Bit-Collapse Theory (BCT)

Falsification Package — Phase-2

Retention-driven bridge, calibrated record channels, and preregistered decision rules

Federico Batori

Working draft for DOI-main package revision

Abstract

This document states the Phase-2 falsification spine for BCT in a form intended to be both operationally disciplined and experimentally discussable. The target is narrow: not a new force, not a modified collapse dynamics, and not a generic decoherence story, but the retention-driven bridge linking verified which-way information to event-like definiteness for a bounded agent. The core claim is that, once which-way information is genuinely present at or above the threshold I_c, interference visibility should become organized by stable record formation in a downstream channel.

Version 4 adds the missing bridge between theory and protocol. It introduces an explicit calibration stage for the record channel, separates the estimation of R, tau_loss, and I_WW from the main visibility scan, states a toy buffer model that generates a threshold-like family in the organizing variable s = R * tau_loss, and turns the primary analyses for Experiments 1 and 2 into preregisterable decision procedures rather than open-ended curve fitting. Experiment 1 remains the main evidential core. Experiment 2 is a structural reinforcement via a two-parameter phase map. Experiment 3 is retained only as a conceptually informative accessibility cross-check and is explicitly de-emphasized as a secondary target.

1. Scope and status

This package tests one specific BCT bridge: whether event-like definiteness for an agent requires not only verified which-way information but also sufficient downstream record stability. The global quantum state may remain unitary throughout. What is under test is the bridge from record formation to experienced definiteness, not a new microscopic dynamics.

Not claimed: a new collapse force, a modified Schrodinger equation, or a universal replacement for all decoherence theory.

Claimed: if a record channel genuinely carries outcome-distinguishing information, event-like behavior should still depend on whether that record persists as a stable accessible trace long enough to cross the relevant threshold for the agent.

Document role: this is the general falsification spine for the main Phase-2 DOI. A specialized thermal route already exists elsewhere and is treated as an external companion, not duplicated here.

2. Canonical variables and frozen discipline

The package uses Phase-2 canonical notation and fixes a frozen operational regime before the main scan.

- AA(tau): accessible algebra of agent A over the integration window tau.
- tau_p: measured persistence or retention time of the record channel.
- tau_loss: characteristic loss or decay timescale of the downstream record buffer.

- tau*: minimal persistence threshold required for a record to count as event-like.
- I_WW: effective which-way information in the record channel.
- I_c: conservative operational threshold for the which-way channel, with default I_c = 1 effective bit.
- R: record-writing or effective refresh rate into the downstream record channel.
- s = R * tau_loss: stability proxy organizing the test wedge.

All claims are made inside a frozen operational regime F_A. Platform-specific tolerances, calibration rules, and confound bounds must be declared before the scan and may not be retuned point-by-point to rescue conclusions.

2.1 Independent calibration principle

The central experimental vulnerability identified in review is the risk that the organizing variable s = R * tau_loss might be estimated from the same phenomena that determine visibility. To block that loophole, Version 4 requires a calibration stage that is independent of the main visibility scan.

Calibration stage A — which-way verification. At each operating point, I_WW must be measured or tightly bounded using a direct record-readout protocol on the record channel. It may not be inferred from the visibility reduction itself.

Calibration stage B — record-channel write and decay. With the same record channel active but with the interference measurement disabled or bypassed, estimate the record-channel acquisition and decay curves. A minimal operational choice is:

D_rec(t) = D_inf * (1 - exp(-R * t)) during write,

D_rec^off(t) = D_0 * exp(-t / tau_loss) after write is switched off.

Here D_rec is a calibrated record distinguishability observable for the buffer alone. The parameters R and tau_loss are extracted from these auxiliary fits before the main visibility scan. Those fitted values, with confidence intervals, are then treated as fixed inputs to the primary analysis.

Calibration stage C — arm-isolation check. Vary tau_loss while the which-way marker is disabled. If visibility changes materially in this control condition, the platform has not isolated the downstream record channel from ordinary arm dephasing, and the main test is not admissible.

3. Claim ladder

The package is easiest to interpret if its claims are ordered by strength.

Minimal falsification claim. If I_WW is independently verified at or above I_c and visibility shows no systematic organization by retention near threshold, the retention-driven bridge fails in that domain.

Signature claim. If a reproducible transition or knee aligns with s = R * tau_loss under confound control and preregistered model comparison, the bridge receives specific support.

Accessibility claim. If matched accessible and inaccessible record conditions diverge near threshold under hardware-level verification, the bridge gains support for an accessibility-sensitive axis beyond bare persistence.

These three claims do not carry equal weight. The core evidential target is the first. The second is stronger but more model-dependent. The third is conceptually valuable but should be treated as a cross-check unless a platform can implement it very cleanly.

4. Toy buffer model and bridge to observables

The package now includes a minimal toy model whose purpose is not to derive a unique microscopic law, but to justify a preregisterable threshold-like family for visibility as a function of s.

Let q(t) be an effective record-support variable for the downstream buffer. Consider the mean-field write-loss equation:

dq/dt = R * (1 - q) - q / tau_loss, with 0 <= q <= 1.

At stationarity,

q_inf(s) = s / (1 + s), where s = R * tau_loss.

Interpretation: q_inf is the asymptotic fraction of stable record support in the channel. Event-like update is expected only when the stationary support exceeds a threshold q* associated with the operational persistence requirement tau_p >= tau*.

The deterministic threshold condition q_inf >= q* is equivalent to s >= s*, where

s* = q* / (1 - q*).

If the threshold crossing is broadened by finite noise, finite readout resolution, or shot-to-shot variability in the channel, the sharp step is naturally smeared into a logistic family in log s. This motivates the primary threshold-like model used in Experiment 1:

V_T(s) = V_min + (V_max - V_min) / (1 + exp[k * (log s - log s_0)]),

with free parameters V_min, V_max, k, and s_0.

A minimal smooth comparison family is

V_S(s) = A * exp(-B * s) + C,

or another explicitly preregistered smooth monotone alternative. The point is not that the logistic is sacred; the point is that the threshold-like family is now tied to a toy buffer model rather than introduced as a purely stylistic curve shape.

5. Shared experimental architecture

All experiments instantiate the same architecture: verified which-way information, an independently characterized record channel, a preregistered organizing variable, and explicit competition against smooth decoherence-only alternatives.

- Step 1. Verify I_WW directly at each operating point. It may not be inferred from nominal settings or from visibility alone.
- Step 2. Characterize the record channel independently through R, tau_loss, and, where possible, tau_p, using auxiliary calibration runs.
- Step 3. Measure visibility V or a platform-equivalent coherence observable in the main scan.
- Step 4. Test whether the data are better organized by the retention proxy s than by unrelated hardware knobs or baseline dephasing alone.
- Step 5. Apply preregistered model competition and declared defeat conditions.

6. Experiment 1 — primary falsification test

Experiment 1 is the main test and should be treated as the primary evidential core of the package. A fixed which-way marking interaction is followed by a downstream leaky record buffer whose loss time tau_loss is scanned across a near-threshold region. The key question is whether visibility V becomes retention-organized once I_WW >= I_c is already verified.

6.1 Primary prediction

For s below threshold, the record decays too quickly to sustain an event-like update and V remains relatively high apart from bounded baseline dephasing.

For s above threshold and tau_p well beyond tau*, V is strongly suppressed in a way consistent with stable event-like record formation.

A reproducible transition region appears when tau_loss crosses the near-threshold band, and the transition aligns with s rather than with arbitrary hardware labels.

6.2 Preregistered model families

Model T (threshold-like): the logistic family V_T(s) defined in Section 4.

Model S (smooth): a smooth monotone family such as V_S(s) = A * exp(-B * s) + C, or another explicitly named smooth alternative fixed before the scan.

6.3 Primary decision rule

The primary metric for Experiment 1 is preregistered model comparison between T and S using one chosen information criterion, for example AIC.

A simple decision policy is:

- support for the retention-driven bridge if Delta_AIC = AIC_S - AIC_T > 4, and the fitted transition point s_0 lies inside the preregistered near-threshold search band;
- no specific support if -4 <= Delta_AIC <= 4;
- defeat of the threshold-like bridge in that domain if Delta_AIC < -4 and the confound controls are satisfied.

The exact numerical cutoff may be platform-adjusted, but one cutoff must be fixed before data collection.

6.4 Defeat conditions for Experiment 1

- F1. No-threshold outcome: no systematic dependence of V on tau_loss in the near-threshold scan while I_WW >= I_c is verified and tau_loss is independently measured.
- F2. Persistence-with-interference: V remains high or recoverable even when the record demonstrably persists well beyond tau* in the same hardware regime.
- F3. Decoherence-only dominance: retention variables add no explanatory power once ordinary arm dephasing and path-noise controls are enforced.
- F4. No s-scaling: any apparent transition fails to align with s when R is varied and instead tracks unrelated hardware settings.

7. Experiment 2 — structural reinforcement in the (R, tau_loss) plane

Experiment 2 maps the same bridge in the two-parameter plane (R, tau_loss). It is not the first experiment to run, but it is the cleanest reinforcement if Experiment 1 already shows a threshold-like signature. The aim is to test whether a phase boundary is organized by s = R * tau_loss rather than by arbitrary device-specific coordinates.

7.1 Primary prediction

A reproducible frontier separates a relatively high-visibility regime from a strongly suppressed regime, and that frontier aligns with s approximately equal to its threshold value.

7.2 Contour-alignment metric

Let x_i = log R_i and y_i = log tau_loss_i for experimentally identified transition points. Let the predicted contour be x + y = log s_0, where s_0 is taken from the preregistered threshold model or from a calibration band fixed in advance.

Define the signed perpendicular distance of point i from the predicted contour by

d_i = (x_i + y_i - log s_0) / sqrt(2).

A simple preregistered alignment score is

D_mean = mean_i |d_i|.

One admissible acceptance rule is: the frontier counts as aligned only if at least 90 percent of transition points satisfy |d_i| <= delta_cont and D_mean <= delta_bar, where delta_cont and delta_bar are fixed before the scan.

7.3 Defeat conditions for Experiment 2

- No reproducible frontier after controls.
- A frontier appears but fails the preregistered contour-alignment rule.
- Strong suppression occurs deep in the low-s regime under verified I_WW >= I_c without a plausible confound explanation.

8. Experiment 3 — accessibility cross-check

Experiment 3 is kept as a conceptually useful but secondary test. Two conditions are matched up to early record creation and then split into an accessible branch and an inaccessible branch. Because hardware-level inaccessibility with matched early record creation is difficult to implement cleanly, this experiment should not be treated as coequal with Experiment 1 unless a concrete platform-specific design exists.

8.1 Primary metric

The primary metric for Experiment 3 is Delta V = V_accessible - V_inaccessible across the same near-threshold scan, with confidence intervals fixed in advance.

8.2 Interpretive status

A null result in Experiment 3 does not by itself kill the retention-driven bridge. It instead pushes the program toward a persistence-only reading. A positive result would be conceptually important, but it should be treated as a cross-check unless the hardware implementation is exceptionally clean.

If no platform-specific implementation is available, Experiment 3 should be read as a proposed extension rather than as part of the minimal falsification package.

9. Controls and non-negotiable confound closure

The package is only as strong as its controls. The following are not optional.

- Arm isolation: changing tau_loss must not measurably alter arm phase noise, losses, alignment, spectral properties, or effective Hamiltonian terms of the interfering system.
- Marker-off control: vary tau_loss with the which-way marker disabled. If baseline visibility shifts materially, the arm-isolation condition fails.
- Record verification at every point: I_WW must be measured or tightly bounded at each scan point.
- Downstream-only discipline: the retention element must be downstream of the coherence-forming step, or back-action must be bounded experimentally.
- Independent calibration: hardware settings must be mapped to tau_loss and, where relevant, to R before the main scan.
- Matched early record creation for Experiment 3: accessible and inaccessible branches must agree before the inaccessibility intervention.

10. Relation to the thermal route

The thermal route is not repeated here because it already exists as a specialized published companion. Conceptually, it is one concrete implementation of the same general architecture: temperature is used only insofar as it modulates tau_loss or tau_p in a record channel. The main DOI should therefore keep the general falsification spine and point to the thermal route externally rather than duplicate it.

11. Clean-loss statement

This package is designed so that a negative result is informative. If visibility is fully explained by independently measured dephasing and does not track retention metrics after the calibration and control steps above, then the retention-driven bridge is falsified as stated in that domain. Conversely, if a threshold-like transition robustly tracks s = R * tau_loss across platforms and survives confound closure, the bridge gains strong support and yields quantitative constraints on tau*.

AI-use note

This draft has undergone AI-assisted editorial and structural drafting support. The author retains responsibility for all claims, definitions, and errors.