Quantum Convergence Threshold (QCT): A First-Principles Framework for Informational Collapse

Author: Gregory P. Capanda
ORCID: 0009-0002-0475-0362
Affiliation: Independent Researcher, Michigan, USA
Journal: Advances in Theoretical and Computational Physics

Abstract

The Quantum Convergence Threshold (QCT) framework presents a deterministic, first-principles model in which wavefunction collapse arises as an informational convergence process rather than a stochastic or observer-dependent event. Collapse occurs when a system's internal informational flux $\Lambda(x,t)$ and temporal resolution Δt exceed its coherence-sustaining capacity Ω, modulated by a dimensionless awareness-threshold function $\Theta(t)$. The formal criterion $\Theta(t)\cdot\Delta t\cdot\Lambda/\Omega \geq 1$ defines the transition point at which quantum superposition becomes informationally unsustainable, forcing a deterministic convergence to a single outcome. This replaces probabilistic postulates with a continuous, measurable buildup of informational density that can, in principle, be monitored experimentally.

QCT introduces explicit, physically motivated terms — informational flux density (Λ), coherence pressure (Ω), and the awareness threshold function (Θ) — into the Schrödinger framework, producing a modified evolution equation that preserves unitarity until convergence becomes unavoidable. The theory predicts measurable threshold signatures in interferometric visibility, hysteresis effects under repeated weak measurement, and scaling relationships between collapse time, decoherence rate, and informational load.

Beyond addressing the measurement problem, QCT connects quantum foundations with quantum information theory and thermodynamics, suggesting a unified informational ontology for the transition from quantum potentiality to classical actuality. Its implications extend to quantum computing, decoherence control, and the emergence of classicality itself, providing a rigorous platform for experimental falsification and theoretical expansion.

1. Introduction

Quantum theory provides an extraordinarily accurate dynamical law for isolated systems — unitary evolution under the Schrödinger equation — yet standard formulations still treat the emergence of definite outcomes as an extra postulate. In practice, we compute amplitudes continuously and then, at the moment of "measurement," we replace superpositions by outcome probabilities according to the Born rule. This split leaves the measurement problem unresolved: what, physically, ends unitary evolution and selects one actual result rather than another?

Deterministic alternatives have historically pursued two main paths. Pilot-wave (Bohmian) models introduce trajectories guided by a nonlocal wave, preserving determinism but at the cost of added ontology and delicate consistency conditions. Objective-collapse models modify the dynamics stochastically (e.g., GRW/CSL), making collapse an intrinsic process but sacrificing determinism and often straining energy conservation or no-signaling. Decoherence theory explains why interference terms become practically unobservable in open systems, but by itself does not explain selection of a single outcome. Many-worlds removes collapse entirely by taking all branches as real, thereby relocating the explanatory burden to probability and ontology.

The Quantum Convergence Threshold (QCT) framework pursues a different route. It treats collapse as a deterministic, threshold-crossing transition driven by informational load that a system accumulates through interaction, internal structure, and entanglement. In QCT, superposition persists until the system's informational flux and temporal resolution outstrip its capacity to sustain coherent alternatives. At that point, the state deterministically resolves to a single branch. The central condition is expressed by a dimensionless collapse index $C(x,t) = \Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega$, and the collapse criterion $C(x,t) \geq 1$. Here $\Lambda(x,t)$ ($\text{bits}\cdot\text{s}^{-1}\cdot\text{m}^{-3}$) is the informational flux density registered by the system, Δt (s) is the temporal resolution window over which coherence must be maintained, Ω ($\text{bits}\cdot\text{s}^{-1}$) is the effective coherence-sustaining capacity (the opposing "pressure" of coherence against decohering load), and $\Theta(t)$ is a dimensionless threshold function capturing system-relative sensitivity to informational overload. In words: when the time-integrated informational load, discounted by the system's coherence capacity and scaled by $\Theta(t)$, exceeds unity, collapse is no longer avoidable.

To connect this threshold picture with standard dynamics, we introduce a minimal coupling to the Schrödinger equation that preserves unitarity away from threshold while allowing deterministic attenuation as convergence builds: $i\hbar\, \partial\psi/\partial t = H\psi - i\hbar\, (\Lambda/\Omega)\, \psi$. The ratio Λ/Ω acts as an internal convergence rate. For $C < 1$ the additional term is negligible and evolution is effectively unitary; near $C \approx 1$ the attenuation becomes appreciable and the state enters a short, deterministic convergence epoch ending in a single outcome. This modification is not a stochastic noise term and does not introduce random jumps; it is a structural, state-dependent drain that reflects the system's informational burden relative to its coherence capacity.

QCT differs from prior approaches in four ways. First, it is deterministic: given Λ, Ω, Δt, and $\Theta(t)$, the onset and timing of collapse follow lawfully, without randomness or hidden trajectories. Second, it is observer-independent: collapse is driven by system-internal informational conditions, not by external observation or consciousness. (Extensions that study conscious modulation via the Quantum Zeno Effect can be layered on top, but the core mechanism does not require it.) Third, it is operational: Λ, Ω, and Δt are tied to measurable properties of interferometers, qubits, and open-system couplings; $\Theta(t)$ is system-relative but can be calibrated empirically. Fourth, it is falsifiable: QCT predicts threshold and hysteresis signatures (for example, sharp visibility knees and path-dependent recovery) that differ from smooth decoherence curves.

This paper develops QCT as a first-principles framework suitable for theoretical analysis and experimental test. Section 2 defines all quantities and physical units used throughout (Λ, Ω, Δt, $\Theta(t)$, C). Section 3 states the logical axioms that ground the model and motivate the threshold form. Section 4 presents the coupled evolution law and shows how the convergence term can be cast in a completely positive, trace-preserving open-system form away from threshold, with deterministic selection at threshold. Section 5 proposes concrete experimental protocols — single-system interferometry with tunable weak which-path taps and engineered dephasing; multipartite entanglement tests with independently controllable Λ and Ω; and gate-level diagnostics on superconducting platforms. Section 6 analyzes theoretical implications for the measurement problem, preferred basis emergence, energy and probability conservation, no-signaling, and the relationship to decoherence and Lindblad dynamics. Section 7 details specific protocol-level predictions (threshold surfaces, hysteresis under repeated weak measurement, and collapse-time scaling), including what would count as a failed prediction. Section 8 concludes and outlines next steps, including extensions to networked thresholds for multipartite systems and calibration strategies for $\Theta(t)$.

Two clarifications frame the scope. First, the “awareness threshold” terminology for $\Theta(t)$ reflects sensitivity to informational overload; it does not smuggle in observers or mental entities. $\Theta(t)$ is a dimensionless response factor that can, in principle, be inferred from system identification and calibration (device complexity, spectral bandwidth of couplings, error-correction overheads). Second, while the present paper focuses on observer-independent collapse, nothing precludes studying how deliberate, repeated interrogation (internal or external) modulates Λ and Δt, thereby delaying or steering collapse (a Zeno-like regime). Those extensions remain optional layers on the same base equation.

In short: QCT reframes collapse as an informational capacity crossing. It preserves the accuracy of unitary quantum dynamics in low-load regimes, explains when and why definiteness emerges without invoking randomness, and yields testable signatures that distinguish threshold-driven convergence from ordinary decoherence.

2. Defined Terms and Physical Units

The Quantum Convergence Threshold (QCT) framework formalizes collapse in terms of measurable, physically motivated quantities. Each represents a distinct aspect of informational dynamics within a quantum system as it transitions from coherent superposition to classical definiteness. This section defines these quantities, their physical dimensions, and their interpretive roles.

2.1 Informational Flux Density — $\Lambda(x,t)$

Unit: bits·s^{-1}·m^{-3}
Definition: Λ(x,t) represents the local rate at which information is registered, exchanged, or encoded within a system's state space per unit volume and per unit time. It quantifies the internal and environmental flow of informational change experienced by the quantum system.

In a physical sense, Λ captures how rapidly distinguishable alternatives in the wavefunction become encoded in the system's degrees of freedom or its environment. This includes both internal correlations (entanglement growth within the system) and external interactions (information leakage into the environment).

Formally, Λ(x,t) may be expressed as a function of the system's reduced density matrix ρ(x,t):

$$\Lambda(x,t) = (d/dt)\ \mathrm{Tr}[\rho(x,t) \log \rho(x,t)^{-1}]$$

This links Λ to the system's information rate, consistent with Shannon information and von Neumann entropy formulations. High Λ indicates rapid information flow and strong coupling; low Λ corresponds to isolated, slow-evolving systems.

2.2 Awareness Threshold Function — Θ(t)

Unit: dimensionless
Definition: Θ(t) quantifies a system's inherent sensitivity to informational overload — its capacity to sustain superposition under increasing informational strain. It is an emergent property dependent on coherence bandwidth, internal entanglement complexity, and structural organization.

Θ(t) acts as a modulating factor between coherence retention and convergence pressure. Systems with high Θ(t) tolerate larger informational flux before collapse; systems with low Θ(t) collapse under smaller loads.

Operationally, Θ(t) can be interpreted as a normalized measure of informational resilience, analogous to a response coefficient in thermodynamic systems. While dimensionless, its calibration can be empirically constrained through experiments measuring coherence duration and interference visibility under controlled Λ and Ω variations.

2.3 Temporal Resolution — Δt

Unit: seconds (s)

Definition: Δt is the characteristic temporal window over which the system maintains phase coherence and within which collapse potential accumulates. It represents the integration period for informational convergence.

In open-system terms, Δt parallels the decoherence time — the interval during which superposition persists before external interactions enforce classicality. However, in QCT, Δt also governs the internal sampling rate of informational change; it sets the clock for when informational flux Λ contributes meaningfully to convergence buildup.

Short Δt values correspond to high-frequency systems where coherence is rapidly renewed; long Δt values indicate systems with extended coherent lifetimes, such as isolated superconducting qubits or atomic interferometers.

2.4 Coherence Pressure — Ω

Unit: bits$\cdot$s^{-1}
Definition: Ω measures a system's capacity to resist decoherence — the effective rate at which it maintains or restores coherence in the face of informational disturbance. It acts as the denominator in the QCT collapse ratio, opposing Λ.

Physically, Ω can be interpreted as the rate of coherence restoration via self-organizing quantum feedback, shielding, or error correction. In open-system models, Ω relates to the system's internal Hamiltonian structure and environmental coupling constants, defining how much informational load the system can sustain without breakdown.

Systems with high Ω (e.g., strongly error-corrected qubits) can maintain superposition longer under heavy informational flux. Low-Ω systems (e.g., macroscopic measurement devices) collapse rapidly due to weak internal coherence resilience.

2.5 Collapse Index — $C(x,t)$

Unit: dimensionless
Definition: The Collapse Index $C(x,t)$ is the dimensionless ratio that quantifies how close a system is to collapse:

$$C(x,t) = \Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega$$

When $C(x,t) < 1$, superposition persists; when $C(x,t) \geq 1$, the system undergoes deterministic convergence into a single outcome.

$C(x,t)$ thus provides a continuous measure of collapse proximity. Unlike probabilistic collapse models, QCT treats $C(x,t)$ as a dynamic state variable that evolves in time with the system's informational structure.

2.6 Collapse Criterion and Threshold Condition

The QCT threshold condition is the central predictive equation of the framework:

$$\Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega \geq 1$$

This expression defines the precise boundary between sustained superposition and deterministic collapse. It establishes a physically grounded convergence law linking information flux, coherence capacity, and temporal evolution.

This formulation provides an experimentally accessible means to test collapse dynamics by modulating Λ and Ω while monitoring coherence observables such as interference visibility or entanglement fidelity.

2.7 Summary of Dimensional Consistency

Each quantity defined in the Quantum Convergence Threshold (QCT) framework has a clear physical meaning and consistent dimensional structure. The informational flux density $\Lambda(x,t)$ carries units of bits per second per cubic meter ($\text{bit}\cdot\text{s}^{-1}\cdot\text{m}^{-3}$) and represents the rate of information encoding or flow within a system. The awareness threshold function $\Theta(t)$ is dimensionless; it expresses the system's relative sensitivity to informational overload and functions as a scaling factor. The temporal resolution Δt, measured in seconds, denotes the time interval over which coherence or collapse is evaluated. The coherence pressure Ω, with units of bits per second ($\text{bit}\cdot\text{s}^{-1}$), quantifies the system's capacity to sustain coherence against informational disturbance. Finally, the collapse index $C(x,t)$ is dimensionless, representing the instantaneous proximity of the system to collapse.

When these quantities are combined, they yield a dimensionless expression:

$$C(x,t) = \Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega.$$

The collapse criterion $\Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega \geq 1$ therefore preserves dimensional coherence and provides a physically interpretable threshold condition. Λ drives convergence, Ω resists it, Δt integrates it over time, $\Theta(t)$ modulates it according to system sensitivity, and $C(x,t)$ measures the resulting proximity to collapse. Together these parameters define a deterministic, law-like transition boundary between superposition and classical definiteness.

3. Logical Foundation and First Principles

This section states the core assumptions of the Quantum Convergence Threshold (QCT) framework, clarifies how they differ from standard interpretive add-ons, and shows how they lead to a falsifiable collapse condition. The aim is to keep the ontology minimal, the assumptions explicit, and the operational content testable.

3.1 Guiding Commitments

1. Ontological minimalism. QCT does not add hidden trajectories, extra worlds, or observer-dependent rules. The quantum state encodes physical dispositions of a system to register distinguishable alternatives over time.

2. Collapse as convergence. Collapse is not external, stochastic, or observer-imposed. It is the lawful end-state of an internal information-balance: when informational loading outpaces coherence-sustaining capacity over a finite window.

3. Determinism up to threshold. Between collapses, dynamics follow standard unitary or open-system evolution. The threshold condition specifies when this evolution can no longer maintain superposition.

4. Operational measurability. Every quantity in the collapse condition has a concrete calibration pathway: Λ via state-distinguishability growth rates, Ω via coherence lifetimes under controlled disturbance, Δt via the experimental integration window, and Θ via system identification (how much flux the system tolerates before interference loss).

3.2 First Principles (Axioms)

Axiom A — Sufficient Definition. A quantum system actualizes a definite outcome only when it has accumulated enough registered information, within a finite window Δt, to define that outcome against alternatives.

Axiom B — Internal Resolution. Collapse does not require an external observer. It is triggered when internal informational convergence outruns coherence support.

Axiom C — Threshold Convergence. There exists a dimensionless collapse index $C(x,t) = \Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega$. When $C(x,t) \geq 1$, the system undergoes deterministic convergence to a single outcome in the relevant basis.

These axioms are deliberately modest: they do not presuppose consciousness, a preferred ontology of particles or fields, or any pilot structure. They only assume that (i) distinguishability accumulates, (ii) coherence is finite, and (iii) there is a finite window over which the competition between the two decides the phase of evolution.

3.3 Collapse Criterion (Restated)

Collapse occurs when
$\Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega \geq 1$.

Interpretation: Λ contributes "pressure to distinguish," Ω contributes "capacity to cohere," Δt is the accumulation window, and Θ is the system's sensitivity coefficient. When the ratio crosses unity, the superposition can no longer be sustained and the state converges to a single outcome.

3.4 Relationship to Standard Dynamics

3.4.1 Closed systems. For perfectly isolated systems, Λ is small and Ω is effectively large, so $C(x,t) \ll 1$ and unitary evolution persists.

3.4.2 Open systems. In realistic settings, dynamics are described by a master equation. QCT is compatible with completely positive trace-preserving (CPTP) evolution below threshold. Collapse corresponds to crossing an operational boundary where the map ceases to be representable as coherent mixing over distinguishable alternatives; the system enters a single-outcome sector.

3.4.3 No contradiction with Wigner's theorem. Wigner's theorem concerns symmetry transformations for closed, purely unitary dynamics. QCT places collapse precisely in the open, thresholded regime where unitary symmetry is not the full description. Below threshold, probabilities are conserved by CPTP dynamics; threshold crossing selects one sector, after which CPTP evolution resumes with the updated state.

3.5 How QCT Avoids the "Observer" Split

The "observer" in QCT is any physical subsystem that registers information and thereby contributes to Λ. No special ontological status is given to human observers or measurement devices. A photodiode, a phonon bath, or an ancilla register contributes to Λ exactly as a person reading a dial would. The same law applies across scales, eliminating the observer–system dualism.

3.6 Preferred Basis by Einselection, Not Postulate

QCT does not stipulate a basis. The effective basis in which outcomes become robust is the one that diagonalizes the system–environment interaction over the window Δt (environment-induced superselection, or einselection). Λ is computed with respect to the rate at which distinguishability in that basis is registered. Thus, the "preferred" basis is emergent from the dynamics that determine Λ and Ω, not inserted by fiat.

3.7 Sketch: Consistency with Probability and Energy Accounting

Probability. Below threshold, evolution is CPTP and probability is conserved. At threshold, a single outcome sector is selected. Post-collapse normalization is standard. Any experimental test of QCT must verify that observed frequencies across many runs match the calibrated mapping from pre-threshold distinguishability to selected outcomes (see Section 6 for the statistical law).

Energy. QCT does not impose an extra Hamiltonian that injects or removes energy. Ω and Λ summarize informational roles of couplings already present in the open-system Hamiltonian and its Lindblad operators. Collapse is the entry into a single outcome sector supported by those same couplings. Energy accounting therefore follows the underlying open-system model; QCT adds a decision surface, not an energy source.

3.8 What $\Theta(t)$ Is — and Is Not

$\Theta(t)$ is not a mystical "awareness constant." It is a dimensionless sensitivity coefficient that reflects how the identified system, with its partition and encoding scheme, tolerates informational flux before interference becomes non-recoverable over Δt. In practice, $\Theta(t)$ is calibrated: prepare the system, vary Λ and Ω under fixed Δt, and empirically mark the knee where visibility or entanglement fidelity shows irreversible loss. That knee fixes $\Theta(t)$ for that partition and protocol.

Because Θ depends on system identification, QCT is explicit about partitions. Change the partition (different encoding, different coarse-graining), and you may change Θ — a feature, not a bug, since interference robustness is known to depend on how information is encoded and where it leaks.

3.9 Why a Threshold Law (and Not "Just" Decoherence)

Standard decoherence explains how interference terms diminish but does not provide a physically grounded, experiment-controllable boundary where a single outcome becomes necessary. QCT adds that boundary. The law $\Theta \cdot \Delta t \cdot \Lambda / \Omega \geq 1$ is the missing criterion that turns gradual suppression into an operational phase change. The distinction is falsifiable: QCT predicts a capacity-linked knee and hysteresis near the knee (Section 7), whereas decoherence-only predicts smooth monotone decay under the same integrated exposure.

3.10 Falsifiability Commitments

1. Knee and hysteresis. If controlled sweeps of Λ and Ω, under fixed Δt and calibrated Θ, never produce a sharp, protocol-invariant knee and the predicted hysteresis loop, QCT is falsified.

2. Threshold-time scaling. Near threshold, the minimum collapse time must scale inversely with Λ/Ω at fixed Θ and Δt. If experiments systematically violate this scaling, QCT is falsified.

3. Multipartite propagation. In entangled systems, QCT predicts structured co-collapse constrained by the shared registration pathway, not arbitrary superluminal effects. If experiments reveal signaling beyond no-signaling bounds under QCT-compatible calibrations, the framework is falsified.

3.11 Summary

From three modest axioms — sufficient definition, internal resolution, threshold convergence — QCT yields a concrete, dimensionless collapse law: $\Theta(t) \cdot \Delta t \cdot \Lambda(x,t) / \Omega \geq 1$. The quantities are calibratable, the basis is emergent by einselection, and the predictions differ from decoherence-only in experimentally testable ways. The logical stance is conservative (no new ontology) yet stronger than interpretation: it provides a threshold surface in parameter space that can, in principle, be confirmed or ruled out.

4. Mathematical Framework for Thresholded Dynamics

4.1 Pre-threshold CPTP evolution

Let $\rho(t)$ satisfy $d\rho/dt = -i[H, \rho] + \Sigma_j (L_j \rho L_j^\dagger - 0.5 L_j^\dagger L_j \rho - 0.5 \rho L_j^\dagger L_j)$, with L_j aligned to the pointer basis and rates consistent with Λ and Ω. Encodings and feedback renormalize rates via Θ.

4.2 Informational bookkeeping and the index C

$C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$ collects (i) how fast bits leak (Λ), (ii) how effectively they become records (Ω), (iii) for how long they accumulate (Δt), and (iv) how hard recording is made by protection (Θ). C is dimensionless; doubling Λ or Θ doubles C; doubling Ω halves C; doubling Δt doubles C.

4.3 Onset mechanics via hazard rates

Define a hazard $h(t) = h0 \cdot s(C(t))$, where $s(C)$ is a steep sigmoid with midpoint at C^*. The onset probability in $[t, t + dt]$ is $h(t)\, dt$. This yields negligible hazard when $C \ll C^*$, a sharp knee near $C \approx C^*$, and near-certain onset shortly after crossing.

4.4 Outcome selection and post-onset state

At onset, select outcomes with $p_i = Tr(\Pi_i \rho_-)$. Set $\rho_+ = \Sigma_i p_i \Pi_i$. This recovers Born weights without extra stochastic postulates.

4.5 Energy and unitarity bounds

If $[H, L_j] = 0$ for dominant channels, $d/dt\ Tr(H \rho) = 0$ pre-threshold. With dissipative components, bound $|dE/dt| \leq \Sigma_j \gamma_j \|[H, L_j]\|$. Report measured drift and controller power used to raise Θ.

4.6 Relation to decoherence-only accounts

Pure decoherence predicts smooth visibility decay with integrated exposure. QCT predicts an additional capacity-linked knee at $C \approx C^*$, hysteresis under cumulative weak taps, and a re-coherence window just below onset that decoherence-only cannot produce once redundancy is supercritical.

5. Experimental Program and Falsification Strategy

This section lays out concrete, device-level tests that can distinguish QCT from decoherence-only accounts and from hidden-variable or many-worlds interpretations. Each protocol specifies the platform, control knobs for the four primitives (Θ, Λ, Ω, Δt), predicted signatures (including "capacity-linked knees," hysteresis, and subcritical re-coherence), and failure modes that would falsify QCT.

5.1 Single-system threshold interferometry

Platforms. Mach–Zehnder or Sagnac photonic interferometers; Talbot–Lau and Kapitza–Dirac matter-wave interferometers; superconducting or NV-center Ramsey interferometers.

Controls.

Λ (informational flux): weak which-path taps (e.g., partial polarizers, low-gain APDs, dispersive cavity probes), tunable tap strength and repetition.

Ω (coherence pressure from the environment): engineered phase noise, temperature, vibrational spectrum, flux noise; independently calibrated dephasing channels.

Δt (coherence window): arm-length or idle-time delay; pulse spacing between weak probes.

Θ (protection factor): dynamical decoupling, error-detect codewords, decoherence-free subspaces; also passive shielding and mode filtering.

Protocol.

1. Establish a high-visibility baseline V_0 at minimal Λ and Ω.

2. Sweep Λ at fixed Ω and Δt; record visibility $V(\Lambda)$.

3. Sweep Ω at fixed Λ and Δt; record $V(\Omega)$.

4. For each configuration, compute $C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$ from independently calibrated Θ, Λ, Ω, Δt.

5. Repeat with cumulative weak-measurement sequences to test rate-independent accumulation.

QCT predictions.

Threshold surface: visibility remains $\approx V_0$ while $C < 1$ (up to smooth dephasing captured in Ω), then exhibits a sharp knee and irreversible drop once C crosses about 1.

Hysteresis under cumulative taps: repeated weak probes increase effective Λ and accumulate into C; fringe loss occurs abruptly just beyond the crossing. Removing the final probe after crossing does not immediately restore V.

Collapse-time scaling: near threshold, the minimal collapse time scales as $\tau_c \approx (\Theta - \Theta_0) / (\Lambda/\Omega)$, holding Δt fixed.

Subcritical re-coherence: if Λ is reduced to push C back to slightly below 1 before full redundancy forms, partial visibility recovers; decoherence-only models predict no such recovery once redundancy is supercritical.

Falsification. Absence of a capacity-linked knee (only smooth exponential decay with integrated exposure), absence of hysteresis, or absence of subcritical re-coherence would disfavor QCT in favor of decoherence-only dynamics.

5.2 Multipartite entanglement: synchronized onset without signaling

Platforms. Polarization-entangled photon pairs (or GHZ states); trapped-ion Bell pairs; superconducting qubit pairs with tunable cross-resonance.

Controls.

Local Λ_A, Λ_B via weak local probes (dispersive readout, Faraday rotation, ancilla-assisted taps).

Local Ω_A, Ω_B via independent engineered dephasing.

Θ_A, Θ_B via local error-detect subspaces or dynamical decoupling schedules.

Space-like separation to enforce relativistic causal independence of settings.

Protocol.

1. Prepare a maximally entangled state in the agreed pointer basis.

2. On side A, raise $C_A = \Theta_A \cdot \Delta t_A \cdot \Lambda_A / \Omega_A$ through unity while keeping $C_B < 1$.

3. Record local marginal statistics and two-party correlations under space-like separation and randomized settings.

4. Swap roles: raise C_B through unity while keeping $C_A < 1$.

QCT predictions.

Synchronized resolution: when C_A crosses unity, both sides' joint state resolves in the shared pointer basis; local marginals remain maximally mixed until local readout.

No signaling: conditional and unconditional marginals on B are statistically independent of A's timing and settings.

Capacity-linked timing: the first side to cross $C \approx 1$ sets the effective onset; the partner's local "record" is fixed at readout without superluminal influence.

Falsification. Any superluminal signaling in marginals; or resolution on the partner side while both $C_A < 1$ and $C_B < 1$ would contradict QCT's thresholded onset.

5.3 Threshold-aware quantum error correction (QEC) on NISQ devices

Platforms. Superconducting transmons (surface codes, repetition codes); trapped ions (Bacon–Shor variants); NV centers with photonic links.

Controls.

Λ: syndrome-extraction frequency and strength; mid-circuit measurement depth.

Ω: engineered noise channels (amplitude, phase, crosstalk) calibrated via randomized benchmarking.

Θ: code distance d, number of stabilizers, dynamical decoupling overlay; logical subspace protection.

Δt: cycle time and idle windows.

Protocol.

1. For a fixed hardware noise Ω, scan (Λ, Θ, Δt) across operating points while monitoring logical error rate p_L and fringe-like coherence proxies.

2. Identify the composite index $C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$ at each operating point from independent calibrations.

3. Track abrupt changes in logical performance as a function of C.

QCT predictions.

Design knee: there exists a narrow band around $C \approx 1$ where the device transitions from recoverable (subcritical) to irrecoverable (supercritical) failure, beyond what standard threshold theorems predict solely from Ω.

Hysteretic failure: repeated weak syndrome extraction can push a code from the safe regime into an irrecoverable regime, even if each tap is below conventional back-action thresholds.

Re-coherence window: backing off Λ to bring C just below unity can restore logical performance if redundancy has not yet saturated.

Falsification. Logical performance depends only on Ω and standard code thresholds, with no capacity-linked knee in C and no hysteresis or re-coherence effects.

5.4 Weak-measurement accumulation and hysteresis

Platforms. Any of the above with high-fidelity, near-QND weak probes.

Controls.

Keep instantaneous Λ per tap small; vary number of taps N and spacing to control effective Δt.

Hold Ω stable; set Θ with or without protective encoding.

Protocol.

1. Apply sequences of N weak taps below individual back-action thresholds.

2. After each sequence, measure interference or state purity.

3. Vary N to sweep cumulative C upward while keeping per-tap disturbance fixed.

QCT predictions.

Accumulation: even with tiny per-tap Λ, cumulative exposure raises C and produces a sharp transition once C crosses ≈ 1.

Hysteresis: once past onset, removing the final tap does not restore coherence; only reducing cumulative exposure or increasing Θ (e.g., by code protection) and doing so before redundancy saturates can recover coherence.

Falsification. If visibility or purity follows only the integrated dephasing expected from Ω with no sharp capacity knee and no hysteresis, QCT's cumulative-threshold claim is falsified.

5.5 Subcritical re-coherence (“new dam upstream”)

Platforms. Same as 5.1–5.3, with programmable control over Λ and Θ in real time.

Protocol.

1. Approach $C \approx 1$ from below, as indicated by near-threshold visibility loss or logical error inflection.

2. Before full redundancy forms, reduce Λ or increase Θ to push C slightly below unity.

3. Observe partial recovery of visibility, purity, or logical performance.

QCT predictions.

Recoverable edge: near the boundary, reducing C below unity yields measurable re-coherence (visibility gain, purity increase, logical error decrease).

Decoherence-only models offer no such recovery once sufficient environmental records already exist.

Falsification. No recovery is observed in the subcritical regime when C is pulled back below 1 despite independent evidence that redundancy had not yet saturated.

5.6 No-signaling and relativistic consistency checks

For all bipartite tests in 5.2, enforce space-like separation and fast random basis choices. Record unconditional marginals and verify independence from the remote party’s settings and timing. QCT predicts strictly no superluminal signaling; any reproducible violation would falsify the framework.

5.7 Parameter extraction and open-data calibration

To convert raw controls into the index C with traceable uncertainties:

1. Calibrate Ω through independent T1/T2, randomized benchmarking, and noise spectroscopy; report Ω with confidence intervals.

2. Calibrate Λ by quantifying information leakage per probe (bits per second per cubic meter) via detector tomography or QND-back-action models tied to measured probe strength and bandwidth.

3. Estimate Θ from protection resources: code distance d, number of stabilizers, leakage rates, and decoupling efficacy; fit Θ as the single free parameter within physically constrained bounds across multiple datasets.

4. Set Δt from controlled idle times or path-length differences.

Model selection. Fit three nested models to the same datasets:
(i) pure decoherence (no threshold),
(ii) decoherence plus hard threshold at C = 1,
(iii) decoherence plus soft logistic activation s(C) with midpoint C*.
Use likelihood-ratio tests and information criteria to determine whether a threshold term is warranted; publish code and data.

5.8 Negative results and falsification criteria

QCT is falsified if, after careful calibration and model selection:

No threshold surface or knee is detected in regimes where C spans well below and above unity,

No hysteresis under cumulative weak measurement is observed,

No subcritical re-coherence occurs when C is pulled back below unity prior to redundancy saturation,

Any superluminal signaling is detected, or

A single parameter set (Θ) cannot consistently account for multiple platforms' knees while Λ, Ω, Δt are independently calibrated.

5.9 Near-term roadmap

Months 0–3: Implement Section 5.1 on an optical Mach–Zehnder with variable weak taps; open-source calibration of Λ and Ω; pre-register analyses.

Months 3–6: Replicate on superconducting Ramsey interferometry; integrate 5.4 accumulation sequences; begin 5.5 re-coherence tests.

Months 6–9: Run 5.2 bipartite tests with space-like separation; publish no-signaling analyses; introduce 5.3 threshold-aware QEC scans.

Months 9–12: Cross-platform meta-analysis and joint fit of Θ; deposit full datasets (controls, raw traces, calibration notebooks, and model-comparison scripts).

This program provides multiple, orthogonal opportunities to prove QCT wrong. If it survives these attempts, the resulting parameterized threshold law would be directly useful for quantum-technology design.

5.10 Power and sensitivity analysis

Define expected knee location C* and slope s'(C*) from pilot runs. For each platform, compute required sample sizes N per condition to detect (i) a visibility drop ΔV at $C \approx C^*$ with significance α and power 1−β, and (ii) hysteresis area A_{hys} in Λ–V trajectories. Use bootstrapped confidence intervals on fitted logistic models V(C) and compare against a nested no-threshold model via likelihood-ratio tests. Report achieved power and sensitivity to Θ misestimation (±10–20%).

5.11 Pre-registration and open science

Pre-register hypotheses (threshold knee at $C \approx 1$, presence of hysteresis and subcritical re-coherence), analysis plans (model comparison: no-threshold vs soft/hard threshold), and exclusion criteria. Release raw data, calibration notebooks for Λ, Ω, Θ, Δt, and code for model fitting and no-signaling checks under a permissive license. Include a replication package so others can recompute C and reproduce all figures from raw traces.

6. Theoretical Implications and Predictions

6.1 Measurement problem and outcome selection

QCT replaces “measurement” as a primitive with a thresholded transition driven by the index

$C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$. While $C < C^*$, the state $\rho(t)$ follows completely positive, trace-preserving (CPTP) dynamics generated by H and $\{L_j\}$. When C reaches the calibrated knee $C^* \approx 1$, an onset hazard h(t) becomes order-one and a single outcome is actualized in the pointer basis $\{\Pi_i\}$. Outcome weights are $p_i = \mathrm{Tr}(\Pi_i \rho_-)$, where ρ_- is the pre-onset state. The post-onset state is $\rho_+ = \Sigma_i p_i \Pi_i$. Thus, selection is not imposed by an observer but by internal informational saturation; probabilities are the Born weights of the state that self-drove C to the threshold.

Prediction P1 (capacity-linked knee). Interference visibility or coherence measures will exhibit a sharp, instrument-calibrated knee at $C \approx C^*$ rather than a purely smooth, exposure-only decay. The knee position shifts in the direction predicted by controlled changes to Θ, Λ, Ω, or Δt.

6.2 Preferred basis and einselection

QCT does not posit a basis by fiat. The pre-threshold Lindblad sector defines pointer projectors $\{\Pi_i\}$ via the dominant environmental and instrument couplings $\{L_j\}$. The same channels that redundantly record information into the environment also maximize Λ/Ω along those directions, pushing C upward fastest in the einselected basis. Consequently, the basis in which collapse occurs is the one already rendered stable by open-system dynamics.

Prediction P2 (basis consistency). The basis that maximizes redundancy (as measured by mutual information with the environment) is the basis in which the knee appears first. Rotating the effective pointer basis (e.g., by engineered dissipation or basis-tracking control) moves the knee accordingly.

6.3 Born rule without ad hoc postulates

Because outcome weights are taken directly from the pre-onset state, $p_i = \mathrm{Tr}(\Pi_i \rho_-)$, QCT does not "insert" the Born rule after the fact. The threshold only governs when sampling occurs; the sampling law is the one already encoded in ρ_-. In repeated-trial ensembles with matched C-trajectories, frequency statistics converge to the Born weights of the pre-onset state.

Prediction P3 (frequency recovery under matched C). Across matched runs with identical ρ_- and C-history up to onset, empirical frequencies of outcomes converge to $\mathrm{Tr}(\Pi_i \rho_-)$ within standard statistical bounds. Drift in frequencies tracks measured drift in the calibrated C-trajectory, not hidden variables.

6.4 Unitarity, CPTP evolution, and energy accounting

Pre-threshold dynamics are CPTP: $d\rho/dt = -i [H, \rho] + \Sigma_j (L_j \rho L_j^\dagger - 0.5 L_j^\dagger L_j \rho - 0.5 \rho L_j^\dagger L_j)$. If $[H, L_j] = 0$ for dominant channels, $\mathrm{Tr}(H \rho)$ is constant; when dissipation is present, energy drift is bounded by the commutator norms and appears in power-budget telemetry (e.g., controller power used to raise Θ). The onset map $\rho_- \rightarrow \rho_+ = \Sigma_i p_i \Pi_i$ is completely positive and trace-preserving. Thus, probability is conserved and energy changes are attributable to explicit couplings, not to the collapse rule itself.

Prediction P4 (energy-drift bounds). Any observed energy drift around onset is captured by measured system–bath/controller couplings. After subtracting those, residual drift falls within pre-declared bounds tied to $\|[H, L_j]\|$.

6.5 No-signaling in multipartite systems

For a bipartite state ρ_{AB} with local pointer projectors $\{\Pi_i^A\}$ and $\{\Pi_k^B\}$, QCT requires that the reduced post-onset state on B be $\rho_{B+} = \mathrm{Tr}_A(\Sigma_i \Pi_i^A \rho_{AB-} \Pi_i^A) = \Sigma_i \mathrm{Tr}_A(\Pi_i^A \rho_{AB-} \Pi_i^A)$. Because Alice's choice of local instrument only changes her $\{L_j^A, \Pi_i^A\}$ and hence her local C_A trajectory, the average (no-postselection) marginal on B is unchanged by Alice's setting. Hence, no controllable superluminal signaling arises. Apparent "synchronous" resolution across spacelike-separated subsystems reflects shared pre-onset correlations and matching pointer structures, not transmitted influence.

Prediction P5 (no-signaling checks). Protocols that vary Λ_A, Ω_A, or Θ_A at site A while monitoring unconditional marginals at site B yield invariance within experimental error. Any conditional differences vanish when postselection is removed.

6.6 Scaling from microscopic to macroscopic

QCT predicts a smooth micro-to-macro crossover governed by Λ, Ω, Θ, Δt:

Microscopic, well-isolated systems: small Λ and large Θ keep C subcritical; long-lived superpositions persist.

Mesoscopic systems with controlled protection: C can be held near but below C^*, enabling prolonged coherence and subcritical re-coherence when Λ is reduced or Θ is increased before onset.

Macroscopic, redundant-record systems: large Λ/Ω drives rapid crossing of C^*, yielding effectively instantaneous classical outcomes.

Prediction P6 (mesoscopic window). In engineered platforms (trapped ions, superconducting qubits, NV centers), a tunable window exists where C can be parked just below C^*, producing (i) extended coherence beyond decoherence-only forecasts and (ii) reversible re-coherence upon reducing Λ or raising Θ before onset.

6.7 Hysteresis and memory effects near threshold

Because the index C integrates over exposure and protection history, near-threshold behavior is path-dependent. Repeated weak which-path taps accumulate Λ and can push C minutely above C^*, producing an abrupt collapse; removing the last weak tap after crossing does not immediately restore coherence. Conversely, carefully timed increases of Θ before C crosses can pull the system back below threshold ("build a new dam upstream").

Prediction P7 (hysteresis loop). Visibility-versus-C trajectories exhibit a clockwise loop under weak-tap accumulation and retraction: the collapse point on the up-sweep occurs at slightly lower nominal C than the recovery point on the down-sweep. Pure decoherence without an explicit threshold lacks this loop.

6.8 Relation to decoherence-only and objective-collapse models

Versus decoherence only. Decoherence explains basis selection and smooth visibility loss with exposure, but it does not single out a unique outcome nor predict a capacity-linked knee, hysteresis, or subcritical re-coherence. QCT adds these threshold phenomena on top of the standard open-system picture.

Versus GRW/CSL. GRW-type models introduce stochastic collapses with new fundamental constants. QCT introduces no new noise source; collapse timing is emergent from Λ, Ω, Θ, Δt that are empirically calibrable on the platform. Consequently, parameter constraints map to device-level quantities rather than universal constants.

Versus Bohmian trajectories. QCT dispenses with trajectories and guiding waves. It inherits decoherence's basis logic and adds a deterministic onset rule tied to informational capacity. Reported tunneling-regime anomalies that challenge trajectory prescriptions are immaterial to QCT.

Prediction P8 (model separation). Side-by-side fits of interferometry data will prefer a thresholded model with a knee and hysteresis (QCT) over a purely smooth model (decoherence-only) by likelihood-ratio tests, while not requiring GRW/CSL-style universal parameters.

6.9 Technology-facing implications

Because C factorizes into Θ, Δt, Λ, and Ω, QCT provides actionable levers for quantum technologies:

Qubit and sensor design. Increase Θ (error-protected encodings, dynamical decoupling, redundancies that do not broadcast to the bath), decrease Λ (shielding, measurement frugality), minimize Ω (noise engineering), and control Δt (faster gates and readouts) to keep C subcritical.

Threshold-aware QEC. Schedules that briefly reduce Λ or raise Θ before heavy-measurement epochs can avoid inadvertent crossing. Conversely, threshold-assisted readout can intentionally cross C^* to harden classical records.

Re-coherence windows. Operate near but below C^* during sensitive logic, then step away from threshold to recover visibility; this behavior has no analog in models without an explicit threshold.

Prediction P9 (performance gains). Controllers synthesized to regulate C in real time will measurably extend algorithmic depth, suppress readout-induced decoherence, and improve metrological sensitivity compared with controllers that optimize only for conventional T1/T2 metrics.

6.10 Limits, edge cases, and falsifiability

QCT is falsified if any of the following are robustly demonstrated:

1. Absence of a knee and hysteresis in platforms where Λ, Ω, Θ, Δt can be independently and precisely swept across the predicted $C \approx C^*$.

2. Controllable superluminal signaling in EPR-style tests under unconditional marginals, despite matched pointer structures.

3. Persistent violations of energy accounting beyond commutator-bound power budgets tied to explicit couplings.

4. Failure of frequency statistics to match $p_i = \mathrm{Tr}(\Pi_i \rho_-)$ across matched C-histories.

Conversely, consistent observation of the capacity-linked knee, hysteresis, subcritical re-coherence, and technology-level gains from C-aware control would strongly support the thresholded-dynamics account.

Section 6 takeaway. QCT delivers a concrete, testable bridge from open-system quantum dynamics to single-outcome classical records: the index C governs when a system resolves, the pointer structure determines where it resolves, and the pre-onset state fixes how often each outcome occurs. The resulting predictions—knee, hysteresis, re-coherence, no-signaling, and scalable control—separate QCT not just from interpretive stories, but from dynamical models that lack an informational threshold.

7. Experimental Protocols and Validation Plan

7.1 Calibration: extracting Λ, Ω, Θ, Δt and the knee C*

Goal: independently calibrate the four drivers of the index $C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$ and locate the empirical knee C^* (expected ≈ 1 within a platform-specific scale).

1. Calibrate Ω (coherence pressure).

Introduce tunable, spectrally characterized dephasing noise with known spectral density S(ω). Extract effective Ω by fitting decay of off-diagonal elements under the Lindblad sector with fixed control settings. Report Ω with confidence intervals and full noise spectra.

2. Calibrate Λ (informational flux).
Implement a weak, parameterized information tap (e.g., partially polarizing pick-off, dispersive probe with known measurement strength, or photon-number-resolving side channel). Define Λ operationally as recorded bits per unit time and volume transferred to controlled ancillae or to the environment. Validate linearity between tap strength and Λ in the weak regime.

3. Calibrate Δt (temporal resolution).
Define Δt as the controllable window during which superposition is maintained between preparation and either readout or intentional protection-release. In pulsed experiments, Δt is the separation between beam-splitter or Ramsey pulses and the final projective step; in continuously driven devices, Δt is an engineered idle interval with fixed Hamiltonian.

4. Calibrate Θ (protection factor).
Estimate Θ by characterizing how error suppression, dynamical decoupling, or logical encoding reduces the effective rate at which recorded bits become redundant environmental records. Operationally: sweep protection settings that do not change Λ (taps) or Ω (bath spectrum) and infer Θ from the shift of the knee in C. Report Θ as a dimensionless multiplier with error bars.

5. Locate C*.
With Λ, Ω, Δt, Θ independently calibrated, sweep C over a range that crosses the knee while monitoring a coherence observable (visibility V, purity, fringe contrast). Identify the inflection point C* via a sigmoid fit to the hazard or visibility curve. Pre-register analysis method.

Deliverables: calibration curves, uncertainties, cross-checks (e.g., swapping baths to keep Λ fixed while changing Ω), and a preregistered plan for subsequent tests.

7.2 Single-system threshold interferometry (specific protocol predictions)

Platforms: Mach–Zehnder or Sagnac photonic interferometers; Talbot–Lau or Kapitza–Dirac matter-wave interferometers; superconducting or NV-center Ramsey interferometers.

Protocol:

1. Prepare a well-defined pointer basis (paths, or $\{|0\rangle, |1\rangle\}$) and measure baseline visibility $V0$ at low Λ and Ω.

2. Sweep Λ at fixed Ω and Δt via weak which-path taps of increasing strength; compute C from calibrated quantities.

3. Sweep Ω at fixed Λ and Δt using engineered dephasing; compute C.

4. For each run, record C, visibility V, purity, and timing.

Specific protocol predictions under QCT:

Threshold surface: V remains near $V0$ for $C < C^*$, then shows a sharp, irreversible drop as C crosses $\approx C^*$.

Hysteresis: cumulative weak taps produce an abrupt loss once C slightly exceeds C^*; removing the last tap after crossing does not immediately restore V.

Collapse-time scaling: near threshold, the minimum onset time scales inversely with Λ/Ω at fixed Θ and Δt.

Decoherence-only alternative: predicts smooth, monotone visibility decay with integrated exposure, no sharp knee, and no hysteresis for identical exposure histories.

7.3 Weak-measurement hysteresis and path dependence

Goal: detect the “memory” of accumulated informational flux.

Protocol:

1. Apply N weak which-path interactions of strength $\{g_1, \ldots, g_N\}$ such that Σg_k increases C toward C^*.

2. After each step, measure visibility and compute current C.

3. Once V drops sharply (onset), step backward by removing the last tap or reducing g_N to restore the pre-onset C.

Predictions: visibility-versus-C traces form a clockwise hysteresis loop. The collapse point on the up-sweep occurs at lower nominal C than the recovery point on the down-sweep. Control runs that keep total integrated exposure the same but distribute it differently produce distinct outcomes (path dependence), inconsistent with simple exposure-only models.

7.4 Multipartite no-signaling with threshold crossing

Goal: verify that QCT's "synchronous" resolutions do not enable controllable superluminal signaling.

Platforms: entangled photon pairs with spacelike-separated stations; superconducting qubit pairs with rapid, isolated readouts.

Protocol:

1. Prepare ρ_{AB} with known correlations.

2. At site A, vary Λ_A, Ω_A, and Θ_A to drive or avoid C_A crossing; at site B keep settings fixed and record unconditional marginals.

3. Ensure spacelike separation between A's setting choice and B's readout in photonic platforms; in solid-state, use strict isolation and fast timing.

Predictions: unconditional marginals at B are invariant under A's setting choices. Apparent "synchrony" appears only in conditioned data and vanishes without postselection. Any residual drift is attributable to measurable cross-couplings, not to action-at-a-distance.

7.5 Subcritical re-coherence window

Goal: demonstrate QCT's distinctive "build a new dam upstream" effect.

Protocol:

1. Park the system near but below threshold: $C \approx C^* - \varepsilon$, with ε small and positive.

2. Before onset, either reduce Λ (weaken taps) or increase Θ (apply protection) to move C farther below C^*.

3. Measure revival of visibility or purity compared to a decoherence-only forecast.

Predictions: a measurable re-coherence window exists below C^*. Decoherence-only models lacking a capacity threshold cannot recover visibility once redundancy is supercritical for the same exposure history.

7.6 Threshold-aware QEC and readout scheduling

Goal: improve algorithmic depth and measurement fidelity by controlling C in real time.

Platforms: superconducting qubits with surface-code or repetition-code primitives; trapped ions with syndrome extraction.

Protocol:

1. Instrument controllers to monitor proxies for Λ, Ω and to switch protection that raises Θ.

2. During heavy-measurement epochs, insert brief "cool-down" intervals that reduce Λ or raise Θ to avoid crossing C^*.

3. For readout hardening, intentionally cross C^* at the end of a cycle to produce stable classical records.

Metrics: logical error rates, algorithmic depth before failure, readout infidelity, and time-integrated C budget. Compare threshold-aware schedules against strong baselines optimized only for T1/T2.

Predictions: measurable gains in depth and readout fidelity; reduced incidence of abrupt, history-dependent failures; reproducible correlation between near-threshold excursions of C and error bursts.

7.7 Collapse-time statistics near threshold

Goal: extract the onset hazard h(t) and its dependence on C.

Protocol:

1. Prepare identical initial states and drive C(t) through the knee with controlled ramps of different slopes.

2. Record the distribution of onset times across repeats.

3. Fit a hazard model $h(t) = h0 \cdot s(C(t))$, where s is a steep sigmoid with midpoint C*.

Predictions: ramp-rate dependence consistent with a thresholded hazard. Fast ramps show a narrow onset-time distribution; slow ramps broaden it in a manner predictable from the C(t) trajectory. Decoherence-only models yield different scaling.

7.8 Cross-platform universality tests

Goal: show the same dimensionless C governs onset across platforms.

Platforms: photons, cold atoms, neutrons, NV centers, superconducting qubits, optomechanics.

Protocol:

1. Perform the 7.2–7.7 batteries on each platform with independent calibrations.

2. Rescale all data by the local C and plot visibility or onset probability versus C.

Predictions: collapse curves from disparate platforms align when expressed in C, revealing a universal knee at $C \approx C^*$ up to platform-specific offsets. Deviations diagnose miscalibration or additional channels.

7.9 Model comparison and preregistered analysis

Goal: distinguish QCT from decoherence-only and GRW/CSL-type models.

Analysis plan:

1. Pre-register observables, fitting functions, priors, and stopping rules.

2. For each dataset, compute Bayes factors and likelihood-ratio tests comparing:
a) thresholded hazard with knee and hysteresis,
b) smooth decoherence-only model,
c) objective-collapse model with universal constants.

3. Penalize model complexity and report posterior predictive checks.

Predictions: datasets featuring knees, hysteresis, re-coherence windows, and ramp-rate hazard scaling favor the thresholded model without requiring universal GRW/CSL parameters.

7.10 Confounds, controls, and falsification criteria

Controls:

Blind and randomized tap sequences to rule out operator bias.

Swap-in baths to vary Ω without changing Λ.

Sham “protection” toggles that alter control power but not Θ.

Hardware-in-the-loop noise-injection tests to verify Ω calibration.

Temperature and vibration logs; isolation checks.

Falsification (any robust finding below invalidates QCT):

1. No knee or hysteresis under independently swept Λ, Ω, Θ, Δt in regimes predicted to cross C*.

2. Controllable changes in unconditional marginals at B driven by A’s settings in spacelike-separated tests.

3. Energy or probability non-conservation not attributable to measured couplings.

4. Failure to recover Born frequencies when C-histories and pre-onset states are matched.

7.11 Reporting standards and reproducibility

Publish full calibration datasets for Λ, Ω, Θ, Δt with uncertainties.

Release raw time-series for C(t), visibility, onset times, and all control settings.

Open-source analysis code and provide dockerized environments.

Provide complete hardware schematics, pulse schedules, and firmware versions.

Archive preregistrations and deviations with justification.

Encourage third-party replication with inter-lab comparison packs.

Section 7 takeaway. This validation program operationalizes QCT's core claims in a platform-agnostic way: isolate the four levers (Λ, Ω, Θ, Δt), sweep the composite index C across a calibrated knee, and look for the predicted package of phenomena—threshold surfaces, hysteresis, subcritical re-coherence, ramp-rate hazard scaling, and no-signaling—while ruling out exposure-only and universal stochastic alternatives.

8. Conclusion, Discussion, Limitations, and Outlook

8.1 What QCT adds beyond decoherence and interpretations

Decoherence explains why interference terms are suppressed in practice but does not specify when a single outcome is selected or why outcome statistics follow the Born rule in each run. QCT contributes a concrete, dimensionless control variable

$C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$

and a thresholded onset mechanism (a knee at $C \approx C^*$) that (i) predicts sharp, history-dependent transitions, (ii) distinguishes capacity-limited collapse from smooth exposure effects, and (iii) recovers $p_i = \mathrm{Tr}(\Pi_i \rho_-)$ at the instant of onset without adding hidden trajectories or branching worlds. Unlike purely interpretive moves, QCT is instrumentable: Λ, Ω, Δt, Θ can be independently calibrated and swept.

8.2 Relation to standard quantum dynamics and CPTP evolution

Pre-threshold evolution is completely positive and trace preserving, governed by

$d\rho/dt = -i [H, \rho] + \Sigma_j (L_j \rho L_j^\dagger - 0.5 L_j^\dagger L_j \rho - 0.5 \rho L_j^\dagger L_j)$,

with L_j aligned to the environmental pointer basis that defines what can be redundantly recorded. Collapse onset is modeled as a hazard process $h(t) = h0 \cdot s(C(t))$ with s a steep sigmoid centered at C^*. For $C \ll C^*$ dynamics remain CPTP; at onset we apply a Born-weighted update $\rho_+ = \Sigma_i p_i \Pi_i$ with $p_i = \mathrm{Tr}(\Pi_i \rho_-)$. This preserves probability, respects positivity, and matches observed statistics.

8.3 Preferred basis and objectivity

The preferred basis emerges from the L_j that dominate environmental coupling (einselection). QCT does not choose a basis by fiat; it exploits the same monitoring channels decoherence identifies. The novelty is that QCT declares a finite capacity to maintain unresolved alternatives in that basis: once C crosses the knee, one record becomes actual and the others are erased from the accessible state description.

8.4 Energy, unitarity, and reversibility bounds

If $[H, L_j] = 0$ for the dominant channels, then $d/dt\, \mathrm{Tr}(H \rho) = 0$ during the pre-threshold CPTP evolution. With dissipative channels present, QCT specifies bounds on energy drift via measured rates γ_j and commutators $\|[H, L_j]\|$. Apparent irreversibility appears only at the thresholded update; below threshold, dynamics are as reversible as the chosen L_j permit. Post-onset "re-coherence windows" are predicted only while C remains subcritical; once redundancy is supercritical, recovery requires reducing Λ or raising Θ before the knee is crossed.

8.5 No-signaling and entanglement

QCT's "synchronous resolution" for entangled subsystems A and B preserves no-signaling because only conditioned statistics show the correlation. Unconditional marginals at B are invariant under A's control of Λ_A, Ω_A, Θ_A, or Δt_A, provided spacelike separation (photonic platforms) or strict isolation (solid-state) is enforced. Practically, QCT instructs experiments to verify that any observed drift in B's marginals arises from measured cross-couplings rather than remote setting choices.

8.6 Multipartite systems and networked thresholds

For a composite $S = S_1 \otimes S_2 \otimes \dots \otimes S_N$ with local controls, define per-subsystem indices
$C_k = \Theta_k \cdot \Delta t_k \cdot \Lambda_k / \Omega_k$
and a network index
$C_{net} = w_1 C_1 + \dots + w_N C_N + w_{corr} C_{corr}$,
where C_{corr} is an empirically calibrated contribution from inter-subsystem couplings (mediated baths, shared buses, common-mode noise). QCT predicts:
(a) local onset if a single C_k exceeds its knee while C_{net} remains subcritical;
(b) coordinated onset if C_{net} crosses its knee even when individual C_k are subcritical;
(c) hysteresis at the network level when cumulative weak cross-couplings raise C_{corr}.
This "networked threshold" generalizes Θ to Θ_{net}, a measurable protection factor of the entire architecture (e.g., logical encoding across qubits or optical modes).

8.7 Relation to objective-collapse (GRW/CSL) and Bohmian mechanics

GRW/CSL introduce universal stochastic parameters (e.g., collapse rate λ, localization width r_C). QCT introduces no universal constant; it predicts platform-specific knees once Λ, Ω, Δt, Θ are calibrated. Where GRW/CSL predict exposure-independent spontaneous hits, QCT predicts thresholded knees and hysteresis under controlled informational loading. Bohmian mechanics posits definite trajectories guided by a pilot wave; QCT dispenses with trajectories and guidance fields and replaces them with capacity-limited resolution in the pointer basis. Distinguishing features: absence of persistent hidden paths, presence of subcritical re-coherence, and threshold-controlled onset.

8.8 Determinism, randomness, and the status of "choice"

Between preparation and onset, evolution is deterministic given H, L_j, and controls. The hazard formalism introduces apparent randomness only in the timing of onset when C(t) approaches the knee. Outcome weights at onset follow $p_i = \mathrm{Tr}(\Pi_i \rho_-)$, not an added stochastic postulate. In this sense QCT is "deterministic until it must choose," and the "choice" is constrained by the Born weights defined by the pre-onset state.

8.9 Engineering consequences and quantum-technology leverage

QCT reframes control strategy from “fight all noise” to “budget the index C.” Practical levers are: reduce Λ during computation (softer taps, better shielding), manage Ω via engineered baths and spectral shaping, shorten Δt between protected refreshes, and raise Θ with dynamical decoupling, error-correcting encodings, and measurement-aware scheduling. Anticipated wins include: increased algorithmic depth before failure, fewer abrupt error bursts (since near-threshold excursions are avoided), and hardened readout by intentionally crossing the knee at the end of a cycle.

8.10 Limitations and open problems

1. Microscopic origin of Θ. We operationalize Θ as a protection factor but have not yet derived it from microscopic control theory or many-body models. A program to relate Θ to code distance, pulse sequences, and control bandwidth is outlined in future work.

2. Exact knee shape. We use a steep but smooth s(C) to fit hazards. The universality class of the knee (logistic vs. other sigmoids) is an empirical question.

3. Full-scale many-body baths. For strongly interacting environments, identifying the dominant L_j and the emergent pointer basis can be nontrivial; QCT inherits those difficulties from decoherence theory.

4. Rigorous no-signaling proofs on all architectures. While the operational protocols are clear, a general theorem for arbitrary network couplings remains to be developed.

5. Integration with cognition. QCT is observer-independent. Optional “volitional” layers (e.g., QZE-style feedback that raises Θ in biological or artificial agents) are reserved for separate work and are not required for the physics here.

8.11 Falsifiable signatures recap

QCT stands or falls on the following package: a calibrated knee in visibility or purity at C ≈ C*, hysteresis under cumulative weak taps, a subcritical re-coherence window, ramp-rate scaling of onset-time statistics consistent with a thresholded hazard, networked knees in multipartite settings, and strict no-signaling in entanglement tests. The absence of these under clean calibrations and controls would falsify the framework.

8.12 Roadmap

Short-term: execute single-platform protocols (Section 7.2–7.7) with preregistration and open data.
Mid-term: demonstrate cross-platform alignment by rescaling to C and extracting consistent C*.
Long-term: build threshold-aware compilers and control stacks that actively budget C in real time, and derive Θ from microscopic models so that protection strategies can be predicted rather than empirically tuned.

Conclusion

The Quantum Convergence Threshold (QCT) framework reframes the measurement problem as a question of informational capacity rather than metaphysical interpretation. It proposes that a quantum system remains coherent until its convergence index, $C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$, crosses a well-defined threshold, at which point collapse becomes an unavoidable, lawful transition. This simple relation unites time resolution, informational flux, environmental decoherence, and system protection into a single dimensionless control parameter.

QCT preserves standard quantum dynamics below threshold and reproduces Born-rule statistics at onset, requiring no hidden variables, branching universes, or observer-dependent postulates. Collapse is not stochastic but structurally determined by informational limits; randomness enters only through the fine structure of near-threshold dynamics. In doing so, QCT grounds wavefunction reduction in measurable, system-specific quantities—transforming collapse from an interpretive gap into a testable boundary condition.

The framework's predictions—threshold knees in interference visibility, hysteresis under cumulative probing, and subcritical re-coherence—are falsifiable through interferometric and entanglement experiments. Beyond foundational insight, QCT introduces practical levers for quantum technologies: by managing Λ, Ω, Δt, and Θ, engineers can deliberately delay or trigger collapse, optimize coherence budgets, and design threshold-aware quantum processors.

In essence, QCT replaces "observation causes collapse" with "informational capacity defines actuality." It retains the rigor of quantum field theory while providing an ontologically consistent, operationally testable account of reality's transition from possibility to fact.

Appendix A — Glossary of Core Terms and Quantities

This glossary consolidates all recurring symbols, parameters, and operators used throughout the Quantum Convergence Threshold (QCT) framework. Each entry includes physical meaning, units, and contextual relevance.

$\Lambda(x, t)$ — Informational Flux Density

• Definition: Rate at which information is registered or exchanged by a quantum system through environmental or internal interactions.
• Units: $bit \cdot s^{-1} \cdot m^{-3}$
• Physical Interpretation: Encodes the inflow/outflow of distinguishable states; high Λ indicates a rapid exchange of information between the system and its surroundings.

Θ(t) — Awareness Threshold Function
• Definition: Dimensionless coefficient quantifying a system's resistance to informational overload; determines the system's capacity to sustain superposition.
• Units: Dimensionless
• Physical Interpretation: Represents an effective "convergence buffer." High Θ corresponds to robust coherence protection (e.g., via isolation, error correction, or symmetry).

Δt — Temporal Resolution Interval
• Definition: Duration over which coherence persistence and informational convergence are evaluated.
• Units: Seconds (s)
• Physical Interpretation: Analogous to the sampling window during which internal quantum evolution remains unperturbed; collapse likelihood increases with longer Δt at constant Λ/Ω.

Ω — Coherence Pressure
• Definition: The decoherence rate associated with environmental disturbance and measurement back-action.
• Units: $bit \cdot s^{-1}$
• Physical Interpretation: Competes with Λ; represents how aggressively the environment fragments phase coherence.

C(x, t) — Collapse Index
• Definition: Dimensionless ratio expressing the balance between informational intake, time accumulation, threshold tolerance, and coherence resistance.
• Formula: $C = \Theta \cdot \Delta t \cdot \Lambda / \Omega$
• Physical Interpretation: The central criterion for collapse; once $C \geq 1$, the system's superposition becomes unsustainable and self-resolves.

$h(t)$ — Collapse Hazard Function
• Definition: Instantaneous collapse probability per unit time, proportional to a sigmoid of $C(t)$.
• Formula: $h(t) = h_0 \cdot s(C(t))$
• Physical Interpretation: Encodes a smooth transition from coherence ($C \ll 1$) to near-certain collapse ($C \geq 1$).

$\tau(t)$ — Informational Divergence
• Definition: Measure of the system's internal informational mismatch or entropy gradient.
• Formula: $\tau(t) = \Sigma\, S(t)(x) \cdot \log[S(t)(x)/A(x)]$
• Units: Dimensionless (bits of divergence)
• Physical Interpretation: Quantifies how far the system's instantaneous state deviates from its equilibrium or attractor distribution.

$R(t)$ — Integrated Coherence Memory
• Definition: Accumulated integral of the coherence rate over time.
• Formula: $R(t) = \int_0^t \rho(\tau)d\tau$
• Units: Dimensionless
• Physical Interpretation: Represents the system's retained "informational inertia." High $R(t)$ corresponds to stronger memory coherence and a higher $\Theta(t)$.

$\rho(t)$ — Coherence Rate
• Definition: Rate at which internal coherence is reinforced through self-consistent evolution or error correction.
• Formula: $d\rho/dt = I(t) - \rho(t)/\tau$
• Units: s^{-1}
• Physical Interpretation: Competes with decoherence pressure; positive $\rho(t)$ sustains wavefunction integrity.

I(t) — Recursive Informational Load / Attentional Intensity
• Definition: Measure of self-referential feedback or recursive monitoring within the system; in cognitive analogs, corresponds to directed attention.
• Units: s^{-1}
• Physical Interpretation: Controls the Quantum Zeno Effect; higher I(t) slows or stabilizes wavefunction evolution.

$\Phi_c(t)$ — Collapse Pressure Field
• Definition: Spatial-temporal gradient of informational divergence; represents the force driving collapse.
• Formula: $\Phi_c = -\nabla \tau$ or $\Phi_c(t) = \eta \cdot (d\tau/dt)$
• Units: $bits \cdot s^{-1} \cdot m^{-1}$
• Physical Interpretation: Analogous to a potential gradient; collapse occurs when Φ_c exceeds a stability limit.

η — Scaling Constant
• Definition: Converts informational gradients into collapse rates.
• Units: Depends on chosen normalization (commonly dimensionless).

Γ[ψ] — Collapse Functional
• Definition: Operator describing decay of coherence due to informational overload.
• General Role: Acts as a nonlinear correction to the Schrödinger equation:

$$i\hbar(\partial\psi/\partial t) = H\psi - i\hbar\Gamma[\psi]$$

• Physical Interpretation: Governs nonunitary transition as $C(x, t) \to 1$.

H — Hamiltonian Operator
• Definition: Governs unitary evolution of the system prior to collapse.
• Relation: Unmodified segment of the Schrödinger equation; QCT corrections are appended as threshold couplings.

$\psi(t)$ — Quantum State Function
• Definition: Standard wavefunction evolving under combined unitary and convergence-influenced dynamics.
• Physical Interpretation: Retains full quantum description until $C(x, t)$ reaches unity, triggering a self-collapse transition.

Π_i — Projection Operator for Outcome i
• Definition: Orthogonal projectors defining possible outcome states.
• Relation to Collapse: At threshold, post-collapse state $\rho_+ = \Sigma_i p_i \Pi_i$, with $p_i = Tr(\Pi_i \rho_-)$.

L_j — Lindblad Operator
• Definition: Decoherence channel operators appearing in the master equation.
• Function: Encapsulates noise, loss, or measurement processes; under QCT, rates depend on Λ and Θ.

$s(C)$ — Sigmoid Transition Function
• Definition: Smooth logistic function regulating the hazard response of collapse onset.
• Typical Form: $s(C) = 1 / (1 + e^{-k(C - 1)})$
• Physical Interpretation: Ensures a continuous crossover from unitary to nonunitary behavior rather than abrupt switching.

C^* — Critical Convergence Point
• Definition: Value of C where collapse hazard reaches 50%.
• Physical Interpretation: Marks the precise balance between coherence retention and informational saturation; typically near $C \approx 1$.

Θ_0, Λ_0, Ω_0 — Baseline Parameters
• Definition: Calibration constants determined empirically for a given system.
• Physical Interpretation: Provide normalization for comparative experiments and theoretical consistency.

Appendix B — Derivations, Computational Notes, and Simulation Framework

---

B.1 Threshold Condition and Collapse Criterion

The Quantum Convergence Threshold (QCT) derives from the balance between four measurable quantities: informational flux Λ, coherence pressure Ω, temporal resolution Δt, and the system's awareness threshold Θ.
The fundamental collapse condition is:

$$C(x, t) = \Theta(t) \cdot \Delta t \cdot \Lambda / \Omega$$

Collapse occurs when $C \geq 1$.
This expresses a simple physical principle: when the cumulative informational intake ($\Lambda \cdot \Delta t$) overcomes the system's tolerance to decoherence (Ω / Θ), superposition becomes dynamically unsustainable.

B.2 Threshold Dynamics and Hazard Rate

To describe how rapidly collapse approaches the threshold, define a hazard function

$$h(t) = h_0 \cdot s(C(t))$$

where h_0 is the baseline hazard and $s(C)$ is a sigmoid response function:

$$s(C) = 1 / [1 + \exp(-k(C - 1))]$$

The probability that collapse has occurred by time t is

$$P(t) = 1 - \exp(-\int_0^t h(\tau)\, d\tau)$$

which yields negligible hazard for $C \ll 1$ and near-certainty when $C \geq 1$. This continuous form avoids unphysical discontinuities and fits time-resolved interferometry data.

B.3 Modified Schrödinger Equation

Below threshold ($C < 1$) evolution is unitary:

$$i\hbar\, (\partial\psi/\partial t) = H\psi.$$

As $C \to 1$, an informational-decay term activates:

$$i\hbar\, (\partial\psi/\partial t) = H\psi - i\hbar\, (\Lambda/\Omega)\, \psi.$$

The second term models amplitude attenuation proportional to the information/decoherence ratio; it preserves energy expectation up to $O(C - 1)$ corrections, remaining trace-preserving until onset.

B.4 Density Matrix Form and Lindblad Coupling

In density-matrix form:

$$d\rho/dt = -i[H, \rho] + \Sigma_j \left(L_j\, \rho\, L_j^\dagger - \tfrac{1}{2} L_j^\dagger L_j\, \rho - \tfrac{1}{2} \rho\, L_j^\dagger L_j \right),$$

where the Lindblad rates γ_j depend on Λ and Ω:

$$\gamma_j \propto \Lambda/\Omega.$$

When $C \geq 1$, the effective rates become divergent, leading to loss of coherence.
The requirement $[H, L_j] \approx 0$ for dominant channels guarantees approximate energy conservation pre-threshold.

B.5 Collapse Outcome and Born-Rule Recovery

At the instant of threshold crossing:

$$p_i = \mathrm{Tr}(\Pi_i\, \rho_-), \quad \rho_+ = \Sigma_i\, p_i\, \Pi_i.$$

This enforces standard Born-rule weighting without additional postulates; randomness reflects microscopic fluctuations of Λ and Ω near $C \approx 1$.

B.6 Energy and Unitarity Bounds

For systems obeying $[H, L_j] = 0$:

$$\frac{d}{dt} \mathrm{Tr}(H\rho) = 0.$$

If dissipative components exist:

$$|dE/dt| \leq \Sigma_j \gamma_j \|[H, L_j]\|,$$

placing a strict bound on energy drift attributable to threshold activation.
Reported experiments should list $\Delta E/E$ and the control power used to raise Θ.

B.7 Relation to Decoherence-Only Models

Pure decoherence predicts monotonic visibility decay with cumulative exposure.
QCT predicts instead:
• a sharp knee near $C \approx 1$,
• hysteresis under cumulative weak taps (collapse delayed then sudden), and
• a re-coherence window just below onset when redundancy becomes supercritical.

B.8 Computational Simulation Outline

A minimal time-step simulation uses four coupled equations:

1. $d\psi/dt = -iH\psi - (\Lambda/\Omega)\psi$

2. $d\rho/dt = I(t) - \rho/\tau$

3. $\Theta(t) = \exp(-1/(R(t)+\varepsilon))$, where $R(t) = \int\rho\, dt$

4. $C(t) = \Theta \cdot \Delta t \cdot \Lambda/\Omega$

Collapse triggers when $C \geq 1$.

Numerically integrate with adaptive Δt to capture near-threshold dynamics and verify energy and trace conservation pre-collapse.

B.9 Toy Model Implementation

Pseudocode sketch (Python-style):

```
for step in range(N):
    Theta = exp(-1 / (R + eps))
    C = Theta * dt * Lambda / Omega
    if C >= 1:
        collapse_flag = True
        break
    psi += (-1j * H @ psi - (Lambda/Omega)*psi) * dt
    rho += (I - rho/tau) * dt
    R += rho * dt
```

Outputs: $\psi(t)$, $\rho(t)$, $C(t)$, and collapse_time. Plot C versus t to show the approach to unity.

B.10 Suggested Empirical Calibration

1. Estimate Λ from entanglement entropy growth ($S \approx \Lambda \cdot \Delta t$).

2. Determine Ω from measured T_2 times ($\Omega \approx 1/T_2$).

3. Extract Θ from decoherence suppression factor of system isolation or feedback.

4. Verify $C \approx 1$ corresponds to interference visibility collapse.

B.11 Outlook for Quantum Control Applications

QCT supplies actionable control parameters:
• Raise Θ (via error correction) to prolong coherence.
• Reduce Λ or Δt to delay collapse.
• Manipulate Ω through environmental engineering to tune threshold behavior.
These levers can be directly integrated into quantum-device feedback protocols.

Appendix C — Networked Collapse, No-Signaling, Basis Selection, and Calibration Details

---

C.1 Multipartite and Networked QCT

Consider N subsystems with local reduced states $\rho_i(t)$, local convergence indices
$C_i(t) = \Theta_i(t) \cdot \Delta t_i \cdot \Lambda_i / \Omega_i$,
and a coupling graph with weights $K_{ij} \geq 0$ quantifying informational interdependence (shared redundancy, entanglement depth, error-correction ties, probe cross-talk).

Define a networked effective index per node: $C_i^{eff}(t) = C_i(t) + \Sigma_{j \neq i} K_{ij} \cdot C_j(t)$.

Global onset criteria (use whichever is stronger for the architecture under test):

• Node-triggered onset: collapse when $\max_i C_i^{eff} \geq 1$.
• Collective onset: collapse when $(1/N) \Sigma_i C_i^{eff} \geq 1$.
• k-core onset: collapse when at least k nodes satisfy $C_i^{eff} \geq 1$.

Interpretation: informational saturation propagates along redundancy channels. In GHZ-like states, K_{ij} are large and symmetric; in weakly entangled clusters, K_{ij} are sparse. This yields concrete, testable scaling: adding redundant observers or copies increases $\Sigma_j K_{ij}$ and can push C_i^{eff} across threshold even if each $C_i < 1$ in isolation.

C.2 No-Signaling: Locality and Marginals

Assume standard microcausality (spacelike separated observables commute) and pre-threshold CPTP dynamics generated by a GKSL map with local Lindblad operators respecting the tensor product structure. Let A and B be spacelike subsystems with joint state ρ_{AB} and local controls that can alter Λ_A, Ω_A, Θ_A, Δt_A but not B's generators.

Claim (no-signaling): changing A's control cannot change B's reduced marginal before classical communication.

Sketch:

1. Pre-threshold, $d\rho_{AB}/dt = L_A \otimes I_B + I_A \otimes L_B + L_{int}$ with L_{int} arranged to preserve microcausality for spacelike separation. The reduced map on B is CPTP and independent of A's choice of local instrument parameters when tracing out A: $Tr_A \circ e^{t(L_A\otimes I + I\otimes L_B + L_{int})} = E_B(t)$, with E_B not depending on A's hidden settings under the microcausal decomposition.

2. At onset in A, the jump map J_A: $\rho_{AB} \to \Sigma_i (\Pi_i^A \otimes I_B) \rho_{AB} (\Pi_i^A \otimes I_B)$ is CPTP and outcome-averaged. B's marginal after the unconditioned jump is $\rho_B' = Tr_A[J_A(\rho_{AB})] = \rho_B$, unchanged.
Therefore, without conditioning on A's communicated outcome, B's statistics are invariant; superluminal signaling is excluded.

C.3 Preferred Basis from Pointer Structure

QCT's outcomes are taken in the decoherence-selected pointer basis. Operationally:

- Let $\{L_j\}$ be the dominant Lindblad operators coupling system to environment or readout.
- The pointer projectors $\{\Pi_i\}$ jointly diagonalize the decoherence functional, i.e., they are eigenprojectors of the dynamical map's decohering part and maximize environmental redundancy ("einselection").
- The onset map $\rho \to \Sigma_i p_i \Pi_i$ with $p_i = Tr(\Pi_i \rho)$ is thus the minimal-disturbance projection consistent with the already-selected basis.

This anchors the "which observable collapses" question in experimentally characterizable system–environment structure.

C.4 CPTP and Energy Bounds (Pre- and Post-Threshold)

Pre-threshold: Evolution under $d\rho/dt = -i[H,\rho] + \Sigma_j (L_j \rho L_j^\dagger - 0.5\{L_j^\dagger L_j, \rho\})$ is completely positive and trace preserving by the GKSL theorem. If $[H, L_j] = 0$ for dominant channels, $d/dt\, Tr(H\rho) = 0$, so energy expectation is conserved up to subdominant commutators.

Onset map: $\rho_+ = \Sigma_i (\Pi_i \rho_- \Pi_i)$ is a convex sum of CP projectors; it is CP and trace preserving. Energy drift at onset satisfies
$|Tr(H\rho_+) - Tr(H\rho_-)| \leq \Sigma_i ||\Pi_i H \Pi_i - H|| \cdot p_i$,
bounded by how non-commuting H is with the pointer projectors; in well-engineered readout bases, this drift is experimentally small and reportable.

C.5 Born Weights via Minimum-Disturbance Projection

Among all states σ that are diagonal in the pointer basis, the state minimizing $D(\sigma \| \rho_-)$ (quantum relative entropy) is $\sigma^* = \Sigma_i \Pi_i \rho_- \Pi_i$ with weights $p_i = Tr(\Pi_i \rho_-)$. Thus, QCT's outcome distribution emerges from a least-change principle: select the diagonal state closest (in information geometry) to the pre-onset state. This reproduces the Born rule without extra stochastic postulates.

C.6 Calibrating Λ, Ω, Θ, Δt in Practice

• Λ (informational flux): Estimate from entanglement entropy production rate $\dot{S}_{env}$, or from the logarithmic rate of redundancy growth in Quantum Darwinism style tomography. Units: bits s^{-1} m^{-3} (coarse-grained to the device mode volume).

• Ω (coherence pressure): Invert measured T_2 processes: $\Omega \approx 1/T_2$ under a single effective dephasing channel; for multi-channel noise, set Ω to the net dephasing rate in the measured pointer basis.

• Δt (temporal resolution): Choose the smallest time bin that preserves stationary Λ and Ω estimates in the regime of interest; in pulsed interferometry, Δt is the pulse separation or interferometer dwell time.

• Θ (threshold factor): Infer from capacity-linked knees in visibility or fidelity versus cumulative exposure, or from the minimal overhead (error correction, dynamical decoupling, encoding) required to shift the knee from $C \approx 1$ to $C \approx 1 + \delta$. Report Θ as a dimensionless multiplier with confidence intervals.

Cross-validation: Ensure the knee predicted by $C = \Theta \, \Delta t \, \Lambda / \Omega \approx 1$ aligns, within error bars, with the empirically observed onset.

C.7 Protocol Library Extensions and Falsifiers

Beyond the single-system and bipartite protocols already specified:

• Multipartite threshold tomography: Prepare W, GHZ, and cluster states of 3–10 qubits. Vary Λ locally on one node and measure onset patterns against the predicted C_i^{eff} structure under known K_{ij} (tunable by adding or removing redundant monitors). QCT predicts topology-dependent knees; decoherence-only predicts smooth decay without sharp graph-linked thresholds.

• Threshold hysteresis loops: Apply many subcritical weak taps (raising Λ slightly) and then one final small increment. QCT predicts abrupt visibility collapse at the final increment and non-immediate recovery upon removing that last tap (history dependence). Decoherence-only predicts history-free monotone decay.

• Re-coherence window: Operate just below threshold while increasing Θ via real-time error correction. QCT predicts partial restoration of visibility when Θ increases at fixed Λ, Ω, Δt; standard models without capacity constraints predict no such knee-adjacent "lift."

Clear falsifier: If carefully calibrated experiments find no capacity-linked knees, no hysteresis under cumulative subcritical exposure, and no sub-threshold re-coherence when Θ is raised, the QCT threshold mechanism is disconfirmed for that platform.

C.8 Constraints from GRW/CSL Bounds

GRW/CSL impose upper bounds on spontaneous objective collapse rates in matter-wave interferometry and optomechanics. QCT is not a universal spontaneous-rate model: its effective "rate" is device- and context-dependent via Λ, Ω, Δt, Θ. To maintain consistency with GRW/CSL exclusions:

• Demonstrate that, for isolated systems where Λ is minimized and Θ is not artificially lowered, C remains well below 1 over the experimental baseline, so no spontaneous onset is predicted.
• Reserve threshold crossing for high-redundancy, measured, or engineered-noise regimes; this aligns with existing null results for unmonitored macroscopic superpositions.

C.9 Re-Coherence Engineering (Pre-Threshold)

Practical levers to move a system from near-onset back into the coherent regime:

• Increase Θ by encoding (e.g., repetition or stabilizer codes), dynamical decoupling, or autonomous feedback;
• Decrease Λ by reducing probe strength, path-information taps, or environmental monitoring bandwidth;

• Decrease Δt by shortening interaction time or gating interferometer dwell time;
• Increase Ω only when calibrated as "pressure" that, in your setup definition, reduces net recording capacity (note: by our convention Ω in the denominator reflects destructuring pressure; verify your mapping).

Rule of thumb: A 10% increase in Θ or a 10% decrease in either Λ or Δt yields roughly a 10% decrease in C, moving the operating point away from the knee.

C.10 Open Mathematical Problems

1. Tight no-go theorems: Prove or disprove that all QCT onset maps respecting pointer structure can be embedded into a continuous-time CP instrument with bounded energy drift.

2. Optimal basis selection: Given H and $\{L_j\}$, characterize the unique pointer basis minimizing post-onset energy disturbance under the minimum-relative-entropy principle.

3. Network thresholds: Derive exact conditions on K_{ij} under which node-triggered and collective onset criteria diverge, and propose critical exponents for large networks.

4. Θ from first principles: Link Θ to code distance, syndrome rate, or Fisher-information curvature for protected subspaces, providing device-independent formulas.

5. Relativistic extension: Formulate a covariant version where C is defined on spacelike hypersurfaces and onset obeys Tomonaga–Schwinger evolution with local CP instruments.

C.11 Reporting Checklist for Experimental Papers

• Specify platform, pointer basis, H, dominant L_j.
• Provide calibrated Λ, Ω, Δt, inferred Θ with confidence intervals.
• Show C versus control parameter and identify the knee (if present).
• Report energy drift bounds and control power used to alter Θ.
• Include no-signaling verification for bipartite tests (marginal invariance).

- Compare to decoherence-only simulations and indicate any hysteresis or re-coherence effects.