within. Nor is waking the $C_{\text{eff}} \rightarrow \infty$ limit—more capacity is more dream; recognition (mokṣa) is a change of standpoint, not a measurement outcome, and not a new state but the seeing of the present one.

The scope of this argument must be drawn as carefully as its content, because it is as easy to over-apply as the strong reading was to assert. What the dream-water argument excludes is intra-level certification of *level-conditions*: the observer’s own boundedness (the strong reading), the global setting–system correlation structure (measurement dependence, which Section 7 concludes can be naturalized but never proven from inside), the dreaming as such. What it leaves untouched is every *intra-level mechanism*: Penrose objective reduction is dream-testable physics—an unrecoverable visibility floor is a perfectly good intra-dream observable, which is exactly why the framework can sponsor the experiment of Part IV—and the argument must never be aimed at the one test the framework itself builds. The two registers kill differently:

induction kills slowly and reversibly—the empirical case against objective reduction would yield to a single well-audited floor—while the transcendental point kills categorically, and only what asks a level to certify its own condition falls to it. The framework’s credibility lives in never letting the second do the first one’s work.

The discrimination that the empirical parts operationalize has a terminus the tradition has already mapped. The recoverability classifier (Section 14.4) is *viveka* performed within the dream: it separates, on the same records, what belongs to the designated observer from what belongs to the world. But the designated observer—apparatus, loop, body—is itself world (Section 6): every part of it can in principle be logged, conditioned on, moved to the far side of the cut. The discrimination can therefore be repeated inward indefinitely and never closes empirically. The tradition’s name for the completed movement is *dr̥g-dr̥śya-viveka*, the discrimination of the seer from the seen: whatever can be witnessed—body, mind, apparatus, record—is thereby the seen, and not finally the seer. The classifier is the operational shadow of that move, the same gesture performed on a designated cut; its terminus—the witness (*sākṣin*) that no log can capture, because every log is its object—is not reached by iteration, since more conditioning is more dream, but by the change of standpoint named above. Within a layer: a real split, a real classifier, real physics. Across all layers: one substrate, and a discrimination that ends not in a measurement but in recognition.

In this sense, physics can be read as the mathematics of the waking dream: not because the waking world is unreal in the trivial sense, but because it is a structured appearance governed by lawful relations within consciousness. IOF does not use this as a physical premise. The physical content remains the finite basis-tracking model and its observer-relative visibility law. The non-dual reading is that even this model belongs to the empirical order: it describes how the knot of apparent separation becomes operationally real for a finite observer within the waking state.

Part VII

Conclusion

We have proposed that what quantum mechanics treats as irreducible randomness can be *read* as the signature of a finite observer embedded in a deterministic world. The word “read” is essential. The mechanism is a control-theoretic constraint: an observer with bounded effective basis-tracking capacity C_{eff} and a nonzero internal instability rate h_{KS} cannot always resolve *why* its own measurement basis took the value it did. The deficit $\kappa = h_{\text{KS}} - C_{\text{eff}} \ln 2$ separates two regimes. In the capacity-wins regime ($\kappa < 0$), textbook fixed-basis quantum mechanics is operationally available. In the chaos-wins regime ($\kappa > 0$), standard quantum mechanics still holds, but raw records can lose contrast because the observer’s reference frame is unresolved. Unitary evolution is untouched; no collapse is added; and the lost contrast returns when the missing reference information is supplied. Apparent collapse becomes an observer-level loss of trackability. We host this mechanism in a deterministic global-history embedding: a single definite history carrying the field $|\psi\rangle$, an ontic state ξ , and the basis θ , with correlations read through a common past.

The visibility law also invites a stronger reading: that the capacity-dependent loss is unrecoverable, a new physical channel beyond standard decoherence. That would make IOF an empirically distinct rival to quantum mechanics. We state that reading precisely in Part III, and existing experiments already answer it. Logged-setting Bell tests, randomized-measurement tomography, and correlation spectroscopy all recover the contrast the strong reading says is lost. The strong reading is excluded. What remains is an interpretation: finite embedded observers sample a single constrained global history, and quantum probability is read as the form taken by dependencies the observer cannot resolve. Alongside it stand a foundations sublayer—the conditional derivation of the binary Born weight from finite-record geometry [3, 5], the same finitude that makes a history read as chance also fixing the one piece of Born structure a finite observer can read from its own records, and no more—and an operational control law whose live experimental use is a discriminating test of Penrose objective reduction [25], not of IOF against quantum mechanics. This is where the framework meets live experiment—not as a rival prediction but as an instrument: because its parameters are built and measured, the same machinery that isolates the observer’s contribution becomes the calibrated subtrahend of an objective-reduction search, making a Penrose floor legible against a controlled background if one exists, and turning a null from an inconclusive absence into a quantitative bound on gravitational collapse if it does not. The framework predicts no floor of its own and so cannot win that test; what it offers is to make the test decidable either way.

The framework, so delimited, is not a physics with a metaphysics laid over it, nor Advaita restated in equations. It grew from a structural resonance between the non-dual analysis of the Advaitic tradition and the control-theoretic structure of bounded observation—each read in the light of the other, as the note on method sets out, with neither reduced to a commentary on the other. On the contemplative side, the three-fold structure of karma supplies a vocabulary for substrate, manifestation, and self-ignorance, and apparent collapse reads as one consciousness appearing as many observers; on the physical side, the same structure is a finite observer’s capacity-limited tracking of its own measurement basis, with apparent collapse an epistemic

update rather than a physical event. That shared structure is what suggested the framework and guided where the formalism should go—yet it cannot establish the physics, and the physics does not prove the metaphysics: the framework’s layers are built to stand or fall separately, whatever one makes of the correspondence that suggested them.

Pursued to its end, that contemplative reading is radical. The mathematics—global-history structure, self-ignorance bounds, sañcita/prārabdhā dynamics—would then not merely describe how quantum systems behave; it would describe how consciousness structures the appearance of multiplicity from unity: the mathematics of a waking dream. EPR correlations are not spooky because there is no space to act across—space is part of the appearance. Measurement is not mysterious because there is no separate observer—the observer is consciousness appearing to observe itself. “Collapse” is not a physical process but, in this contemplative reading, a change in how consciousness’s own appearance is registered within the dream.

This is speculation beyond physics proper, and we mark it as such. It shows only how far the interpretive resonance may extend: not that physics confirms Vedantic ideas, but that physics may provide one mathematical language for exploring what Ramana was pointing to. The framework gives the structure; the contemplative tradition gives the substrate. Together: physics as a rigorous description of how the One appears as many, and why the many cannot know they are One—held by self-ignorance—until they wake.

What, then, is the resonance worth, once the empirical books are settled? Stated exactly: *the mapping is public; the weight is private*. The structural correspondences are checkable claims—that the layer-relative ontic/epistemic labels of Section 6 factor as the tradition’s vyāvahārika/pāramārthika distinction, that the recoverability classifier performs the gesture of viveka on a designated cut (Section 14.4), that the exclusion of Section 12 and the dream-water argument of Part VI share one logical shape—and anyone may verify the isomorphism without adopting either side of it. What no third party can be compelled to share is the significance assigned to it. Nor is that subjectivity a defect imported from physics’ standards: first-person confirmation (anubhava) is the contemplative science’s native and final mode of validation (Part I), exactly as third-person measurement is physics’. One structure, two sciences, each confirming in its own register—that is the precise sense in which they meet, and the only sense this paper claims.

One feature of the correspondence, however, resists deflation as authorial selection. A mapping built by choosing flattering parallels would align the two registers where they affirm; this one aligns most sharply where they refuse. The strong reading of the visibility law is rejected twice, by each register on its own authority and for its own reasons: by physics, empirically, through logged-settings interferometry and clock spectroscopy (Section 12)—data with no stake in any tradition—and by the tradition, on principle, through the teaching that no dream experiment certifies the dreaming (Part VI). Neither rejection borrows from the other, and the claim they jointly eliminate is the one reading that would have set the two registers at odds. Convergence at a refusal is the one kind of agreement that selective mapping cannot manufacture: parallels can be chosen; independent vetoes that land on the same point cannot. It proves neither science. It is, however, the reason the resonance—proven or not—cannot simply be ignored. And it is the separation of the books that gives the observation any force: resonance between blended registers would be circularity; resonance between registers kept rigorously apart—one answerable to interferometers, the other to inquiry—is consilience.

What the framework asks for is scrutiny, not assent—and the scrutiny begins within: the strong reading its own law suggests is stated precisely enough to be tested, and existing data answer it. The measurement problem has outlasted a century of proposed mechanisms; perhaps what it demands is exactly this discipline: say precisely what an observer's limits would change, check whether physics already knows the answer, and keep separate books for what is physics, what is foundations, and what is the lens through which one chooses to see the whole.

Appendices

A Derivation of Tracking Timescales (τ_{fill} , t_{break})

This appendix derives the tracking timescales from a high-rate, converse-type bound from information theory for tracking a chaotic source: the bound gives the minimum information rate required for a given distortion, so the resulting variance law is a lower envelope, attained only by near-optimal tracking, and is used in the main text as the minimal effective closure (Section 10.4). We derive τ_{fill} (capacity-wins convergence time) and t_{break} (chaos-wins failure time).

A.1 Formula for Tracking Timescale

For a chaotic source with Kolmogorov–Sinai entropy rate h_{KS} (in nats/s), the information rate required to track its state with a target mean-squared error (distortion) D_{target} over a time horizon t is given by the high-rate bound. (In systems with a single positive Lyapunov exponent, such as the kicked rotor, h_{KS} coincides with that exponent; we use h_{KS} as the general entropy-rate measure of instability.)

$$R(D_{\text{target}}; t) \approx h_{\text{KS}}t + \frac{1}{2} \ln \left(\frac{\sigma_0^2}{D_{\text{target}}} \right) \quad [\text{nats}] \quad (52)$$

where σ_0^2 is the prior variance (the uncertainty at $t = 0$). An observer with an effective basis-tracking channel of capacity C_{eff} (in bits/s) can gather $C_{\text{eff}} \cdot t \cdot \ln(2)$ nats of task-relevant information in time t . Setting $R = C_{\text{eff}} \cdot \tau \cdot \ln(2)$ and $t = \tau$ yields:

$$C_{\text{eff}} \cdot \tau \cdot \ln(2) = h_{\text{KS}}\tau + \frac{1}{2} \ln \left(\frac{\sigma_0^2}{D_{\text{target}}} \right) \quad (53)$$

Solving for τ :

$$\tau = \frac{\ln(\sigma_0^2/D_{\text{target}})}{2(C_{\text{eff}} \ln 2 - h_{\text{KS}})} \quad (54)$$

Here, $D_{\text{target}} = -2 \ln(V_{\text{target}})$ is the target error variance corresponding to a target visibility V_{target} .

Capacity-wins regime ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$): The denominator is positive, and we denote this convergence timescale τ_{fill} . The observer converges to accurate tracking.

Chaos-wins regime ($h_{\text{KS}} > C_{\text{eff}} \ln 2$): With appropriate sign conventions ($\sigma_{\text{target}}^2 > \sigma_0^2$, uncertainty growing from initial calibration to tolerance threshold), the loss timescale is:

$$t_{\text{break}} = \frac{\ln(\sigma_{\text{target}}^2/\sigma_0^2)}{2(h_{\text{KS}} - C_{\text{eff}} \ln 2)} \approx \frac{1}{h_{\text{KS}} - C_{\text{eff}} \ln 2} \quad (55)$$

where the approximation assumes the log factor is $O(1)$.

A.2 Steady-Cycle Prior for Kicked Rotor (Capacity-Wins)

For a cyclic experiment like the kicked rotor in the capacity-wins regime, the prior uncertainty σ_0^2 should be self-consistently derived from the experiment’s dynamics.

In steady state, the uncertainty at the end of one cycle equals the uncertainty at the start of the next. Balancing chaotic amplification ($\sigma^2 \rightarrow \sigma^2 e^{2h_{\text{KS}}T_{\text{kick}}}$) against task-relevant information gathering ($C_{\text{eff}} \cdot T_{\text{kick}} \cdot \ln(2)$ nats), the self-consistency condition yields a simplified formula for the convergence timescale:

$$\tau_{\text{fill}} = \frac{h_{\text{KS}}T_{\text{kick}}}{C_{\text{eff}} \ln 2 - h_{\text{KS}}} \quad (56)$$

This formula applies specifically to kicked-rotor experiments in the capacity-wins regime ($C_{\text{eff}} \ln 2 > h_{\text{KS}}$).

A.3 Sensitivity Analysis

Figure 5 shows how the convergence time (54) depends on the effective basis-tracking capacity C_{eff} for several internal-instability rates h_{KS} , holding the tracking-task parameters ($\sigma_0, V_{\text{target}}$) fixed. The timescale is governed by C_{eff} : at the baseline capacity $C_{\text{eff}} = 2600$ bits/s the observer sits deep in the capacity-wins regime ($\tau_{\text{fill}} \approx 0.89$ ms), reducing capacity to $C_{\text{eff}} = 600$ bits/s raises it to a few milliseconds, and as capacity approaches the boundary $C_{\text{eff}} \ln 2 = h_{\text{KS}}$ the convergence time diverges. The strong dependence on C_{eff} at fixed dynamics is what makes capacity the primary experimental knob.

A.4 Conclusion of Analysis

The systematic parameter exploration confirms that the convergence timescale τ_{fill} is robustly in the low-millisecond regime for the mesoscopic laboratory parameters explored here. The baseline calculation yields $\tau_{\text{fill}} \approx 0.89$ ms. This represents the convergence timescale in the capacity-wins regime—the time for an observer to spin up from initial uncertainty to accurate basis tracking. This reveals the framework’s epistemic picture: observers with these parameters occupy the Goldilocks zone—sufficient capacity to track their basis (not perfect, not chaos-dominated)—where ideal quantum visibility is recovered.

Regime interpretation: These parameters ($C_{\text{eff}} = 2600$ bits/s, $h_{\text{KS}} \approx 1$ nats/s) fall in the *capacity-wins regime* where $C_{\text{eff}} \ln 2 \gg h_{\text{KS}}$. After the transient spin-up time τ_{fill} , the observer maintains good knowledge of the measurement basis and visibility approaches the quantum ideal: $V_{\text{IOF}} \approx 0.999-1$. Laboratory systems naturally occupy this epistemic Goldilocks zone where bounded but capable observers approximately track their measurement basis—explaining why QM works so remarkably well in controlled settings.

Overlap with Penrose OR: For the capacity-wins laboratory parameters above ($C_{\text{eff}} = 2600$ bits/s, $h_{\text{KS}} \approx 1$ nats/s) the convergence time τ_{fill} is sub-millisecond, and there is *no* overlap with the Penrose window $\tau_{\text{OR}} \sim 10-100$ ms. The two timescales coincide only in a *chaos-wins regime* with low effective tracking capacity (for example $C_{\text{eff}} \sim 10$ bits/s, $h_{\text{KS}} \sim 50$ nats/s),

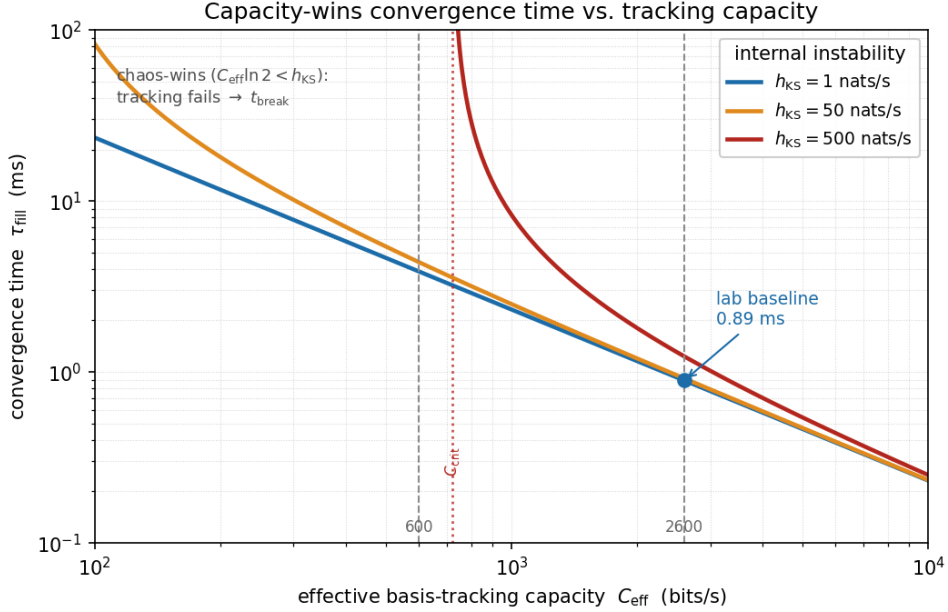


Figure 5: Capacity-wins convergence time τ_{fill} (54) versus effective basis-tracking capacity C_{eff} , for three internal-instability rates h_{KS} (tracking-task parameters held fixed). At laboratory capacities ($C_{\text{eff}} = 600$ and 2600 bits/s, dashed) the observer is deep in the capacity-wins regime; the baseline ($h_{\text{KS}} = 1$ nats/s, $C_{\text{eff}} = 2600$ bits/s) gives $\tau_{\text{fill}} \approx 0.89$ ms. As C_{eff} falls toward the boundary $C_{\text{eff}} \ln 2 = h_{\text{KS}}$ (e.g. $C_{\text{eff}} \approx 721$ bits/s for $h_{\text{KS}} = 500$ nats/s), τ_{fill} diverges; below it the system is chaos-wins and tracking fails, governed instead by t_{break} .

where $t_{\text{break}} \sim 50\text{--}70$ ms. The coincidence is numerical, not mechanistic: the scales come from unrelated principles—gravitational instability on one side, information-theoretic tracking failure on the other—so it is no evidence of a shared mechanism. What it marks is the regime where the framework’s control variables become a useful experimental axis; the raw-record suppression there is recoverable reference-frame physics, and the live discrimination along the mass-geometry axis is of Penrose objective reduction, not of IOF (Section 14).

A.5 Sanity Checks

- **High-Rate Validity:** The capacity-wins derivation is valid when $C_{\text{eff}} \ln 2 > h_{\text{KS}}$. For these parameters, $C_{\text{eff}} \ln 2 \approx 1802$ nats/s and $h_{\text{KS}} \approx 1$ nats/s, so this condition is strongly satisfied.
- **Wrapped-Gaussian validity:** The factor $V = \exp(-\sigma^2/2)$ is the exact first circular moment of a wrapped-Gaussian basis distribution (Part III), not a small-angle approximation; the modeling assumption is that the posterior stays first-harmonic-dominated. For these capacity-wins parameters σ_θ is in any case small ($V_{\text{IOF}} \approx 1$).

A.6 Canonical Clock Conventions

The IOF dynamics define several timescales that differ by order-unity factors, and the same 10–100 ms scale can be labeled in more than one way. Table 2 fixes the notation for the whole

corpus.

Symbol	Definition	Role
τ_κ	$\frac{1}{\kappa}$	Amplitude e-folding time ($\sigma \propto e^{\kappa t}$); the corpus's reference clock.
τ_{var}	$\frac{1}{2\kappa}$	Variance e-folding time ($\sigma^2 \propto e^{2\kappa t}$).
t_*	$\frac{1}{\kappa} \ln \frac{\sigma_*}{\sigma_0}$	Operational threshold-crossing time; = τ_κ up to an $\mathcal{O}(1)$ log factor.
t_{break}	$\propto \frac{1}{\kappa}$	Measured visibility-breakdown time of the unconditioned record; an instance of t_* .
τ_{SK}	threshold-crossing t_*	"Self-knowledge" scale used in the biological supplement [16]; a threshold variant of τ_κ , <i>not</i> the bare $1/\kappa$ clock.

Table 2: Canonical timescales of the IOF dynamics. The amplitude e-folding time $\tau_\kappa = 1/\kappa$ is the reference clock; threshold-crossing times differ from it by $\mathcal{O}(1)$ logarithmic factors $\ln(\sigma_*/\sigma_0)$. The perceptual/biological scale written τ_{SK} in the biological supplement is a threshold-crossing variant t_* .

References

- [1] John S. Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3): 195–200, 1964. doi: 10.1103/PhysicsPhysiqueFizika.1.195.
- [2] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961. doi: 10.1147/rd.53.0183.
- [3] Aernoud Dekker. The Born Rule from Finite Observation: A Conditional Derivation of the Binary Born Form. 2026. doi: 10.17605/OSF.IO/U5RDE. OSF Preprints.
- [4] Aernoud Dekker. Bandwidth-Limited Quantum Control: A finite-rate phase-reference control law and its operational benchmark. 2026. doi: 10.17605/OSF.IO/G5WRH. OSF Preprints.
- [5] Aernoud Dekker. Why Fisher Geometry Gives Binary Born but Not Quantum Phase: A Limitative Theorem on Markov-Invariant Record Geometry. 2026. doi: 10.17605/OSF.IO/N3SU9. OSF Preprints.
- [6] Tim Palmer. Superdeterminism without Conspiracy. *Universe*, 10(1):47, 2024. doi: 10.3390/universe10010047.
- [7] Stephen D. Bartlett, Terry Rudolph, and Robert W. Spekkens. Reference frames, superselection rules, and quantum information. *Reviews of Modern Physics*, 79:555–609, 2007. doi: 10.1103/RevModPhys.79.555.
- [8] Nicholas Harrigan and Robert W. Spekkens. Einstein, incompleteness, and the epistemic view of quantum states. *Foundations of Physics*, 40:125–157, 2010. doi: 10.1007/s10701-009-9347-0.
- [9] Gerard 't Hooft. *The Cellular Automaton Interpretation of Quantum Mechanics*, volume 185 of *Fundamental Theories of Physics*. Springer, 2016. doi: 10.1007/978-3-319-41285-6.
- [10] Sabine Hossenfelder and Tim Palmer. Rethinking Superdeterminism. *Frontiers in Physics*, 8:139, 2020. doi: 10.3389/fphy.2020.00139.
- [11] Michael J. W. Hall. Local Deterministic Model of Singlet State Correlations Based on Relaxing Measurement Independence. *Physical Review Letters*, 105(25):250404, 2010. doi: 10.1103/PhysRevLett.105.250404.
- [12] Girish N. Nair and Robin J. Evans. Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM Journal on Control and Optimization*, 43(2):413–436, 2004.
- [13] Sekhar Tatikonda and Sanjoy Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004.
- [14] Wojciech H. Zurek and Juan Pablo Paz. Decoherence, Chaos, and the Second Law. *Physical Review Letters*, 72(16):2508–2511, 1994. doi: 10.1103/PhysRevLett.72.2508.
- [15] Wojciech H. Zurek. Decoherence, Chaos, Quantum-Classical Correspondence, and the Algorithmic Arrow of Time. *Physica Scripta*, T76:186–198, 1998. doi: 10.1238/Physica.Topical.076a00186. Also *Acta Phys. Pol. B* 29(12), 3689 (1998).

- [16] Aernoud Dekker. Biological Observers: Supplement to the Ignorant Observer Framework. 2026. doi: 10.17605/OSF.IO/CQHNS. OSF Preprints.
- [17] Gregor Weihs, Thomas Jennewein, Christoph Simon, Harald Weinfurter, and Anton Zeilinger. Violation of bell’s inequality under strict einstein locality conditions. *Physical Review Letters*, 81:5039–5043, 1998. doi: 10.1103/PhysRevLett.81.5039.
- [18] Tiff Brydges, Andreas Elben, Petar Jurcevic, Benoît Vermersch, Christine Maier, Ben P. Lanyon, Peter Zoller, Rainer Blatt, and Christian F. Roos. Probing rényi entanglement entropy via randomized measurements. *Science*, 364:260–263, 2019. doi: 10.1126/science.aau4963.
- [19] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16:1050–1057, 2020. doi: 10.1038/s41567-020-0932-7.
- [20] C. W. Chou, D. B. Hume, M. J. Thorpe, D. J. Wineland, and T. Rosenband. Quantum coherence between two atoms beyond $Q = 10^{15}$. *Physical Review Letters*, 106:160801, 2011. doi: 10.1103/PhysRevLett.106.160801.
- [21] E. R. Clements, M. E. Kim, K. Cui, A. M. Hankin, S. M. Brewer, J. Valencia, J.-S. Chen, D. B. Hume, D. R. Leibbrandt, D. J. Wineland, et al. Lifetime-limited interrogation of two independent $^{27}\text{Al}^+$ clocks using correlation spectroscopy. *Physical Review Letters*, 125:243602, 2020. doi: 10.1103/PhysRevLett.125.243602.
- [22] Yoon-Ho Kim, Rong Yu, Sergei P. Kulik, Yanhua Shih, and Marlan O. Scully. Delayed “choice” quantum eraser. *Physical Review Letters*, 84:1–5, 2000. doi: 10.1103/PhysRevLett.84.1.
- [23] Xiao-Song Ma, Johannes Kofler, Angie Qarry, Nuray Tetik, Thomas Scheidl, Rupert Ursin, Sven Ramelow, Thomas Herbst, Lothar Ratschbacher, Alessandro Fedrizzi, Thomas Jennewein, and Anton Zeilinger. Quantum erasure with causally disconnected choice. *Proceedings of the National Academy of Sciences*, 110:1221–1226, 2013. doi: 10.1073/pnas.1213201110.
- [24] S. Merlet, J. Le Gouët, Q. Bodart, A. Clairon, A. Landragin, F. Pereira Dos Santos, and P. Rouchon. Operating an atom interferometer beyond its linear range. *Metrologia*, 46:87–94, 2009. doi: 10.1088/0026-1394/46/1/011.
- [25] Aernoud Dekker. A Capacity-Calibrated Protocol for Testing Penrose Objective Reduction. 2026. doi: 10.17605/OSF.IO/2QJNE. OSF Preprints.
- [26] Matthew F. Pusey, Jonathan Barrett, and Terry Rudolph. On the reality of the quantum state. *Nature Physics*, 8:475–478, 2012. doi: 10.1038/nphys2309.
- [27] Swami Sarvapriyananda. Lectures on Advaita Vedanta. Vedanta Society of New York (e.g. the Katha Upanishad series). https://www.youtube.com/playlist?list=PLDqahtm2vA72naWj1foEqGFQiN_bRI5my.
- [28] Michael James. Happiness and the Art of Being. Translations of and commentary on the teachings of Bhagavan Sri Ramana Maharshi. <https://happinessofbeing.com>.

End of Document

*Ten monks crossed the river wide,
Each counted nine on the other side.
“One has drowned!” they wept in fear—
Till shown: the missing monk was here.*