Bit-Collapse Theory (BCT)

Manifesto (Phase 2) — Canonical front-door statement

Federico Batori
April 2026

Purpose. A short front-door statement for the Phase-2 package.
Method. Crisp claims, explicit scope, and explicit pointers to companion notes.
Status. Interpretational / ontological framework plus a killable empirical wedge.

0. One-sentence thesis

BCT defines the experienced state of agent A as an equivalence-class label $[\Psi]_A$ over global states that are indistinguishable relative to A's physically accessible algebra $AA(\tau)$ and the values of A-accessible stable records. In this framework, "collapse" is the discrete agent-relative update of that label when stable record structure changes.

1. Why this exists

Decoherence and related approaches explain why interference becomes practically unavailable in reduced descriptions and why macroscopic records can become robust and redundant. But a conceptual gap remains.

- We speak of definite outcomes while also accepting global unitarity.
- We speak of the same observer over time while the underlying microphysical realization may vary drastically.
- We often let semantic language sneak into what should be a physical statement about access, records, and stability.

BCT is aimed at one missing piece: a non-elastic, physically grounded notion of observer-definiteness and experienced identity that can be tied to explicit defeat conditions.

2. What BCT claims

2.1 Identity and experienced state

BCT treats an agent as a finite physical system with bounded access and finite integration. The experienced state of agent A is not a primitive intuition and not a raw microstate. It is the equivalence-class label $[\Psi]_A$ induced by $AA(\tau)$ together with the values of stable-for-A records.

2.2 Collapse as class-label update

BCT does not introduce an objective collapse force. Global Schrödinger dynamics may remain unitary throughout. "Collapse" is reinterpreted as the discrete update of $[\Psi]_A$ when stable accessible record structure changes.

2.3 Outcomes as record-defined macro-conditions

An outcome is not treated as a primitive "measurement result." It is a record-defined macro-condition: a stable, accessible record structure that refines the experienced state of the agent. Distinct micro-realizations need not imply distinct experienced states unless they split the relevant stable record structure for that agent.

2.4 Probability as constrained support over record-defined outcomes

In Phase 2, probabilities are assigned to record-defined macro-conditions, not to vague branch labels. Once the support assignment over those macro-conditions satisfies standard operational consistency constraints, the Born form follows by Gleason-style uniqueness. BCT's added content is not new probability mathematics, but the ontology of what the outcome projectors represent.

2.5 A killable empirical wedge

BCT also commits to a narrow empirical bridge: conditioned on verified which-way information at or above threshold, event-like definiteness should depend on sufficiently stable downstream record formation. This is the retention-driven bridge, not a general claim about all decoherence phenomena.

3. What BCT does not claim

- It does not modify Schrödinger dynamics.
- It does not claim that consciousness or semantics causes collapse.
- It does not promise a single omniscient experienced narrative valid for all agents at all stages.
- It does not present itself as a universal theory of everything.
- It does not claim that decoherence, Quantum Darwinism, or Spectrum Broadcast Structures by themselves already *are* BCT.

4. The core move, stated plainly

Two global physical states may differ microscopically while remaining the same experienced state for agent A, provided that:

- the differences are not physically accessible through $AA(\tau)$, and
- the values of stable-for-A records are the same.

This is the central BCT move. It allows many microstates to remain compatible with one experienced label. The theory becomes deliberately non-elastic at exactly this point: observer-definiteness is fixed by operational access and stable records, not by intuitive language.

5. Outcomes, support, and the Born route

In the Phase-2 canon, outcomes are represented as record-defined macro-conditions R_i, with associated macro-subspaces $\mathcal{H}_{R_i}$ and projectors Π_i. At the adopted coarse-grained resolution, these macro-conditions are treated as mutually exclusive and exhaustive.

The support assignment $\mu(\Pi \mid \Psi)$ is constrained by the usual operational requirements: non-negativity, normalization, additivity on orthogonal alternatives, non-contextuality at fixed record condition, and unitary covariance. Under those constraints, the Born form is uniquely selected. BCT's distinctive claim lies in the interpretive reading of those projectors: they represent stable, accessible, record-defined macro-conditions tied to experienced definiteness for an agent.

6. Relation to Many-Worlds and Wigner-type scenarios

BCT is compatible with global unitarity, so the real disagreement with Many-Worlds is not "unitarity versus collapse." The disagreement is about definiteness.

BCT denies the silent inference:

$$\text{many micro-alternatives} \;\Rightarrow\; \text{many experienced subjects.}$$

For BCT, distinct experienced labels for a given agent require distinct stable record conditions for that agent. This is why BCT allows "many microstates, one experienced label" until record structure actually splits.

The same logic reduces pressure in Wigner-type scenarios. Different agents may induce different equivalence relations because they do not have the same physical access or the same stable record channels at the same stage. Cross-agent agreement is therefore a problem of record propagation and access, not a demand for one omniscient experienced narrative.

7. Falsifiability: what would make BCT lose

Phase 2 commits in advance to defeat conditions along a specific bridge:

measurement-like interaction $\rightarrow$ record formation $\rightarrow$ record persistence $\rightarrow$ experienced definiteness.

The narrow test sector uses the following variables:

$$I_{\mathrm{WW}}, \quad I_c, \quad \tau_p, \quad \tau_{\mathrm{loss}}, \quad R, \quad s := R\tau_{\mathrm{loss}}.$$

The core prediction is not that any loss of visibility confirms BCT. The core prediction is narrower: once verified which-way information is genuinely present at or above I_c, visibility should become organized by downstream record stability.

BCT is willing to lose in that domain if carefully controlled tests show, for example:

- no systematic retention-organization near threshold under verified $I_{\mathrm{WW}} \geq I_c$;
- persistence-with-interference, namely high or recoverable visibility despite records that demonstrably persist beyond the relevant threshold;

- decoherence-only dominance, where ordinary arm dephasing fully explains the data and retention variables add no explanatory power;
- failure of s-scaling when R is varied.

A negative result is informative. The bridge is structured so that it can fail cleanly rather than survive by post-hoc parameter rescue.

8. Phase-2 discipline

Phase 2 is defined by four upgrades.

1. **Non-elastic identity.** Experienced definiteness is fixed by an explicit equivalence relation grounded in access and stable records.
2. **Record-defined outcomes.** Outcomes are stable macro-conditions, not primitive measurement labels.
3. **Conditional Born route.** Probability is treated as a support assignment over record-defined outcomes under standard operational constraints.
4. **Defeat conditions upfront.** The empirical wedge is preregisterable, calibration-sensitive, and designed to die cleanly if the bridge fails.

9. Document map

A sensible reading order for the DOI-main package is:

1. *Axiom 0 for Bit-Collapse Theory (BCT) and a Retention-Threshold Falsification Program*
2. *BCT Glossary (Canonical Terms)*
3. *Born-Rule Constraints from Stable Records and Accessible Algebra*
4. *Falsification Package — Phase-2*
5. *Bit-Collapse Theory (BCT) vs Many-Worlds (MWI)*
6. *Appendix B — Objections and Structural Responses*
7. *Why τ^* Is Not an Anthropic Parameter*

10. Closing

BCT is a bet on a specific kind of clarity: make identity physical, make collapse discrete but agent-relative, make outcome structure record-defined, and make defeat conditions explicit. If the retention-driven bridge fails under the declared controls, BCT should be considered wrong in that domain or in need of structural replacement, not merely rhetorical patching.

AI-use note. This document benefited from AI-assisted editorial and structural drafting support. The author remains solely responsible for all claims, definitions, and errors.