BCT Level-3 — Revised Working Draft

Coherence-Filtered Typicality, White-Rabbit Suppression, and Open Proof Obligations

Federico Batori
Working draft for collaboration — revised editorial version
Built on Phase-2 / Axiom0 commitments

Abstract

This document presents a revised working version of the BCT Level-3 program. Level-3 is not a new project phase. It is an exploratory theoretical layer built on the Phase-2 / Axiom0 framework and aimed at a narrower question: under a substrate-indifferent generative ontology, why should coherent, record-stable, law-like subject-histories dominate over arbitrary or pathological continuations?

The paper adopts four methodological constraints. First, subjects are modeled as record-chains rather than as primitive selves or unconstrained observers. Second, typicality claims are evaluated under an Occam-weighted generative measure rather than by raw multiplicity. Third, coherence is treated as a continuation constraint on stabilized records, not as a psychological or narrative notion. Fourth, all substantive claims are relative to a frozen version of the framework, so that post-hoc retuning counts against the program rather than in its favor.

Within that disciplined setting, the paper does three things. It states the Level-3 program in a version-frozen form, gives a rigorous partial result for generic single-step white-rabbit deviations under a local step-regularity assumption, and makes the remaining bottlenecks explicit. In particular, it extracts a quantitative sparse-irregular-step consequence already available from the V0 coherence budget, distinguishes density control from adversarial-time local closure, and sharpens the open tasks concerning predictor rigidity, family counting, and computable approximations.

This draft does not claim a full solution to the selection problem, a derivation of law-like physics, or a completed suppression theorem for all anomaly classes. Its current status is explicit: some results are proved, some are conditional, and several proof obligations remain open. That structure is a feature, not an embarrassment. Level-3 is presented here as a disciplined exploratory program, not as a finished framework.

Status note and taxonomy

BCT uses **Phase** for project stages and **Level** for theoretical layers.

- **Phase 1**: initial exploratory stage.
- **Phase 2**: Axiom0-anchored operational / canonical stage.
- **Level-3**: exploratory theoretical layer built on Phase-2 commitments.

The label **Level-3** is a local roadmap label. It does not imply that “Level-1” and “Level-2” exist as separately formalized public modules.

1. Scope and stance

This is not a Phase-2 physics-interface paper. It is a generative / ontological layer intended to motivate, and later reconnect with, physics-facing BCT constraints.

Three negative clarifications are important from the outset.

1. The paper does **not** assume a primitive classical measurement cut.
2. The paper does **not** privilege a specific physical substrate.
3. The paper does **not** claim that Level-3 has already derived law-like physics or fully solved white-rabbit suppression.

The target is narrower:

> Given a substrate-indifferent generative ontology, can one define a disciplined version of record-based typicality under which coherent, law-like subject-histories dominate over arbitrary continuations?

That question is exploratory but not unconstrained. The program is only meaningful if it can be version-frozen, defeated, and improved without hidden retuning.

1.1 Compatibility with Phase-2 canonical language

To keep terminology stable across BCT documents, this draft remains aligned with Phase-2 canonical usage:

- **Agent A**: an operationally defined physical system with finite resources and a finite integration window τ.
- **Accessible algebra** $\mathcal{A}_A(\tau)$: the observables physically accessible to A over τ.
- **Record** R: a physically instantiated correlation or finite record variable that can constrain later updates.
- **Persistence / loss scales**: τ_p, τ_{loss}, and threshold τ^* remain operational parameters.
- **Which-way threshold**: I_c or I_{WW} remains an operational threshold in Phase-2 contexts.

Level-3 does not replace this vocabulary. It extends it by asking what generative / typicality structure could make record-stable, law-like histories dominate.

1.2 Relation to adjacent programs

Any substrate-indifferent generative stance inherits familiar difficulties: a measure problem, reference-machine sensitivity, and a white-rabbit problem. The BCT-specific constraint is **non-elastic subjecthood**: a subject is a record-chain, and coherence is continuity of stabilized records under a fixed representational regime.

Measures that ignore record-based continuity, or require uncontrolled redefinition to preserve conclusions, count against the program rather than in its favor.

2. Core premises used in this paper

This paper works with three Level-3 premises.

- **P1. Unitary informational primitive.** There exists a single primitive domain I: a structured informational space.
- **P2. Substrate indifference.** If two implementations realize the same informational configuration up to structure-preserving isomorphism, they are identical-in-kind for BCT purposes.
- **P3. Exhaustive binary generation as a presentation.** An exhaustive binary generation program may be treated as an equivalent presentation of I, not as a second ontological layer "running inside" physics.

P3 is a Level-3 ontological stipulation, not a theorem. It is also the most controversial premise in the present paper. The draft does not claim to settle general objections to pancomputational or exhaustive-generation ontologies here. It only states the stipulation under which the current Level-3 program is explored.

3. Primitive, structure, and the status of the wave

Definition 3.1. Unitary structured informational space

Let I be a domain of informational configurations equipped with:

1. a compatibility / coherence structure; and
2. composition rules for combining compatible configurations.

No commitment is made here to a unique mathematical implementation. Hilbert-space objects may later appear as effective coordinates rather than primitive ontology.

A BCT-native move at this level is that the "wave" is not primitive. For a subject-configuration C, define the coherence region

$$W(C) \subseteq I$$

as the set or substructure of configurations coherent with C given its records and continuation constraints. A wavefunction-like object ψ_C is then understood as a chart, compression, or coordinate description of $W(C)$, not as the ontological base layer.

4. Subjects as record-chains

Definition 4.1. Record

A **record** is a finite informational structure that can:

1. be re-accessed,
2. constrain future updates, and
3. participate in a continuity condition for a subject-chain.

In Level-3 language, a record is a stability / constraint artifact within I. In Phase-2 language, it remains compatible with operational record language and finite accessibility.

Definition 4.2. Subject / experience chain

A **subject** is modeled as a chain of record-states

$$S := (R_0 \to R_1 \to R_2 \to \cdots),$$

where each R_t is a record-state and "$\to$" denotes an admissible update preserving the relevant continuation constraints.

Principle 4.3. Coherence as continuation constraint

Coherence is the requirement that admissible continuations preserve the non-contradictory integration of stabilized records, maintain a stable mapping between predictive structure and encountered inputs, and do not permit arbitrary rewriting of the past.

Coherence is therefore a record-chain continuation condition, not a psychological, narrative, or merely aesthetic notion.

5. Measure: Occam as a structural choice

In an exhaustive generative domain, raw multiplicity is not enough. One needs a measure.

This program adopts an Occam-weighted generative measure as a standard candidate under substrate indifference, prefix coding, and coding-theoretic compression principles. Without such weighting, incoherent fragments and arbitrary glitches can dominate and no stable notion of typicality survives.

Definition 5.1. Occam-weighted generative measure (sketch)

Let π be a finite binary description in a chosen reference machine U. Assign weight

$$w(\pi) \propto 2^{-|\pi|},$$

and induce weight on generated histories or structures in I by summing over descriptions that generate them.

The present paper does not claim that every reasonable measure must take exactly this form. It claims a narrower point: without a principled Occam-like weighting, the Level-3 program lacks a disciplined notion of dominance.

6. Main conjecture: coherence-selection typicality

Conjecture 6.1. Coherence-selection typicality (fixed version)

Fix a frozen version

$$V := (C, pred, D, X),$$

where:

- C is a record representation / codec family,

- $pred$ is a continuation predictor class,
- D is a discrepancy functional,
- X is a contradiction / instability penalty.

Under an Occam-weighted generative measure and the coherence discipline induced by V, the dominant weight on subject-chains concentrates on histories that maximize longitudinal coherence, thereby yielding law-like regularities as coherence constraints.

This is the central conjecture of the current Level-3 program. It is **not** fully proved in this draft.

The fixed-version condition is part of the claim, not a cosmetic add-on. Without it, the entire program collapses into uncontrolled retuning.

7. Coherence functional, filtering, and defeat conditions

7.1 Coherence as a functional template

Definition 7.1. Coherence functional template

Let $S = (R_t)_{t\geq 0}$ be a subject-chain. Define

$$K(S) := \sum_{t\geq 0} \big(\alpha D(R_{t+1}, pred(R_t)) + \beta X(R_t)\big),$$

with fixed $\alpha, \beta > 0$.

This is a schema, not yet a unique formalism. It becomes non-elastic only after fixing the record representation family, predictor class, discrepancy functional, and contradiction penalty.

7.2 Coherence-filtered Occam weight

Let π range over a prefix-free set of finite binary descriptions with

$$w(\pi) := 2^{-|\pi|}.$$

A hard coherence filter accepts only histories below a threshold θ:

$$\mathbf{1}_{coh}(\pi) := \mathbf{1}[K(S_\pi) \leq \theta].$$

A soft alternative is

$$\varphi_\lambda(\pi) := e^{-\lambda K(S_\pi)}.$$

Definition 7.2. Coherence-filtered Occam weight

Define the unnormalized weight of a subject-chain S by

$$W_{coh}(S) \propto \sum_{\pi : S_\pi = S} 2^{-|\pi|}\Phi(\pi),$$

where $\Phi(\pi)$ is either the hard filter or the soft penalty.

Depending on the setup, W_{coh} may be a semimeasure until normalized on the coherent subset. Dominance claims in this paper use relative weight.

7.3 Frozen-version discipline and defeat conditions

All major claims are relative to a fixed version $V = (C, pred, D, X)$ and bounded admissible perturbations. If qualitative conclusions require unbounded re-specification of these ingredients, the version is defeated.

The following conditions are currently best understood as **methodological defeat conditions**. They are important, but they are not all yet operationalized as full technical falsifiers.

- **Defeat 1. Reference-machine pathology.** If qualitative dominance claims depend on arbitrary re-description or encoding beyond bounded constants irrelevant to long-run dominance, the version fails.
- **Defeat 2. Measure arbitrariness.** If no principled Occam-like measure, or measure class, can be justified without smuggling in target physics, the version fails.
- **Defeat 3. Coherence under-determination.** If coherence constraints remain so elastic that radically different law-classes retain comparable weight under admissible changes of V, the version fails.
- **Defeat 4. White-rabbit dominance.** If a nontrivial class of glitch-rich histories remains dominant under W_{coh}, the version fails.

Defeat 3 and Defeat 4 are already part of the methodological logic of the program, but they are not yet fully reduced to closed technical criteria.

8. Heuristic failure-mode map

Heuristic Claim 8.1. White-rabbit suppression dichotomy

For a fixed frozen version V, a white-rabbit deviation must lose by at least one of the following routes:

1. **Occam loss:** its description complexity rises enough to suppress its weight;
2. **Coherence loss:** it violates the coherence budget or incurs downstream coherence penalties;
3. **Version defeat:** if neither loss occurs without retuning V, then the version is defeated.

This is a failure-mode map, not a general theorem.

9. Technical V0 instantiation and partial rigorous result

The rest of the draft fixes one narrow version V_0 so that suppression becomes a mathematically meaningful claim.

9.1 Formal setup (V0)

Fix:

- record length $n \in \mathbb{N}$,
- chain length $T \in \mathbb{N}$,
- memory depth k,

- a restricted MDL / Kolmogorov-style predictor $pred_k$,
- normalized Hamming discrepancy D.

A record-state is a binary string $R \in \{0,1\}^n$, and a subject-chain is

$$S = (R_0, R_1, \ldots, R_T) \in (\{0,1\}^n)^{T+1}.$$

Define the V0 coherence functional by

$$K_{V_0}(S) := \sum_{t=0}^{T-1} \alpha D(R_{t+1}, pred_k(t)) + \sum_{t=0}^{T} \beta X(R_t),$$

where $X(R_t)$ is a contradiction / instability penalty and $\alpha, \beta > 0$ are fixed.

Write "θ_K-coherent" for chains satisfying

$$K_{V_0}(S) \le \theta_K.$$

For the Occam term, use a prefix-free universal machine U and the induced semimeasure m_U. Then

$$W_{coh}(S) \propto m_U(S)\Xi(S),$$

where $\Xi(S)$ is either a hard coherence threshold or a soft coherence penalty.

9.2 White-rabbit deviation (single-step, V0)

Fix a θ_K-coherent "law-like" chain S_{law} and choose a time $t^* \in \{1, \ldots, T\}$. A chain S_{WR} is a (t^*, ε) white-rabbit deviation if it agrees with S_{law} except at time t^* and

$$D(R^{WR}_{t^*}, pred_k(t^* - 1)) \ge \varepsilon.$$

Define the implied average per-step discrepancy budget

$$\theta_D := \frac{\theta_K}{T\alpha}.$$

This quantity is only an average budget. It does **not** by itself yield a local step-regularity bound at an arbitrary chosen time.

9.3 Local regularity assumption for the generic-glitch theorem

The main quantitative suppression theorem requires an explicit local assumption.

Local step-regularity assumption at t^*. At the chosen glitch time,

$$D(R^{law}_{t^*}, pred_k(t^* - 1)) \le \theta_{step},$$

for some $0 < \theta_{step} < \varepsilon \le 1/2$.

This assumption is not currently derived from average coherence alone. Its status in this draft is explicit.

9.4 The generic-glitch theorem

Call the glitch at time t^* **generic** relative to S_{law} if its conditional Kolmogorov complexity is large:

$$K(R_{t^*}^{WR} \mid R_0^{law}, \dots, R_{t^*-1}^{law}) \geq nH_2(\varepsilon) - c_1 \log n.$$

Theorem 9.1. Occam suppression for generic single-step glitches

Let S_{law} be a θ_K-coherent chain under V_0. Let S_{WR} be a (t^*, ε) white-rabbit deviation with a generic glitch, and assume local step-regularity at t^*:

$$D(R_{t^*}^{law}, pred_k(t^* - 1)) \leq \theta_{step}, \qquad 0 < \theta_{step} < \varepsilon \leq 1/2.$$

Then

$$\frac{m_U(S_{WR})}{m_U(S_{law})} \leq 2^{-nH_2(\varepsilon) + nH_2(\theta_{step}) + c_2 \log(nT)}.$$

In particular, for fixed $\varepsilon > \theta_{step}$, the ratio tends to zero exponentially in n.

Interpretation. This is the main rigorous partial result of the present draft. It applies to generic single-step deviations under an explicit local regularity assumption.

9.5 Sparse irregular-step control

The average budget does not yield local regularity at every step. But it does imply a weaker and real consequence: sufficiently bad steps must be sparse.

For any threshold $\tau > 0$, define the irregular-step set

$$B_\tau(S) := \{t \in \{0, \dots, T-1\} : D(R_{t+1}, pred_k(t)) > \tau\}.$$

Lemma 9.2. Sparse irregular-step bound

Assume $X(R_t) \geq 0$ for all t. If S is θ_K-coherent under V_0, then for every $\tau > 0$,

$$|B_\tau(S)| \leq \frac{\theta_K}{\alpha \tau}.$$

Equivalently,

$$\frac{|B_\tau(S)|}{T} \leq \frac{\theta_D}{\tau}.$$

So average coherence does not deliver a local bound at every step, but it does bound the density of irregular steps.

Corollary 9.3. Most-time generic suppression

Let S_{law} be θ_K-coherent under V_0. Fix $\tau > 0$, and choose a time $t^* \notin B_\tau(S_{law})$. Then local regularity holds with $\theta_{step} = \tau$, and every generic (t^*, ε) white-rabbit glitch with $\varepsilon > \tau$ obeys

$$\frac{m_U(S_{WR})}{m_U(S_{law})} \le 2^{-nH_2(\varepsilon)+nH_2(\tau)+c_2\log(nT)}.$$

This does not close the average-to-local problem in the adversarial-time sense. But it does show that generic suppression applies automatically outside a quantitatively controlled sparse exceptional set.

9.6 The simple-glitch gap

The generic theorem does not cover anomalous but low-complexity glitches.

Example 9.4. Simple glitch far from prediction

Suppose the pre-glitch chain supports a simple alternating pattern so that the predictor expects 0^n at time t^*. If the realized glitch is 1^n, then the Hamming discrepancy is maximal but the string 1^n itself has low Kolmogorov complexity.

So the genericity assumption fails even though the deviation is maximally far from the predictor.

This is a real gap, not a cosmetic corner case.

9.7 Predictor rigidity as a conditional closure route

Definition 9.5. Predictor rigidity

A predictor class $pred$ is (δ,λ)-rigid for a law-like chain S_{law} if, for any deviation at time t^* with

$$D(R^{WR}_{t^*}, pred(t^*-1)) \ge \varepsilon > \delta,$$

the updated predictor $pred'$, re-fitted on the post-glitch history, satisfies

$$K(pred') - K(pred) \ge \lambda n(\varepsilon-\delta) - O(\log(nT)).$$

Informally: a predictor cannot cheaply absorb a sufficiently far-from-law event.

At present, rigidity is not claimed in full generality. The real technical target is narrower: prove such a bound for an explicit restricted predictor subclass with a canonical code.

Proposition 9.6. Gap closure under predictor rigidity

Assume the V0 predictor is (δ,λ)-rigid and the law-chain is locally regular at the glitch time:

$$D(R^{law}_{t^*}, pred(t^*-1)) \le \delta.$$

Then for any (t^*,ε) white-rabbit deviation with $\varepsilon>\delta$, one of the following holds:

1. $W_{coh}(S_{WR})=0$, or 2.

$$\frac{m_U(S_{WR})}{m_U(S_{law})} \le 2^{-\lambda n(\varepsilon-\delta)+c_6\log(nT)}.$$

This closes the simple-glitch gap only **conditionally**.

9.8 Aggregate suppression and its genuine limit

Per-chain suppression does not automatically imply concentration against exponentially large deviation families.

Theorem 9.7. Aggregate suppression under an entropy-gap condition

Let $\mathcal{F}_{n,T}$ be a family of candidate white-rabbit chains such that, for each $S \in \mathcal{F}_{n,T}$,

$$\frac{m_U(S)}{m_U(S_{law})} \le 2^{-cn + c_0 \log(nT)}$$

for some constants $c > 0$, c_0, uniformly in S. Assume moreover that

$$\log_2 |\mathcal{F}_{n,T}| \le \beta n + o(n), \qquad \beta < c.$$

Then

$$\frac{\sum_{S \in \mathcal{F}_{n,T}} m_U(S)}{m_U(S_{law})} \to 0$$

up to the same polynomial corrections.

This is a correct aggregation lemma. It is **not** a proof of full concentration for the most natural exponentially large families.

For the family of all single-step ε-glitches at fixed t^*, the family size typically saturates the same entropy scale appearing in the generic-glitch theorem. That is why the counting problem remains real.

10. What is proved, what is conditional, what is open

10.1 Proved in this draft

- A disciplined Level-3 framework with a frozen-version rule and explicit defeat logic.
- A concrete V0 instantiation of the coherence / Occam machinery.
- A quantitative Occam-suppression theorem for generic single-step white-rabbit glitches under local step-regularity.
- A sparse-irregular-step bound derived from the average V0 coherence budget.
- A most-time version of the generic suppression result.
- A correct aggregation lemma showing when per-chain suppression survives summation over a family.

10.2 Conditional in this draft

- Suppression of simple glitches is conditional on predictor rigidity.
- The generic-glitch theorem at an adversarially chosen time still depends on a local regularity assumption not derived from average coherence alone.
- The broader coherence-selection typicality conjecture remains only partially grounded.

10.3 Open proof obligations

The open tasks are best separated into narrower components.

- **O1a. Canonical-code injection lemma.** For an explicit frozen predictor subclass, prove that post-glitch refitting which absorbs a far-from-law deviation must pay a lower-bounded increase in canonical description length.
- **O1b. Rigidity theorem for the chosen subclass.** Deduce the linear lower bound needed for predictor rigidity.
- **O2a. Density control from average coherence.** This is essentially extracted in the present draft through sparse irregular-step control.
- **O2b. Adversarial-time local closure.** Derive, justify, or axiomatize a local regularity bound that applies at a chosen glitch time, not merely outside a sparse exceptional set.
- **O3. Multi-step white rabbits.** Extend the present arguments to multiple and possibly correlated glitches.
- **O4. Reference-machine robustness.** Show sign-stability of suppression gaps under structurally neutral machine changes up to lower-order terms.
- **O5. Dominance of coherent histories.** Go beyond anomaly suppression and show that coherent law-like chains retain non-negligible weight.
- **O6. Computable approximations.** Replace uncomputable Kolmogorov / MDL ingredients with computable proxies and quantify the degradation.
- **O7. Natural-family counting gap.** Close the entropy gap for exponentially large deviation families without hiding the combinatorics.

For bookkeeping, O7 itself is best understood as splitting across distinct subfamilies: generic glitches at regular times, generic glitches at irregular times, simple low-complexity glitches, and multi-step or structurally correlated deviations.

11. Discussion

This draft is deliberately narrow. It does not claim that Level-3 has derived quantum formalism from scratch, solved all measure problems, or established a universal white-rabbit suppression theorem.

Its current value lies elsewhere: it converts a vague program into a frozen-version framework with explicit defeat logic, one nontrivial rigorous partial theorem, one extracted intermediate consequence from average coherence, and a clean map of the remaining bottlenecks.

That is enough to justify circulation for targeted collaboration.

12. Conclusion

This revised working draft does four things.

1. It states the Level-3 program in a disciplined frozen-version form.
2. It proves a genuine partial suppression result for generic single-step deviations under an explicit local assumption.
3. It extracts a weaker but rigorous most-time control from average coherence.
4. It sharpens the remaining proof obligations instead of hiding them inside vague future work.

The immediate mathematical tasks are now more precise than before: prove or refute predictor rigidity for a restricted predictor subclass, and close the remaining adversarial-time and family-counting gaps without violating the frozen-version discipline.

That is the present status of Level-3: not a completed derivation of law-like reality, but a disciplined exploratory program with one real theorem, one real intermediate consequence, and clearly named bottlenecks.

Appendix A. Bridge note: record-broadcast and physics-facing anchors

The current technical V0 machinery does not formally integrate a Quantum Darwinism / Spectrum Broadcast Structures record-broadcast condition. That bridge remains important, but it should not be confused with the present theorem set.

A natural future bridge condition would require two jointly satisfied constraints for subject-definite recordhood:

1. **Public-record condition:** the record is redundantly encoded and publicly recoverable in the relevant environment fragments.
2. **Chain-stabilization condition:** the record persists long enough, and coherently enough, to stabilize subject-chain continuation.

This appendix is therefore bridge-oriented, not part of the current V0 proof core.

Appendix B. Optional future fork: adaptive predictor path

One possible future extension is an explicitly new frozen version V_1, in which one evaluates not only a history S but also a predictor path $(p_t)_{t=0}^{T}$, with an explicit retuning cost between successive predictor states.

Schematically,

$$K_{V_1}(S,(p_t)) = \sum_{t=0}^{T-1} \alpha D(R_{t+1},p_t) + \eta \sum_{t=0}^{T-1} \Delta(p_{t+1},p_t) + \sum_{t=0}^{T} \beta X(R_t) + \gamma G(S).$$

This adaptive fork is not part of the current V0 core. It is mentioned only to show a disciplined extension path for internalizing the D1 / D2 tradeoff without covertly altering the current version.

AI-use note

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