The Actualization State

Artist's Proof 01

Foundations

The spine — viability geometry, agency, coupled corridors

Table of Contents

Paper B — Selection as Irreversible Exclusion

Paper C — Agency as Constrained Control

Paper D — Coupled Viability: Structural Conditions for Multi-Agent Persistence Under Irreversible Dynamics

Appendices

Kill Switch Ledger

P0.1 — Before the Beginning

Close your eyes for a moment. Try to imagine nothing. Not darkness — darkness is something. Not silence — silence is something. Not empty space — space is something. Nothing.

The absence of everything, including the absence itself.

You cannot do it. Your mind keeps producing something to fill the void. That inability is not a failure of imagination. It is the first clue.

Begin with nothing. Not empty space, not vacuum fluctuation, not a quantum field in its ground state. Nothing. No topology, no dimension, no time, no observer. The empty set: Ø.

Ø is not a place. It has no properties to describe. But it is not incoherent. Mathematics begins with the empty set and builds everything from it.

Set theory constructs the integers, the reals, topology, and eventually the structures that physicists use to describe the universe.

The question is not whether Ø is real — you cannot answer that.

The question is whether the transition from Ø to structure has a shape — and if that shape leaves traces in what we observe.

P0.2 — The Fracture

You have seen symmetry break. A glass falls from a table. Before it falls, all directions are equally possible. After it falls, one direction is actual. The glass cannot un-fall.

The first distinction is binary. From Ø, two values in perfect balance: 1:1. Not numbers yet — just the minimal possible differentiation. A fracture in undifferentiated potential.

One side is absence (0), the other is presence (1). But perfect balance is not structure.

Structure requires the smallest possible perturbation — a deviation from symmetry so slight that it could not be smaller and still exist. Call it ε.

The fracture is not 0 and 1 alone; it is 1:1 + 1×ε. This is the axiom from which the argument proceeds. It is not a physical event.

It is a structural observation: the simplest thing that can happen to nothing is that it becomes two things, and the simplest thing that can happen to two things in balance is that the balance breaks.

Call the 0-side orientation. It is the residue of what was not chosen, the background against which presence is defined. Call the 1-side actualization.

It is the fact of record — that something, rather than nothing, has been committed. The perturbation ε is what makes the difference between potential and record.

Without it, the two sides are indistinguishable and no structure exists. Crucially, the break does not distinguish two pre-existing sides. There are no sides before ε.

The break creates the sides by breaking the symmetry that made them indistinguishable. “Orientation” and “actualization” are consequences of the break, not preconditions of it.

Crucially, actualization is not merely a label applied to one side of the fracture. It is a dimension — a degree of freedom as real as any spatial direction that will later emerge.

If the manifold that forms has three spatial dimensions and one temporal dimension, actualization is the fifth: the dimension of possibility from which records are written into the four.

The canvas is not less real than the painting; it is what makes the painting possible.

The 0-side (orientation) and the 1-side (actualization) do not sit inside the manifold. The manifold sits inside them. Every record is written from the actualization dimension into the manifold.

This observation is structural, not formal; it is developed operationally in Paper A (where AS quantifies movement along this dimension) and formally in the dimensional analysis of AP10.

The fracture is silent. No energy is released, because energy has not yet been defined. No observer records it, because recording requires structure that does not yet exist.

The symmetry breaks and there is no sound. This is the silent pop.

What follows — the expansion of structure, the differentiation of forces, the emergence of spacetime — is the big bang. The silent pop precedes it: the break that makes the bang possible.

P0.3 — The First Force

The fracture is not passive. It does something. You know this from experience — every time equilibrium breaks, motion follows.

If structure can emerge from potential, then the first question is: what mediates between them? What is the interaction between actualized structure and the undifferentiated background from which it emerged?

Gravity has a unique property among the known forces. It is universal — it couples to all energy, not just to specific charges. It is unscreenable — there is no gravitational insulator.

And it is always attractive — it draws structure together rather than separating it into types.

These properties make gravity the only known interaction that could plausibly serve as the first mediator between differentiated structure and undifferentiated background.

This is not a derivation.

It is a structural observation: if you need one force to emerge first, and that force must couple to everything that exists simply by virtue of its existence, gravity is the only candidate in the known inventory.

Whether this observation is deep or coincidental is precisely the kind of question that cannot be settled by argument.

P0.4 — Accumulation

You have never undone a moment. Not one.

Once the fracture has occurred and structure begins to actualize, the process has a direction. Records form. Alternatives are excluded. Irreversibility accumulates.

This is the upper half of the hourglass: potential converting into record, the 0-side draining into the 1-side.

The formal version of this process is Actualization State increasing under decohering dynamics (Paper A, Theorem T1). But the intuition precedes the formalism. The universe, once it has begun to differentiate, does not spontaneously undifferentiate.

Records, once formed, do not unform. The arrow is structural, not thermodynamic — though thermodynamics inherits it.

During accumulation, the available space for new records is vast. Branching is cheap. Alternatives proliferate. The viability kernel (Paper A, Definition D7) is large relative to the occupied state.

Agency, in the control-theoretic sense of Paper C, is near its maximum. There is room to move.

P0.5 — Saturation

Everything fills up. Your hard drive. Your patience. The universe.

Accumulation cannot continue without limit. Every record consumes capacity. Every actualization forecloses alternatives. The viability kernel shrinks. The no-return surface (Paper A, Definition D9) advances inward.

Saturation is the state in which the capacity for new record-structured branching approaches zero. The system has committed nearly all of its available degrees of freedom.

New differentiation requires old structure to be recycled — but recycling requires energy that is itself subject to the same capacity constraints.

Black holes are the extreme expression of saturation. They represent states of maximal gravitational commitment — configurations from which no further internal differentiation is accessible to any external agent.

In the language of Paper A, they are deep within the capture basin: states from which exit is impossible under all admissible controls.

They are not reset buttons. They are endpoints of the accumulation process.

P0.6 — The Turn

Here is where the narrative enters territory you cannot test — yet. Hold it lightly.

The turn is the most speculative element of this narrative. It is included because the question it addresses — what happens when accumulation is complete? — is unavoidable once you take the argument seriously.

It is not included because there is evidence for it.

A companion document, Artist Proof 03: The Loop Hypothesis, develops this speculation into a formal conjecture with explicit falsification conditions. What follows here is the intuition that preceded that conjecture.

At saturation, two things are true. First, all capacity has been consumed: no further branching is possible.

Second, the structure that has been built is real — it consists of irreversible records that cannot be undone.

The question is whether there exists any admissible transformation that restores capacity without violating the irreversibility of existing records.

Paper A addresses this in Section A6 as an optional module. The formal conditions are: no reversal of realized selections, no bypass of the selection mechanism, and restoration of the effective record-algebra dimensionality.

Conformal rescaling — a transformation that is insensitive to absolute scale — is one candidate satisfying these conditions at extreme dilution.

In general relativity, there exists a structural correspondence between the interior geometry of a collapsing configuration at maximum compression and the geometry of an expanding configuration at its origin.

This correspondence is not a temporal sequence but a geometric identity: the two descriptions may refer to the same structure read from different sides.

Whether this identity is physically realized is an empirical question addressed in the companion document.

The intuitive image is the bottom of the hourglass. Sand has accumulated. The bulb is full.

But the bottom of the hourglass is also the top of the next one — not because the glass has been turned, but because the geometry at maximum compression is structurally identical to the geometry at the origin of expansion.

The old records remain as boundary conditions. The capacity renews.

The structure continues, with the prior record intact.

Whether this actually happens is not a question this argument can answer.

It is flagged here because the structure of the argument makes the question well-posed, and because intellectual honesty requires acknowledging the places where the intuition reaches beyond what the formalism can support.

P0.7 — The Loop

If the structural identity holds, the process is not cyclic in time but identical in geometry: compression ≡ origin. Each side inherits the record structure of the other as a boundary condition. Nothing is erased.

The loop is not a repetition; it is a structure with memory, read differently from each side of the identity.

The most provocative reading of this structure is that a universe is operationally defined by its record structure. The records produced by actualization constitute the only evidence that anything happened at all.

A universe with no records is indistinguishable from Ø. A universe with records is, precisely and only, those records.

Terms such as “witness” or “observe,” if used elsewhere in this narrative, mean record formation only — not consciousness, inner experience, or subjective awareness. The spine invokes none of those concepts.

This is where the artist’s intuition ends and the physicist’s discipline begins. The preceding paragraphs are a story — a structural story, constrained by the mathematics that follows, but a story nonetheless.

Stories do not have truth values. They have coherence, and they have consequences.

The consequences of this story are the four papers that follow.

P0.8 — A Conjecture on Energy and Actualization

The following conjecture is retained for historical completeness. It is not a current claim of the argument.

Subsequent work (AP03: The Loop Hypothesis) indicates it is likely incorrectly formulated: systems at maximum compression represent states of maximal coarse-grained entropy, not minimal energy contribution.

The conjecture is included because it was the original compact expression of the argument's intuition, and because intellectual honesty requires preserving the record of what was thought before it was corrected.

The simplest such relationship would be: $E = mc^2 \times AS$, where $AS \in [0, 1]$ is the Actualization State defined in Paper A.

At $AS = 0$, no record structure exists and the system contributes nothing to the energy budget of actualized reality. At $AS = 1$, the system is maximally actualized and its full mass-energy is committed.

The original intuition was that reality is not given but earned, one irreversible record at a time. This intuition survives even though this particular formulation does not.

The conjecture does not appear in, and is not referenced by, any of the four formal papers. The spine is unaffected by its status.

P0.9 — Bridge to the Spine

The preceding sections describe an intuition. The following four papers formalize a set of consequences that are consistent with that intuition, but do not depend on it.

No definition, theorem, proposition, or falsifier in Papers A through D requires anything from Paper 0. The spine is self-supporting.

Paper A defines Actualization State as an operational measure of record-structured irreversibility. It proves that AS increases under decohering dynamics, establishes no-return surfaces from bounded capacity, and specifies falsifiable experimental tests.

It depends on nothing outside standard quantum mechanics and viability theory.

Paper B characterizes selection — the transition from multiplicity to definiteness — as a costly, rate-limited exclusion process. It derives structural requirements and a falsifiable gravitational rate bound.

It depends on Paper A and nothing else.

Paper C develops agency as a control-theoretic quantity: the fraction of the viability kernel reachable from where you currently stand under admissible control. It formalizes drift, fatigue, coupling, and exit as consequences of irreversibility.

It depends on Papers A and B and nothing else.

Paper D extends coupling to multi-agent systems operating in shared constraint environments. It derives structural filtering, hierarchy, cooperation, and deterrence as geometric consequences.

Every power structure you have ever encountered — every hierarchy, every alliance, every threat — has this geometry underneath it of irreversible drift. It depends on Papers A, B, and C and nothing else.

Each paper is independently falsifiable. You can kill any one of them. Each contains explicit conditions under which it fails.

The dependency chain is one-way: failure of D does not invalidate C, failure of C does not invalidate B, and failure of B does not invalidate A.

The papers stand or fall on their own logic, independent of the narrative that motivated them.

In the symbolic notation that motivates the formal development:

A record is the actualization state of an irreversible symmetry-breaking event applied to the vacuum.

— where $\varnothing_0$ is the undifferentiated potential of P0.1 and Crack is the symmetry-breaking fracture of P0.2. This notation is evocative, not formal; Paper A defines all quantities operationally.

End of Paper 0. Non-Falsifiable · Structural Narrative · Complete

Paper A

Actualization State An Operational Measure of Record-Structured Irreversibility Reference Document · Canonical

Paper A is reproduced in full on the following pages. It is the foundation of the spine. It depends on nothing outside standard quantum mechanics and viability theory. All subsequent papers inherit from it.

Actualization State (AS) An Operational Measure of Record-Structured Irreversibility

A0 — Front Matter

A0.1 — Title Block ****&**** Abstract

Title

Actualization State (AS): An Operational Measure of Record-Structured Irreversibility

You are reading this sentence. That is a record. Photons hit your retina, neurons fired, a pattern was recognised. The event cannot unhappen.

This paper builds a tool to measure how far that process has gone — and proves that, under the right conditions, it can only go in one direction.

Abstract

Actualization State (AS) is an operational measure of irreversible record formation in quantum systems.

AS is defined relative to physically realisable coarse-grainings induced by system–environment interactions and quantifies the degree to which mutually exclusive classical alternatives have become durably encoded. Irreversibility here is not about entropy.

It is about reachability — the boundary beyond which you cannot get back, no matter what you do.

The paper establishes criteria under which AS is well defined, operationally invariant, and falsifiable, and the proof shows that AS is monotonic under decohering, record-forming dynamics within a precisely delimited scope.

The paper further introduces a domain-neutral no-return theorem showing that bounded maintenance capacity generically induces irreversible loss of reachability, independent of quantum mechanics or gravity.

Together these results provide a falsifiable, interpretation-agnostic framework that isolates irreversible record formation as a measurable physical process, independent of collapse, gravity, or consciousness. You need no interpretation of quantum mechanics to use this tool.

You need only the measurements. No collapse mechanism, gravitational hypothesis, or cosmological assumption is invoked.

The argument isolates the definitional and theorem layers required for any subsequent theory of selection or definiteness.

A0.2 — What This Paper Does and Does Not Do

It does: Define AS as a physically meaningful measure of irreversible record formation. Prove AS increases under decohering dynamics (Theorem T1). Establish no-return surfaces from bounded capacity (Theorem T2).

Require operational invariance — and die if that requirement fails (Kill Switch F0).

It does not: Propose a collapse mechanism. Derive the Born rule. Appeal to gravity or cosmology. Solve the measurement problem. Explain consciousness.

Sections A0–A3 are self-contained. Sections A4–A5 add independently falsifiable postulates. If A4–A5 fail, A0–A3 are untouched.

A1 — Problem Statement

A1.1 — The Actualization Problem

You have never experienced a superposition. Every moment of your life has been definite — this room, this chair, this sentence.

Yet quantum mechanics says that before measurement, systems exist in superpositions of all possible outcomes. Something bridges the gap between “all possible” and “one actual.” That bridge is the subject of this paper.

Quantum theory describes closed systems by unitary evolution in Hilbert space. Experiments, however, report records: mutually exclusive, persistent, classical facts. Between these descriptions lies a structural gap.

Standard measurement language attempts to bridge this gap using observers, projections, or epistemic updates.

These notions do not specify a physical transition; they describe when an agent updates a description, not when a system becomes unable to support alternatives.

Decoherence explains the suppression of interference, but by itself it does not quantify how much irreversible structure has formed, nor does it specify when alternative histories cease to be operationally recoverable.

What is missing is a quantity that refers only to physically accessible degrees of freedom, distinguishes loss of coherence from mere ignorance, and measures the accumulation of durable record structure prior to any claim about outcome definiteness.

That quantity is Actualization State (AS).

Note. This argument is intentionally minimal. It does not ask why the universe permits records, only when they become irreversible.

It does not explain the ‘feltness’ of outcomes, only the structural conditions under which multiple outcomes cease to be simultaneously accessible.

By isolating the transition from quantum coherence to classical record, the paper provides a shared phenomenological target for any deeper theory of definite outcomes.

A1.2 — What Is New: Positioning Relative to Existing Notions

Actualization State is not a redefinition of decoherence, entropy, or thermodynamic irreversibility. The following distinctions are structural.

AS vs. Decoherence

Decoherence is a dynamical process that suppresses interference between alternatives. AS is an operational quantity that measures the extent of record-structured commitment resulting from such decoherence.

The two are distinct: decoherence can occur without significant growth of AS, and AS can increase even when total entropy change is negligible.

A concrete demonstration of their independence is given in the worked comparison below.

AS vs. Entropy

Entropy quantifies total uncertainty or mixedness, including contributions from unobserved degrees of freedom. AS deliberately discards such contributions and tracks only inter-sector branching relative to the physically realizable record algebra.

A system may have high entropy and low AS, or low entropy and high AS. The von Neumann entropy comparison below makes this independence explicit.

AS vs. Quantum Darwinism

Quantum Darwinism (Zurek) quantifies the redundancy with which information is imprinted in environmental fragments. AS measures the informational richness of committed classical branching, not the number of copies of that information.

The two quantities are operationally independent: each can be maximized or minimized independently of the other, as the worked comparison below demonstrates.

AS vs. Consistent Histories

The history-based representation AS_h (Section A2.4) is restricted to single-time record histories under complete decoherence. This is a deliberate narrowing relative to the full consistent-histories framework.

The full framework permits multi-time, multi-branching history sets; AS_h does not. AS_h serves as an equivalence bridge to the primary definition, not as a replacement for consistent histories.

Worked Comparison: Where AS and Quantum Darwinism Redundancy Diverge

The preceding distinctions are structural, but their force is best seen in a concrete system where AS and Quantum Darwinism redundancy move independently.

The following two cases share identical quantum dynamics; they differ only in the number of environmental fragments and the number of pointer sectors.

Case 1: High redundancy, zero AS. A qubit S with pointer basis $\mathcal{O} = \{|0\rangle\langle 0|, |1\rangle\langle 1|\}$ is prepared in the pure pointer state $|0\rangle$.

The environment consists of N = 1000 fragments, each of which independently records that the system is in sector $|0\rangle$.

The Quantum Darwinism redundancy is $R\delta \approx 1000$: the classical information "the system is in $|0\rangle$" is broadcast across a thousand environmental fragments, and any small fraction of the environment suffices to determine the system state.

However, the sector weights are $p_0 = 1$, $p_1 = 0$. The Shannon entropy $H(\{p_i\}) = 0$ and therefore AS = 0. No branching exists.

The environment has recorded a single, definite outcome with extreme redundancy, but there is no record-structured irreversibility to measure.

Case 2: Zero redundancy, maximal AS. A four-level system S with pointer basis $\mathcal{O} = \{\Pi_1, \Pi_2, \Pi_3, \Pi_4\}$ is prepared in an equal superposition and then fully dephased by coupling to a single environmental fragment E.

The sector weights are $p_i = 1/4$ for all i. Quantum Darwinism redundancy is $R\delta = 1$: only one fragment carries the classical information, and loss of that fragment destroys access.

But AS = H({1/4, 1/4, 1/4, 1/4}) / log 4 = log 4 / log 4 = 1. Maximal record-structured branching exists across four mutually exclusive alternatives.

In Case 1, the system is classically definite and robustly broadcast but has no actualization structure. In Case 2, the system is maximally branched but fragile in the Darwinist sense.

AS and redundancy are therefore not merely different in definition; they are operationally independent quantities that can be maximized or minimized independently of each other.

AS vs. von Neumann entropy. A similar divergence arises with respect to von Neumann entropy S(ρ).

Consider a single pointer sector Π_1 of rank $d_i = 100$, with the system state confined entirely to that sector in a maximally mixed intra-sector state.

The von Neumann entropy is S(ρ) = log 100, which is large. But AS = H({1}) / log 1 = 0, since all weight resides in one sector.

Conversely, a two-sector system with rank-1 projectors and equal weights $p_0 = p_1 = 1/2$ has S(ρ) = log 2 and AS = 1. Von Neumann entropy tracks total uncertainty including intra-sector degeneracy; AS tracks only inter-sector branching.

They answer different questions.

Operator Horizon vs. the Second Law

The Second Law expresses the typical growth of entropy under macroscopic dynamics.

The Operator Horizon introduced in Section A3 instead defines irreversibility as a boundary of operational accessibility: a geometric limit in state space beyond which recovery is impossible given bounded control capacity.

Irreversibility here is not a statement about likelihood or typicality, but about reachability under admissible operations. This notion applies equally to quantum, classical, and abstract control systems and is independent of thermodynamic assumptions.

A1.3 — Scope Clarification

The paper does not propose a mechanism of collapse, derive the Born rule, or assume any cosmological or gravitational hypothesis.

It isolates the minimal definitional and theorem-level structure required to make irreversible record formation a well-defined, operationally testable concept.

Any subsequent theory of outcome selection or definiteness must be built on this foundation—or fail against it.

A2 — Definitions

What follows are the tools. Each definition names a specific thing and says exactly what it does. If you lose track, come back here. The definitions do not move.

A2.1 — D1: Physically Realizable Coarse-Graining $\mathcal{O}$ Let $\mathcal{H}_s$ be the system Hilbert space and $\mathcal{H}_e$ its environment.

A physically realised coarse-graining $\mathcal{O}$ is a finite set of mutually orthogonal projectors $\mathcal{O} = \{\Pi_i\}$ satisfying all of the following conditions:

$\mathcal{O}$ is not observer-chosen. It is selected by the physics of coupling. You do not choose what gets measured. The interaction chooses.

The critical point: $\mathcal{O}$ is not your choice. It is nature's choice. The physics of the interaction determines what gets measured. You do not get to pick the basis. The coupling picks it for you.

Computability remark.

In practice, the physically realizable coarse-graining is identified as the stable algebra generated by the interaction Hamiltonian's pointer observables — for instance, via the predictability sieve (Zurek, 1993) or stability analysis under the system–environment coupling.

Definition D5 (Operational Invariance) then tests robustness across any co-admissible candidates that survive this selection.

The identification of the pointer algebra for a given Hamiltonian is a research problem, not a closed algorithm; D5 converts this openness into a falsifiable condition rather than leaving it as an ambiguity.

A2.2 — D2: Dephasing Map $\Delta\mathcal{O}$

Given a density matrix ρ on $\mathcal{H}_s$, define the dephasing map relative to $\mathcal{O}$ as $\Delta\mathcal{O}$ removes quantum interference between record sectors while preserving classical probabilities.

Critical clarification. Pay attention to this — it is where most confusion enters. $\Delta\mathcal{O}$ does not measure ignorance.

It enforces projection onto the record algebra, isolating entropy attributable to irreversible branching rather than lack of knowledge. This step prevents conflating classical uncertainty with physical actualization.

A2.3 — D3: Actualization State — Primary Definition

Statement of the Definition

Let ρ be the reduced density matrix of a system after tracing over inaccessible degrees of freedom. Let $\mathcal{O} = \{\Pi_i\}$ be a physically realizable coarse-graining selected by system–environment interaction (Definition D1).

Define the dephasing map relative to this record algebra: $\Delta\mathcal{O}(\rho) \equiv \Sigma_i \Pi_i \rho \Pi_i$. The Actualization State (AS) is defined as where S_eff is the effective record entropy defined below.

Effective Entropy When record sectors Π_i have rank greater than one, the dephased entropy decomposes as where $p_i = Tr(\Pi_i \rho)$ and σ_i is the normalised intra-sector state. AS tracks inter-sector branching only.

What happens inside each sector is invisible to AS — deliberately so. The effective entropy entering AS is defined as

with intra-sector entropy discarded by construction. For rank-1 sectors, $S(\Delta\mathcal{O}(\rho)) = H(\{p_i\})$ and no distinction arises. A formal derivation is provided in Appendix A.

Rank-1 Simplification For rank-1 sectors (pure pointer states), the definition simplifies to: where H is the Shannon entropy and $N = |\mathcal{O}|$ is the number of record sectors.

Normalization and Physical Bounds

Let $\mathcal{O} = \{\Pi_i\}$ be the physically realizable record algebra, with total record-algebra dimension $d\mathcal{O} \equiv \Sigma_i \operatorname{rank}(\Pi_i)$.

Define: $S_{max}(\mathcal{O}) = \log d\mathcal{O}$, $S_{min}(\mathcal{O}) = 0$, where the minimum corresponds to support on a single record sector.

Crucial clarification. $S_{min}(\mathcal{O}) = 0$ is attained whenever the system's accessible state is pure and confined to a single record sector, even when Π_i has rank greater than one.

Internal degeneracy or unmonitored degrees of freedom within a sector do not contribute to actualization. AS therefore vanishes upon complete selection into a single record sector, independent of that sector's internal rank.

Interpretation: The AS Inversion

Why Dephasing-Relative Entropy (and Not Purity Alone)

Purity $\mathrm{Tr}(\rho^2)$ conflates two distinct situations: a system that was never coherent but is classically mixed due to ignorance, and a system that was coherent and has irreversibly decohered into record-consistent alternatives.

Both can share the same purity. Only the second represents actualization.

By applying $\Delta\mathcal{O}$ before evaluating entropy, the definition isolates suppression of interference relative to physically realized records. Ignorance alone does not count as actualization. Formally, AS is a dephasing-relative entropy of coherence in the $\mathcal{O}$-algebra.

Normalization Discipline

Normalization is tied to $d\mathcal{O}$, not the abstract Hilbert-space dimension. S_{min} and S_{max} are fixed by what the environment can physically record.

This prevents artificial inflation or suppression of AS by adding unmonitored degrees of freedom.

Operational Meaning

AS answers one question, and only one: To what extent has the system developed record-structured commitment across the alternatives that its physical environment can distinguish?

Think of it this way. A spinning coin has AS = 1 — maximum branching, both faces equally possible. A landed coin has AS = 0 — one face, no alternatives.

AS measures how much spinning is left. Not which face will land. Just: how much spinning.

It does not answer: how fast this occurred, what caused it, or whether you should care, whether it is fundamental or emergent. Those questions belong to later sections. Notice what AS does NOT tell you.

It does not tell you which outcome will happen. It does not tell you when.

It tells you how much branching exists right now. That is all. Crucially, AS is computed from the dephased state $\Delta\mathcal{O}(\rho)$, not from the physical state ρ directly.

A system may have AS = 1 before decoherence has physically occurred, because the sector weights of the dephased state already distribute maximally.

AS measures branching structure, not decoherence progress; the latter is tracked by the no-return surface (D13).

Clarification. AS measures inter-sector branching structure in the record algebra. It does not by itself measure operational irreversibility, which is separately determined by loss of recoverability (Definition D13, Section A4.1).

A system may have AS = 1 before decoherence is physically complete, because the sector weights of the dephased state already distribute maximally.

AS tracks branching potential; the no-return surface tracks irreversibility.

AS is defined as a scalar. But the space it measures — the space of possible record configurations — is a genuine degree of freedom.

In the full corpus (AP10), it is the fifth dimension: three spatial, one temporal, one actualization.

In the full framework (AP10), the actualization dimension is identified as the fifth degree of freedom, alongside three spatial and one temporal. The manifold lives in 3+1 dimensions; the quantum sector lives in the fifth.

Every record is written from the actualization dimension into the manifold.

AS measures how far along this dimension a system has progressed: from pure possibility (AS = 0 in a single sector) through maximal branching (AS = 1).

This interpretation is not required by the definitions of A0–A3, which are self-contained. It is offered as structural orientation for readers approaching AP01 within the broader corpus.

The fifth dimension is prior to the manifold — it is the precondition for the manifold's existence — but it is not less real for being prior.

It is the dimension from which physics is written.

Scope and Kill-Switch

AS is defined relative to a coarse-graining. Absolute, basis-free AS is meaningless. Physical legitimacy is enforced by Definition D5 (Operational Invariance): if AS varies beyond experimental tolerance across physically realizable $\mathcal{O}$, the argument fails.

This is the built-in termination condition.

Worked Example: Two-Qubit Dephasing

Numbers make abstractions real. Follow this example and you will understand AS better than any amount of definition-reading can provide.

Consider two qubits, S_1 and S_2, each coupled to an independent environment fragment, with pointer basis $\mathcal{O} = \{\Pi_{00}, \Pi_{01}, \Pi_{10}, \Pi_{11}\}$ where $\Pi_{ij} = |ij\rangle\langle ij|$ and $i, j \in \{0, 1\}$.

The record algebra has $d_t = 4$ sectors.

Initial state: $|\psi(0)\rangle = (|00\rangle + |01\rangle + |10\rangle + |11\rangle)/2 \otimes |E_0\rangle$. The reduced state after dephasing is $\Delta_t(\rho_s) = \mathrm{diag}(¼, ¼, ¼, ¼)$. Sector weights: $p_{ij} = ¼$ for all (i,j).

Effective entropy: S_eff = H({¼, ¼, ¼, ¼}) = log 4. Normalization: S_max = log 4. Therefore AS = log 4 / log 4 = 1. Maximal branching across four record sectors.

Now suppose only S_1 has decohered while S_2 remains coherent.

The accessible state is $\rho_s = ½(|0\rangle\langle 0| \otimes |+\rangle\langle +|) + ½(|1\rangle\langle 1| \otimes |+\rangle\langle +|)$ where $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. After dephasing in the four-sector basis: $p_{00} = p_{01} = p_{10} = p_{11} = ¼$.

AS = 1 again.

The sector weights are identical despite partial decoherence. This illustrates the key point: AS tracks the branching structure of the dephased state, not the physical progress of decoherence.

The distinction between these two situations is captured not by AS but by the no-return surface (D13): in the first case the system has crossed it; in the second it has not.

Finally, suppose decoherence has completed but one sector has been depopulated by dissipation, yielding weights $p_{00} = ½$, $p_{01} = ¼$, $p_{10} = ¼$, $p_{11} = 0$.

Then S_eff = H(½, ¼, ¼, 0) = ½ log 2 + ½ log 4 = ½ log 2 + log 2 = $^3/_2$ log 2, and AS = ($^3/_2$ log 2) / log 4 = $^3/_4$ = 0.75. Dissipation has reduced branching.

This is consistent with Proposition T1b: non-unital dynamics can decrease AS.

A2.4 — D4: History-Based Representation of AS (AS_h)

Let {α} denote a set of coarse-grained histories defined by $\mathcal{O}$, with decoherence functional D(α, β).

Define the history-based representation of Actualization State as where $p_a = D(\alpha, \alpha)$, $H(\{p_a\})$ is the Shannon entropy, and N is the number of admissible histories.

$AS_h = 0$: a single trivial history (no branching). $AS_h = 1$: maximal record-structured branching. Under complete decoherence in the $\mathcal{O}$-algebra, the dephased entropy reduces to $H(\{p_i\})$.

In this regime, AS_h coincides with the primary AS definition up to normalization. Formal conditions for equivalence are established in Appendix A.

If you are reading for the structural argument, the primary definition is all you need. The history representation exists for readers who work in the consistent-histories formalism.

Conventions on Approximation and Operational Criteria

Throughout this work, qualitative terms such as effective, operational, or inaccessible are shorthand for quantitatively defined conditions.

Effective orthogonality. Two states ρ and σ are effectively orthogonal iff $\frac{1}{2}\|\rho - \sigma\|_1 \geq 1 - \varepsilon$, for fixed operational tolerance $\varepsilon > 0$. Equivalently, they are operationally distinguishable to within ε of perfect discrimination.

Operational inaccessibility. A property is operationally inaccessible iff no admissible CPTP map Λ acting on the accessible system (optionally with fresh ancilla) can alter the reduced state by more than ε in trace distance: $\|\Lambda(\rho_s) - \rho_s\|_1 \leq \varepsilon$ for all admissible Λ.

The parameter ε represents experimental resolution and is held fixed within any given analysis.

A2.5 — D5: Operational Invariance Test (Kill Switch F0)

Setup

Let $\mathcal{O}_1 = \{\Pi_i^{(1)}\}$ and $\mathcal{O}_2 = \{\Pi_j^{(2)}\}$ be two physically realizable coarse-grainings of the same experimental system, each admissible under Definition D1 for the same preparation and control protocol.

Let ρ be the reduced state inferred from experimentally accessible data for that protocol, and let $\delta_{exp} > 0$ denote the experimentally justified tolerance.

Definition D5 (Operational Invariance Requirement)

AS is operationally invariant iff for all admissible pairs ($\mathcal{O}_1$, $\mathcal{O}_2$) and all experimentally accessible ρ,

If two coarse-grainings are both physically realized by the same system–environment coupling and apparatus constraints, they must not yield incompatible AS values beyond experimental tolerance.

Falsifier F0 (Global Kill Switch)

If there exists any experimentally realizable system and any co-admissible pair ($\mathcal{O}_1$, $\mathcal{O}_2$) such that repeated trials yield $|AS(\rho; \mathcal{O}_1) - AS(\rho; \mathcal{O}_2)| > \delta_{exp}$ persistently, then AS is not a well-defined operational quantity and the argument is falsified.

That is a global termination condition.

Notes on Scope

Read F0 again. It is the most important sentence in this paper. If two legitimate ways of measuring the same system give different AS values beyond experimental tolerance, the ENTIRE programme is dead.

Not just this paper. Everything built on it. Every subsequent proof. Every ethical conclusion. All of it.

That is what honest argument looks like — it hands you the tools to destroy it.

D5 does not require invariance under arbitrary refinements, coarse mathematical partitions, or observer-chosen bases. D5 requires robustness only across physically realized record algebras within the same experimental context.

Failure of later dynamical postulates does not affect D5; conversely, failure of D5 invalidates the entire AS program.

Worked Example: Operational Invariance in Circuit QED

Consider a transmon qubit dispersively coupled to a microwave cavity.

The system–environment interaction selects the charge-parity basis as the pointer algebra: $\mathcal{O}_1 = \{|g\rangle\langle g|, |e\rangle\langle e|\}$, where $|g\rangle$ and $|e\rangle$ are the ground and excited states of the transmon.

This is a physically realizable coarse-graining in the sense of D1: it is selected by the dispersive Hamiltonian $H\_int = \chi a^\dagger a \sigma\_z$, which entangles photon number with qubit state.

Now consider a second coarse-graining arising from the same physical setup.

If the cavity is driven to produce a coherent state that rotates conditionally on the qubit state, the environment effectively records in a rotated basis.

However, for the dispersive interaction, the conditional phase shift $\varphi = \chi t$ on the cavity field produces pointer states that remain the energy eigenstates $|g\rangle$, $|e\rangle$ regardless of drive parameters.

Any co-admissible coarse-graining $\mathcal{O}_2$ arising from the same dispersive coupling must therefore coincide with $\mathcal{O}_1$ up to a relabeling of sectors.

Operational invariance test. Prepare the qubit in state $|\psi\rangle = \alpha|g\rangle + \beta|e\rangle$ and allow decoherence via photon-number-dependent dephasing. The reduced state is $\rho_s = |\alpha|^2|g\rangle\langle g| + |\beta|^2|e\rangle\langle e|$.

$AS(\rho; \mathcal{O}_1) = H(|\alpha|^2, |\beta|^2) / \log 2$. Because $\mathcal{O}_2$ coincides with $\mathcal{O}_1$ for this coupling, $AS(\rho; \mathcal{O}_2) = AS(\rho; \mathcal{O}_1)$ exactly.

D5 is satisfied with $\delta\_exp = 0$.

The example is deliberately simple: in circuit QED, the dispersive coupling uniquely determines the pointer basis, so co-admissible coarse-grainings are trivially equivalent. That is the point.

Operational invariance is most easily verified in systems where the coupling Hamiltonian strongly constrains the pointer algebra.

The interesting tests of D5 arise in systems with richer coupling structures — and those are the systems that will either confirm or kill the argument, where multiple pointer candidates compete and small differences in AS across co-admissible bases can be measured against δ_{exp}.

A3 — Theorems: Irreversibility and No-Return

The definitions are set. Now the proofs. What follows cannot be argued with — it can only be tested.

A3.1 — Theorem T1: Monotonicity of AS Under Decohering Dynamics

Setup

Let $\rho(t)$ be the reduced state of a system evolving under a completely positive trace-preserving (CPTP) dynamical semigroup $\{\mathcal{E}_t\}_{t \geq 0}$ with generator $\mathcal{L}$.

Let $\mathcal{O} = \{\Pi_i\}$ be a physically realizable record algebra (Definition D1), and let $\Delta\mathcal{O}(\rho)$ denote the associated dephasing map.

Theorem T1 (Statement) Here is the central result of Paper A. Everything before it was preparation. Everything after it is consequence.

The result says: under three precisely stated conditions, the branching can only grow. The coin cannot un-land. The ink cannot un-dry. The record cannot un-write. Possibility becomes fact, and the transition is one-way.

Not because physics prohibits reversal — but because the conditions that create records are the conditions that make the inequality hold.

Actualization State is monotonic nondecreasing along the evolution $\rho(t)$ provided the following minimal sufficient conditions hold:

(1) Decoherence relative to the record algebra. Interference between record sectors is not regenerated: $(d/dt)\, C\mathcal{O}(\rho(t)) \leq 0$, where $C\mathcal{O}$ is any coherence monotone that vanishes on $\Delta\mathcal{O}(\rho)$.

Equivalently, off-diagonal terms in the $\mathcal{O}$-basis decay monotonically.

(2) Closure of the record algebra. $\Delta\mathcal{O} \circ \mathcal{E}_t = \mathcal{E}_t \circ \Delta\mathcal{O}$ for all $t \geq 0$. This ensures that populations in the record sectors evolve autonomously once decohered.

(3) Unital (mixing) dynamics on the record algebra. The sector weights $p_i(t) = \mathrm{Tr}(\Pi_i\, \rho(t))$ evolve under a doubly stochastic map: $p(t) = M(t)\, p(0)$, where M preserves the uniform distribution.

In plain language: the mixing is fair. No sector is favoured. The dynamics spread probability out, not funnel it in.

Conclusion Under conditions (1)–(3), That is, Actualization State is monotonic along decohering, record-forming dynamics.

Scope of the theorem. Theorem T1 is a conditional statement. Its domain of applicability is exactly the set of dynamics satisfying conditions (1)–(3).

Dynamics outside this domain — including dissipative, non-unital, or feedback-controlled evolution — may cause AS to decrease.

This does not contradict the theorem; it indicates the dynamics are not purely record-forming in the sense defined above. The explicit scope boundary and converse are given below.

Note on novelty and scope. Conditions (1)–(3) are sufficient but not necessary. The monotonicity of Shannon entropy under doubly stochastic mixing is a standard result (Schur convexity), and this paper does not claim otherwise.

What is new is not the mathematical inequality but its physical application: the identification of conditions under which doubly stochastic mixing is the correct effective description of decoherence in a physically realized record algebra, the isolation of inter-sector branching entropy (via the dephasing map $\Delta\mathcal{O}$) from thermodynamic entropy, and the explicit rate bound (T1a) and converse (T1b) that turn the monotonicity into a diagnostic tool for identifying record-forming dynamics.

The theorem is a known inequality applied in a new physical context; the contribution is the context, not the inequality.

You are seeing a known mathematical tool applied to a new physical question — and the answer it gives is decisive.

Lemma T1.1 (Doubly stochastic mixing). Under conditions (2) and (3), the sector weights p(t) evolve under a doubly stochastic matrix M(t) for all $t \geq 0$. The proof is short, and it is the engine that drives everything.

Proof. By condition (2), $\Delta_{\mathcal{O}} \circ \mathcal{E}_t = \mathcal{E}_t \circ \Delta_{\mathcal{O}}$, so the evolution commutes with dephasing. Therefore the diagonal elements evolve autonomously: there exists a linear map $M(t)$ such that $p(t) = M(t)\, p(0)$.

The map is stochastic because $\mathcal{E}_t$ is trace-preserving: $\Sigma_i\, p_i(t) = 1$ for all t.

Unitality (condition 3) means $\mathcal{E}_t(I/d) = I/d$. Applying $\Delta_{\mathcal{O}}$ to both sides and using condition (2): $\Delta_{\mathcal{O}}(\mathcal{E}_t(I/d)) = \mathcal{E}_t(\Delta_{\mathcal{O}}(I/d)) = \mathcal{E}_t(I/d) = I/d$. In terms of sector weights, the uniform distribution is fixed: $M(t)(1/N, \ldots, 1/N)^T = (1/N, \ldots, 1/N)^T$.

A stochastic matrix that preserves the uniform distribution is doubly stochastic. □

The lemma is short. Its consequence is not. Once you know the mixing is doubly stochastic, Shannon's inequality does the rest.

The monotonicity of AS follows from a chain of implications, each link forged by standard mathematics.

With Lemma T1.1 established, the monotonicity of Shannon entropy under doubly stochastic mixing is a standard result (Schur-convexity of −H). The theorem follows. The proof is standard mathematics applied to a new physical context.

What is new is not the inequality but the recognition that record-forming dynamics satisfy exactly the conditions that make the inequality hold.

Interpretation

When interference between record-distinguishable alternatives is suppressed, the record algebra is dynamically closed, and record-sector probabilities mix without coherent backflow, the informational richness of committed classical branching cannot decrease.

This monotonic increase defines the arrow of actualization. You have lived inside this arrow your entire life. Every moment has been forward.

Every record has been permanent. The theorem says why: under the conditions that create records, the branching can only grow.

Explicit Scope Boundary

Outside these three conditions, the guarantee is void. AS may decrease under cooling, decay, relaxation, feedback control, or any dynamics that funnel probability into fewer sectors. The theorem does not claim universality. It claims precision.

These cases do not contradict the theorem; they lie outside its scope by construction.

Why the Scope Is Exact

Pay attention to this. It is the difference between a real theorem and a hand-wave.

AS measures branching, not definiteness. Decoherence creates branching — AS increases. Selection resolves branching — AS decreases. T1 applies only to the creation phase.

The theorem is not claiming that AS always increases everywhere forever. It is claiming something much more precise: under exactly these three conditions, AS cannot decrease. If you violate the conditions, the guarantee is void.

The conditions are not technicalities. They are the physics. And the physics is testable.

Quantitative Sharpening: Rate Bounds and Converse

Theorem T1 establishes monotonicity but does not quantify the rate at which AS approaches its equilibrium value, nor does it characterize conditions under which AS must decrease.

The following two results sharpen T1 in both directions.

Corollary T1a (Rate of convergence). Under conditions (1)–(3), let the induced classical dynamics on the sector weights be governed by a continuous-time doubly stochastic rate matrix W, so that dp/dt = Wp.

Let $\lambda_2 < 0$ denote the second-largest eigenvalue of W (the spectral gap). Then the deviation of AS from its equilibrium value AS_eq = 1 satisfies where C depends on the initial condition.

Derivation: for any initial distribution p(0), your deviation from uniformity satisfies $\|p(t) - \pi\|_1 \leq \sqrt{N} \cdot \exp(\lambda_2 t)$ (by standard contraction bounds for reversible Markov chains).

The Shannon entropy H(p) is Lipschitz in the L_1 norm on the probability simplex: $|H(p) - H(q)| \leq \|p - q\|_1 \cdot \log N$ (continuity bound for entropy; Cover and Thomas 2006).

Therefore $|AS(t) - 1| = |H(p(t)) - \log N| / \log N \leq \|p(t) - \pi\|_1 \leq \sqrt{N} \cdot \exp(\lambda_2 t)$, where the constant $C = \sqrt{N}$ absorbs the dimension dependence.

The rate of convergence to maximal branching is therefore controlled by the spectral gap of the mixing dynamics, not by any property intrinsic to the AS definition.

For the symmetric d-sector model with uniform inter-sector mixing rate w, the spectral gap is $\lambda_2 = -dw$, and the convergence time to AS ≈ 1 scales as $\tau \sim 1/(dw)$.

This connects AS growth directly to the physical decoherence rate of the record algebra.

The faster the environment records the system, the faster AS climbs. You can measure this. The spectral gap is a physical quantity. The convergence rate is a prediction.

Proposition T1b (Converse: conditions for AS decrease). If condition (3) is violated—specifically, if the induced classical dynamics on sector weights are governed by a stochastic matrix M that is not doubly stochastic, with stationary distribution $\pi \neq$ uniform—then there exist initial distributions $p(0)$ for which AS is strictly decreasing.

In particular, if $p(0)$ majorizes π, the Shannon entropy $H(p(t))$ may increase, but if π majorizes $p(0)$, then $H(p(t))$ can decrease toward $H(\pi) < \log d$.

Physically, Proposition T1b corresponds to dissipative dynamics that preferentially funnel population into a subset of record sectors (e.g., amplitude damping, spontaneous decay into a ground-state sector).

Such dynamics violate the unital condition (3) and drive AS downward.

This converse confirms that condition (3) is not merely a technical convenience but a physical requirement: monotonicity of AS is a signature of symmetric, environment-driven mixing, not of dissipative relaxation.

When AS decreases, something is funnelling probability into fewer branches — cooling, decay, relaxation. When AS increases, the environment is writing records.

The direction tells you which process is dominant.

Summary. T1, T1a, and T1b together establish that AS monotonicity is the exact fingerprint of doubly stochastic record-forming dynamics. T1 gives the direction, T1a gives the rate, and T1b gives the converse.

No further characterization of the monotonicity is needed or claimed.

A3.2 — The Operator Horizon: No-Return as an Inequality

Setup

Let $x(t) \geq 0$ be a scalar variable representing the degree of maintained structure of a system—displacement from its unmaintained equilibrium baseline.

Assume deterministic dynamics: where a is an intrinsic decay/drift rate toward equilibrium, u is a control/maintenance input, and $u_{max} \geq 0$ is a hard upper bound on control capacity.

Theorem T2 (Operator Horizon)

You have felt this theorem in your body. Every system with limited resources — every body, every business, every civilisation — has a point beyond which no strategy can save it.

The theorem names that point. Define the operator horizon

If at some time t_0 the system satisfies $x(t_0) > x_h$, then for all admissible controls u(t),

Once you cross the horizon, x(t) decreases no matter what you do. Maximum effort slows the decline but cannot reverse it. Recovery is impossible by control alone.

Proof

From the dynamics, $dx/dt = -a x + u \leq -a x + u_{max}$.

If $x > u_{max}/a$, then $-a x + u_{max} < 0$, hence $dx/dt < 0$. At $x = x_h$, maximal control yields $dx/dt = 0$; continuity implies monotone approach to x_h from above. □

Interpretation

x_h is a capacity boundary, not a physical wall. It is the maximum sustainable structure given maximum maintenance effort.

Beyond x_h, the system decays toward the horizon regardless of strategy: irreversibility arises from insufficiency of admissible control, not from prohibition of reverse dynamics.

Generalizations

Nonlinear decay: If $dx/dt = -g(x) + u$ with $g(0) = 0$, $g(x)$ increasing, the horizon is implicitly defined by $g(x_h) = u\_max$.

Your body operates under nonlinear decay — the maintenance cost of health increases with age, and the horizon shifts. The qualitative no-return result is unchanged.

Time-dependent capacity: If a(t) or u_max(t) varies, $x_h(t) = u\_max(t)/a(t)$ defines a time-dependent horizon; exceeding the instantaneous horizon for sufficient time yields practical no-return behavior.

You know this. You have watched a garden become overgrown past the point where you could maintain it. You have watched debt compound past the point where income could service it.

You have watched a body deteriorate past the point where medicine could restore it.

The mathematics is confirming what your experience already knows: there is a line, and once you cross it, effort is not enough.

Classical analogy: Consider a leaking bucket with a bounded refill rate. The horizon is the maximum fill level sustainable at maximum pumping. Once above it, the bucket drains regardless of effort.

A3.3 — No-Return Surfaces and Operational Irreversibility

Setup

The scalar horizon is the simple case. Real systems have many dimensions. The generalisation uses viability theory (Aubin, 1991) — the mathematics of survival under constraint.

Definition D6: State Space and Admissible Dynamics. Let $X \subseteq \mathbb{R}^n$ be the state space. Let admissible controls satisfy $u(t) \in U$, where U is compact.

Dynamics: $dx/dt = f(x, u)$, $u \in U$, with f locally Lipschitz in x. Let $R \subset X$ be the recoverable (safe) set.

Definition D7: Viability Kernel. $\mathrm{Viab}(R) \equiv \{ x_0 \in R \mid \exists\, u(\cdot) \in U \text{ s.t. } x(t; x_0, u) \in R\ \forall t \geq 0 \}$.

States from which the system can be kept inside R indefinitely using admissible control.

Definition D8: Capture Basin. $\mathrm{Cap}(R) \equiv \{ x_0 \in X \mid \forall\, u(\cdot) \in U, \exists\, t \geq 0: x(t; x_0, u) \notin R \}$.

States from which exit from R is inevitable under all admissible controls.

Definition D9: No-Return Surface. $\Sigma_h \equiv \partial\, \mathrm{Viab}(R)$. This is the geometric generalization of the scalar horizon x_h.

Proposition P3.3

Assume the standard viability conditions (local Lipschitz continuity of f, compact U).

Then: (1) Viab(R) consists of states from which at least one admissible control avoids loss indefinitely; Cap(R) consists of states from which all admissible controls lead to loss in finite time.

- Crossing Σ_h transfers the system from a region where recovery is reachable to a region where it is not. Caveat: Viab(R) and Cap(R) partition X up to the boundary Σ_h.

Boundary states may be marginal.

Definition D10: Operational Irreversibility (Quantified). A state x_0 is operationally irreversible with respect to R iff $x_0 \notin$ Viab(R). The reverse transition to R does not exist under admissible controls.

Operational irreversibility depends on reachability under constraints, not on microscopic time-reversal symmetry. The distinction matters. A shattered vase is not irreversible because physics forbids reassembly.

It is irreversible because you do not have the resources, precision, or time to reassemble it. Irreversibility is about what you can do, not about what nature prohibits.

Consistency with A3.2

A3.2 is recovered as the special case: $n = 1$, $f(x, u) = -ax + u$, $R = [0, x_h]$. Then $Viab(R) = [0, x_h]$, $\Sigma_h = \{x_h\}$.

The scalar horizon is exactly the one-dimensional no-return surface.

You now have the complete no-return geometry. The scalar horizon (T2) is the simple case. The viability kernel (D7) is the general case.

The no-return surface (D9) is the boundary between where you can still recover and where you cannot.

Every system you have ever cared about — your body, your relationships, your work — has this geometry. The mathematics names what your experience already knows.

Robustness Notes

With stochastic forcing, replace viability by almost-sure or probabilistic viability; Σ_h becomes a probabilistic boundary. The concept is unchanged; only the quantifier changes. Viab(R) and Σ_h are computable via Hamilton–Jacobi reachability methods.

A4 — Quantum Mechanical Instantiation

Sections A0–A3 are complete. They stand alone. What follows adds independently falsifiable postulates — each one an invitation to destroy a specific claim. If any postulate in this section falls, everything above it survives.

You lose the extension, not the foundation.

Sections A0–A3 establish definitions and theorems that are self-contained: they depend only on operational definitions, standard quantum mechanics, and viability theory. Nothing in A0–A3 requires the content of A4 or A5.

The section introduces postulates that extend the argument to address selection and gravitational rate constraints.

These postulates are independently falsifiable: each has explicit conditions under which it fails (F1–F3, G1–G3), and failure of any postulate here does not invalidate the definitions, theorems, or irreversibility results of A0–A3.

The transition from theorem to postulate is marked explicitly at each point of introduction.

A4.1 — Operational Irreversibility in Open Quantum Systems

The viability geometry of Section A3 now meets quantum mechanics. The abstract becomes concrete. You will recognise the structure.

Setup: Accessible Dynamics

You have access to the system but not the environment. That restriction is the source of irreversibility.

Let the total Hilbert space factor as $\mathcal{H} = \mathcal{H}_s \otimes \mathcal{H}_e$, with joint state evolving unitarily under H_{se}.

The accessible system state — the one you can actually measure — is $\rho_s(t) = \mathrm{Tr}_e[U(t)\, \rho_{se}(0)\, U^\dagger(t)]$.

Admissible operations are restricted to system-local CPTP maps — the operations you can actually perform on the system without accessing the environment $\Lambda: \mathcal{B}(\mathcal{H}_s) \to \mathcal{B}(\mathcal{H}_s)$, optionally employing fresh ancilla, but with no access to the original environment E.

Definition D11: Coherence-Recoverable States. A system state ρ_s is coherence-recoverable relative to $\mathcal{O}$ iff there exists an admissible CPTP map Λ such that $\|\Lambda(\rho_s) - \rho\_coh\|_1 \leq \varepsilon$, for some state ρ_coh satisfying $\Delta\mathcal{O}(\rho\_coh) \neq \rho\_coh$, and for fixed operational tolerance $\varepsilon > 0$.

Definition D12: Recoverable Set and Viability Kernel. $K\varepsilon(\mathcal{O}) := \{ \rho_s \mid \rho_s$ is coherence-recoverable relative to $\mathcal{O} \}$. $K\varepsilon(\mathcal{O})$ is the viability kernel of coherence under admissible control.

States outside $K\varepsilon(\mathcal{O})$ are operationally irreversible with respect to $\mathcal{O}$.

Mapping to viability theory (Definitions D6–D9). The abstract viability framework of Section A3.3 instantiates in the quantum setting as follows.

State space X: the convex set of density operators on $\mathcal{H}_s$, equipped with the trace-norm topology.

Admissible controls: the set of system-local CPTP maps (optionally with fresh ancilla), representing all operations an agent can perform on S alone. Dynamics f(x, u): the discrete-time or continuous-time evolution generated by applying admissible operations.

Recoverable set R: the set of states from which coherence between record sectors can be restored (Definition D11).

Viability kernel Viab(R): exactly $K\varepsilon(\mathcal{O})$ — the set of states for which there exists an admissible control strategy that maintains coherence-recoverability indefinitely.

Capture basin Cap(R): the set of states from which, under any admissible control strategy, the system eventually becomes operationally irreversible. No-return surface Σ_h: the boundary $\partial K\varepsilon(\mathcal{O})$, separating recoverable from irreversibly decohered states.

The scalar horizon x_h = u_max/a of Theorem T2 is the one-dimensional special case of this general geometric construction.

Proposition P4.1: Tracing Induces Loss of Recoverability

If system–environment interaction produces correlations such that distinct record sectors become correlated with orthogonal (or trace-distance-separated) environment states, then for sufficiently small ε, $\rho_s(t) \notin K\varepsilon(\mathcal{O})$.

Once which-record information is encoded in inaccessible degrees of freedom, no admissible operation on S alone can restore coherence between record sectors.

You have felt this. Once the words leave your mouth, you cannot un-say them.

The environment has recorded them — in the other person's memory, in the vibrations of the air, in the electromagnetic radiation that left the room at the speed of light.

No operation on your mouth alone can undo what the environment now holds.

Proof sketch. Orthogonality of environment states implies non-injectivity of the reduced dynamics on $\mathcal{H}_s$.

Multiple globally distinct states map to the same ρ_s, and no CPTP map on S can reconstruct the lost phase information. ■

Definition D13: Operational No-Return Surface (Quantum). The operational no-return surface relative to $\mathcal{O}$ is the boundary $\partial K\varepsilon(\mathcal{O})$.

Crossing this boundary transfers the system from a region where coherence is recoverable to a region where it is not. Irreversibility is identified with loss of reachability. Not with energy dissipation. Not with entropy increase.

With the fact that you cannot get back, not with violation of microscopic reversibility.

Interpretation

Operational irreversibility in quantum mechanics arises not from non-unitarity, but from restricted access.

Tracing over inaccessible degrees of freedom removes states from the recoverable set K_ε, producing a no-return surface in state space exactly analogous to the operator horizons of Section A3.

Irreversibility is a geometric property of admissible control, not a statement about time reversal at the microscopic level.

A4.2 — Objective Actualization Channels and Selection Dynamics

Impossibility of Deterministic Selection by Linear CPTP Dynamics

Proposition. No deterministic, linear CPTP map acting on the system state can transform a diagonal mixture over record sectors into a single realized sector in individual runs.

Linear CPTP evolution preserves convex mixtures: $\mathcal{E}(\Sigma_i\, p_i\, \rho_i) = \Sigma_i\, p_i\, \mathcal{E}(\rho_i)$. Any linear CPTP map that acts identically on each component preserves the mixture structure and cannot yield single-outcome definiteness in individual realizations.

Any mechanism that resolves a decohered mixture — any mechanism that produces the definiteness you experience into a single realized branch must involve either a stochastic unraveling or an explicitly nonlinear effective evolution at the level of single trajectories.

Read that again. Standard quantum mechanics — linear, deterministic, trace-preserving — cannot produce definiteness in individual runs. Something else is required. The proof is not complicated. Linear maps preserve mixtures.

If you feed a mixture in, you get a mixture out. Something nonlinear or stochastic must intervene for one outcome to become THE outcome.

Definition D14: Decoherence vs. Selection. Decoherence suppresses interference between record-distinguishable alternatives and yields a stable diagonal mixture in the record algebra $\mathcal{O}$.

Selection is the further transition from a diagonal mixture to a single realized branch. Decoherence is sufficient for irreversibility. Selection is required for definiteness. You live in a definite world. Something selects.

These are distinct processes. You experience them as distinct. The formalism confirms your experience.

Operational Meaning of “Realized Branch”

A branch i is realized iff, for all subsequent times and for any sequence of admissible operations confined to $\mathcal{H}_s$, the accessible state behaves as if it had been prepared in the conditional state $\rho_i \equiv (\Pi_i\ \rho_s\ \Pi_i)/p_i$ at the onset of selection, in the sense that for all admissible POVMs M on S, $|\mathrm{Tr}(M\ \rho_s(t)) - \mathrm{Tr}(M\ \rho_i)| \leq \varepsilon$ for operational tolerance ε.

Decoherence separates the branches. Selection picks one. You have felt both — the moment when the options became clear (decoherence) and the moment when you chose (selection). The physics mirrors the experience.

Postulate P: Objective Actualization Channel

The paper posits an objective actualization channel acting on the reduced state after decoherence has established operational irreversibility.

The ensemble evolution of the reduced state is written schematically as where $D\mathcal{O}$ is the standard pointer-dephasing channel and $A\mathcal{O}$ is a selection channel responsible for definiteness. The master equation governs ensemble dynamics.

Single-run definiteness requires a stochastic or nonlinear unraveling of $A\mathcal{O}$.

Structural Requirements on the Selection Channel

(S0) Activation Condition. $A\mathcal{O} \approx 0$ until the decoherence condition (D13) holds within tolerance ε. Selection activates only after branches are operationally distinct.

(S1) Record-algebra locality. $A\mathcal{O}(\rho_s) = A\mathcal{O}(\Delta\mathcal{O}(\rho_s))$. Selection never creates interference.

(S2) Sector fixed points. $A\mathcal{O}(\Pi_i\, \rho_s\, \Pi_i) = 0$ for all i. Once a branch is realized, dynamics cease.

(S3) Contractivity (single-run resolution). Under the stochastic unraveling, the Shannon entropy $H(\{p_i\})$ is a supermartingale: it decreases along individual trajectories almost surely.

(S4) Born boundary condition. For an ensemble of identical preparations at the onset of selection, the distribution of realized branches converges to $\{p_i\}$. S4 is a boundary condition, not a derivation.

Notice what this does and does not say. It does not explain why the Born rule holds. It says: whatever selection is, it must produce Born statistics at the ensemble level. The constraint is structural.

The explanation is someone else's problem.

Five structural requirements. Not a mechanism — an interface. Whatever selection IS, it must satisfy these five constraints. The constraints are testable. The mechanism is nature's business.

Observational consequence. Any selection channel satisfying S0–S4 produces, at the single-trajectory level, dynamics that are not reproducible by any linear CPTP map acting on $\mathcal{H}_s$.

This is the empirical signature of selection: post-decoherence statistics that violate linearity.

Concretely, monitoring a single system through the selection process should reveal either diffusive wandering or discrete jumps in sector weights—neither of which is consistent with a Lindblad master equation applied after dephasing is complete.

This prediction is the target of test R5 (A5.2).

Toy Model: Stochastic Selection on a Qubit

Abstract requirements are convincing only when you can see them work in a concrete case. Here is the simplest possible example — a single qubit. Watch how all five requirements are satisfied simultaneously.

The structural requirements S0–S4 constrain the selection channel but do not uniquely determine it.

To demonstrate that the requirements are jointly satisfiable and to anchor the postulate in a concrete mathematical object, we exhibit a minimal toy model: diffusive selection on a single qubit.

Setup. Let the system be a qubit with pointer basis $\mathcal{O} = \{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. After decoherence is complete, the reduced state is $\rho_s = \mathrm{diag}(p, 1-p)$, with $p \in [0, 1]$.

The selection channel acts on the single free parameter p via the Itô stochastic differential equation where W(t) is a standard Wiener process and $\gamma > 0$ is the selection rate parameter.

The diffusion coefficient $\sigma(p) = \sqrt{\gamma} \cdot p(1 - p)$ vanishes at the boundaries $p = 0$ and $p = 1$, making both endpoints absorbing states.

Verification of S0–S4.

(S0) Activation condition: γ is set to zero until the decoherence threshold (D13) is satisfied. Prior to activation, p evolves only under standard quantum dynamics.

(S1) Record-algebra locality: The SDE acts entirely on the diagonal sector weight p. No off-diagonal (coherence) terms are generated. The map is record-algebra local by construction.

(S2) Sector fixed points: At $p = 0$ and $p = 1$, the diffusion coefficient vanishes identically: $\sigma(0) = \sigma(1) = 0$. Both sector-pure states are absorbing. Once a branch is realized, dynamics cease.

(S3) Contractivity (full Itô calculation): Define $H(p) = -p \log p - (1-p) \log(1-p)$. The selection SDE is $dp = \sqrt{\gamma} \cdot p(1-p)\, dW$. Apply Itô's formula: $dH = H'(p)\, dp + ½\, H''(p)\, (dp)^2$.

Compute: $H'(p) = -\log p + \log(1-p) = \log[(1-p)/p]$. $H''(p) = -1/p - 1/(1-p) = -1/[p(1-p)]$.

The quadratic variation is $(dp)^2 = \gamma\, p^2(1-p)^2\, dt$. Substituting: $dH = \log[(1-p)/p] \cdot \sqrt{\gamma} \cdot p(1-p)\, dW + ½ \cdot [-1/(p(1-p))] \cdot \gamma\, p^2(1-p)^2\, dt$.

The drift term simplifies: $½ \cdot [-1/(p(1-p))] \cdot \gamma\, p^2(1-p)^2 = -(\gamma/2)\, p(1-p)$. Therefore: $dH = -(\gamma/2)\, p(1-p)\, dt + \sqrt{\gamma} \cdot p(1-p) \cdot \log[(1-p)/p]\, dW$.

The drift term $-(\gamma/2)\, p(1-p)$ is strictly negative for all $p \in (0, 1)$, with equality only at the absorbing boundaries $p = 0$ and $p = 1$. The dW term is a martingale (zero expectation).

Therefore $E[dH] = -(\gamma/2)\, p(1-p)\, dt \leq 0$, establishing that H is a supermartingale.

Shannon entropy decreases along individual trajectories almost surely, with rate proportional to γ. The mathematics just proved that the branching resolves. Not on average. Not in expectation. Along every single trajectory.

The coin lands. and to the product $p(1-p)$, which is maximal at $p = ½$ and vanishes at the absorbing boundaries.

(S4) Born boundary condition: The process $dp = \sqrt{\gamma} \cdot p(1-p)\, dW$ is a martingale, since the drift term is zero: $E[p(t)] = p(0)$ for all t.

Because p converges almost surely to {0, 1}, the probability of absorption at $p = 1$ is exactly $p(0)$, and absorption at $p = 0$ occurs with probability $1 - p(0)$.

The ensemble statistics of realised branches therefore reproduce the Born weights without any additional assumption. No extra postulate. No interpretation.

The martingale property of the process — the same mathematics that governs fair games in probability theory — forces the Born rule. The universe plays fair.

Connection to the gravitational limiter. Postulate G (A4.3) constrains the selection rate: $\gamma \le \Delta E_G/\hbar$.

In the toy model, the mean time to absorption scales as $\tau \sim 1/\gamma$, so the gravitational bound imposes $\tau \ge \hbar/\Delta E_G$.

For gravitationally indistinguishable records ($\Delta E_G = 0$), $\gamma = 0$ and no selection occurs.

What the toy model does not determine. The diffusive form of the SDE is a choice, not a derivation.

A jump (Poisson) unraveling would also satisfy S0–S4, as would other diffusion coefficients of the form $\sigma(p) = f(p)$ with $f(0) = f(1) = 0$.

The toy model demonstrates joint satisfiability of the structural requirements and provides a concrete anchor for falsification predictions (e.g., the distinction between diffusive and jump statistics in R5), but it does not claim to be the unique or correct selection dynamics.

The postulate specifies an interface; the toy model is one implementation.

You do not need to know which implementation nature uses. You need to know that any implementation satisfying S0–S4 reproduces what you observe.

This underdetermination is a feature of the postulate's generality, but it has experimental consequences.

Diffusive unravelings predict continuous, Brownian-like wandering of sector weights before absorption, yielding a characteristic 1/f noise spectrum in single-trajectory monitoring. Jump unravelings predict sudden, discrete transitions between sector weights, yielding telegraph noise.

These are operationally distinguishable in systems where single trajectories can be monitored (e.g., continuously measured superconducting qubits or fluorescence-monitored trapped ions).

Test R5 is designed, among other things, to make this distinction. Until such data exist, the postulate correctly remains agnostic about which unraveling nature selects.

Explicit Distinction from Spontaneous Collapse Models

The selection postulate (Postulate P) shares surface features with spontaneous collapse models such as GRW (Ghirardi, Rimini, Weber) and CSL (Continuous Spontaneous Localization) but differs from them in structure, scope, and commitment.

The following distinctions are precise.

Basis of collapse. GRW and CSL collapse the wavefunction in the position basis by construction. The localization operators are spatial: they multiply the wavefunction by a Gaussian centered at a random point in physical space.

The selection postulate instead acts on whatever record algebra $\mathcal{O}$ is selected by the system–environment coupling.

In systems where the pointer basis is not position (superconducting qubits, photon polarization, spin ensembles), GRW/CSL and Postulate P make different predictions. That is the content of falsifier F1 and test R1.

Order of operations. In GRW/CSL, collapse is a universal, always-on process that competes with unitary evolution at all times.

The selection postulate requires an activation condition (S0): selection is negligible until decoherence has rendered record sectors operationally distinct. This ordering is falsifiable via test R5.

If collapse signatures appear before decoherence is complete, the selection postulate fails; if they appear only after, GRW/CSL's always-on character is unnecessarily strong.

Rate structure. GRW introduces a universal collapse rate λ (one collapse per particle per ~10^{16} seconds) and a localization width a ~ 10^{-7} m as free parameters.

CSL replaces the discrete hits with continuous diffusion but retains the same two parameters. The selection postulate introduces no universal rate.

Instead, the rate is bounded from above by the gravitational self-energy distinguishability of the record sectors (Postulate G).

This means the selection rate is system-dependent and vanishes for gravitationally indistinguishable records, a prediction that GRW/CSL do not make.

Ontological commitment. GRW/CSL modify the Schrödinger equation at the fundamental level: they add a stochastic, nonlinear term to the dynamical law governing all matter. The selection postulate does not modify quantum mechanics.

It posits an additional channel that acts on the reduced state after decoherence, without specifying whether this channel is fundamental, emergent, or an effective description of deeper physics.

The limitation is deliberate, not an oversight: the postulate specifies an interface, not an ontology.

Summary of experimental discriminants.

Three tests distinguish the selection postulate from GRW/CSL: (R1) whether selection targets the pointer basis or always position; (R2) whether selection occurs between gravitationally indistinguishable records; and (R5) whether selection requires prior decoherence or operates independently of it.

Agreement on all three would be surprising and informative; disagreement on any one is decisive.

Relation to Actualization State

Selection reduces AS. Decoherence increases AS by creating record-structured branching (A3.1). Selection decreases AS by collapsing that branching into a single realized history. There is no contradiction: AS measures branching richness, not outcome definiteness.

Decoherence opens the fan. Selection closes it. AS tracks the fan.

Falsifiability

F1 (Pointer failure): Selection does not respect the record algebra $\mathcal{O}$.

F2 (Born violation): Ensemble statistics of realized branches deviate from $\{p_i\}$.

F3 (Context dependence): Selection depends on observer intervention rather than objective dynamics.

Failure of the postulate does not invalidate AS, T1–T2, or A4.1.

A4.3 — Physical Constraints on Selection Rates (Gravity as Limiter)

Definition D15: Gravitational Self-Energy Distinguishability. Let two decohered record sectors i and j correspond to mass–energy densities $\rho_i(x)$ and $\rho_j(x)$. Define the gravitational self-energy difference

ΔE_G = 0: the two records are gravitationally indistinguishable. Larger ΔE_G: stronger gravitational distinguishability.

Definition D16: Inter-Sector Selection Rate. Let τ_{ij} be the characteristic time for an individual realization to become operationally indistinguishable from the conditional state within tolerance ε. Define the inter-sector selection rate $\lambda_{ij} \equiv 1/\tau_{ij}$.

Postulate G: Gravity-Limited Selection Rate

The objective selection rate between two record sectors is bounded above by their gravitational distinguishability:

The bound is a limiting inequality, not an equality. Selection may be slower. Selection cannot be faster without invoking a coupling stronger than gravity.

Physical motivation. This bound is not derived from first principles; it is a postulate. But it is not arbitrary.

Two record sectors with distinct mass–energy distributions source distinct gravitational fields. If these sectors are in superposition, the gravitational field itself is in a superposition of distinguishable configurations.

The gravitational self-energy difference ΔE_G quantifies the degree to which the two field configurations are distinguishable: it is the interaction energy of the difference mass distribution with itself.

The energy–time relation $\Delta E \cdot \Delta t \geq \hbar$ then implies that the minimum time required for any physical process to resolve this distinguishability is $\Delta t \sim \hbar/\Delta E_G$.

The bound $\lambda_{ij} \leq \Delta E_G/\hbar$ therefore says: selection cannot resolve two record sectors faster than the gravitational field configurations sourced by those sectors can be distinguished.

The bound is not a derivation of gravitational collapse. It is a consistency constraint.

Whatever mechanism performs selection, it cannot outrun gravitational distinguishability unless it couples to the system more strongly than gravity does.

The bound is testable (R3) and its failure would indicate either that selection couples to a non-gravitational degree of freedom or that the energy–time reasoning does not apply to the selection process.

Think about what this means for everyday objects. A cat in a box has enormous gravitational self-energy difference between “alive” and “dead” configurations — different mass distributions, different gravitational fields.

The bound says selection happens almost instantly. You never see a cat in superposition because gravity resolves it before you could notice.

An electron spin has essentially zero gravitational self-energy difference between “up” and “down” — same mass, same distribution. The bound says selection is negligible. Electrons remain in superposition indefinitely. One inequality.

Both the classical and the quantum world explained.

Distinction from Diósi–Penrose

The Diósi–Penrose proposal argues that gravity causes collapse due to spacetime ambiguity in superposed mass configurations. The present framework makes a weaker claim: gravity limits the rate of selection.

The mechanism of selection is not specified; gravity provides only a ceiling on how fast it can proceed.

This is the difference between “gravity collapses the wavefunction” and “whatever performs selection cannot do so faster than gravitational distinguishability allows.”

Special Cases

Gravitationally indistinguishable records. $\Delta E\_G = 0 \Rightarrow \lambda_{ij} = 0$. Gravity forbids objective selection between them. A superposition of such records can persist indefinitely unless another interaction provides a limiter.

Macroscopic spatial superpositions. For spatially separated mass distributions, ΔE_G grows with mass and separation, yielding $\tau_{ij} \geq \hbar/\Delta E\_G$. Macroscopic records resolve quickly but not instantaneously.

Falsifiability

G1: Selection occurs faster than $\Delta E_G/\hbar$ for gravitationally distinguishable records.

G2: Selection occurs between records with $\Delta E_G = 0$.

G3: Selection rates scale universally with non-gravitational parameters across macroscopic records.

Failure here invalidates only the gravitational-limiter hypothesis, not the selection postulate or prior results.

A5 — Experimental Regimes and Falsification Paths

A5.1 — Orientation: Exclusions Before Fits

Up to A4.3, the argument specifies what must be true if the theory is correct. A5 specifies how it can fail, and how that failure would be observed. Principles: Qualitative exclusions before quantitative fits.

Absence tests before rate tests. Pointer-basis dependence before gravitational scaling. Operational signatures before interpretation.

If any regime below fails, the corresponding theoretical component is dead—cleanly and locally.

A5.2 — Test Map (R0–R5)

R0 — Operational Invariance of AS (Global Kill Switch)

Prepare a system with two physically realizable coarse-grainings $\mathcal{O}_1$ and $\mathcal{O}_2$. Compute AS(ρ; $\mathcal{O}_1$) and AS(ρ; $\mathcal{O}_2$). Prediction: $|AS(\rho; \mathcal{O}_1) - AS(\rho; \mathcal{O}_2)| \leq \delta\_exp$. Falsifier F0: Persistent disagreement beyond tolerance → entire framework fails.

Priority ordering. The tests are listed in order of logical dependency but should be executed in order of discriminatory power.

R0 (operational invariance) is the global kill switch and must be tested first: if R0 fails, the entire framework is dead and no further test is meaningful.

If R0 passes, R2 (gravitational null case) is the most discriminating test of Postulate G, because it probes the sharpest prediction — zero selection rate for gravitationally indistinguishable records.

R5 (order of operations) tests the selection postulate directly. R3 (rate bound) provides quantitative constraint. R1 and R4 test secondary predictions.

In short: R0 first; if R0 passes, R2 and R5 next; then R3; then R1 and R4.

R1 — Pointer Basis vs. Position Basis

Setup: Systems where the environment-selected pointer algebra $\mathcal{O}$ is not position. Concrete examples: superconducting qubits (e.g., flux-tunable transmons), cavity QED (e.g., circuit QED, Schuster et al. 2007), collective spin ensembles.

Prediction: Selection targets $\mathcal{O}$, not position. Falsifier F1: If definiteness consistently appears in position despite $\mathcal{O} \neq$ position, the selection postulate fails.

R2 — Gravitational Null Case ($\Delta E\_G = 0$)

Setup: Decohered records that differ only by internal degrees of freedom with identical mass distributions. Candidates: nuclear spin states, photon polarization states, hyperfine states with identical spatial profiles.

Prediction: $\Delta E\_G = 0 \Rightarrow \lambda_{ij} = 0$. No selection dynamics beyond standard decoherence. Falsifier G2: Observation of objective selection between gravitationally indistinguishable records.

This does not forbid decoherence between such states via other interactions (e.g., EM); it forbids only objective selection on gravitational timescales.

R3 — Upper-Bound Rate Test (Speed Limit)

Setup: Mesoscopic/macroscopic superpositions with controlled mass distributions (levitated nanospheres, optomechanical resonators). Compute: $\Delta E\_G \Rightarrow \tau\_min = \hbar/\Delta E\_G$. Prediction: $\tau_{ij} \geq \tau\_min$. Falsifier G1: Selection faster than $\hbar/\Delta E\_G$.

Concrete estimates: For a spherical nanoparticle of radius R = 100 nm composed of a high-density material (tungsten, $\rho \approx 19$ g/cm^3), with two record sectors separated by $\Delta x \sim R$, the gravitational self-energy scales as $\Delta E\_G \propto \rho^2 R^5$.

This yields $\tau\_min \sim 1$–10 seconds, which is within reach of current cryogenic optical/magnetic levitation platforms and proposed microgravity free-fall experiments.

For comparison, silica nanoparticles ($\rho \approx 2$ g/cm^3) at the same radius give $\tau\_min \sim 10^2$–10^3 seconds, pushing toward the boundary of current coherence times. High-density materials are therefore strongly preferred for near-term tests.

R4 — Born Boundary Condition

Repeated preparations of identical decohered mixtures $\{p_i\}$. Prediction: The final ensemble of realized branches converges to $\{p_i\}$. Intermediate-time biases allowed; only asymptotic ensemble constrained. Falsifier F2: Systematic deviation from Born weights.

R5 — Order-of-Operations Test

Continuously tune environmental coupling to control degree of decoherence. Prediction: Selection is negligible until decoherence has rendered pointer sectors operationally distinct (D13). Falsifier F3: Detection of selection signatures prior to operational irreversibility.

This kills models where collapse is invoked to cause decoherence.

A5.3 — Operational Signature of Selection

Selection corresponds to nonlinear or stochastic dynamics at the single-trajectory level after decoherence is complete, producing effects not reproducible by any linear CPTP map acting on $\mathcal{H}_s$.

Detectable signatures include: single-trajectory anomalies (jump or diffusion statistics inconsistent with any linear Lindbladian after dephasing), irreversible loss of interference-revival capacity even under idealized system-only control, and telegraph-like stabilization (once a branch is realized, subsequent measurements behave as if prepared in the conditional state within operational tolerance).

A5.4 — What Counts as Confirmation vs. Survival

Passing a test does not confirm the argument. It only allows it to survive. Confirmation would require joint success across multiple regimes. Even then, what is established is structure, not interpretation.

Failure, by contrast, is immediate and final.

A5.5 — Timeline to Falsification

The following estimates reflect the state of experimental capability as of 2025 and are intended as orientation, not prediction.

R0 (Operational Invariance): Near-term (0–2 years). Circuit QED platforms already produce systems with multiple co-admissible pointer bases. Verification of AS invariance across such bases requires only measurement and classical post-processing of existing data.

R1 (Pointer Basis vs. Position): Near-term (0–3 years). Superconducting qubit and cavity QED experiments routinely prepare states where the pointer basis is energy or charge, not position.

Checking whether definiteness tracks the pointer basis requires monitoring which observable resolves first under controlled decoherence.

R5 (Order of Operations): Near-term to medium-term (1–5 years). Requires continuously tunable environmental coupling with single-trajectory readout. Superconducting qubits with adjustable coupling to engineered reservoirs are the most promising platform.

The key observable is whether selection-like signatures (telegraph stabilization, non-Lindbladian statistics) appear only after decoherence is complete.

R4 (Born Boundary Condition): Medium-term (2–5 years). Requires large ensembles of identically prepared, fully decohered systems with high-fidelity single-shot readout. Trapped-ion and superconducting qubit arrays are approaching the required scale and fidelity.

R2 (Gravitational Null Case): Medium-term to long-term (3–10 years). Requires decohered records that differ only by gravitationally indistinguishable internal degrees of freedom (e.g., nuclear spin states with identical spatial profiles).

The challenge is isolating such systems from all non-gravitational decoherence sources long enough to confirm the absence of selection.

R3 (Upper-Bound Rate Test): Long-term (5–15 years). Requires maintaining spatial superpositions of high-density nanoparticles (~100 nm, tungsten or osmium) for seconds in a decoherence-free environment, then measuring whether selection occurs faster than $\hbar/\Delta E_G$.

Cryogenic levitation and proposed space-based platforms (e.g., MAQRO mission concept) are plausible but not yet operational at the required scale.

Ordering logic. Tests are listed in order of experimental accessibility, not theoretical importance. R0 and R1 can be performed with existing hardware. R5 and R4 require modest extensions. R2 and R3 require dedicated experimental programs.

A negative result on R0 terminates the entire framework before any other test is needed.

A5.6 — Stop Condition

Definitions are operational. You can measure every one of them. Theorems are scoped. Each tells you exactly where it applies and where it does not. Postulates are isolated. Kill one and the rest survive.

Bounds are testable. You can check them with existing equipment. Falsifiers are explicit. You know exactly what would kill each claim.

Nothing further can be settled by argument. The mathematics has spoken. The experiments are specified. The kill switches are published. What remains is nature's answer. The argument has told you everything it can tell you.

It has defined the measuring tool. It has proven the monotonicity. It has established the no-return surface. It has characterised selection. It has formalised agency. It has derived the consequences of coupling.

Every claim is stated. Every kill switch is published. Every experiment is specified.

What remains is nature's answer. The argument cannot give you that. Only measurement can.

Only contact with reality can determine whether the definitions are operationally invariant, whether the gravitational bound holds, whether selection respects the pointer basis, whether the Born statistics emerge.

The argument waits. It has placed itself at the mercy of experiment. That is the only place an honest argument belongs.

A6 — Optional Module: The Turn

Status: Optional module. Not load-bearing. Included for conceptual completeness; failure leaves the entire laboratory program intact.

A6.1 — Capacity Saturation

Let K be the viability kernel of maintainable structure (A3.3). Capacity saturation obtains when the accessible state space for creating new durable record sectors has measure zero under admissible controls.

Operationally: further decoherence may occur, but no new independent records can be written.

Capacity saturation corresponds to thermodynamic heat death in the limit where all free-energy gradients are exhausted. It is not identical to heat death in general: saturation may occur locally or structurally before global thermal equilibrium.

A6.2 — Restoration Without Reversal

Any admissible Turn must satisfy: (1) No reversal: previously realized selections are not undone. (2) No selection bypass: A4.2 remains valid locally. (3) Capacity restoration: the effective record algebra regains room for new, independent branches.

This is an interface specification, not a dynamical law.

A6.3 — Conformal Rescaling

At extreme dilution, dynamics become insensitive to absolute scale. A conformal identification can map a capacity-saturated configuration to an initial configuration with renewed branching capacity without reversing causal order. Existence proof, not assertion of actuality.

“Records are compressed, not erased” means: distinguishable record sectors at late times map to a lower-resolution effective record algebra under the conformal identification, preserving orthogonality relations and causal precedence while reducing accessible distinguishability.

A6.4 — Relation to Heat Death

Black holes are capacity sinks, not reset buttons. They localize saturation and demonstrate no-return boundaries (A3.3). Any global Turn, if it exists, must respect the same non-reversal constraint.

A6.5 — Why This Module Is Optional

The core program answers how irreversibility arises, how definiteness arises, and how fast definiteness may arise. Module T addresses only whether global capacity can ever be restored.

Failure of Module T leaves the entire laboratory program intact.

Appendices

Appendix E — Glossary of Terms and Notation

Actualization State (AS). An operational scalar $\in [0, 1]$ measuring the degree of inter-sector branching in the record algebra. AS = 0: all weight in one record sector.

AS = 1: maximal branching across all record sectors. Defined in D3 (A2.3).

Admissible operation. A CPTP map acting on the accessible system, optionally with fresh ancilla, but with no access to the original environment.

The set of admissible operations defines what an agent can do, and thereby what counts as operationally irreversible.

Capture basin, Cap(R). The set of states from which exit from the recoverable set R is inevitable under all admissible controls. Defined in D8 (A3.3).

Coherence-recoverable state. A system state from which coherence between record sectors can be restored by admissible operations within tolerance ε. Defined in D11 (A4.1).

Coarse-graining, physically realizable ($\mathcal{O}$). A finite set of mutually orthogonal projectors selected by the physics of system–environment coupling, not by observer choice. Defined in D1 (A2.1).

Dephasing map, $\Delta\mathcal{O}$. The map $\Delta\mathcal{O}(\rho) \equiv \Sigma_i \Pi_i \rho \Pi_i$ that removes quantum interference between record sectors while preserving classical probabilities. Defined in D2 (A2.2).

Effective entropy, S_eff. The Shannon entropy $H(\{p_i\})$ of the sector weights, with intra-sector entropy discarded. This is the entropy entering the AS definition. Defined in A2.3.

Falsifier (F0, F1, F2, F3, G1, G2, G3). An experimentally observable condition whose occurrence would invalidate a specific component of the argument.

F0 is global (kills AS itself); F1–F3 target the selection postulate; G1–G3 target the gravitational limiter. Listed in A5.2.

Gravitational self-energy distinguishability, ΔE_G. The Newtonian self-energy of the difference mass distribution between two record sectors. Defined in D15 (A4.3). Appendix C provides explicit forms.

No-return surface, Σ_h. The boundary of the viability kernel. States beyond this surface cannot return to the recoverable set under any admissible control. Defined in D9 (A3.3) and D13 (A4.1).

Operator horizon, x_h. The scalar specialization of the no-return surface: $x_h \equiv$ u_max/a, the maximum sustainable structure given maximum maintenance effort. Defined in T2 (A3.2).

Operational invariance. The requirement that AS values computed from different physically realizable coarse-grainings of the same system agree within experimental tolerance. Defined in D5 (A2.5). Violation triggers the global kill switch F0.

Operational irreversibility. A state is operationally irreversible with respect to a recoverable set R iff it lies outside the viability kernel of R.

Irreversibility is defined by loss of reachability under admissible control, not by entropy increase or violation of time-reversal symmetry. Defined in D10 (A3.3).

Record algebra. The algebra generated by the physically realizable coarse-graining $\mathcal{O}$ = $\{\Pi_i\}$. States in this algebra are diagonal in the record basis. The record algebra defines what the environment can physically distinguish.

Selection. The transition from a diagonal mixture over record sectors (post-decoherence) to a single realized branch (definiteness). Distinguished from decoherence in D14 (A4.2).

The mechanism is specified by Postulate P; the rate is bounded by Postulate G.

Viability kernel, Viab(R). The set of states from which the system can be kept inside the recoverable set R indefinitely using admissible control. Defined in D7 (A3.3).

ε (operational tolerance). A fixed positive parameter representing experimental resolution. All operational definitions (effective orthogonality, inaccessibility, recoverability) are quantified relative to ε. Defined in A2.4.

Appendix A — Equivalence of AS Representations

A.1 Purpose

This appendix establishes the precise conditions under which the primary $AS(\rho; \mathcal{O})$ coincides with the history-based representation $AS_h(D)$. No equivalence is assumed in the main text.

A.2 Objects and Restrictions

The histories $\{\alpha\}$ are taken to correspond one-to-one with the record sectors $\{\Pi_i\}$ at a single time slice. No multi-time or branching-tree histories are included. Under this restriction: $N = d\mathcal{O}$.

This restriction is explicit and intentional.

A.3–A.4 Decoherence Condition and Block Entropy

Assume complete decoherence: $D(\alpha, \beta) \approx 0$ for $\alpha \neq \beta$. Under this condition, the dephased state has block form $\Delta\mathcal{O}(\rho) = \Sigma_i\, p_i\, \sigma_i$.

The von Neumann entropy decomposes exactly as $S(\Delta\mathcal{O}(\rho)) = H(\{p_i\}) + \Sigma_i\, p_i\, S(\sigma_i)$. No assumption of maximal mixing within sectors is made.

A.5–A.6 Equivalence Statement

AS uses $S\_eff(\rho; \mathcal{O}) \equiv H(\{p_i\})$ with normalization $AS(\rho; \mathcal{O}) = H(\{p_i\}) / \log d\mathcal{O}$.

A.7 Non-Equivalence Regimes

The equivalence fails when: decoherence is incomplete, histories span multiple time slices, or one attempts to include intra-sector entropy as “actualization.” In these regimes AS(ρ; $\mathcal{O}$) remains well-defined; $AS_h(D)$ ceases to be a faithful representation.

AS(ρ; $\mathcal{O}$) is preferred in all ambiguous cases because it requires only the reduced state and record algebra, not a full history space.

Appendix B — Viability Theory Background

Dynamics: $dx/dt = f(x, u)$, $u \in U$. Viability kernel: $\mathrm{Viab}(K) = \{x_0 \in K \mid \exists\, u(t): x(t) \in K\ \forall t \geq 0\}$.

Capture basin: $\mathrm{Cap}(K^c) = \{x_0 \mid \forall\, u(t), \exists\, t: x(t) \notin K\}$.

No-return surface: $\Sigma\_NR = \partial\, \mathrm{Viab}(K)$. The partition holds up to boundary sets of measure zero. For further background, see Aubin (1991), Viability Theory.

Appendix C — Gravitational Self-Energy

C.1 Definition For two record sectors i, j with mass densities $\mu_i(x)$, $\mu_j(x)$:

C.2 Weighted Multi-Sector Form

C.3 Special Cases

Identical mass distributions: $\Delta E_G = 0$. Rigid sphere (radius R, displacement $\Delta x \ll R$): $\Delta E_G \sim Gm^2/R \cdot (\Delta x/R)^2$ (order estimate).

C.4 Positivity By construction, $\Delta E_G \geq 0$.

Appendix D — Experimental Feasibility Estimates

D.1 High-Density Nanoparticle Regime

Consider a spherical nanoparticle of radius $R = 100$ nm composed of a high-density material (tungsten or osmium, $\rho \approx 19–22$ g/cm^3). For comparison, silica has $\rho \approx 2$ g/cm^3.

With two record sectors separated by $\Delta x \sim R$, and using $\Delta E_G \sim Gm^2/\Delta x$ with $m \propto \rho R^3$, the self-energy scales as $\Delta E_G \propto \rho^2 R^5$.

Tenfold density increase yields ~100× increase in ΔE_G. Resulting timescale: $\tau_{min} \sim \hbar/\Delta E_G$ yields $\tau_{min} \sim 1–10$ s for $R \sim 100$ nm high-density particles.

D.2 Experimental Platforms

Timescales in the second-to-tens-of-seconds range are within reach of: cryogenic optical or magnetic levitation of heavy nanoparticles (e.g., Delić et al., 2020; Tebbenjohanns et al., 2021), hybrid optomechanical traps with active feedback cooling (e.g., Aspelmeyer group, Vienna), and space-based or microgravity proposals (long-coherence free-fall platforms, e.g., MAQRO mission concept).

D.3 Interpretation

The purpose of this estimate is to demonstrate that the relevant regime is experimentally accessible. Observation of selection faster than the bound falsifies the gravity-limiter hypothesis.

Absence of selection constrains gravity’s relevance without undermining the core AS framework.

Appendix F — Worked Example: AS Calculation for a Dephasing Qubit

This appendix provides a complete, explicit AS calculation for the simplest nontrivial case, intended as a pedagogical anchor.

F.1 Setup

System: a single qubit S with Hilbert space $\mathcal{H}_s = \mathbb{C}^2$, coupled to an environment E. Pointer basis (selected by coupling): $\mathcal{O} = \{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. Record algebra dimension: $d\mathcal{O} = 2$.

F.2 Initial State

$|\psi(0)\rangle = (|0\rangle + |1\rangle)/\sqrt{2} \otimes |E_0\rangle$. The system is in a coherent superposition. The reduced state is $\rho_s(0) = |+\rangle\langle+|$. After dephasing: $\Delta\mathcal{O}(\rho_s(0)) = \mathrm{diag}(½, ½)$. Sector weights: $p_0 = p_1 = ½$.

F.3 Before Decoherence

Off-diagonal elements are present.

However, $S\_eff = H(\{p_i\}) = H(½, ½) = \log 2$. Normalization: $S\_max = \log 2$. Therefore AS = log 2 / log 2 = 1. Even before decoherence is complete, the sector weights already distribute maximally.

AS measures the branching structure of the dephased state, not whether decoherence has physically occurred.

F.4 After Decoherence

Environment records which-path: $|\psi\rangle \rightarrow (|0\rangle|E_0\rangle + |1\rangle|E_1\rangle)/\sqrt{2}$ with $\langle E_0|E_1\rangle \approx 0$. Reduced state: $\rho_s = \mathrm{diag}(½, ½)$. Off-diagonals have been physically suppressed.

AS = 1. Same numerical value, but now the system has crossed the no-return surface: coherence is not recoverable. Operational irreversibility has been established.

F.5 After Selection

Selection resolves the mixture into sector $|0\rangle$ (say). Now $p_0 = 1$, $p_1 = 0$. $H = 0$. $AS = 0$. One story remains. The system has transitioned from maximal branching to definiteness.

F.6 Summary

AS tracks branching richness, not definiteness. It rises during decoherence (branching phase) and falls during selection (definiteness phase).

The worked example illustrates that AS ≈ 0 and AS ≈ 1 are both physically meaningful endpoints of different processes, not a hierarchy of “more actualized” vs. “less actualized.”

Paper A — Canonical Reference Locked · Execution-Complete

Paper B

Selection as Irreversible Exclusion

Depends on Paper A

Paper A measured the branching. It proved the branching can only grow under the right conditions. It established the point of no return.

But it left one question unanswered — the question that haunts every interpretation of quantum mechanics.

If definiteness occurs in individual runs — and it does, every experiment ever conducted says so — what must selection be?

Not what might it be. What must it be. What structural requirements does any selection mechanism have to satisfy? What must it cost? How fast can it act?

And is there a universal bound on that speed?

No collapse mechanism is proposed. No interpretation is invoked. No result of Paper A is rederived. All hypotheses are independently falsifiable. Failure of any does not invalidate Paper A.

B0.1 — Dependency Statement

This work is a strict continuation of Paper A and assumes as established: the operational definition and validity of Actualization State (AS), operational irreversibility as loss of reachability under admissible control, the existence of no-return surfaces induced by bounded capacity, and the separation between branching (AS increase) and definiteness.

For reference, the record-sector algebra $\mathcal{R}$ is the algebra generated by the projectors $\{\Pi_i\}$ of a physically realizable coarse-graining $\mathcal{O}$ (Paper A, Definition D1), representing the set of operationally accessible record observables.

All norms and maps in this paper that reference $\mathcal{R}$ act on this algebra.

No construct from Paper A is redefined or rederived here.

B0.2 — Purpose

Paper A establishes irreversibility without definiteness: after decoherence and loss of recoverability, multiple mutually exclusive record sectors can persist simultaneously in the reduced description.

This paper addresses the remaining physical question:

If definiteness occurs, what must selection be, given the constraints already established?

A second question follows necessarily:

What physical resources must be expended to enforce such selection?

This work does not argue that selection must exist. It characterizes the structure and constraints of selection if it exists at all.

B0.3 — Hard Non-Claims

The paper does not: • redefine Actualization State, • propose a collapse mechanism, • derive or assume the Born rule, • invoke observers, consciousness, or epistemic update, • introduce agency, decision-making, or control,

• claim gravity causes selection.

Failure of this paper does not invalidate Paper A. B1 — The Definiteness Problem (Reframed)

B1.1 — What Remains After Paper A

After the results of Paper A, the following are established:

(1) Interference is suppressed between record-distinguishable alternatives (Paper A, T1).

(2) Recoverability is lost once record information is encoded in inaccessible degrees of freedom (Paper A, D13).

(3) Actualization State increases during the branching phase, quantifying record-structured multiplicity (Paper A, T1).

Yet none of these imply that, in an individual experimental trial, only one record persists.

A diagonal reduced state of the form where $\{\Pi_i\}$ are the record-sector projectors defined in Paper A (Definition D1) and $\{p_i\}$ are the corresponding diagonal coefficients of the reduced state inherited from the decoherence process (no probabilistic interpretation is assumed here), is fully consistent with all results of Paper A.

B1.2 — Why Decoherence Is Not Definiteness

Decoherence explains why interference terms become inaccessible. It does not explain why alternatives are excluded.

Operationally:

Decoherence answers: why alternatives cannot interfere.

Definiteness asks: why alternatives are no longer reachable.

These are distinct constraints. Paper A resolves the first and intentionally stops there.

B1.3 — Individual Realizations (Operational Definition)

An **individual realization** (or individual run) is defined as:

a single experimental trial producing a definite, time-ordered record stream in the environment, which subsequently constrains all future accessible system behavior.

This definition is purely operational and refers only to record structure.

B1.4 — Selection as Irreversible Exclusion

If definiteness exists, it must correspond to a physical exclusion process acting after irreversibility is established, because all subsequent experimental outcomes in the record stream depend causally on which sector obtains (B1.3).

Selection is defined here as:

The transition of the system state into a restricted reachable region of state space (under admissible control) in which only one record sector remains reachable. Equivalently, selection is the irreversible removal of alternative record sectors from operational accessibility in an individual realization.

Once selection has occurred, no admissible system-local operation can restore reachability of excluded sectors.

B1.5 — Consequence for Actualization State

Selection has a precise consequence for AS:

Decoherence increases AS by creating record-structured branching (Paper A, T1).

Selection reduces the **accessible AS** of an individual realization by restricting reachability to a single record sector.

This does not imply erasure of environmental records. It reflects the collapse of future operational accessibility, not the destruction of past structure.

B1.6 — Selection Cost (Foreshadowing)

Exclusion is not free.

Any process that removes reachability of alternatives must expend physical resources to enforce that restriction. This is generically as the **cost of selection**: the minimal physical resource expenditure required to enforce irreversible exclusion.

This cost need not be thermal energy; it may appear as time, interaction strength, or consumption of distinguishability capacity. Its precise form depends on the limiting interaction and is quantified below.

B2 — The Nonlinearity Requirement and Selection Cost

B2.1 — Linearity Constraint

Deterministic linear completely positive trace-preserving (CPTP) dynamics acting on the reduced system state preserve convex structure. Consequently, linear ensemble evolution cannot, by itself, enforce single-sector definiteness in individual realizations.

Formally, for any deterministic linear CPTP map $\mathcal{E}$:

Linearity preserves convex mixtures. No such map can select a single component from a diagonal mixture in individual runs. This is a structural consequence of linearity and does not depend on interpretation.

B2.2 — Ensemble Linearity vs. Trajectory Resolution

The implication of the linearity constraint is precise:

Ensemble-level evolution may remain linear and CPTP.

Selection, if it occurs, must act at the trajectory level, resolving individual realizations via stochastic or effectively nonlinear dynamics.

There is no contradiction with quantum linearity at the ensemble level.

B2.3 — Quantifying Selection Deviation

Selection is a trajectory-level phenomenon. Its signature is not a deviation of the ensemble state — ensemble consistency is required (B3.5) and the ensemble state is preserved by construction.

The signature of selection is that individual trajectories resolve to outcomes that no deterministic CPTP map could produce from the same initial state.

Let $\mathcal{E}_{\mathrm{ens}}$ denote the ensemble-level CPTP evolution acting on the reduced state.

Let Φ denote a stochastic selection process that, for an initial state ρ, produces a random trajectory with realization-dependent final state $\rho^{W}(\rho)$ indexed by trajectory ω. Ensemble consistency (B3.5) requires $\mathbb{E}[\rho^{W}] = \mathcal{E}_{\mathrm{ens}}(\rho)$.

The quantity that distinguishes selection from deterministic CPTP evolution is not the ensemble mean but the trajectory spread. Define the selection deviation: where $\|\cdot\|_{\mathscr{R}}$ is an operational norm restricted to the record-sector algebra $\mathcal{R}$.

This is the expected squared deviation of individual trajectory outcomes from the ensemble mean, measured in the record algebra. For any deterministic CPTP map, all trajectories produce the same output, so $\delta_{\mathrm{sel}} = 0$.

For selection, trajectories resolve to different record sectors, so $\delta_{sel} > 0$.

All subsequent results hold for any choice of norm satisfying: (i) contractivity under admissible system-local CPTP maps, and (ii) sensitivity to distinguishability between record sectors.

Compatibility with ensemble consistency. δ_{sel} measures the spread of trajectory outcomes, not the deviation of their mean. Ensemble consistency (B3.5) constrains the first moment: $\mathbb{E}[\rho^W] = \mathcal{E}_{ens}(\rho)$. The selection deviation constrains the second moment.

These are independent: a process can have zero mean deviation and nonzero trajectory spread. Selection is precisely such a process.

Verification (qubit toy model). For the selection SDE $dp = \sqrt{\gamma} \cdot p(1-p)\, dW$ (Paper A, Section A4.2), trajectories resolve to $p = 0$ or $p = 1$ with probabilities $1-p_0$ and p_0 respectively.

The ensemble mean is preserved: $\mathbb{E}[\rho^W] = \mathrm{diag}(p_0, 1-p_0)$. The trajectory variance is $\delta_{sel} = p_0(1-p_0) > 0$ for any nontrivial initial mixture, confirming that selection produces nonzero δ_{sel} while satisfying ensemble consistency.

Selection requires $\delta_{sel} > 0$. If $\delta_{sel} = 0$ for all accessible states, every trajectory produces the same output as the ensemble map, and no resolution to individual sectors has occurred.

B2.4 — Definition: Selection Cost

The cost of selection is defined as the minimal physical resource expenditure required to produce nonzero trajectory variance ($\delta_{sel} > 0$) sufficient to resolve individual realizations into single record sectors.

This cost may be expressed as: a time scale (rate of exclusion), a coupling strength to enforcing interactions,

or a resource budget required to maintain exclusion.

Later sections identify universal constraints on this cost.

B2.5 — Falsifier B2: Pre-Irreversibility Selection

If exclusion signatures appear before the system has crossed the no-return surface (D13), the entire selection model is dead. Selection must wait for irreversibility. If it does not wait, the model is wrong.

Selection cannot precede irreversibility. B3 — Structural Requirements on Selection Dynamics

Any admissible selection dynamics must satisfy a minimal set of structural requirements implied jointly by Paper A and Sections B1–B2. These requirements are necessary, not sufficient.

Failure of any requirement falsifies the selection hypothesis without affecting the irreversibility results of Paper A.

B3.1 — Post-Irreversibility Activation

Selection may act only after operational irreversibility is established.

Let $K_\varepsilon(\mathcal{O})$ denote the ε-recoverable set defined in Paper A (Definition D12).

For all states $\rho \in K_\varepsilon(\mathcal{O})$, admissible dynamics must satisfy: That is, no selection deviation is permitted while recovery remains reachable under admissible control. Selection dynamics may become active only once $\rho \notin K_\varepsilon(\mathcal{O})$.

B3.2 — Record-Algebra Locality

Selection must act only on degrees of freedom that distinguish record sectors, and only during active selection.

Let $\Delta\mathcal{O}$ denote the dephasing map onto the record algebra $\mathcal{O}$. During active selection (i.e., after $\rho \notin K_\varepsilon(\mathcal{O})$):

This condition is consistent with B3.1: by the time selection activates, off-diagonal terms in the record basis are already operationally inaccessible.

Selection may not generate or reintroduce interference, nor act on unmonitored intra-sector degrees of freedom.

B3.3 — Absorbing Record Sectors

Selection is an absorbing process. Once a record sector Π_i is realized, sector membership must remain fixed under subsequent selection dynamics. Formally:

up to operational tolerance set by experimental resolution. This condition enforces irreversible confinement to the realized sector while allowing arbitrary intra-sector evolution.

B3.4 — Contractivity of Multiplicity

Selection resolves multiplicity; it must not amplify it.

Let $\{p_i(t)\}$ denote the diagonal coefficients of the reduced state in the record-sector basis, treated here as record weights. Shannon entropy

is used strictly as a multiplicity measure, not as a thermodynamic or epistemic entropy.

Any admissible selection dynamics must be such that, along individual trajectories generated by those dynamics, $H(\{p_i(t)\})$ is a supermartingale:

with strict decrease during active selection. The expectation is taken with respect to the trajectory measure induced by the selection dynamics.

This is a requirement on the class of admissible dynamics, not a property derived from a specific generator: any candidate selection process whose trajectories amplify multiplicity is excluded.

No assumption about the Born rule is made here; H is used purely as a multiplicity measure.

Note that contractivity of H along trajectories is a consequence of the trajectory-level character established in B2.2: the stochastic process Φ must resolve the mixture, which requires H to decrease along individual realizations.

B3.5 — Ensemble Consistency

While individual trajectories resolve to single sectors, the ensemble description must remain consistent with linear evolution. Averaging over all trajectory realizations must reproduce the ensemble map:

ensuring compatibility with standard ensemble-level quantum predictions. This requirement constrains the first moment of the trajectory distribution.

It does not constrain the second moment: trajectory spread ($\delta_{sel} > 0$, Section B2.3) is fully compatible with ensemble consistency. Selection is characterized by the combination of preserved ensemble mean and nonzero trajectory variance.

B3.6 — Boundary Condition on Outcomes (BC1)

The paper does not derive outcome statistics. However, admissible selection dynamics must produce a well-defined distribution over realized record sectors.

The analysis restricts to the class of selection dynamics that are Born-consistent: under the trajectory measure induced by the dynamics, the marginal distribution over realized sectors converges to the diagonal weights $\{p_i\}$ inherited from decoherence.

This is a defining constraint of the model class studied here, not a derived result. The existence and properties of Born-inconsistent selection dynamics are a separate question outside the present scope.

Born-consistency is experimentally testable: repeated preparations of identically decohered systems must yield realized-sector frequencies converging to $\{p_i\}$. Persistent deviation falsifies the Born-consistent class, not selection itself.

B3.7 — Summary of Structural Requirements

Selection dynamics, if they exist, must be:

Post-irreversibility — inactive while recovery remains reachable.

Record-local — acting only on the record algebra during active selection.

Absorbing — once a sector is realized, sector membership remains fixed.

Contractive — monotonically reducing multiplicity along trajectories.

Ensemble-consistent — preserving linear ensemble evolution.

Any candidate process violating these conditions is not a physically admissible form of selection under the argument established by Paper A.

Worked Example: Structural Requirements Applied to Qubit Selection

Paper A (Section A4.2) defines a toy model of stochastic selection on a qubit with pointer basis $\mathcal{O} = \{|0\rangle\langle 0|, |1\rangle\langle 1|\}$ and selection dynamics $dp = \sqrt{\gamma} \cdot p(1-p)\, dW$.

The verification shows that this toy model satisfies all five structural requirements of B3.

B3.1 (Post-irreversibility activation): The selection rate γ is set to zero while the system remains within the recoverable set $K\varepsilon(\mathcal{O})$. The SDE activates only after decoherence has rendered pointer sectors operationally distinct.

Prior to activation, $\delta_{\text{sel}} = 0$ identically.

B3.2 (Record-algebra locality): The SDE acts entirely on the diagonal sector weight p. No off-diagonal coherence terms are generated or accessed. The process is record-algebra local by construction: $\Phi(\rho) = \Phi(\Delta\mathcal{O}(\rho))$.

B3.3 (Absorbing sectors): At $p = 0$ and $p = 1$, the diffusion coefficient $\sigma(p) = \sqrt{\gamma} \cdot p(1-p)$ vanishes identically. Both sector-pure states are absorbing fixed points. Once realized, sector membership is permanent.

B3.4 (Contractivity): By the full Itô calculation (Paper A, Section A4.2, requirement S3): $dH = -(\gamma/2)\, p(1-p)\, dt + \sqrt{\gamma} \cdot p(1-p) \cdot \log[(1-p)/p]\, dW$.

The drift term $-(\gamma/2)\, p(1-p)$ is strictly negative for $p \in (0, 1)$. H is a supermartingale with strict decrease during active selection, as required.

B3.5 (Ensemble consistency): The SDE $dp = \sqrt{\gamma} \cdot p(1-p)\, dW$ is a martingale: $\mathbb{E}[p(t)] = p(0)$ for all t. Averaging over trajectories reproduces the ensemble state $\rho_{\text{ens}} = \text{diag}(p(0), 1-p(0))$ at all times.

The ensemble map is linear and CPTP.

This confirms that the structural requirements B3.1–B3.5 are jointly satisfiable. The toy model is not the unique solution; it is a proof of existence.

Any candidate selection dynamics must pass all five requirements to be admissible.
B4 — Universal Rate Constraints on Selection

Selection, if it exists, cannot occur arbitrarily fast.

This section establishes necessary upper bounds on the rate at which admissible selection dynamics may act, without introducing new physics and without exceeding the scope of Paper A.

B4.1 — Selection Rate as an Operational Quantity

For two record sectors i and j, define the **inter-sector selection time** τ_{ij} as the minimal duration required, in an individual realization, for the system's accessible behavior to become operationally indistinguishable from confinement to sector i rather than j, within experimental tolerance.

The corresponding selection rate is:

This rate is operationally measurable: it characterizes how rapidly exclusion is enforced between competing record sectors.

B4.2 — Requirements on a Universal Rate Limiter

Any candidate universal rate limiter on selection must satisfy the following constraints:

Universality. The bound must apply across all macroscopic records, independent of composition or charge.

Context Independence. The bound must not depend on observer intervention, measurement choice, or apparatus-specific tuning.

Discriminatory Relevance. The bound must couple directly to the physical features that distinguish record sectors.

These constraints do not uniquely determine a limiter, but they strongly restrict admissible candidates.

B4.3 — Gravity as a Candidate Universal Limiter (Hypothesis)

Among known interactions, gravity satisfies all three requirements above: it is universal, unscreenable, and directly sensitive to mass–energy configuration, which distinguishes macroscopic records.

The hypothesis: that gravity provides a universal upper bound on selection rates. This is an empirical claim about known interactions, not a proof of uniqueness, and it does not assert that gravity causes selection.

B4.4 — Gravitational Distinguishability of Record Sectors

The gravitational self-energy distinguishability ΔE_G between record sectors i and j is defined in Paper A (Definition D15, Appendix C).

It measures the Newtonian self-energy of the difference mass distribution between two record sectors and is zero when the sectors are gravitationally indistinguishable.

The definition, explicit integral form, and positivity proof are given in Paper A and are not repeated here.

B4.5 — Rate Inequality

The gravity-limited selection bound (Paper A, Postulate G) states:

Paper A introduces this bound as a physical constraint.

Here we elevate it to a candidate universal rate limiter by demonstrating that gravity satisfies the universality, context-independence, and discriminatory relevance requirements of B4.2 — requirements that were not articulated in Paper A.

The experimental consequences of this elevation are developed in B5.

The bound is limiting, not exact. Selection may be slower; it may not be faster without invoking a coupling stronger than gravity to mass–energy distinguishability.

B4.6 — Null Case (Conditional)

Under the gravity-limited hypothesis, if two record sectors are gravitationally indistinguishable: $\Delta E^{G-ij} = 0$

then the gravity-constrained contribution to the selection rate vanishes:

If no alternative limiter satisfying the requirements of B4.2 applies, such superpositions persist indefinitely. If a non-gravitational mechanism consistent with B4.2 exists, it would supply an independent rate bound not covered by the present hypothesis.

This constitutes a testable conditional prediction of the gravity-limited framework.

B4.7 — Consistency with Prior Results

The gravity-limited bound is consistent with all earlier sections:

it applies only after operational irreversibility (B3.1),

it constrains rates, not outcome statistics (B3.6),

it preserves ensemble linearity (B3.5),

and it does not explain decoherence or branching (Paper A).

Gravity here functions solely as a rate limiter, not a causal mechanism.

B4.8 — Falsifiers (Rate-Level)

The gravity-limited hypothesis is falsified if any of the following are observed:

FG1: Selection occurs with $\lambda ij > \Delta EG/\hbar$.

FG2: Selection occurs between records with $\Delta EG = 0$ in the absence of any alternative limiter satisfying B4.2.

FG3: Selection rates scale universally with non-gravitational parameters across macroscopic records.

Failure here invalidates the limiter hypothesis only; it does not invalidate selection as defined in this paper, nor irreversibility as defined in Paper A.

B4.9 — Closing

If selection occurs, it is constrained by physical limits on how rapidly alternatives can be distinguished. Definiteness cannot emerge arbitrarily fast; it can emerge no faster than a universal interaction can discriminate between competing records.

B5 — Experimental Regimes and Discriminating Tests

This section translates the structural and rate constraints of Sections B1–B4 into experimentally discriminable regimes.

The aim is not parameter fitting, but to specify what observations would count as confirmation, survival, or falsification of selection as defined in this paper.

B5.1 — Principle of Test Construction

Experiments testing selection must satisfy three criteria:

Post-Irreversibility Regime. Decoherence and loss of recoverability must already be established (Paper A). Tests conducted while the system remains within $K\varepsilon(\mathcal{O})$ are irrelevant to selection.

Trajectory Sensitivity. The experiment must probe single-run behavior or trajectory-level signatures, not ensemble averages alone.

Rate Sensitivity. The experiment must be capable of resolving timescales comparable to the predicted selection rate λ_{ij}^{-1}.

Only experiments satisfying all three can meaningfully constrain selection dynamics.

Test Map Summary

The following tests are ordered by what they target, not by experimental accessibility. Each test is independently meaningful.

BT1 — Order-of-Operations (B5.5). Target: selection itself. Falsifies: selection as defined in B1.4. Method: continuously tune decoherence and check whether selection signatures appear before the system exits $K\varepsilon(\mathcal{O})$.

Platform: superconducting qubits with tunable coupling to measurement cavity.

BT2 — Active Selection Signature (B5.2). Target: existence of selection. Falsifies: selection is present in the tested regime. Method: compare single-trajectory statistics against all linear Lindblad models fitted to the same decoherence data.

Observable: telegraph noise or diffusive wandering inconsistent with any CPTP unraveling. Platform: continuously monitored superconducting qubits or trapped ions with fluorescence readout.

BT3 — Null-Rate Regime (B5.3). Target: gravity-limited hypothesis. Falsifies: FG2. Method: prepare decohered superpositions with $\Delta E_G = 0$ and monitor for selection. Observable: persistent multiplicity (no single-sector stabilization) vs. rapid selection.

Platform: nitrogen-vacancy centers in diamond, nuclear spin states with identical mass distributions.

BT4 — Rate-Bound Regime (B5.4). Target: gravity-limited hypothesis. Falsifies: FG1. Method: create spatial superpositions of mesoscopic masses, measure selection timescale, compare against $\tau_{min} = \hbar/\Delta E_G$. Observable: selection faster or slower than bound.

Platform: levitated nanoparticles (tungsten, $R \approx 100$ nm, $\tau_{min} \sim 1$–10 s) in cryogenic vacuum.

BT5 — Born Boundary Condition (B3.6). Target: Born-consistency of selection. Falsifies: the Born-consistent model class. Method: large ensembles of identically prepared, fully decohered systems with single-shot readout.

Observable: realized-sector frequencies deviating from $\{p_i\}$ beyond statistical tolerance. Platform: trapped-ion arrays or superconducting qubit arrays.

B5.2 — Signature of Active Selection

Selection is operationally distinct from decoherence. An **active selection signature** is any trajectory-level behavior, occurring after operational irreversibility, that:

cannot be reproduced by any system-local linear CPTP evolution consistent with the independently characterized decoherence dynamics of the system, and

enforces persistent confinement to a single record sector under all admissible system-local controls.

Examples of admissible signatures include:

irreversible loss of interference revival capacity despite full system-only control,

stochastic stabilization of record-sector behavior inconsistent with linear Lindblad dynamics fitted to the same decoherence data,

telegraph-like trajectory behavior resolving into a single sector with no subsequent switching within accessible timescales.

Absence of such signatures implies absence of selection in the tested regime.

B5.3 — Null-Rate Regime (Gravitational Degeneracy)

Consider record sectors that are operationally decohered but gravitationally indistinguishable: $\Delta EG = 0$.

Under the gravity-limited hypothesis, the gravity-constrained contribution to the selection rate vanishes. The argument therefore predicts one of two outcomes:

Persistent multiplicity: no selection signatures appear within experimentally accessible timescales; or

Non-gravitational selection: selection occurs at a slower rate governed by an alternative limiter satisfying the requirements of B4.2.

Concrete example: a nitrogen-vacancy (NV) center in diamond prepared in a superposition of spin states $|m_s = +1\rangle$ and $|m_s = -1\rangle$.

These states have identical mass distributions ($\Delta E_G = 0$) but are operationally distinguishable via microwave spectroscopy. After environmental decoherence has suppressed spin coherence, the reduced state is $\rho = \frac{1}{2}|+1\rangle\langle+1| + \frac{1}{2}|-1\rangle\langle-1|$.

Under the gravity-limited hypothesis, no gravitational contribution to selection exists. If single-trajectory monitoring reveals rapid stabilization to one spin state inconsistent with any Lindblad model of the decoherence dynamics, the gravity-limited hypothesis is falsified.

Observation of rapid selection in this regime falsifies the gravity-limited hypothesis.

B5.4 — Rate-Bound Regime (Macroscopic Distinguishability)

For record sectors with significant gravitational distinguishability $\Delta EG \gg \hbar/T$, where T is the duration over which the experiment maintains sensitivity to trajectory-level behavior, the gravity-limited hypothesis predicts an upper bound:

Experiments in this regime can test whether observed selection times: respect the bound (hypothesis survives),

approach the bound (gravity-limited selection likely active), or

violate the bound (hypothesis falsified).

Candidate systems and estimates. A tungsten nanoparticle ($R = 100$ nm, $\rho \approx 19$ g/cm^3) in spatial superposition with separation $\Delta x \sim R$ yields $\Delta E_G \sim Gm^2/R$, giving $\tau_{min} = \hbar/\Delta E_G \sim 1$–$10$ seconds.

This is within reach of cryogenic levitation experiments maintaining coherence for seconds (cf. Paper A, Appendix D).

For silica ($\rho \approx 2$ g/cm^3), $\tau_{min} \sim 10^2$–10^3 seconds, at the boundary of current coherence times. High-density materials are strongly preferred for near-term tests.

Observation of selection with $\tau < \tau_{min}$ falsifies the gravity-limited hypothesis (FG1). Absence of selection within accessible timescales is consistent with the hypothesis but does not confirm it.

B5.5 — Order-of-Operations Test

Selection must not precede irreversibility.

Experiments that continuously tune environmental coupling can test whether selection signatures appear only after the system exits the recoverable set $K\varepsilon(\mathcal{O})$.

Specifically, if trajectory-level confinement to a single record sector is observed while the system remains within $K\varepsilon(\mathcal{O})$ as defined in Paper A (Definition D12), then selection as defined in B1.4 is falsified.

This test targets selection itself, not merely the gravity-limited hypothesis.

B5.6 — Outcome Classification

Experimental outcomes partition as follows:

No selection observed: selection absent in the tested regime.

Selection observed, rate indeterminate: selection present; gravity-limited hypothesis neither confirmed nor falsified.

Selection observed within bound: selection present and consistent with the gravity-limited hypothesis.

Selection observed faster than bound: gravity-limited hypothesis falsified.

Selection observed in null-rate regime without alternative limiter: gravity-limited hypothesis falsified or incomplete.

No outcome retroactively rescues the hypothesis.

B5.7 — Scope Closure

This paper establishes: what selection must be if it exists, what it must cost, how fast it may occur,

and how it can be falsified.

It does not determine whether selection actually occurs in nature. That question is empirical. B6 — Conclusions and Program Status

This paper has treated selection as a physical exclusion process constrained by irreversibility, control limits, and rate bounds, without invoking interpretation, agency, or collapse mechanisms. The results can be summarized as follows.

B6.1 — What Has Been Established

If selection exists, it must satisfy all of the following:

1. Post-Irreversibility Constraint. Selection cannot act before operational irreversibility is established. Any exclusion prior to exit from $K\varepsilon(\mathcal{O})$ falsifies selection as defined here.

2. Trajectory-Level Character. Selection must act at the level of individual realizations while preserving linear ensemble evolution.

3. Record-Algebra Locality. Selection may act only on degrees of freedom that distinguish record sectors and may not reintroduce interference.

4. Absorbing Dynamics. Once a record sector is realized, sector membership is fixed under subsequent selection dynamics.

5. Contractivity of Multiplicity. Selection must monotonically reduce record-sector multiplicity along individual trajectories.

6. Cost and Rate Constraints. Selection requires physical resources and cannot occur arbitrarily fast.

7. Universal Rate Limiter (Hypothesis). Gravity provides a candidate universal upper bound on selection rates, expressible through gravitational distinguishability ΔEG, and is falsifiable by explicit rate tests.

Each condition is necessary. None are assumed to be sufficient.

B6.2 — What Has Not Been Assumed

This paper has not: • assumed that selection must occur,

• derived outcome statistics or the Born rule, • specified a concrete dynamical generator, • invoked observers, consciousness, or epistemic update, • claimed gravity causes selection,

• or extended irreversibility beyond what is established in Paper A.

Failure of any hypothesis in this paper leaves the foundations of Paper A intact.

B6.3 — Status of the Gravity-Limited Hypothesis

The gravity-limited hypothesis introduced in Section B4 is empirically motivated, dimensionally consistent, and experimentally falsifiable. It stands or falls entirely on observation. Its failure would constrain the space of admissible selection mechanisms, not rescue them.

B6.4 — Programmatic Closure

Together with Paper A, this work completes the physics-level characterization of selection:

Paper A establishes irreversibility without definiteness.

Paper B establishes definiteness as costly, rate-limited exclusion, if it exists.

No further progress on selection can be made by argument alone. The remaining uncertainty is empirical.

B6.5 — Forward Dependency

If selection is absent or constrained, the remaining question is not about definiteness but about structure: how does behavior unfold within a single realized record sector under irreversible constraint?

That question concerns control under irreversibility, not the emergence of definiteness. It is addressed in Paper C, where agency is treated as constrained dynamics downstream of the physics established here.

End of Paper B.

Paper C

Agency as Constrained Control

Depends on Papers A and B

You are an agent. You make choices. You maintain yourself against decay. You navigate a space of possibilities that narrows with every irreversible step. You have a budget that depletes.

You face drift that never stops. And somewhere ahead of you, invisible but real, is a boundary beyond which no choice you make can save you.

Everything you just read is geometry. Not philosophy. Not metaphor. Geometry — measurable, computable, falsifiable.

This paper strips the philosophy out of agency and replaces it with a number. The number measures the fraction of survivable states you can still reach from where you stand.

That number is more honest than any definition philosophy has ever produced, because it does not care about your intentions.

It cares about your position in the state space and the size of your control set. The rest is arithmetic.

Dependent on:

Paper A — Actualization State (AS): An Operational Measure of Record-Structured Irreversibility

Paper B — Selection as Irreversible Exclusion: Rates, Costs, and Constraints on Definiteness Abstract

This paper develops a control-theoretic account of agency under irreversible physics. Your agency, measured as a number.

Building on Papers A and B, agency is defined as a geometric property: the fraction of survivable states you can reach from where you currently stand, using whatever control you have, within the record sector you actually occupy.

No new physical assumptions are introduced.

The results: if you stop maintaining, agency decays. At the no-return boundary, agency hits zero. High-variance or misaligned strategies waste the budget faster than steady ones. None of this is surprising.

All of it is now proven. Control fatigue, noise, coupling, and exit are defined as consequences of constrained reachability rather than psychological or normative phenomena.

The paper establishes necessary conditions for persistence of controlled behavior under irreversibility and provides falsifiers for the control framework. What remains unresolved is empirical — and that is exactly as it should be.

The geometry is proven. Which systems instantiate it is nature's answer: which real systems instantiate these constraints and how closely.

C0 — Scope

C0.1 — Dependency Statement

This work depends explicitly and exclusively on the physical results established in **Paper A** and **Paper B**.

It assumes as given:

irreversibility as loss of reachability under admissible control (Paper A),

the existence of no-return surfaces induced by bounded capacity (Paper A),

selection, if it exists, as a costly, rate-limited exclusion process acting after irreversibility (Paper B).

Paper C requires only that selection produces confinement to a single record sector; it does not depend on the mechanism, rate, or statistics of selection.

No physical construct is redefined or rederived here.

Cross-reference note. The viability kernel Viab(R) used throughout this paper (Paper A, Definition D7) corresponds, in the quantum setting, to the recoverable set $K_\varepsilon(\mathcal{O})$ (Paper A, Definition D12).

Paper C operates entirely within a single realized record sector, so the relevant constraint set R is the set of states accessible to the system after selection, not the full quantum state space.

C0.2 — Purpose

Paper C addresses a question that is not physical in origin, but structural in consequence:

Given irreversible physics and costly definiteness, how can controlled behavior persist within a single realized record sector?

Agency is treated not as intention, belief, or choice, but as a **control property** — a number you can compute of a system evolving under irreversible constraints.

C0.3 — Hard Non-Claims

The paper does not: • introduce new physical laws, • modify or reinterpret quantum mechanics, • explain why selection occurs, • invoke psychology, motivation, ethics, or meaning,

• provide prescriptions or normative guidance.

Failure of Paper C does not invalidate Papers A or B.

C1 — Agency as a Geometric Control Quantity

C1.1 — Definition of Agency

Within a single realized record sector, define **agency** as:

The fraction of the viability kernel reachable from the current state under admissible control.

Let x(t) denote the system state confined to a realized record sector.

Let Viab(R) be the viability kernel defined in Paper A relative to admissible controls, and let Reach(x) be the set of states reachable from x under those controls.

Define where μ is the natural volume measure induced by the state-space metric and Viab(R) has finite positive measure.

The normalization ensures $\mathcal{M} \in [0, 1]$, with $\mathcal{M} = 1$ when the entire viability kernel is reachable and $\mathcal{M} = 0$ at the no-return surface where no viable future remains.

Monotonicity and regularity assumptions. The measure μ is monotone with respect to set inclusion: if $S_1 \subseteq S_2$ then $\mu(S_1) \leq \mu(S_2)$.

The analysis assumes Reach(x) varies continuously with x in the Hausdorff metric on compact subsets of the state space, ensuring $\mathcal{M}$ is continuous.

These are standard regularity conditions in viability theory (Aubin, 1991) and are not additional physical assumptions.

C1.2 — Control Authority

Let admissible controls $u(t) \in U$ be bounded by physical and energetic constraints. Control authority is determined by:

Bandwidth: the maximal rate at which control can counteract irreversible drift,

Reachability: the remaining volume of Viab(R) accessible from x(t),

Slack: the time-to-boundary from x(t) under zero control (defined formally in C8.1).

Limit condition. By definition of $\Sigma NR = \partial$ Viab(R) and continuity of $\mu(\mathrm{Reach}(\cdot) \cap \mathrm{Viab}(R))$ as a function of state (guaranteed by the Hausdorff regularity assumption above):

At the boundary, only a single future trajectory remains.

C2 — Drift as a Consequence of Irreversibility

C2.1 — Irreversible Drift

For open systems with nonzero irreversible drift, ordered states decay toward loss of structure in the absence of sustained control.

This follows directly from bounded control capacity and the operator-horizon result of Paper A (Theorem T2); it is not an independent axiom.

C2.2 — Baseline Dynamics

Absent control ($u = 0$), the system evolves as where $f(x)$ is the irreversible drift field pointing toward an attractor of structural loss (equilibrium, failure, or saturation).

The scalar decay model of Paper A ($dx/dt = -ax + u$) is a special case of this general form.

Proposition C2.1 (Agency decay under drift — the mathematical expression of what you already know: everything falls apart without maintenance). Let $x(t)$ evolve under $dx/dt = f(x) + u$ with $u(t) \in U$, and suppose the drift field f points inward toward an attractor $x^* \notin \mathrm{Viab}(R)$.

If $|f(x)| \geq a\|x - x^*\|$ for some $a > 0$ (linear lower bound on drift), and if $\mu(\mathrm{Reach}(x) \cap \mathrm{Viab}(R))$ is Lipschitz in x with constant L, then along any trajectory with $|u(t)| \leq u\_max$:

When the system is far from the attractor ($\|x - x^*\| > u\_max/a$), the right-hand side is strictly negative: agency decreases regardless of control.

This reproduces the operator horizon result of Paper A (Theorem T2) in the agency framework and quantifies the rate of agency loss beyond the horizon.

Proof sketch. $d\mathcal{M}/dt = (d/dt)[\mu(\mathrm{Reach}(x) \cap \mathrm{Viab}(R))]/\mu(\mathrm{Viab}(R))$. By Lipschitz continuity, $|\Delta\mu| \leq L|\Delta x|$. The state velocity is $|dx/dt| = |f(x) + u| \leq |f(x)| + |u|$.

The drift pushes toward x (reducing Reach), while control pushes away (expanding it). Net rate: $d\mathcal{M}/dt \leq L(|u| - |f(x)|)/\mu(\mathrm{Viab}(R)) \leq L(u\_max - a\|x - x\|)/\mu(\mathrm{Viab}(R))$. □ C3 — Necessary Conditions for Agency Preservation

C3.1 — Continuous Control Cost

For open systems with $f(x) \neq 0$ away from fixed points, maintaining distance from ΣNR requires continuous expenditure of control effort. Except at exact fixed points of f, no finite intervention permanently arrests drift.

Such fixed points, if they exist, may themselves lie outside Viab(R) or require sustained control to reach; their existence does not generally provide a cost-free maintenance strategy.

C3.2 — Variance-Conditioned Control Effectiveness

Proposition (conditional). For admissible control systems in which instantaneous control cost c(u) is convex in |u|, low-variance control trajectories preserve $\mathcal{M}$(x) more effectively than high-variance or impulsive strategies with the same mean control effort.

Proof sketch. For convex c, Jensen's inequality gives $\mathbb{Z}[c(u)] \geq c(\mathbb{Z}[u])$.

Variable control with fixed mean effort therefore incurs greater cumulative cost than constant control at the mean level, depleting the control budget B(t) (defined in C5.1) more rapidly and thereby reducing reachable viability volume. □

Corollary C3.1a (Maintenance condition). For the scalar system dx/dt = −ax + u with a > 0 and u ∈ [0, u_max], the agency $\mathcal{M}$(x) is maintained (d$\mathcal{M}$/dt = 0) if and only if u = ax, i.e., control exactly balances drift.

This requires x ≤ u_max/a = x_h (the operator horizon). For x > x_h, no admissible control can maintain $\mathcal{M}$, and d$\mathcal{M}$/dt < 0 strictly.

The maintenance condition is the agency-framework restatement of Paper A's Theorem T2: the horizon is the boundary between maintainable and inevitably decaying agency.

C4 — No-Return Geometry Within a Realized Sector

C4.1 — Horizon Geometry

The operator horizon from Paper A (Theorem T2, Definition D9) applies strictly within a realized record sector. Crossing this boundary removes states from Viab(R).

C4.2 — Ruin as Absorbing State

Ruin is defined as x ∉ Viab(R)

Once this occurs, recovery is impossible under admissible control. Ruin is a geometric property of state space, not a subjective condition.

Worked Example: No-Return Geometry in a 2D Linear System

Consider a two-dimensional system with state $x = (x_1, x_2) \in \mathbb{R}^2$, drift $f(x) = (-a_1x_1, -a_2x_2)$ with $a_1, a_2 > 0$, and control $u = (u_1, u_2) \in [0, u_1^{max}] \times [0, u_2^{max}]$.

The constraint set is $R = \{x : x_1 \geq 0, x_2 \geq 0\}$.

The viability kernel is the rectangle $\mathrm{Viab}(R) = [0, x_{1h}] \times [0, x_{2h}]$ where $x_{ih} = u_i^{max}/a_i$ is the per-axis operator horizon.

The no-return surface Σ_NR is the boundary of this rectangle: any state with $x_1 > x_{1h}$ or $x_2 > x_{2h}$ is in the capture basin and will be driven to the boundary regardless of control.

Agency computation. For a state $x = (x_1, x_2)$ inside Viab(R), the reachable set within Viab(R) is $\mathrm{Reach}(x) \cap \mathrm{Viab}(R) = [0, \min(x_1 + u_1^{max}/a_1, x_{1h})] \times [0, \min(x_2 + u_2^{max}/a_2, x_{2h})]$ (for steady-state reachability).

The normalized agency is $\mathcal{M}(x) = \mu(\mathrm{Reach}(x) \cap \mathrm{Viab}(R))/\mu(\mathrm{Viab}(R))$. At the origin, $\mathcal{M} = (u_1^{max}/a_1)(u_2^{max}/a_2)/(x_{1h} \cdot x_{2h}) = 1$ (full viability kernel reachable).

At the corner (x_{1h}, x_{2h}), Reach shrinks to the single point, and $\mathcal{M} \to 0$.

This example illustrates three features: (i) the no-return surface is axis-separable in the linear case; (ii) agency varies continuously from 1 to 0 across the viability kernel; (iii) position within Viab(R) determines how much future flexibility remains, independent of the system's history.

C5 — Control Budgets and Fatigue

C5.1 — Control Budget Define the control budget:

Experimental instantiation: E. coli chemotaxis in a microfluidic gradient chamber. The agency framework maps directly to a concrete biological system. State space: cell position and internal chemotactic signalling state within the gradient.

Viability kernel: the region of the chamber where nutrient concentration supports growth (positions with [nutrient] > threshold).

Drift: diffusion and fluid flow carry cells toward the nutrient-depleted zone ($f(x) \neq 0$). Control: chemotactic swimming ($u \in U$, bounded by flagellar motor torque and tumble rate).

Operator horizon: the position beyond which maximum chemotactic swimming cannot overcome the flow rate — x_h = u_max/a where a is the effective advection rate.

Budget: internal energy reserves (ATP, proton-motive force) that deplete with swimming effort. Fatigue: as reserves approach zero, flagellar motor stalls; control ceases. Noise: Brownian motion and tumble stochasticity consume control bandwidth without expanding reachable set.

Ruin: cell exits the viable nutrient zone; no admissible swimming restores it.

Slack: time-to-washout at current position and flow rate if swimming stops.

Falsifier (C10.1, condition 1): if a non-motile mutant (u_max = 0) maintains its position in the gradient without external intervention, agency as defined here is falsified — reach grew without control expenditure.

Falsifier (C10.1, condition 4): if a motile cell with finite ATP reserve persists indefinitely in a persistent nutrient gradient flow, the survival time bound (Theorem C5.1) is falsified.

Every construct in Paper C — drift, control, horizon, budget, fatigue, noise, slack, ruin — maps to a measurable variable in this system.

The experiments are within the capability of standard microfluidics laboratories. where $c(u)$ is the instantaneous control cost and $B_0 > 0$ is the initial budget. Admissible control requires $B(t) \geq 0$.

C5.2 — Control Fatigue

Control fatigue occurs as $B(t) \to 0$. High-frequency or high-magnitude control accelerates depletion of $B(t)$, reducing $\mathcal{M}(x)$.

By Proposition C3.2, impulsive strategies with convex cost deplete the budget strictly faster than steady control at the same mean effort.

Theorem C5.1 (Survival time bound). Let $x(t)$ evolve under $dx/dt = f(x) + u$ with $u(t) \in U$ and control cost $c(u) \geq c\_min > 0$ for all $u \neq 0$. Let $B(t) = B_0 - \int_0^t c(u(s))\, ds$ be the control budget.

Define the survival time T as the first time at which either $B(T) = 0$ or $x(T^*) \notin Viab(R)$. Then:

Proof. Since $c(u) \geq c\_min$ for any nonzero control, the budget depletes at rate $dB/dt = -c(u) \leq -c\_min$ whenever the system is actively controlled.

If the system requires continuous control to remain in Viab(R) (i.e., f(x) points outward at x for all x on the trajectory), then control must be nonzero for the entire survival period, giving $B(T) = B_0 - \int_0^T c(u)\, ds \leq B_0 - c\_min \cdot T^*$.

Setting $B(T) = 0$ yields $T \leq B_0/c\_min$. □

The bound is tight for constant minimal-cost control. It establishes that finite budgets imply finite survival: no system with bounded resources can maintain agency indefinitely against persistent drift.

The bound does not depend on the drift rate a, only on the control cost floor.

Faster drift depletes the budget faster (higher u needed), but the absolute bound is set by the budget-to-cost ratio.

Worked example (scalar system). For dx/dt = −ax + u with a = 1, u_max = 2, c(u) = u (linear cost), B_0 = 10, and initial state x_0 = 1.5 (inside Viab(R) = [0, 2]): maintaining x = 1.5 requires u = ax = 1.5, costing c = 1.5 per unit time.

Survival time: T* = B_0/c = 10/1.5 ≈ 6.67 time units.

After budget depletion, u = 0 and x decays exponentially toward 0. The system crosses into ruin when it can no longer reach any target within Viab(R). C6 — Noise and Silence

C6.1 — Noise

Noise is defined as exogenous or stochastic input to the system dynamics that is not under admissible control and that consumes control bandwidth without increasing Reach(x) ∩ Viab(R).

Formally, noise is any perturbation ξ(t) added to the drift field, f(x) → f(x) + ξ(t), where ξ is not an element of the admissible control set U.

Proposition C6.1 (Noise-induced agency decay). Let the system evolve under dx/dt = f(x) + u + ξ(t) where ξ is a zero-mean stochastic forcing with $\mathbb{E}[\xi] = 0$ and $\mathbb{E}[|\xi|^2] = \sigma^2$.

If the system must expend additional control Δu to compensate for ξ, then your effective budget depletion rate increases: dB/dt = −c(u + Δu) ≤ −c(u) − $\alpha\sigma^2$ for some α > 0 depending on the convexity of c. Consequently, noise reduces your survival time: T_noisy ≤ B_0/(c_min + $\alpha\sigma^2$) < T_quiet.

Noise taxes your control budget without expanding reachable viability volume.

C6.2 — Silence

Withholding response (u(t) = 0) is an admissible control action. When the drift field f(x) is slow or favorable (directed away from Σ_NR), silence preserves control budget at no agency cost.

This is not inaction in the colloquial sense; it is the optimal control policy when the marginal agency cost of intervention exceeds the agency cost of drift.

Formally, silence is preferred when $c(u)/|\partial\mathcal{M}/\partial u| > |d\mathcal{M}/dt|_{u=0}$, i.e., when the budget cost per unit of agency preservation exceeds the drift-induced agency loss rate.

In noise-dominated regimes, silence may also prevent noise-amplifying feedback loops in which control effort introduces additional disturbance. C7 — Coupling and Rescue

C7.1 — Coupled Systems and Agency Transfer

When systems are coupled, their drift fields combine and control capacities load jointly. **Agency transfer** occurs when, under coupled dynamics:

indicating expansion of reachable viability for system A at the expense of system B. The total agency of the coupled system is not conserved.

This is not an assumption; it follows from the geometry of Viab(R) under coupling.

C7.2 — Rescue Instability (Sufficient Condition)

“Rescue” is coupling a stabilized system A to a divergent system B to offset B’s drift using A’s control capacity. Let |·| denote the norm induced by the coupled dynamics on the joint state space.

A **sufficient condition** for joint loss of viability is: $|f_a| + |f^b| > |u_a|_{\max} + |u^b|_{\max}$

Under this condition, the total drift magnitude exceeds the total available control, and the coupled system approaches ΣNR faster than either system in isolation.

This is a sufficient, not necessary, condition; directional alignment of drift and control fields may permit stability even when this scalar inequality holds.

Non-conservation: two examples. (1) Cooperative coupling. Two scalar systems with drift $a = 1$, $u_{max} = 1$ each, coupled so that each contributes control to the other.

If the coupling allows total control capacity to be shared: effective u_max per system = 2, x_h doubles for both. $\mathcal{M}$_A + $\mathcal{M}$_B increases. Coupling creates agency.

(2) Parasitic coupling. System A (a = 1, u_max = 2) is coupled to system B (a = 3, u_max = 0).

B diverts A's control capacity: effective u_max for A drops to 1, while B still cannot sustain itself (3 > 1). Both systems lose agency: $\mathcal{M}$_A + $\mathcal{M}$_B decreases. Coupling destroys agency.

These examples demonstrate that agency transfer is not zero-sum. The coupling topology and the relative drift-to-control ratios determine whether joint agency expands, contracts, or redistributes. No conservation law governs total agency.

C8 — Slack and Robustness

C8.1 — Slack

Slack is defined as the minimum time to reach ΣNR under zero control: where ϕ_t is the uncontrolled flow generated by f(x).

Slack measures time-to-boundary, not Euclidean distance, and is the operationally relevant quantity for assessing control margin. Greater slack increases the time window available for corrective control and absorbs perturbations.

Proposition C8.1 (Scalar Slack-Agency Correspondence). For the scalar system dx/dt = −ax + u with u ∈ [0, u_max], the slack s(x) = x/a (time to reach x = 0 under zero control) and the horizon x_h = u_max/a satisfy: $\mathcal{M}$(x) is monotonically increasing in s(x) for x ∈ [0, x_h].

Greater slack implies greater agency. You have felt this — the difference between having three months of savings and having three days.

At s = 0 (boundary), $\mathcal{M}$ = 0. At s = x_h/a (maximal slack at origin), $\mathcal{M}$ = 1. Slack is the operationally measurable proxy for agency in systems where direct computation of Reach(x) ∩ Viab(R) is intractable.

In higher dimensions or systems with non-convex constraints, slack is a necessary but not sufficient condition for agency: a state may have large time-to-boundary under zero control yet be surrounded by regions from which no viable trajectory exists (geometric cul-de-sacs).

C8.2 — Redundancy

A system has **redundancy r ≥ 1** with respect to a target state $x \in \mathrm{Viab}(R)$ if there exist at least r distinct admissible control trajectories reaching x from the current state while remaining within Viab(R).

Redundancy reduces sensitivity of viable trajectories to perturbations in f(x) and u(t). Higher redundancy increases robustness at the cost of efficiency, since maintaining multiple viable pathways consumes control capacity that could otherwise extend reachability.

C9 — Exit as a Control Outcome

C9.1 — Withdrawal

When $\mathcal{M}(x(t))$ decreases monotonically under all admissible controls in a coupled system, decoupling preserves more reachable viability volume than continued coupling.

This holds when the coupled drift exceeds the joint control capacity (C7.2), so that decoupling removes the excess drift load. Exit is therefore a control outcome implied by reachability geometry, not a prescription.

C9.2 — Agency-Dissipative Environments

An environment is **agency-dissipative** if, for all admissible controls:

Persistence in such an environment strictly reduces reachable viability volume. This is a geometric characterization, not a recommendation.

Proposition C9.1 (Decoupling condition). Let systems A and B be coupled with joint dynamics.

Decoupling is agency-preserving for A if and only if the coupled drift acting on A exceeds A’s isolated drift: $|f_{coupled,A}(x)| > |f_A(x)|$.

That is, decoupling is preferred when the coupling increases the effective drift on A beyond what A experiences in isolation.

You know this. The relationship that costs more energy to maintain than it provides is a relationship that increases your drift. The mathematics says: leave. Not because leaving is morally right.

Because the geometry of your viability kernel contracts while you stay. This is a necessary and sufficient condition for decoupling to instantaneously increase $d\mathcal{M}_A/dt$.

It does not account for future recoupling opportunities or transient effects. C10 — Falsifiability and Closure

C10.1 — Falsifiers

Paper C is falsified if:

FC1**:** $\mathcal{M}(x)$ increases without corresponding control expenditure (violates C5.1).

FC2**:** Irreversible loss of reachability is reversed without external intervention violating the admissibility constraints of Paper A.

FC3**:** Stable control persists beyond ΣNR under admissible control (violates C4.2).

FC4 (Free lunch): A system maintains $\mathcal{M}(x) > 0$ indefinitely with B_0 finite and no external resource input, in the presence of persistent nonzero drift. This violates Theorem C5.1.

FC5 (Resurrection): A system recovers $\mathcal{M}(x) > 0$ after reaching $\mathcal{M} = 0$ (ruin) without external intervention that violates the admissibility constraints of Paper A. This violates C4.2.

C10.2 — Closure

Paper C introduces no new physics. It applies the irreversible and selective constraints of Papers A and B to controlled dynamics within a realized record sector.

Identification of concrete systems instantiating these constraints—biological, engineered, or otherwise—is treated separately. No extension of this argument is possible without new physical assumptions.

Experimental Instantiation

The definitions and propositions of Paper C are abstract control-theoretic structures. They become empirically meaningful when instantiated in concrete systems.

Two candidate systems are outlined below, one biological and one engineered, to demonstrate that the argument makes operationally testable predictions.

System 1: Bacterial chemotaxis. A bacterium navigating a nutrient gradient instantiates the scalar control model. State: nutrient concentration at cell location. Drift: diffusion-driven nutrient depletion ($a > 0$).

Control: flagellar motor switching ($u \in \{run, tumble\}$). Budget: metabolic energy store (ATP). Viability kernel: nutrient concentrations supporting growth. No-return surface: starvation threshold below which metabolic shutdown is irreversible.

Testable prediction: survival time scales with initial metabolic reserve divided by maintenance metabolic rate (Theorem C5.1). Noise: Brownian rotational diffusion acts as stochastic forcing ξ, taxing the control budget (Proposition C6.1).

Observable: mean survival time decreases with increasing environmental noise, controlling for nutrient availability.

System 2: Autonomous robotic navigation. A battery-powered robot avoiding obstacles instantiates the 2D control model. State: (position, battery level) $\in \mathbb{R}^2 \times \mathbb{R}_+$. Drift: gravitational or terrain slope.

Control: motor torque ($u \in U$, bounded by motor capacity). Budget: battery charge (B_0). Viability kernel: states from which the robot can reach a charging station before battery depletion.

No-return surface: states where remaining battery is insufficient to reach any charger under optimal control. Testable prediction: the robot's reachable viable set shrinks monotonically as battery depletes (Proposition C2.1).

Slack: time-to-boundary under zero motor input = coasting distance / terrain slope. Observable: optimal control policies should exploit silence (zero motor) on favorable slopes, consistent with C6.2.

These instantiations are not metaphors. Each maps the abstract quantities ($\mathcal{M}$, B, s, Σ_NR) to physically measurable variables with quantitative predictions. Failure of the predictions in either system falsifies the corresponding propositions of Paper C.

C11 — Structural Closure

Together, the trilogy establishes a layered, one-way dependency chain:

Paper A: irreversibility as loss of reachability under bounded control. Defines Actualization State, proves monotonicity under decohering dynamics, and establishes no-return surfaces. Independent of Papers B and C.

Paper B: selection as costly, rate-limited, irreversible exclusion of alternative record sectors, if it exists. Derives structural requirements and a falsifiable gravitational rate bound. Depends on Paper A; independent of Paper C.

Paper C: agency as normalized reachable viability volume under your constrained control within a single realized record sector.

Establishes propositions on agency decay under drift (C2.1), survival time bounds (C5.1), noise-induced depletion (C6.1), non-conservation under coupling (C7), and slack-agency correspondence (C8.1). Provides worked examples and experimental instantiations in biological and engineered systems.

Depends on Paper A; uses outcome of Paper B but not its mechanism.

Failure of Paper C does not invalidate Paper B. Failure of Paper B does not invalidate Paper A. Each layer is independently falsifiable.

What remains is empirical: which systems realize these structures, and how closely.

End of Paper C. Paper C — Canonical Reference Locked · Execution-Complete

Paper D Coupled Viability

Structural Conditions for Multi-Agent Persistence Under Irreversible Dynamics

Depends on Papers A, B, and C

Paper D extends coupling to multi-agent systems. It depends on Papers A, B, and C and nothing else. Failure of Paper D does not invalidate Papers A, B, or C.

[Paper D content follows — see standalone document]

Paper D

Coupled Viability: Structural Conditions for Multi-Agent Persistence Under Irreversible Dynamics

Dependent on: Paper A — Actualization State Paper B — Selection as Irreversible Exclusion Paper C — Agency as Constrained Control

D0 — Dependency, Scope, and Non-Overlap

D0.1 — Dependency Statement

Paper D depends explicitly and exclusively on the results of Papers A, B, and C.

It assumes as given:

Actualization State as an operational measure of record-structured irreversibility (Paper A).

Selection as costly, rate-limited exclusion to definiteness, if it exists (Paper B).

Agency as normalized reachable viability volume under constrained control within a single realized record sector (Paper C).

No construct from Papers A, B, or C is redefined or rederived.

Failure of Paper D does not invalidate any prior paper.

D0.2 — Purpose

Paper D addresses: Given multiple agents, each described by Paper C's formalism, operating within shared constraint environments under irreversible physics, what are the structural conditions for persistent joint dynamics, and what forms of emergent order are admissible?

This is a question about the geometry of coupled viability kernels under drift.

It is not a question about society, cooperation, or morality.

D0.3 — Positioning Relative to Existing Literature

Multi-agent viability theory exists. Aubin, Bayen, Saint-Pierre, and colleagues have developed the mathematics of viability kernels for coupled systems, differential games, and multi-agent control.

The contribution of Paper D is not in proving new viability theorems.

It is in applying viability theory to the specific structure of irreversibility (Paper A), selection (Paper B), and agency (Paper C).

The results are constraints derived from physical irreversibility and record structure, not abstract control theory.

Paper D is not evolutionary game theory. It does not invoke fitness, replication, or selection pressure.

Paper D is not multi-agent reinforcement learning. It does not invoke reward signals, policy gradients, or learning.

Paper D is not mechanism design. It does not invoke incentive compatibility, revelation principles, or social welfare functions.

Paper D is viability geometry applied to physically irreversible, record-structured, agency-bearing coupled systems.

D0.4 — Hard Non-Claims

The paper does not:

Introduce new physical laws.

Modify or reinterpret quantum mechanics.

Invoke psychology, motivation, ethics, value, meaning, or consciousness.

Assume rationality, optimization, or fitness maximization.

Model communication, signaling, or strategic negotiation.

Derive evolutionary fitness.

Propose normative guidance or prescriptions.

Claim emergent structures are designed, intended, or purposeful.

Claim that social structures emerge from viability constraints alone.

Failure of Paper D does not invalidate Papers A, B, or C.

D0.5 — Loaded Terms: Geometric Definitions

Several terms in Paper D carry normative or sociological connotations in ordinary language. Each receives a strict geometric definition at first appearance. No connotation beyond the definition is implied.

“Cooperation” — a geometric condition where mutual record externalities expand joint viability. No intention, reciprocity, or payoff implied.

“Hierarchy” — asymmetric coupling where higher-capacity agents’ record externalities dominate the constraint landscape of lower-capacity agents. A consequence of scale asymmetry.

“Deterrence” — a coupling configuration where the cost of unilateral decoupling exceeds the cost of continued coupling for both agents. A viability geometry.

“Impedance” — the ratio of control authority to drift rate: $Z = u_{max} / a$. Two agents are impedance-matched when their operator horizons are comparable.

“Resonance” — frequency and phase compatibility between coupled control strategies. Constructive resonance expands joint viability; destructive resonance contracts it.

D1 — Shared Constraint Environments

D1.1 — Shared Viability Domain

When multiple agents operate within a common physical environment, their individual viability kernels may overlap.

The joint state space is the product of individual state spaces.

The joint dynamics are defined by the individual drift fields, the individual control sets, and the coupling terms that transmit the effect of one agent's actions onto another's drift.

The joint viability kernel is the set of joint states from which there exists a joint admissible control strategy that maintains all agents within their individual viability constraints for all future time.

The shared viability domain is the projection of the joint viability kernel onto the shared constraint dimensions. It is the region of the common environment in which joint persistence is geometrically possible.

Structural assumption: Agents share constraint dimensions. They operate in a common physical environment whose state is affected by the actions of all agents. This is the defining condition for being in a shared constraint environment. If agents' state spaces are fully orthogonal (no shared dimensions), they are uncoupled and Paper D does not apply.

This assumption is not a theorem. It is a scope condition. Paper D analyzes systems that satisfy it.

D1.2 — Constraint Coupling

Paper C (C7) treats coupling as direct energy and control transfer between two systems.

Paper D introduces a second coupling mode: constraint coupling. When Agent A's actions modify the shared environment in a way that alters Agent B's drift field, control set, or viability kernel, the agents are constraint-coupled.

No direct energy exchange is required. The coupling operates through the shared constraint landscape.

Example: Robot A occupies the charging station. Robot B's admissible trajectories contract (it cannot charge). No energy flowed from A to B. But B's reachable set changed because A's action modified the shared environment.

Scope: Pairwise coupling as base analysis. Network effects modeled as cascades through pairwise links. Testable with 3-agent systems (minimal non-trivial network).

D1.3 — Record Externalities (Geometric Exclusion Principle)

Definition (Record-Writing Action):

An irreversible action by Agent A whose recorded environmental change lies in the shared constraint coordinates (e), and which modifies B's admissible dynamics $f_B(\cdot; e)$ or admissible control set $R_B(e)$.

The shared constraint coordinates are the dimensions of the environment state space that appear in B's dynamics or constraints.

Geometric Exclusion Principle:

For coupled agents with $K_A \cap K_B \neq \emptyset$, if Agent A performs a record-writing action that changes the shared constraint coordinates on which B's viability depends, then K_B changes, and $\mu(K_B)$ changes generically (Corollary D1.3).

Non-degeneracy assumption:

The map $e \mapsto K_B(e)$ is non-degenerate: the viability kernel boundary ∂K_B depends smoothly on e, and the constraint surface intersects the kernel boundary transversally.

This excludes pathological cases where the environmental change lies entirely within the kernel interior, producing no boundary change.

Proof:

(1) A's record-writing action irreversibly modifies the shared constraint coordinates $e \to e'$ (by definition of record-writing action and Paper B's irreversibility).

(2) B's viability kernel is a function of the shared constraint coordinates: $K_B = K_B(e)$.

Since B's admissible dynamics or control set depend on e (by definition), and the kernel boundary depends smoothly on e (by the non-degeneracy assumption), changing e changes the set of states from which B can persist.

(3) Under the non-degeneracy assumption, $K_B(e') \neq K_B(e)$. The transversality condition ensures the boundary moves under perturbation of e.

The change may be positive (expansion) or negative (contraction), depending on the direction of $e \to e'$ relative to B's constraint surface.

(4) Sign classification: if e' tightens B's constraints (reduces B's control set or increases B's drift), K_B contracts (negative externality). If e' loosens B's constraints, K_B expands (positive externality).

(5) $\mu(K_B)$ changes generically: by the transversality theorem, the set of e' for which $K_B(e') \neq K_B(e)$ but $\mu(K_B(e')) = \mu(K_B(e))$ (volume-preserving deformations) has measure zero in the space of admissible environmental changes.

This is the content of Corollary D1.3. □

Corollary D1.3 (Genericity of Non-Neutrality):

In smooth families of coupling maps, the set of record-writing actions that produce exactly zero change in $\mu(K_B)$ has measure zero. Neutral externality requires parameter-level fine-tuning. This holds under the same non-degeneracy assumption stated above.

Falsifier D1 (No Free Survival):

If Agent A exerts a negative record externality on Agent B (measured as decrease in $\mu(K_B)$), and Agent B increases its agency $\mathcal{M}_B$ (Paper C measure) without: (a) severing coupling, (b) increasing its control budget $u_{B,max}$, or (c) receiving compensating positive externalities from a third agent, the argument is falsified.

Scope boundary. The measure-zero neutrality result (Corollary D1.3) depends on smoothness of the coupling map and regularity (C^2) of the viability kernel boundary.

If the coupling map is non-smooth or the kernel boundary contains cusps, corners, or discontinuities, the transversality argument may fail and non-neutral actions could have positive measure.

This is an explicit scope limitation: Paper D's genericity claims apply to smooth families of coupling maps with regular kernel boundaries.

A further exception arises in systems with continuous symmetries (e.g., rotational invariance) where record-writing actions correspond to symmetry operations that preserve the viable volume by construction. Outside of such symmetry-protected subspaces, non-neutrality is generic.

D2 — Composition of Agency

D2.1 — Joint Agency and Non-Additivity

Paper C established that agency is non-conservative under coupling (C7.1). Paper D extends this to N agents.

Definition (Joint Agency):

Joint agency $\mathcal{M}_{joint}$ is defined as the viability volume of the joint state space under joint admissible controls, normalized by the total joint viability kernel.

The measure $\mathcal{M}$ is inherited from Paper C's Definition D5, applied to the product state space $X_1 \times X_2 \times \ldots \times X_N$.

Proposition D2.1: Joint agency is non-additive. $\mathcal{M}$_joint ≠ Σ $\mathcal{M}$_i in general.

The non-additivity term depends on: (a) alignment of individual drift fields, and (b) compatibility of individual control sets.

Joint agency is superadditive ($\mathcal{M}_{joint} > \Sigma \mathcal{M}_i$) when drift fields are anti-aligned (agents face complementary threats) and control sets are compatible.

Joint agency is subadditive ($\mathcal{M}_{joint} < \Sigma \mathcal{M}_i$) when drift fields are co-aligned (agents face the same threat simultaneously) or control sets conflict.

Proof sketch (by construction):

Superadditive witness: Two scalar systems, each with drift $a = 1$ and $u_{max} = 1$. Uncoupled, each has $\mathcal{M}_i$ = viability volume of its individual kernel.

Coupled cooperatively (Paper C, C7.1, cooperative example), shared control capacity yields effective $u_{max} = 2$ per system.

The joint viability kernel expands: $\mathcal{M}_{joint} > \mathcal{M}_1 + \mathcal{M}_2$. This is the cooperative coupling case from Paper C.

Subadditive witness: System A ($a = 1$, $u_{max} = 2$) coupled parasitically to system B ($a = 3$, $u_{max} = 0$). B diverts A's control capacity.

Effective u_{max} for A drops to 1. The joint viability kernel contracts: $\mathcal{M}_{joint} < \mathcal{M}_1 + \mathcal{M}_2$. This is the parasitic coupling case from Paper C.

Since both strict inequality directions are realizable, $\mathcal{M}_{joint} \neq \Sigma\, \mathcal{M}_i$ in general. The sign depends on drift alignment and control compatibility. □

D2.2 — Impedance Matching

Definition:

Impedance $Z_i \equiv u_{i,max} / a_i$, where $u_{i,max}$ is maximum control authority and a_i is drift rate. This mirrors the operator horizon (Paper A, Theorem T2).

Drift timescale:

$\tau_i = 1/a_i$. The characteristic time for Agent i's drift to carry it a significant distance toward its no-return surface.

When $Z_i \neq Z_j$, three distinct failure modes arise:

Failure mode A (primary): Deadline mismatch.

The low-Z agent has short slack (small τ, little time-to-ruin). Help must arrive before the no-return surface is crossed.

If the high-Z agent's response delay exceeds the low-Z agent's remaining slack, the help arrives after no-return. Control effort applied after no-return produces zero viability gain.

This is the primary failure mode because no-return geometry makes time windows hard and asymmetric: once missed, the effect is permanently zero.

Failure mode B: Absorption bottleneck.

Even if help arrives in time, the low-Z agent may not be able to convert it into viability. If the limiting constraint is u_{max}, then energy that does not increase u_{max} does not change Z.

Resources delivered to a system whose bottleneck is not resources produce zero reachability gain.

Failure mode C (secondary): Budget drain on the high-Z agent.

To rescue a short-slack agent, the high-Z agent must accelerate its response. The high-Z agent's budget drains faster.

This is secondary because the drain is usually caused by failure mode A: the high-Z agent spends, but the effect arrives too late or cannot be converted.

Proposition D2.2 (qualitative):

Coupling efficiency between agents i and j degrades as impedance ratio $|Z_i/Z_j|$ deviates from unity. Waste increases with mismatch. The primary waste mechanism is control effort applied outside the low-Z agent's viable intervention window.

Conjecture D2.2 (quantitative):

For a minimal interface model with transfer delay τ, conversion factor κ, and slack s, coupling efficiency η is bounded by the fraction of control effort deliverable within the slack window.

A candidate functional form is $\eta \leq 1/(1 + |Z_i/Z_j - 1|)$.

This requires derivation in the specific interface model and is not claimed as a universal result.

D2.3 — Resonance and Phase

System class:

Linear periodically forced agents with scalar state x_i, symmetric pairwise coupling κ, identical drift a, and sinusoidal control $u_i(t) = U \sin(\omega_i t + \varphi_i)$.

Stability condition: $a > \kappa$ (drift exceeds coupling strength; if $\kappa \geq a$ the sum mode is unstable and both agents diverge regardless of phase).

Theorem D2.3 (Toy Model):

In the above system class, the measure of the joint viability set (the set of joint initial conditions from which both agents persist indefinitely) is maximized when $\omega_1 = \omega_2$ and $\varphi_1 - \varphi_2 = 0$ (in-phase resonance).

The joint viability set contracts monotonically as $|\varphi_1 - \varphi_2|$ increases from 0 to π.

Proof:

Dynamics: $dx_i/dt = -a \cdot x_i + u_i(t) + \kappa \cdot x_j$ ($i \neq j$), with viability constraint $x_i(t) \geq 0$. For identical frequencies $\omega_1 = \omega_2 = \omega$, define $\Delta\varphi = \varphi_1 - \varphi_2$. Define sum mode $S = x_1 + x_2$ and difference mode $D = x_1 - x_2$. These decouple:

By trigonometric identities: $\sin(\omega t + \varphi_1) + \sin(\omega t + \varphi_2) = 2 \cdot \cos(\Delta\varphi/2) \cdot \sin(\omega t + (\varphi_1+\varphi_2)/2)$, and $\sin(\omega t + \varphi_1) - \sin(\omega t + \varphi_2) = 2 \cdot \cos(\omega t + (\varphi_1+\varphi_2)/2) \cdot \sin(\Delta\varphi/2)$.

The effective control amplitude on the sum mode is $2U \cdot \cos(\Delta\varphi/2)$. The effective control amplitude on the difference mode is $2U \cdot |\sin(\Delta\varphi/2)|$.

The viability constraint $x_i \geq 0$ translates to $S \geq |D|$ (both components are non-negative if and only if their sum exceeds the absolute value of their difference).

Joint viability therefore requires: (a) S remains large, which requires maximum control amplitude on the sum mode, and (b) |D| remains small, which requires minimum forcing on the difference mode.

Condition (a) is optimized when $\cos(\Delta\varphi/2) = 1$, i.e., $\Delta\varphi = 0$. Condition (b) is optimized when $\sin(\Delta\varphi/2) = 0$, i.e., $\Delta\varphi = 0$. Both conditions are simultaneously optimized at $\Delta\varphi = 0$ (in-phase resonance).

Monotonic contraction: as $|\Delta\varphi|$ increases from 0 to π, $\cos(\Delta\varphi/2)$ decreases monotonically from 1 to 0 (sum-mode control weakens) and $|\sin(\Delta\varphi/2)|$ increases monotonically from 0 to 1 (difference-mode forcing strengthens).

Both effects reduce the set of initial conditions from which both agents persist.

At $\Delta\varphi = \pi$: $\cos(\pi/2) = 0$ (zero sum-mode control) and $|\sin(\pi/2)| = 1$ (maximum difference-mode forcing). This is the worst case. □

Conjecture D2.3 (General):

For broader classes of periodically controlled coupled agents, joint viability is generically maximized under frequency commensurability and phase alignment.

The empirical signature is: perturbing a persistent coupled system's phase relationship should contract the joint viability margin (measured as minimum slack over one control cycle).

Falsifier for Conjecture D2.3:

If a persistent coupled system is shown to have maximum joint viability at anti-resonant phase ($\varphi_1 - \varphi_2 = \pi$) under standard coupling, the conjecture is falsified.

D3 — Stable Configurations Under Drift

D3.1 — Compositional Equilibrium

A compositional equilibrium (CE) is a joint state-control configuration in which all agents maintain positive agency ($\mathcal{M}_i > 0$ for all i) indefinitely, given their joint drift field and joint control constraints.

Mathematical obstacle. The two-agent sinusoidal proof (above) succeeds because the viability constraint is linear in the sum/difference decomposition and the control is purely periodic.

The general case resists proof for three identified reasons: (i) nonlinear coupling terms (e.g., multiplicative or saturating interactions) break the sum/difference decoupling that enables the trigonometric argument; (ii) non-convexity of the joint viability kernel under general control policies means the viability volume does not decompose into independent modal contributions; (iii) for non-periodic control strategies, the phase relationship is not a well-defined scalar parameter, and the optimisation landscape may have local maxima at non-zero phase offset.

Any proof of the general conjecture must either restrict the coupling class (e.g., to affine or monotone coupling) or establish a variational principle on the joint viability volume that is monotone in phase alignment.

Until such a proof exists, the conjecture is supported by the two-agent result and by the specified falsifier.

CE does not invoke rationality.

CE does not invoke payoff maximization.

CE is a geometric fixed-point condition: the joint system remains within the interior of the joint viability kernel.

CE ≠ NE: The Operator-Required Charger

System: Two robots, each with battery $b_i(t) \in [0, 100]$. Ruin at $b_i \leq 10$.

Drift: Battery drains at 12 units/hour (always on).

Coupling condition: To charge, the other robot must crank (operate the charger). Charging rate: +30 units/hour. Cranking cost: additional −4 units/hour (total −16/hour for the cranking robot). If no one cranks, no one charges.

Compositional Equilibrium: Alternation.

Hour 0–1: R1 charges, R2 cranks. R1: 100 → 100 (capped). R2: 100 → 84.

Hour 1–2: R2 charges, R1 cranks. R2: 84 → 100 (capped). R1: 100 → 84.

Cycle repeats. Both oscillate between 84 and 100. Both remain far above ruin threshold (10). Persistence is indefinite. This is CE.

Nash-type defection: Refuse to crank.

At any cranking turn, defection costs only −12/hour instead of −16/hour. Net gain: +4 units/hour. Strict local improvement.

If R1 defects every cranking turn but accepts cranking from R2:

Hour 0–1: R1 charges, R2 cranks. R1: 100, R2: 84.

Hour 1–2: R1 refuses. Nobody charges. R1: 88, R2: 72.

Hour 2–3: R1 charges, R2 cranks. R1: 100, R2: 56.

Hour 3–4: R1 refuses. R1: 88, R2: 44.

Hour 4–5: R1 charges, R2 cranks. R1: 100, R2: 28.

Hour 5–6: R1 refuses. R1: 88, R2: 16.

Hour 6–6.375: R2 crosses ruin threshold while cranking. R2 dead.

Hour 6.375–13.5: R1 cannot charge (no cranker). Drifts to ruin at 12/hr. R1 dead.

Result: CE (alternation) produces indefinite survival. The Nash move (refuse to crank) is a strict local improvement (+4/hour). The Nash move kills the partner at t ≈ 6.375 hours.

Without a partner, the defector cannot charge and dies at t ≈ 13.5 hours.

Conclusion: CE is not Nash-stable under immediate-payoff incentives. The Nash-type unilateral improvement exits the joint viability kernel. Viability ≠ utility. A rational agent can calculate its way into extinction by ignoring the operator horizon.

D3.2a — Necessary Conditions for Persistence (under alignment)

Assumption (Monotonic Alignment):

All agents face drift toward the same boundary of the viability kernel. No agent's drift partially compensates another's drift without control expenditure.

Assumption (Regularity):

The viability kernel boundary is smooth (C^2). No agent's state is exactly on a cusp or non-differentiable point of the kernel boundary.

Proposition D3.2a (Necessary Conditions):

Under the monotonic-alignment and regularity assumptions, a multi-agent configuration that persists indefinitely must satisfy:

(N1) Aggregate control capacity exceeds aggregate drift (joint operator horizon condition).

(N2) Impedance compatibility: agents' drift timescales are close enough that help can arrive within the low-Z agent's slack window.

(N3) No agent's record externalities push another past its no-return surface faster than that agent can compensate (measured as: drift increase from externality ≤ compensating control authority).

(N4) Joint control budget is sufficient to maintain all agents above ruin (total energy expenditure ≤ total available budget over any finite horizon).

Failure of any condition implies at least one agent reaches its no-return surface in finite time (under the stated assumptions).

Proof of necessity (sketch):

N1 violated (aggregate drift exceeds aggregate control): Under monotonic alignment, all drift vectors point toward the same boundary. The joint system's state moves toward Σ_{NR} at net rate $\Sigma a_i - \Sigma u_{i,max} > 0$.

Since the viability kernel has finite diameter, the boundary is reached in finite time bounded by $\text{diam}(K) / (\Sigma a_i - \Sigma u_{i,max})$. At least one agent exits.

N2 violated (impedance incompatibility): Under monotonic alignment, the low-Z agent's drift carries it toward Σ_{NR} with slack $s_{low} = d(x_{low}, \Sigma_{NR}) / a_{low}$.

If the high-Z agent's minimum response delay $\tau_{response} > s_{low}$, the intervention arrives after no-return. By Paper A's absorbing-state result, the low-Z agent cannot recover.

Its survival time T is bounded by s_{low}.

N3 violated (externality exceeds compensation): Agent i's record-writing action increases j's effective drift by Δa_j. If $\Delta a_j > u_{j,max} - a_j$ (the remaining control margin), then j's net drift becomes positive toward Σ_{NR}.

By the same finite-diameter argument as N1, j reaches the boundary in finite time.

N4 violated (budget exhaustion): Each agent's control expenditure rate is at least a_i (the maintenance cost from Paper C, Corollary C3.1a).

If $\Sigma a_i >$ total budget rate, the aggregate budget depletes to zero in finite time. Once budget is exhausted, all agents are subject to uncontrolled drift and reach Σ_{NR} in finite time. □

Scope of necessity:

Under non-aligned drift (where agents' drift fields partially cancel), configurations may persist while violating N2 (because natural cancellation reduces the effective impedance mismatch). The necessity claim is strictly conditional on the alignment assumption.

D3.2b — Sufficient Conditions for Persistence (without alignment)

Proposition D3.2b (Sufficient Conditions):

The following are sufficient conditions for a multi-agent configuration to persist under irreversible drift, with no alignment assumption required:

(S1) Each agent independently satisfies its single-agent viability condition (Paper C): $u_{i,max} > a_i$ and budget > maintenance cost.

(S2) All pairwise record externalities are non-negative (no agent's actions contract any other agent's kernel).

(S3) Coupling is impedance-compatible (D2.2).

(S4) Joint control budget exceeds joint maintenance cost.

Under these conditions, all agents persist. No agent's kernel contracts due to coupling, and each has sufficient resources to maintain itself.

Proof sketch:

By S1, each agent satisfies Paper C's single-agent viability condition independently: $u_{i,max} > a_i$ and budget exceeds maintenance cost.

By S2, no agent's record externalities contract any other agent's kernel, so coupling does not degrade any individual viability condition. By S3, coupling is impedance-compatible, so control transfers between agents arrive within viable intervention windows.

By S4, the joint budget sustains the joint maintenance cost indefinitely. Each agent therefore remains within its individual viability kernel for all time, and the joint state remains within the joint viability kernel. □

Note: D3.2a and D3.2b are categorically different claims. D3.2a tells you what persistent configurations must look like (under alignment). D3.2b tells you what configurations will definitely persist. Neither subsumes the other.

A configuration can satisfy D3.2b without violating D3.2a, but violating D3.2a does not imply violating D3.2b, because D3.2a requires alignment assumptions that D3.2b does not.

D3.3 — Instability and Cascade Failure

When Agent i exits its viability kernel ($\mathcal{M}$_i = ∅), and the coupling was such that i's control actions were partially compensating j's drift, then j's effective drift increases by the lost coupling contribution.

Proposition D3.3:

Failure of Agent i propagates to Agent j if the removal of i's control contribution increases j's effective drift beyond j's remaining control margin ($u_{j,max}$).

If the new effective drift exceeds $u_{j,max}$, j is pushed beyond its operator horizon and cascades toward its own no-return surface.

Containment: Cascade stops at Agent j if j has sufficient slack (s_j from C8.1) to absorb the shock before the increased drift pushes it past Σ_{NR}.

Falsifiable test: 3-agent system: A coupled to B coupled to C. Terminate A. Measure survival time T_B and T_C.

The cascade condition predicts: if removal of A's contribution increases B's drift past $u_{B,max}$, T_B is finite. If B's failure increases C's drift past $u_{C,max}$, T_C is finite.

Propagation stops when the drift increase at an agent is less than that agent's remaining control margin.

D4 — Emergent Order Without Design

D4.1 — Null Model and Order Metric

Before claiming emergent order, Paper D establishes what the absence of order looks like.

Null model:

A population of agents with random (uncorrelated) control policies under the same drift field, coupling topology, initial condition distribution, and environment noise. Only survivors are analyzed.

The null model must match every confound except control policy coordination.

Order metric (Slack Correlation):

For surviving agents, measure individual slack $s_i(t)$ = time-to-ruin if controls are frozen at time t.

Compute pairwise cross-correlation $\rho_{ij} = \mathrm{corr}(s_i(t), s_j(t))$ over time.

Random survivors: $\rho \approx 0$ (independent fluctuations).

Coordinating agents: ρ significantly positive (slack levels move together).

Statistical test:

Compute $\bar{\rho}_{obs}$ = mean pairwise correlation among survivors in the observed system.

Build null distribution $\{\bar{\rho}_{null}\}$ from N simulation runs with random control policies and survivorship filtering.

Compute empirical p-value: $p = (1 + \#\{\text{null runs with } \bar{\rho} \geq \bar{\rho}_{obs}\}) / (1 + N)$.

Declare order if $p < 0.05$.

Structural precondition: Overlap ratio $O(t) = \mu(\cap K_i) / \mu(\cup K_i)$ measures geometric capacity for coordination. Slack correlation measures actual coordination. Paper D uses both.

Falsifier D4 (Order Indistinguishable from Noise):

If observed persistent configurations cannot be statistically distinguished from the null ($p \geq 0.05$ for empirical slack correlation), D4 is falsified. Observable: empirical p-value.

D4.2 — Structural Filtering of Configurations

Under irreversible drift, configurations that violate the necessary conditions of D3.2a (under the stated alignment and regularity assumptions) are eliminated.

Survivors are biased toward configurations satisfying these conditions—not because they were selected for, but because everything else exited the viability kernel.

Proposition D4.2:

Under irreversible drift and the monotonic-alignment and regularity assumptions of D3.2a, the long-run support of persistent multi-agent configurations is contained in the set of configurations satisfying the necessary conditions N1–N4. No optimization, fitness function, or teleology is required.

Conjecture D4.2:

Under additional assumptions (ergodicity, stationary drift, specified policy update process), the distribution of persistent configurations converges to the set of compositional equilibria.

This conjecture requires explicit specification of the stochastic process and is not claimed as a theorem.

D4.3 — Hierarchy as Constraint Geometry

When agents have asymmetric capacity (different Z values), stable configurations generically exhibit hierarchical structure: higher-capacity agents' record externalities dominate the constraint landscape of lower-capacity agents.

Regularity assumption:

The coupling map is smooth and the constraint surface is non-degenerate (as in D1.3). Agent impedances are distinct: $Z_i \neq Z_j$.

Proposition D4.3:

In a coupled system with asymmetric agent impedances under the above regularity assumption, persistent configurations exhibit hierarchical coupling: the high-Z agent's record externalities alter $\mu(K_j)$ for the low-Z agent more than the reverse.

Proof sketch:

The high-Z agent has a larger viability kernel (higher u_max/a ratio means more reachable viable states, by Paper A's operator horizon).

Under the regularity assumption (smooth coupling, non-degenerate constraint surface), the kernel expansion projects non-trivially onto the shared constraint coordinates — that is, the additional reachable states include states that differ in the shared dimensions, not only in private dimensions.

A larger footprint in the shared constraint coordinates means that the high-Z agent's record-writing actions generically (on a set of full measure in the space of coupling parameters) modify more of the shared constraint space than the low-Z agent's actions.

By the Geometric Exclusion Principle (D1.3), this produces larger changes in the partner's kernel.

The hierarchy is geometric, not intentional. □

Falsifier D4.3 (Hierarchy Inversion):

If, in a system with $Z_i \gg Z_j$ (impedance ratio > 10:1) satisfying the regularity assumption, the low-Z agent's record externalities dominate the constraint landscape of the high-Z agent ($\Delta\mu(K_i)$ from j's actions > $\Delta\mu(K_j)$ from i's actions, measured over equivalent action magnitudes), D4.3 is falsified.

Observable: $\Delta\mu(K_i)$ and $\Delta\mu(K_j)$ per unit record-writing action.

D4.4 — Cooperation and Deterrence as Structural Outcomes

Proposition D4.4a (Cooperation):

Cooperative equilibria exist when mutual record externalities expand each agent's viability kernel more than coupling cost contracts it. The operator-required charger (D3.1) is an instance. Observable: $\mathcal{M}_{joint} > \Sigma\ \mathcal{M}_i$.

Falsifier D4.4a (Cooperation Nonexistence):

If, in every tested coupled system where mutual record externalities are positive (each agent's actions expand the other's kernel), $\mathcal{M}_{joint} \leq \Sigma \mathcal{M}_i$ (joint agency never exceeds the sum of individual agencies), D4.4a is falsified.

Observable: $\mathcal{M}_{joint}$ and $\Sigma \mathcal{M}_i$ computed from the joint and individual viability kernels.

Proposition D4.4b (Deterrence):

Deterrence equilibria exist when unilateral decoupling cost exceeds continued coupling cost for both agents. Observable: for each agent, $\mathcal{M}_i$(coupled) > $\mathcal{M}_i$(decoupled). Neither agent can improve its viability by exiting the coupling.

This is a geometric fixed-point, not a threat.

Falsifier D4.4b (Deterrence Exit):

If an agent in a coupled system with $\mathcal{M}_i$(coupled) > $\mathcal{M}_i$(decoupled) for all i can unilaterally decouple and increase its agency ($\mathcal{M}_i$(after decoupling) > $\mathcal{M}_i$(coupled)), the characterization of the configuration as a deterrence equilibrium is falsified.

Observable: $\mathcal{M}_i$ before and after decoupling.

Both are geometric. Neither is normative.

D5 — Experimental Instantiations and Falsifiers

D5.1 — Worked Examples

System 1: Microbial ecology (chemostat).

Shared viability domain: nutrient-population configuration space. Record externalities: waste products altering pH/nutrient availability (irreversible environmental modification). Impedance matching: metabolic rate compatibility between species. Slack: time-to-washout at current dilution rate and population density.

Cascade failure: trophic coupling propagation. Each construct maps to a measurable variable with a quantitative prediction.

System 2: Operator-required charger (two robots).

Shared viability domain: joint (position, battery) space with shared charging infrastructure. Record externalities: station occupation. Coupling condition: charging requires partner's cranking. Impedance matching: battery capacity and discharge rate compatibility.

Slack: time-to-ruin at current battery level and discharge rate. CE ≠ NE demonstration: Section D3.1. Each construct maps to a measurable variable with a quantitative prediction.

D5.2 — Falsifiers

Global Falsifier F0 (Kill Switch):

If a multi-agent system persists indefinitely (survival time $T = \infty$ for all agents) while violating all necessary conditions N1–N4 of D3.2a, under a configuration satisfying the monotonic-alignment and regularity assumptions, Paper D is falsified.

Observable: survival time T for each agent; verification of N1–N4 violation; verification that alignment and regularity assumptions hold.

Falsifier D1 (No Free Survival):

Defined in D1.3. Observable: μ(K_B) and $\mathcal{M}$_B before and after A's record-writing action.

Falsifier D2.1 (Additivity Under Coupling):

If joint agency equals the sum of individual agencies ($\mathcal{M}$_joint = Σ $\mathcal{M}$_i) in a coupled system with non-zero coupling terms (non-orthogonal shared constraint coordinates), D2.1 is falsified.

Observable: $\mathcal{M}$_joint and Σ $\mathcal{M}$_i in coupled vs uncoupled configurations.

Falsifier D2.2 (Impedance-Independent Efficiency):

If coupling efficiency (measured as viability transfer per unit control effort) does not degrade as impedance ratio $|Z_i/Z_j|$ deviates from unity, D2.2 is falsified.

Observable: viability transfer rate at impedance ratios 1:1, 2:1, 5:1, and 10:1 under matched conditions.

Falsifier D2.3 (Anti-Resonant Optimality):

Defined in D2.3. Observable: joint viability margin at phase offsets 0, π/4, π/2, 3π/4, π.

Falsifier D3.3 (Cascade Non-Propagation):

Defined in D3.3. Observable: survival times T_B and T_C after termination of A in a 3-agent chain.

Falsifier D4.2 (Persistent Violators):

If a multi-agent configuration persists indefinitely while violating one or more of N1–N4, under a system satisfying the monotonic-alignment and regularity assumptions, D4.2 is falsified.

This differs from F0 (which requires violation of all four conditions).

D4.2 claims the long-run support is contained in the satisfying set; a single persistent violator of any condition falsifies it. Observable: persistence time T and verification of individual N1–N4 conditions for each surviving configuration.

Falsifier D4 (Order Indistinguishable from Noise):

Defined in D4.1. Observable: empirical p-value for slack correlation. If $p \geq 0.05$ for all candidate systems, D4 is falsified.

Falsifier D4.3 (Hierarchy Inversion):

Defined in D4.3. Observable: $\Delta\mu(K_i)$ and $\Delta\mu(K_j)$ per unit record-writing action in impedance-asymmetric systems.

Falsifier D4.4a (Cooperation Nonexistence):

Defined in D4.4. Observable: $\mathcal{M}_{joint}$ and $\Sigma \mathcal{M}_i$ in systems with mutually positive externalities.

Falsifier D4.4b (Deterrence Exit):

Defined in D4.4. Observable: $\mathcal{M}_i$ before and after unilateral decoupling.

Every proposition has at least one testable falsifier with a specified observable. Falsifiers are independent of A, B, C. Failure of any proposition leaves all prior papers intact.

D5.3 — Scope Closure

Paper D establishes: what multi-agent composition must look like under the trilogy's constraints, what persistent configurations require, what destroys them, and how to test these claims.

It does not determine whether specific configurations are realized in nature. That question remains empirical.

D6 — Structural Closure

Paper A: Irreversibility as loss of reachability. Independent of B, C, D.

Paper B: Selection as costly exclusion, if it exists. Depends on A. Independent of C, D.

Paper C: Agency as constrained control. Depends on A; uses outcome of B. Independent of D.

Paper D: Coupled viability under multi-agent constraint. Depends on A, B, C. Extends coupling (C7), introduces shared constraint environments, derives structural filtering, hierarchy, cooperation, and deterrence as geometric consequences.

The one-way dependency is preserved.

Failure of D does not invalidate C, B, or A.

Each layer adds structure. None adds physics.

Appendix E — Exploratory: Reclamation and Renewal (Non-Load-Bearing)

Paper A's optional module (A6) addresses capacity saturation and restoration for single systems. This appendix extends that to coupled systems: joint saturation, partial reclamation, the multi-agent loop.

It inherits the speculative status of A6. Explicitly non-load-bearing. No proposition in the main body depends on it. Included for structural completeness and intellectual honesty.

End of Paper D.

All stated proofs in this document follow from the definitions and assumptions declared locally. All propositions have specified observables and testable falsifiers. All conjectures are fenced. Papers 0, A, B, C, D

Series: The 420 Code Artist Proof 01 — The Physics of the Operator Medium: Foundational Physics / Viability Geometry

Artist: G STUDIO **G** Published for free forever Kill Switch Ledger

The following ledger maps every falsifiable claim in AP01 to the corpus-wide kill switch numbering system. Each kill switch has a unique identifier (KS-N), a status, and a specified observable.

Statuses: CLOSED (proven within the argument), LIVE-EMPIRICAL (testable by experiment), LIVE-HARD (open theoretical problem).

All kill switches in AP01 are LIVE-EMPIRICAL in principle: each has an operational observable and becomes directly testable once a concrete instantiation (physical or engineered system) is specified. Paper A Kill Switches

KS-V.1 (F0) — AS operational invariance. Global kill switch. If co-admissible coarse-grainings yield incompatible AS values beyond tolerance, the entire framework fails. Status: LIVE-EMPIRICAL. Test: R0.

KS-V.2 (F1) — Pointer-basis targeting. If selection targets position rather than the environment-selected pointer algebra, the selection postulate fails. Status: LIVE-EMPIRICAL. Test: R1.

KS-V.3 (F2) — Born violation. If ensemble statistics of realised branches deviate systematically from the diagonal weights $\{p_i\}$, the selection postulate fails. Status: LIVE-EMPIRICAL. Test: R4.

KS-V.4 (F3) — Context dependence. If selection depends on observer intervention rather than objective dynamics, the selection postulate fails. Status: LIVE-EMPIRICAL. Test: R5.

KS-V.5 (G1) — Selection rate exceeds gravitational bound. If selection occurs faster than $\hbar/\Delta E_G$ for gravitationally distinguishable records, the gravity limiter fails. Status: LIVE-EMPIRICAL. Test: R3.

KS-V.6 (G2) — Selection in gravitationally degenerate regime. If objective selection occurs between records with $\Delta E_G = 0$, the gravity limiter fails. Status: LIVE-EMPIRICAL. Test: R2.

KS-V.7 (G3) — Non-gravitational rate scaling. If selection rates scale universally with non-gravitational parameters across macroscopic records, the gravity limiter fails. Status: LIVE-EMPIRICAL. Test: R3. Paper B Kill Switches

KS-V.8 (B2) — Pre-irreversibility selection. If exclusion signatures appear before operational irreversibility is established (Paper A, D13), selection as defined in Paper B is falsified. Status: LIVE-EMPIRICAL. Test: BT1. Paper C Kill Switches

KS-V.9 (FC1) — Agency increase without control. If reachable viability volume increases without corresponding control expenditure, Paper C is falsified. Status: LIVE-EMPIRICAL. Test: C10.1.

KS-V.10 (FC2) — Irreversible loss reversed. If irreversible loss of reachability is reversed without external intervention violating admissibility constraints, Paper C is falsified. Status: LIVE-EMPIRICAL. Test: C10.1.

KS-V.11 (FC3) — Stable control past no-return surface. If stable control persists beyond the no-return surface under admissible control, Paper C is falsified. Status: LIVE-EMPIRICAL. Test: C10.1.

KS-V.12 (FC4) — Free lunch. If a system maintains positive agency indefinitely with finite budget and persistent nonzero drift, the survival time bound (Theorem C5.1) is falsified. Status: LIVE-EMPIRICAL. Test: C10.1.

KS-V.13 (FC5) — Resurrection. If a system recovers positive agency after reaching ruin without inadmissible external intervention, Paper C is falsified. Status: LIVE-EMPIRICAL. Test: C10.1. Paper D Kill Switches

KS-V.14 (FD0) — Multi-agent persistence violating all necessary conditions. If a multi-agent system persists indefinitely while violating all necessary conditions N1–N4 under monotonic-alignment and regularity assumptions, Paper D is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.15 (FD1) — No free survival. If Agent B increases agency despite negative record externality from Agent A, without severing coupling, increasing control budget, or receiving compensating externalities, the Geometric Exclusion Principle is falsified.

Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.16 (FD2.1) — Additivity under coupling. If joint agency equals sum of individual agencies in a non-trivially coupled system, non-additivity (Proposition D2.1) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.17 (FD2.2) — Impedance-independent efficiency. If coupling efficiency does not degrade as impedance ratio deviates from unity, impedance matching (Proposition D2.2) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.18 (FD2.3) — Anti-resonant optimality. If a persistent coupled system has maximum joint viability at anti-resonant phase under standard coupling, the resonance conjecture (D2.3) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.19 (FD3.3) — Cascade non-propagation. If failure of Agent A does not propagate to Agent B despite coupling that exceeds B's remaining control margin, cascade failure (Proposition D3.3) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.20 (FD4) — Order indistinguishable from noise. If persistent configurations cannot be statistically distinguished from random survivors ($p \geq 0.05$ for slack correlation), emergent order (D4) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.21 (FD4.2) — Persistent violators. If any configuration persists while violating any of N1–N4 under alignment and regularity assumptions, structural filtering (Proposition D4.2) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.22 (FD4.3) — Hierarchy inversion. If the low-impedance agent's record externalities dominate the high-impedance agent's constraint landscape (at impedance ratio greater than 10:1), hierarchy as constraint geometry (Proposition D4.3) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.23 (FD4.4a) — Cooperation nonexistence. If joint agency never exceeds sum of individual agencies in any system with mutually positive externalities, cooperation as structural outcome (Proposition D4.4a) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2.

KS-V.24 (FD4.4b) — Deterrence exit. If an agent in a deterrence equilibrium can unilaterally decouple and increase its agency, the deterrence characterisation (Proposition D4.4b) is falsified. Status: LIVE-EMPIRICAL. Test: D5.2. Summary

Total kill switches: 24 (KS-V.1 through KS-V.24). All LIVE-EMPIRICAL. Global kill switch: KS-V.1 (F0). If KS-V.1 fires, the entire framework is dead and no further test is meaningful.

Kill switch numbering begins at KS-V.1 to avoid conflict with existing corpus assignments (KS-1 through KS-49, assigned in AP05–AP22).

A corpus-wide renumbering will assign final numbers once all 22 Artist Proofs have been reviewed. Conditionality Footer

Conditional on: Nothing external. AP01 is self-contained. It depends only on standard quantum mechanics (unitary evolution, CPTP maps, decoherence) and viability theory (Aubin, 1991).

No result in AP01 depends on the axiom system {S, B, R, C}, on the Embedding Hypothesis (EH), on the Quadratic Regularity Assumption (QRA), or on any other Artist Proof.

Conditioned upon by: Subsequent Artist Proofs may inherit the operational definitions, irreversibility results, and viability geometry established here.

Load-bearing inheritance: Actualization State (D3), monotonicity under decohering record-forming dynamics (T1, within scope), Operator Horizon / no-return structure (T2; D9/D13).

Optional inheritance (explicitly postulate-level here): selection channel (A4.2) and gravitational rate limiter (A4.3) are referenced only where later proofs explicitly require them.

Kill switches: KS-V.1 through KS-V.24 (all LIVE-EMPIRICAL). See Kill Switch Ledger above.

Status: Publication-ready. Locked.

Last Page

Series	The 420 Code
Volume	Notebook Ø.1 The Premise
Edition	Artist Proof 01
Title	The Actualization State
Artist	G
Date	February 2026

This work is published for free forever at www.the420code.org.

No paywall. No institutional affiliation. No gatekeepers.

The spine stands or falls on its own logic. Your job is to test it.

Edition 02 will present the same spine in prose.

Edition 03 will present the same spine in metaphor.

Edition 04 will present the same spine in rhyme.

One spine. Four languages. Same architecture. Same truth.

STUDIO G