BCB Inevitability Programme: Records, Commitment, TPB, and Restricted Structural Uniqueness

---

Executive Summary

Question. How much of physics is forced by the demand that experiments have repeatable outcomes?

Axioms. Six — empirical meaning, stable records, finite distinguishability, microscopic reversibility, low-energy locality, consistency — none of which a working theory of physics can deny.

What is proven (Theorems A–F, Theorem U):

- Repeatable experiments ⇒ physical records ⇒ coarse-graining ⇒ irreversibility ⇒ minimum entropy cost per bit
- Reversible substrate + irreversible records ⇒ a Ticks-Per-Bit conversion invariant (TPB) with finite admissibility window
- Finite observational resolution ⇒ EFT-type local description (under standard Wilsonian conditions)
- Stable solitonic excitations in 3D ⇒ a Skyrme-type stabiliser, unique within its derivative class
- Redundancy of microdescription ⇒ gauge symmetry with compact Lie group
- **Restricted uniqueness (Theorem U):** within fold-based EFTs with stable records, the BCB architecture is forced

Axiom translation to gauge papers (Lemmas H3.1–H3.9):

- Axiom 4 ⇒ BC1 (distinguishability continuity) — formally derived
- Axiom 3 + observational equivalence ⇒ BC2 (entropy minimisation) — operationally grounded
- Axiom 6 ⇒ BC3 (anomaly cancellation) — formally derived
- Axioms 1 + 3 + Čencov monotonicity ⇒ FIM (Fisher–Rao geometry) — formally derived

What remains open (H1–H3):

- H1: Derive TPB window from substrate microphysics (eliminates parametric freedom)
- H2: Solve fold eigenproblem → three generations spectrally forced (eliminates generational ambiguity)

- H3: Complete axiom translation → SU(3)×SU(2)×U(1) internal to the programme (eliminates gauge ambiguity)

Conditional structural uniqueness (Theorem SU): Under closure of H1–H3, any admissible theory in the defined class is equivalent to the BCB architecture up to field redefinitions, operator basis choice, gauge redundancy, and UV completion above cutoff. No additional structural freedom remains.

Abstract

This document asks a simple question: *how much of physics is optional?*

Most theories of physics begin with assumptions — a set of equations, a choice of symmetry group, a menu of particles — and then check whether those assumptions match experiment. The BCB programme inverts that process. It starts from the most minimal possible demand — that a theory of physics must be capable of producing repeatable experimental results — and asks what structures are *forced* by that demand alone.

The answer, developed through a chain of theorems in this document, is that a surprising amount of familiar physics is not a choice but a consequence. If experiments must have outcomes, something physical must store those outcomes: records. If the underlying microdynamics is probability-consistent (normalisation-preserving) and invertible at the substrate level, then records are irreversible at the macroscopic level — they cost entropy. The ratio of substrate ticks to irreversible record commitments defines a conversion invariant (Ticks Per Bit, or TPB) that any such theory must contain. The requirement that observers have finite resolution forces an EFT-type local description at accessible scales under standard Wilsonian conditions (locality, cluster decomposition, and finite resolution). And if the theory's stable structures are solitonic (fold-like), then the mechanism that prevents them from collapsing is essentially unique.

The document also identifies, with equal precision, what *cannot* yet be proven — most importantly, why the specific symmetry group of the Standard Model (rather than some other group) is the one that emerges. This gap is not hidden; it is stated as an open problem with a defined path to resolution.

The result is a restricted uniqueness theorem: within a well-defined class of theories, the architectural elements of BCB — records, commitment, TPB, entropy-weighted time, Skyrme-stabilised folds, and SMEFT-compatible effective field theory — are not modelling choices. They are structural necessities. The document maps every logical dependency, flags every open problem, and draws a clear line between what is proven and what remains to be done.

What this document adds to the BCB Lagrangian paper

The companion paper — *BCB Fold v3: A Complete Lagrangian Field Theory* — constructs BCB as a calculable quantum field theory. It writes down the explicit action, computes electron mass, proton structure, Higgs VEV, strong coupling, hypercharges, and three-generation structure, and achieves a 60–67% reduction in free parameters relative to the Standard Model. That paper answers the question: *what does BCB predict?*

This document answers the question the Lagrangian paper cannot: *why this theory and not some other?*

The Lagrangian paper takes records, commitment, entropy costs, TPB, fold stability, Role-4, and EFT structure as its starting ingredients. This document derives every one of those ingredients from six axioms that no empirically meaningful theory can deny. Specifically:

- **TPB** moves from a defined bookkeeping variable to a proven structural necessity — any admissible theory with reversible microdynamics and irreversible records must contain an equivalent object (Theorem C). A referee reading the Lagrangian paper can ask "why do I need TPB?" After this document, the burden shifts: TPB (or an equivalent conversion invariant) becomes a structural requirement of any admissible record-forming framework in the defined class.
- **The Skyrme stabiliser** moves from "chosen because it works" to "proven unique within its derivative class" (Theorem E, Proposition E.2). The Lagrangian paper shows folds are stable. This document proves no other stabiliser could do the job.
- **Gauge symmetry** moves from a constructive argument (Fisher geometry on CP^2 gives you the gauge group) to an abstract derivation (any theory with redundant microdescriptions must have gauge symmetry with a compact Lie group — Theorems F1–F2).
- **EFT locality** moves from an adopted methodology to a motivated consequence of finite observational resolution under standard conditions (Proposition D).
- **The formal theory class $\mathcal{T}$** is entirely new. The Lagrangian paper never defines what BCB is competing against. This document defines the precise class of admissible theories and proves results within it — which is the only context in which "uniqueness" means anything.
- **The restricted uniqueness theorem (Theorem U)** does not exist in the Lagrangian paper at all. It assembles the full chain and proves: any theory in the defined class must contain a TPB-like invariant, an entropy-weighted time mapping, a Skyrme-like stabiliser, and a SMEFT-compatible EFT expansion.

The Lagrangian paper is the calculable theory — it computes, predicts, and can be tested against experiment. This document is the structural argument that the theory's architecture is forced within its class. Together, they establish that BCB is both calculable and structurally necessary within the defined theory class.

Relationship to the BCB Fold Lagrangian Paper

This document is the companion to *Bit Conservation and Balance Fold v3: A Complete Lagrangian Field Theory Unifying Particle Physics and Gravity from Information-Theoretic*

Principles (the Lagrangian paper). The two documents answer different questions, and neither is complete without the other.

What the Lagrangian paper does

The Lagrangian paper constructs BCB as a calculable quantum field theory. It writes down the explicit action $S = \int d^4x \sqrt{-g}\, \mathscr{L}_{BCB}$, specifies every field, coupling, and interaction term, and then computes: electron mass from fold geometry, proton mass and radius from energy minimisation, Higgs VEV from a Planck-rooted derivation chain, strong coupling α_s from CP^2 curvature, hypercharges uniquely from anomaly cancellation plus fold stability, and the three-generation structure from a conditional eigenvalue theorem. It provides Feynman rules, gauge-fixing procedures, Faddeev–Popov ghosts, renormalisation prescriptions, and a systematic EFT expansion. It achieves a 60–67% reduction in free parameters relative to the Standard Model, with roadmaps toward 90–93%.

This is a major achievement. It transforms BCB from a conceptual framework into a theory that can be coded, simulated, loop-corrected, and compared quantitatively with every measurement in the Particle Data Group tables. It answers the question: *what does BCB predict, and how do you calculate it?*

What the Lagrangian paper cannot do

The Lagrangian paper is vulnerable to a specific class of objections that no amount of successful prediction can address: *why this Lagrangian?*

A referee can always ask: why did you choose these fields? Why this stabiliser? Why this time-mapping? Why this commitment structure? Why folds rather than strings, or loops, or causal sets? The Lagrangian paper can answer "because it works" — it produces the right masses, the right couplings, the right number of generations. But "because it works" is empirical adequacy, not structural necessity. It is the same justification the Standard Model uses for its 25 free parameters. A theory that works is valuable. A theory that *must* work — that is the only theory compatible with the requirement that experiments have outcomes — is something categorically stronger.

What this document does that the Lagrangian paper cannot

This document answers the "why this Lagrangian?" question. It does not start from the BCB action and check its predictions. It starts from six axioms that no empirically meaningful theory can deny, and derives the BCB architecture as a consequence. The difference is the difference between *consistency* and *inevitability*.

Specifically, the Inevitability Programme upgrades each foundational element of the Lagrangian paper from an assumption to a theorem:

Element	Status in Lagrangian paper	Status in this document
Records and commitment	Assumed as starting point	Derived from Axiom 1 (Theorem A, Lemma A1, Corollary A2)
Entropy cost per bit (ε_bit)	Defined and used	Proven as structural necessity (Theorem B)
Ticks Per Bit (TPB)	Defined as a conversion ratio	Proven to exist in any admissible theory (Theorem C)
TPB admissibility window	Assumed finite	Motivated from operational requirements (Proposition C1)
EFT structure	Adopted as methodology	Motivated by finite observational resolution under standard conditions (Proposition D)
Skyrme-type stabiliser	Chosen because it works	Proven unique within minimal derivative class (Theorem E, Proposition E.2)
Entropy-weighted time (Role-4)	Written into the Lagrangian	Shown to be a canonical consequence of reversible substrate + irreversible records
Gauge symmetry (abstract form)	Motivated by Fisher geometry	Derived from redundancy of microdescription (Theorems F1–F2)
Gauge group (specific: SU(3) × SU(2) × U(1))	Assumed as starting point	Derived in companion gauge papers from related axioms (BC1–BC3 + FIM) + three minimal empirical anchors; axiom translation to Axioms 1–6 identified as remaining task
SMEFT compatibility	Demonstrated by operator crosswalk	Expected as natural consequence of EFT locality + gauge structure

The Lagrangian paper shows that BCB *is* consistent. This document shows that BCB *must be* — within the defined theory class.

How they work together

The relationship is architectural. This document establishes *why* the structure exists; the Lagrangian paper establishes *what* the structure predicts. A reader who encounters only the Lagrangian paper will see a powerful and predictive theory but may reasonably ask whether its foundations are arbitrary. A reader who encounters only this document will see a rigorous inevitability argument but will not know whether it leads to a theory that actually matches experiment. Together, they form a complete package: structural necessity (this document) plus empirical adequacy (Lagrangian paper).

A third component — the gauge paper trilogy (*Minimal Internal Symmetry Theorem*, *Distinguishability Conservation and Gauge Structure*, *Closing the Interfaces*) — addresses the specific gauge group derivation that neither this document nor the Lagrangian paper fully resolves on its own. Theorems F1–F2 of this document derive abstract gauge structure; the gauge papers derive the specific group SU(3) × SU(2) × U(1) from a closely related axiom framework

(BC1–BC3 + FIM). The remaining task — translating between the two axiom sets — is identified in Section H3 as a dictionary problem, not a conjecture.

The documents share open problems at different levels. The independent TPB window derivation (H1) and three-generation spectral proof (H2) correspond to the "Critical Calculations Needed for Rigor" in Section 16.0 of the Lagrangian paper. The gauge group integration (H3) corresponds to the open problems listed in Section 8.3 of the gauge papers (Yang–Mills dynamics, symmetry breaking, generation structure, coupling constants). In each case, the Lagrangian paper frames them as calculations that would strengthen predictions; this document frames them as theorems that would close the gap between restricted and full uniqueness; the gauge papers frame them as extensions of the foundational infrastructure. They are the same problems viewed from different angles.

Why this matters for BCB

The structural results derived in this document are not abstract philosophical curiosities. They align directly with the architecture of BCB.

BCB already contains a reversible substrate ordering parameter, irreversible record commitments, a conversion invariant between substrate progression and commitment accumulation (TPB), an entropy-weighted calibration of operational time, a Skyrme-type stabiliser in its fold sector, a compact Lie gauge structure, and an effective field theory expansion at accessible scales.

What this document shows is that these are not arbitrary design choices. They are precisely the structures that any admissible theory with stable records and reversible microdynamics must contain. The structural features proven inevitable here — records, entropy cost, a substrate-to-record conversion invariant (TPB), entropy-weighted time calibration, fold stabilisation, and compact Lie gauge symmetry — are not merely compatible with BCB. They are the defining ingredients of BCB. The inevitability programme therefore does not merely tolerate BCB; it structurally selects its architecture within the admissible theory class.

This does not yet prove that BCB is the only possible theory in all of physics. But it does show that BCB inhabits the uniquely admissible region of theory space defined by the axioms — which is a much stronger position than simple empirical fit. The Lagrangian paper shows BCB works. This paper shows why something like BCB must exist. The overlap is not accidental.

Table of Contents

A. The Admissible Theory Class

Without defining the class of theories over which uniqueness is claimed, the word "unique" is meaningless. This section defines the battlefield.

Definition (Admissible Empirical Theory Class $\mathcal{T}$)

A theory $T \in \mathcal{T}$ is admissible if it supplies:

Axiom 1 (Repeatable empirical content). There exist physically realisable procedures whose outcomes can be compared across times, places, and agents. This is the minimal requirement for a theory to make contact with experiment.

Axiom 2 (Stable records). There exist physical degrees of freedom that can encode outcomes in a way that remains available to later interactions. Without this, Axiom 1 is vacuous — comparison requires something to compare against.

Axiom 3 (Finite distinguishability). For any finite region and finite resources, the number of mutually distinguishable record states is finite. This excludes theories that require infinite precision to extract finite empirical content.

Axiom 4 (Microscopic reversibility). The fundamental microdynamics is invertible (or measure-preserving), even if emergent dynamics is irreversible.

Axiom 5 (Low-energy locality). Interactions at accessible scales admit a local effective description (EFT locality).

Axiom 6 (Consistency and stability). The effective description admits a positive-norm state space (no negative-residue propagating modes below cutoff) and a stable — or sufficiently long-lived metastable — vacuum.

Convention (Observational equivalence quotient). Throughout, theories are identified up to observational equivalence on the admissible protocol class: if two internal descriptions yield identical outcome distributions for all admissible protocols, they represent the same physical theory. Equivalently, internal structure that is provably unobservable under the admissible protocol class is quotiented out as redundancy. This convention implements the non-redundancy principle used in Proposition H3.3 without introducing an additional axiom.

BCB + Role-4 + TPB + folds is a member of $\mathcal{T}$. The uniqueness ambition is: within $\mathcal{T}$, BCB-type structure is forced (up to field redefinitions and equivalent formulations). Uniqueness is always understood up to field redefinitions, redundancy equivalences, and EFT operator-basis choices related by equations of motion.

Axiom 4: Anticipated Challenges and Defence

Axiom 4 carries more weight than the others in the theorem chain — Corollary A2 depends on it entirely — and it is the axiom most likely to be challenged by a referee. The defence has two layers.

First layer (probability consistency). In any framework that assigns probabilities to experimental outcomes, the evolution rule must preserve normalisation and positivity of probabilities across time. In quantum mechanics this is unitarity; in classical statistical mechanics it is measure-preservation; in more general probabilistic theories it is a normalisation-preserving evolution on states. Axiom 4 is therefore not an additional dynamical hypothesis but a consistency constraint: without probability-preserving micro-evolution, the theory cannot maintain a coherent link between preparation procedures and outcome frequencies, undermining Axiom 1.

Second layer (classical substrate). Even in a pre-quantum or substrate-level formulation where the language of unitarity does not directly apply, microscopic reversibility is equivalent to the requirement that the fundamental dynamics preserves phase-space volume (Liouville's theorem for Hamiltonian systems, or measure-preservation more generally). This is the standard assumption underlying statistical mechanics. Theories that violate it at the fundamental level require an external entropy sink with no physical realisation — which, again, conflicts with the self-contained character required by Axiom 1.

A referee who rejects Axiom 4 is effectively proposing that fundamental information destruction is physically realisable. The burden of proof falls on that claim, not on the assumption of reversibility.

B. Theorem Chain: Empirical Meaning Forces Records, Commitment, and Entropy

Theorem A (Empirical Meaning ⇒ Record Primitive)

Statement. If a theory supports repeatable empirical facts (Axiom 1), then it must contain a primitive that functions as a record (Axiom 2): a physical degree of freedom whose state can correlate with outcomes and remain accessible for later comparison.

Proof sketch. Repeatability means that an outcome at stage 1 can be compared against an outcome at stage 2. Comparison requires some physically instantiated carrier of the earlier outcome — otherwise the earlier outcome is not operationally defined. A mental record held by an observer is itself a physical state of a physical system (the observer's brain, a notebook, a detector register). Therefore, any empirically meaningful theory must instantiate records as physical degrees of freedom.

Note on strength. This theorem is deliberately close to tautological. That is a feature, not a weakness. The foundations of an inevitability argument should be claims that no reasonable interlocutor can deny. Theorem A says: if your theory makes predictions that can be checked, something physical must store the prediction until it is checked. Denying this is denying that experiments have outcomes.

Lemma A1 (Record ⇒ Many-to-One Coarse-Graining)

Statement. A stable record requires a many-to-one map from microstates to macrostates.

Proof sketch. Stability means: small perturbations in microscopic configuration do not erase the recorded macrostate. Consider a record encoding a single macrostate $r \in R$. If r corresponded to a unique microstate $\omega \in \Omega$ (a one-to-one map), then any perturbation $\delta\omega$ that moves the system to $\omega' \neq \omega$ would generically change the macrostate, since the map has no redundancy to absorb the perturbation. Therefore, stable records require that each macrostate r corresponds to a set of microstates — a coarse-graining map:

$$\pi : \Omega \to R$$

where Ω is the substrate microstate space and R is the finite set of record macrostates (finite by Axiom 3). π must be many-to-one on typical trajectories; otherwise records are unstable against thermal noise.

Quantitative refinement. The size of the preimage $|\pi^{-1}(r)|$ determines the noise tolerance of the record. A record with $|\pi^{-1}(r)| = 1$ is destroyed by any perturbation. A record with $|\pi^{-1}(r)| = k$ tolerates perturbations that remain within the k-element basin. The minimum basin size for stability against thermal fluctuations at temperature T is set by the Boltzmann ratio: $|\pi^{-1}(r)|$ must be large enough that the free-energy barrier between basins exceeds kT by a margin sufficient for the desired record lifetime. This connects record stability directly to thermodynamic quantities.

Corollary A2 (Many-to-One Coarse-Graining + Invertible Microdynamics ⇒ Emergent Irreversibility)

Statement. Under invertible microdynamics (Axiom 4), stable records imply emergent irreversibility at the record level.

Proof sketch. Let microdynamics be invertible: $U : \Omega \to \Omega$ is a bijection. The coarse-grained record dynamics would be:

$$\tilde{U} := \pi \circ U \circ \pi^{-1}$$

But π^{-1} is not well-defined as a function (it is a set-valued map, since π is many-to-one). Therefore $\tilde{U}$ is not well-defined as an invertible map on R. Concretely: if two microstates ω_1, ω_2 both map to record state r, their images $U(\omega_1)$ and $U(\omega_2)$ may map to different record states under π, or the same one. The coarse-grained dynamics is generically many-to-one — multiple micro-histories merge into the same record trajectory. This is emergent irreversibility: information about the microstate is lost at the record level, even though it is preserved at the substrate level.

This is the core commitment move. A record is a coarse-grained sink. Once the system enters the basin of a record macrostate, the micro-level information about which specific microstate it occupied is no longer accessible at the record level. The commitment is irreversible *at the level at which records are read*, even though the underlying dynamics is reversible.

Connection to the second law. This corollary is the information-theoretic core of the second law of thermodynamics, stated without reference to thermodynamic potentials or equilibrium assumptions. The increase of entropy is a consequence of coarse-graining under invertible dynamics — not an independent law.

Theorem B (Irreversible Commitment ⇒ Minimum Entropy Cost per Bit)

Statement. If a system performs a stable two-class commitment (a 1-bit record), then the entropy production satisfies a lower bound:

$\Delta S \geq \varepsilon_{bit} > 0$

where ε_{bit} is the minimum entropy cost per committed record, bounded below by ln 2 (in units with $k_B = 1$) in the standard cyclic (write–erase–reset) setting, and bounded below by a strictly positive constant in one-way commitment settings. The exact constant depends on how the committed macrostate is stabilised and how the unchosen macrostate's accessibility is removed.

Proof sketch (Landauer-style, standard case). A stable 1-bit record maps a larger set of micro-configurations into one of two macrostates {0, 1} that remain distinguishable under noise. Making that mapping robust against perturbations requires a logically irreversible reduction in accessible microstate measure associated with committing to one of two robustly distinguishable macrostates, yielding $\Delta S \geq \ln 2$. Any physical implementation can exceed this bound; none can go below it while maintaining stability and reset capability.

Extension to one-way commitment (BCB-native case). The standard Landauer bound assumes that the record system must support cyclic operation — write, read, erase, repeat. In the BCB substrate architecture, commitment events are one-way: once a bit is committed, it is never erased at the substrate level. This relaxes the Landauer bound because the erasure step (which is what actually forces the ln 2 cost in Landauer's original argument) is absent.

However, the entropy cost does not vanish. Even without erasure, the act of commitment requires the system to transition from a state where both macrostates are accessible (pre-commitment) to a state where only one is occupied (post-commitment). This transition requires the system to shed the phase-space volume associated with the unchosen macrostate.

More generally, the entropy production is bounded below by the logarithmic contraction of accessible microstate volume induced by the commitment map — i.e. by the reduction in the measure of the accessible set under coarse-graining. For symmetric binary commitment this reduces to the familiar ln 2 scale; for asymmetric or one-way commitments the bound remains strictly positive but implementation-dependent.

Lemma B1 (One-way commitment entropy lower bound). Let $\Omega\_pre \subset \Omega$ denote the accessible microstate set prior to commitment and $\Omega\_post \subset \Omega$ the accessible set conditioned on the committed macrostate (the basin selected by the commitment). Let μ be the invariant micro-measure under reversible dynamics. Then any physically realisable one-way commitment that maps Ω_pre into Ω_post with stability against perturbations requires an entropy production bounded below by the logarithmic contraction of accessible measure:

$$\Delta S \geq \ln(\mu(\Omega\_pre) / \mu(\Omega\_post))$$

For symmetric binary commitment with $\mu(\Omega\_post) = ½ \mu(\Omega\_pre)$, this reduces to $\Delta S \geq \ln 2$. For asymmetric basins the bound remains strictly positive whenever $\mu(\Omega\_post) < \mu(\Omega\_pre)$. Deriving $\mu(\Omega\_post)$ for the substrate's record basins from the specific basin geometry is a closure subtask; the structural guarantee $\varepsilon\_bit > 0$ holds independently of that calculation.

What this establishes: Record formation has a nonzero minimum entropy cost. The natural scale is ln 2 for symmetric commitments. The precise bound for asymmetric one-way commitments is implementation-dependent, but $\varepsilon\_bit > 0$ is guaranteed by the phase-space contraction inherent in any genuine many-to-one map. This is the thermodynamic price of empirical meaning.

C. TPB as a Derived Quantity

Definition (Commitment Density and TPB)

Let $N\_bit(\tau)$ count the number of committed record-events along the reversible substrate parameter τ. Define the **commitment density** along the substrate parameter:

$$\rho\_commit(\tau) := d\, \mathbb{E}[N\_bit] / d\tau$$

This is the expected number of irreversible commitments per substrate tick — a purely combinatorial quantity requiring no notion of time.

The **Ticks Per Bit** ratio is its reciprocal:

$$TPB(\tau) := d\tau / d\, \mathbb{E}[N\_bit] = \rho\_commit(\tau)^{-1}$$

This is the average number of substrate ticks required to produce one committed record. Both quantities are defined entirely in terms of τ and event counts. No time variable is assumed or required.

Note on emergent time. Any experimentally relevant "rate" enters only after introducing the emergent monotone mapping $dt\_phys = f(s)\, d\tau$ (Role-4, Section G). The experimentally observable commitment frequency is then:

$$d\,\mathbb{E}[N_{bit}] / dt_{phys} = (1 / f(s)) \cdot \rho_{commit}(\tau) = 1 / (f(s) \cdot TPB(\tau))$$

Prior to that mapping, all statements are expressed purely in terms of τ-progress and commitment counts. The ordering of commitments along the reversible substrate progression is well-defined without time; the calibration of that ordering into an operational time variable is emergent.

Theorem C (Empirical Meaning + Reversible Substrate ⇒ Existence of a TPB-like Conversion Invariant)

Statement. In any theory satisfying Axioms 1–4 (empirical meaning, stable records, finite distinguishability, microscopic reversibility), there exists a statistical conversion invariant relating micro-parameter progress to record increments, equivalent to a TPB up to monotone reparametrisations of τ.

Proof sketch. From Theorem A, records exist. From Lemma A1, records are coarse-grained. From Corollary A2, record formation is irreversible. From Theorem B, each record event costs entropy $\varepsilon_{bit} > 0$. Because record formation events occur irregularly along the reversible micro-evolution parametrised by τ, any attempt to relate micro-parameter progress to macro-commitment count requires a conversion invariant. This is the commitment density $\rho_{commit}(\tau)$ and its reciprocal TPB(τ):

$$\rho_{commit}(\tau) = d\,\mathbb{E}[N_{bit}] / d\tau\ , \quad TPB(\tau) = \rho_{commit}(\tau)^{-1}$$

Reparametrising τ by any monotone function τ' = g(τ) rescales the numerator but cannot eliminate the need for the ratio itself: without such a conversion, the ordering of commitments along the reversible substrate progression is well-defined, but its calibration into an operational time variable requires a conversion invariant (TPB) and an entropy-weighted mapping (Role-4). A theory that lacks this conversion invariant cannot make predictions about how commitment events distribute along the micro-evolution.

Therefore, any admissible theory in the sense of $\mathcal{T}$ that has a reversible substrate and stable records must contain a TPB-like object. The object is forced by the structure of the problem; naming it is convention.

Significance. This is a genuine inevitability result. TPB is not a modelling choice imposed on the theory — it is a structural necessity derived from the requirement that the theory be empirically meaningful. Any framework that claims to describe physics with reversible microdynamics and irreversible record formation either contains an equivalent of TPB (possibly under a different name) or is incomplete.

Proposition C1 (Operational Bounds Motivating a TPB Admissibility Window)

Statement. In any record-forming substrate, operational requirements of (i) record resolvability against noise and (ii) maintenance of stable composite structures motivate bounds $0 < TPB_{min} \leq TPB \leq TPB_{max} < \infty$. The numerical values are architecture-dependent and are a target for derivation in the closure programme (Section H1).

Motivation.

Lower bound ($TPB \geq TPB\_min$). Each committed record must be resolved against the substrate dynamics. Resolution requires a minimum number of substrate ticks to distinguish the committed state from uncommitted fluctuations — this is a minimal resolvability requirement: commitments must be distinguishable from reversible substrate fluctuations over a finite number of substrate updates. If TPB falls below this minimum, the substrate cannot accumulate enough ticks to certify that a commitment has occurred. The resulting "records" would be indistinguishable from noise and would violate Axiom 2 (stability).

Upper bound ($TPB \leq TPB\_max$). If TPB is too large — too many substrate ticks elapse between commitments — then the commitment density (commitments per substrate tick) drops below the threshold needed to sustain coherent physical structures. Specifically, any composite structure that requires ongoing record maintenance (a particle, a bound state, a measurement apparatus) will lose coherence if the tick-to-commitment ratio exceeds the number of substrate ticks over which environmental perturbations accumulate to the point of disrupting the structure. The maximum TPB is set by the ratio of the perturbation-accumulation scale (in substrate ticks) to the minimum commitment count needed for structural maintenance.

The finite window [TPB_min, TPB_max] is therefore a consequence of the joint requirements of record resolution (lower bound) and structural coherence (upper bound). Its width is determined by the substrate architecture, not by modelling choice.

D. EFT Locality from Finite Observers

Proposition D (Wilsonian EFT Locality Under Standard Conditions)

Statement. If observers have finite spatiotemporal resolution and finite energy (Axiom 3), and under the additional conditions stated below, then the predictive content of a theory at accessible scales can be represented by a local effective action expanded in operators ordered by dimension.

Additional conditions (explicit). Proposition D assumes, beyond Axioms 1–6: (i) an approximate scale separation between accessible excitations and integrated-out structure, and (ii) cluster decomposition at accessible scales. These are not independent modelling choices but empirical regularities of the regimes where modern physics is tested; their full derivation from Axioms 1–6 alone is not attempted here, though partial justification is given below.

Proof sketch. Finite resolution implies that observers cannot probe arbitrarily short-distance structure. Degrees of freedom below the resolution scale are inaccessible and must be integrated out. Under the following conditions the resulting effective description is a local EFT:

1. **Translation invariance at scales ≫ cutoff.** Axiom 5 (low-energy locality) guarantees that the physics at accessible scales does not depend on absolute position. This ensures that the effective action can be written in terms of local fields.

2. **Cluster decomposition.** Axiom 6 together with empirical repeatability (Axiom 1) motivates cluster decomposition at accessible scales: distant experiments yield independent results. This forces the effective action to factorise at large separations. (This motivation falls short of a derivation; cluster decomposition is treated as an empirical input consistent with, but not strictly entailed by, the axioms.)
3. **Operator ordering by dimension.** Finite energy (Axiom 3) means that higher-derivative operators, which probe shorter distances, are suppressed by powers of (energy/cutoff). The effective action admits an expansion in operators of increasing dimension, with coefficients suppressed by corresponding powers of the cutoff scale.

This is the standard Wilsonian argument, but here it enters as part of the inevitability chain rather than as a methodological choice. Under the standard assumptions of scale separation and cluster decomposition, EFT locality is not a preference — it is a consequence of finite access to the substrate. (This does not exclude the possibility of nonlocal UV completions above the cutoff; the claim is about the effective description at accessible scales.)

Logical character. Proposition D does not claim that locality is derived from Axioms 1–6 alone. It claims that under empirically observed scale separation and cluster decomposition — conditions satisfied in all tested regimes — the effective description takes EFT form. Proposition D has a different logical character from Theorems A–C: those results follow from Axioms 1–6 alone, while Proposition D additionally requires these empirical inputs. The programme treats them as conditions rather than axioms: they are testable, they hold in every regime where modern physics has been checked, and they could in principle be weakened or derived in future work.

Connection to the theorem chain. Theorems A–C establish that any admissible theory contains records, commitment, entropy costs, and a TPB conversion invariant. Proposition D establishes that these structures must be describable by a local EFT at accessible scales. The remaining question is: what does that EFT look like?

E. Fold Stability: The Skyrme-Type Stabiliser is Forced

Theorem E (Static Finite-Energy Solitons in 3D Require a Higher-Derivative Stabiliser)

Statement. Consider a 3D static energy functional for a fold field Ψ containing only a quadratic gradient term and a potential term:

$$E_0[\Psi] = \int d^3x \, [\, a \, |\nabla\Psi|^2 + b \, V(\Psi) \,]$$

with $a, b > 0$. Then for scaling $\Psi_\lambda(x) = \Psi(\lambda x)$, E_0 cannot be bounded away from collapse in general. Adding a quartic-gradient term:

$$E[\Psi] = \int d^3x \, [\, a \, |\nabla\Psi|^2 + b \, V(\Psi) + c \, |\nabla\Psi|^4 \,]$$

with $c > 0$ produces coercivity under collapse sequences.

Proof. Under scaling $\Psi_\lambda(x) = \Psi(\lambda x)$ in 3D:

- Gradient term: $\int|\nabla\Psi_\lambda|^2 \, d^3x = \lambda \cdot \int|\nabla\Psi|^2 \, d^3x$ (scales as λ^1)
- Potential term: $\int V(\Psi_\lambda) \, d^3x = \lambda^{-3} \cdot \int V(\Psi) \, d^3x$ (scales as λ^{-3})
- Quartic-gradient term: $\int|\nabla\Psi_\lambda|^4 \, d^3x = \lambda^3 \cdot \int|\nabla\Psi|^4 \, d^3x$ (scales as λ^3)

For E_0 (without the quartic term), the energy as a function of λ is:

$$E_0(\lambda) = a\lambda \cdot I_2 + b\lambda^{-3} \cdot I_V$$

where $I_2 = \int|\nabla\Psi|^2$ and $I_V = \int V(\Psi)$. This has a single extremum and no barrier preventing $\lambda \to \infty$ (collapse to a point) or $\lambda \to 0$ (dispersal). This is Derrick's theorem: no stable static solitons exist with only quadratic gradient and potential terms in 3D.

With the quartic term:

$$E(\lambda) = a\lambda \cdot I_2 + b\lambda^{-3} \cdot I_V + c\lambda^3 \cdot I_4$$

As $\lambda \to \infty$ (collapse), the $c\lambda^3 \cdot I_4$ term dominates and drives $E \to \infty$. Collapse is excluded. As $\lambda \to 0$ (dispersal), the $b\lambda^{-3} \cdot I_V$ term dominates and drives $E \to \infty$. Dispersal is excluded. The energy functional is coercive: it grows without bound under both collapse and dispersal sequences, forcing the existence of a finite-energy minimum.

Proposition E.2 (Partial Uniqueness of the Stabiliser)

Statement. Within the class of stabilisers that are (i) local, (ii) rotationally invariant, (iii) quartic in first derivatives, and (iv) positive definite modulo total derivatives, the stabilising operator is unique up to linear combinations that are equivalent on-shell or differ by redundancies.

Proof sketch. The requirement is a term that scales as λ^3 under $\Psi_\lambda(x) = \Psi(\lambda x)$ to dominate the collapse direction. Among local, rotationally invariant expressions quartic in first derivatives of a field on $\mathbb{R}^3$, the independent structures are (after removing total derivatives and using equations of motion):

- $|\nabla\Psi|^4$ (the Skyrme invariant)
- Contractions involving $\partial_i\Psi \cdot \partial_j\Psi$ with index structures that reduce to the above on-shell

Any positive-definite quartic-gradient invariant that provides the required λ^3 scaling is a positive linear combination of these structures, which are equivalent up to field redefinitions and equations of motion.

Conclusion. Within the minimal derivative order required to defeat Derrick scaling, the stabilising sector is essentially unique up to redundancy. This is a defensible uniqueness result for the fold stabilisation mechanism.

F. Gauge Symmetry: What Can Be Proven and What Remains Open

This section requires the most intellectual honesty in the programme. The chain from records to gauge structure can be made rigorous at the abstract level; the derivation of the specific Standard Model gauge group remains an open problem.

Theorem F1 (Redundancy of Description ⇒ Gauge Structure)

Statement. If the physically distinguishable state is an equivalence class of microdescriptions, and local comparisons of states require choosing local representatives, then the theory admits a principal-bundle description with a structure group G acting as redundancy — i.e. gauge symmetry.

Proof sketch.

Step 1 (Equivalence classes define fibres). Let M denote the space of physically distinguishable states (record macrostates in the BCB context). Let Ω denote the space of microdescriptions. The physical-equivalence relation ~ partitions Ω into equivalence classes: for each physical state $m \in M$, the fibre $F_m = \{\omega \in \Omega : \pi(\omega) = m\}$ is the set of all microdescriptions consistent with that physical state. This defines a surjection $\pi : \Omega \to M$ with fibres F_m.

Step 2 (Local representatives define sections). To do local computations (write equations of motion, compute observables), one must choose a representative $\omega \in F_m$ for each physical state m in some neighbourhood $U \subset M$. This choice is a local section $\sigma : U \to \Omega$ with $\pi \circ \sigma = id_U$. The choice is not unique — any other section σ' related by a fibre-preserving transformation is equally valid.

Step 3 (Transitions between local representatives define a connection). When two local sections σ_U and σ_V overlap on $U \cap V$, they are related by a transformation $g_{UV} : U \cap V \to G$, where G acts on the fibres. Consistency across triple overlaps requires $g_{UV} \cdot g_{VW} = g_{UW}$ (the cocycle condition). This is the definition of a principal G-bundle over M. Comparing sections at nearby points (parallel transport) requires specifying how representatives change along paths in M — this is a connection on the bundle. The connection is a gauge field.

Step 4 (Physical observables are gauge-invariant). Quantities that are independent of the choice of section — i.e. invariant under G-transformations on the fibres — are the physically meaningful observables. This is gauge invariance, derived here not as a postulate but as a consequence of the redundancy inherent in describing physical states via microdescriptions.

Connection to BCB. In the BCB context, the fibres are the sets of substrate microstates that correspond to the same record macrostate (the many-to-one map π of Lemma A1). The gauge group G is the group of transformations that permute microstates within a fibre without changing the macrostate. The gauge field arises from the need to compare records at different spacetime points, which requires a convention for how to identify microstates across fibres — i.e. a connection.

Theorem F2 (Continuity + Composition ⇒ Lie Group)

Statement. If the redundancy transformations (i) compose (group multiplication), (ii) have identity and inverses, and (iii) vary continuously with configuration, then the structure group G is a Lie group.

Proof sketch. This is a direct application of the solution to Hilbert's fifth problem (Montgomery–Zippin, 1952): a locally compact topological group that acts faithfully and continuously on a manifold is a Lie group. The conditions are satisfied by construction: the fibre transformations compose (by definition of group action on fibres), continuity is required for local sections to be smoothly defined, and local compactness and manifold regularity are motivated by finite distinguishability (Axiom 3) and the requirement that redundancy transformations act smoothly on the admissible state manifold.

Unitarity/compactness refinement. Axiom 6 (consistency and stability) strongly motivates restriction to compact gauge groups: non-compact gauge groups generically produce negative-norm states or vacuum instabilities in quantum realisations. This restricts G to a compact Lie group in any ghost-free, stable quantum theory.

What Is Established

Theorems F1 and F2 together prove: any admissible theory in $\mathcal{T}$ that has a redundancy of microdescription admits a gauge structure with a compact Lie group. This is a general structural result — it establishes the *form* of gauge symmetry without specifying the *content* (which group).

What the Companion Gauge Papers Accomplish

Three companion papers — *The Minimal Internal Symmetry Theorem*, *Distinguishability Conservation and Gauge Structure*, and *Closing the Interfaces* — address the content question in a related axiom framework (BC1–BC3 + FIM). Their combined results are:

Gauge connections are forced (not postulated). The foundations paper proves that local distinguishability conservation (BC1) over spacetime requires a gauge connection on the internal state bundle. Any two admissible connections differ by gauge freedom. This is the same result as Theorem F1 of this document, proved independently from information-geometric axioms.

Internal dimension is bounded. A curvature-controlled distinguishability capacity bound on the internal Fisher manifold constrains the admissible internal dimension to $n \leq 3$. In the companion gauge papers, the distinguishability capacity functional $C_\varepsilon(n)$ for $\boldsymbol{\mathcal{M}}_{int} = CP^{(n-1)}$ with Fisher/Fubini–Study geometry exhibits a sharp decrease beyond $n = 3$: the available capacity drops by approximately 23% at the $n = 3 \to 4$ transition and approximately 38% at $n = 4 \to 5$, while the minimum anchor-imposed capacity remains fixed — a scissors argument that closes at n = 3. Detailed calculations appear in the companion gauge trilogy.

The specific gauge group is derived. Three minimal empirical anchors — (A) stable three-body confined bound states exist, (B) a fundamental two-level internal sector with continuous symmetry exists, (C) quantum phases are unobservable — combined with the geometric classification (admissible internal manifolds are CP^{n-1} with $n \leq 3$, acted on by SU(n)) and entropy minimisation, force the gauge group uniquely to SU(3) × SU(2) × U(1).

Abelian uniqueness is proven beyond phase redundancy. Given the derived chiral matter content, any BCB-admissible Abelian charge must be proportional to hypercharge. Independent dark-photon-type U(1) sectors are excluded at the fundamental level.

Chirality is derived from information conservation. Bidirectional SU(2) coupling to spinorial matter produces overconservation of distinguishability (two independently conserved currents where BC1 requires one). The chiral restriction to a single Weyl sector is structurally forced. The identification with S_L requires one bit of empirical input — analogous to choosing which direction is "future" in thermodynamics.

The Remaining Integration Task

The gauge papers use axioms BC1–BC3 + FIM. This document uses Axioms 1–6. These are not identical axiom sets, but they overlap substantially:

Gauge paper axiom	Inevitability Programme counterpart
BC1 (Distinguishability conservation)	Axiom 4 (Microscopic reversibility) + record/coarse-graining structure (Theorem A, Lemma A1)
BC2 (Entropy minimisation)	Structural parsimony driving TPB admissibility (Proposition C1) and EFT truncation
BC3 (Anomaly exclusion)	Axiom 6 (Consistency and stability), specialised to gauge sector
FIM (Fisher information manifold)	The specific geometric realisation of the distinguishability structure derived abstractly by Theorems F1–F2

The gap is therefore not mathematical — the proofs of gauge group uniqueness exist in the companion papers. The gap is showing that this document's Axioms 1–6 *imply* the gauge papers' axioms BC1–BC3 + FIM, which would make the gauge group derivation a downstream consequence of the same inevitability chain. This is an axiom-translation problem, not a theorem problem. The honest statement is now:

The abstract structure of gauge symmetry is forced by Theorems F1–F2 of this document. The specific identity of the Standard Model gauge group SU(3) × SU(2) × U(1) is derived in the companion gauge papers from a related axiom framework (BC1–BC3 + FIM), conditional on three minimal empirical anchors strictly weaker than Standard Model assumptions. Full integration into the Inevitability Programme requires an explicit demonstration that Axioms 1–6 entail BC1–BC3 + FIM. This is a dictionary task with a defined scope, not an open conjecture.

The Three Empirical Anchors

Any derivation of specific gauge content from abstract principles must eventually make contact with observation. The gauge papers identify the minimum empirical input required — three observable facts, each strictly weaker than assuming the Standard Model:

Anchor A (Confinement). There exist stable bound states composed of exactly three fundamental constituents, where no individual constituent can be isolated. This is an observable fact about hadrons. It does not assume "quarks," "colour charge," or SU(3).

Anchor B (Two-level internal sector). There exists at least one fundamental two-state internal degree of freedom acted on by a continuous symmetry preserving Fisher distinguishability. This is an observable fact about the doublet structure of fundamental matter. It does not assume "SU(2)," "W bosons," or a specific gauge structure.

Anchor C (Universal phase redundancy). The quantum mechanical description of any physical system admits a global phase transformation $|\psi\rangle \to e^{i\theta}|\psi\rangle$ that leaves all observable predictions invariant. This is a structural fact about the Born rule. It does not assume electromagnetism, U(1) gauge theory, or the existence of photons.

These anchors are the gauge-sector analogue of this document's Axioms 1–6: minimal, defensible, and strictly weaker than the conclusions they support. Their integration into the Inevitability Programme's axiom chain is part of the H3 translation task.

G. The Uniqueness Theorem: What Can Be Claimed Now

Definition (Fold-Based EFT)

A fold-based EFT is a local effective theory containing at least one field Ψ whose finite-energy configurations support topological or boundary-stabilised soliton sectors used to model particle-like excitations.

Theorem U (Restricted Uniqueness of the BCB Architecture)

Statement. Fix the admissible theory class $\mathcal{T}$ and restrict to the subclass $\mathcal{T}_{fold} \subset \mathcal{T}$ of EFTs that:

1. Realise records via an explicit commitment-count $N_{bit}(\tau)$ with reversible micro-parameter τ;
2. Contain a stable finite-energy fold sector in 3D;
3. Are ghost-free below cutoff;
4. Are renormalization-closed at a stated truncation order;
5. Admit an SMEFT-legible operator basis at dimension ≤ 6.

Then any T ∈ $\mathcal{T}$_fold must contain:

- (a) A TPB-like conversion invariant (by Theorem C),
- (b) A canonical entropy-weighted time mapping of Role-4 type (by Corollary A2 + Theorem B + the bridge equation),
- (c) A Skyrme-like stabiliser term, unique up to redundancy within the minimal derivative order (by Theorem E and Proposition E.2),
- (d) An EFT operator expansion living inside SMEFT classes (by Proposition D + condition 5).

Moreover: Within minimal-derivative stabilisers and fixed symmetries, the stabilising operator class is unique up to field redefinitions (Proposition E.2). The TPB admissibility window is a structural necessity (Proposition C1), and — as a substrate-defined constraint — is expected to be invariant under EFT-scale RG running, in the sense that RG flow does not redefine the substrate's commitment bounds; it only changes how close a given effective configuration sits to those bounds.

On the scope of $\mathcal{T}$_fold. The restriction to $\mathcal{T}$_fold is phenomenological, not ontological. $\mathcal{T}$_fold contains theories whose empirically accessible excitations behave as stable, localised, finite-energy structures describable within a low-energy effective field theory.

At accessible energies, any theory whose experimentally observed excitations behave as approximately localised, finite-energy particle states must admit an effective description in which those excitations are represented by stable configurations in a local field expansion. Whether the underlying ontology consists of strings, loops, or other extended objects is immaterial at this level; the claim concerns the structure of the effective description, not the microscopic substrate.

This does not imply that all ultraviolet completions reduce to folds. Some UV frameworks (e.g. extended-object theories) may not admit a soliton description in their fundamental formulation. Theorem U therefore makes a conditional claim: within the class of empirically admissible theories whose accessible particle-like excitations admit a stable, localised EFT representation, the BCB architectural elements are forced. Whether every UV completion that matches observed particle phenomenology necessarily flows into this effective class is a separate dynamical question not addressed here.

Proof of Theorem U

The proof assembles the theorem chain:

1. T ∈ $\mathcal{T}$_fold satisfies Axioms 1–6 (by membership in $\mathcal{T}$) and conditions 1–5 (by membership in $\mathcal{T}$_fold).
2. By Theorem A, T contains records. By Lemma A1, records are coarse-grained. By Corollary A2, record formation is irreversible. By Theorem B, each record event costs ε_bit > 0.

3. By condition 1 and Theorem C, T contains a TPB-like conversion invariant. This establishes (a).
4. By Corollary A2 and Theorem B, the irreversible commitment process defines a thermodynamic arrow. Given a reversible substrate ordering parameter τ and an irreversible commitment count $\mathbb{E}[N_{bit}(\tau)]$, any operational "time" variable that preserves the commitment ordering is defined only up to monotone reparametrisation $t \mapsto \tilde{t}(t)$ with $d\tilde{t}/dt > 0$; no unique time coordinate exists at the structural level. Role-4 selects a canonical representative of this equivalence class by imposing the entropy-weighted calibration $dt_{phys} = f(s)\, d\tau$, which ties the operational time parameter to local entropy weighting. This is a gauge-fixing (a normalisation choice within the equivalence class of monotone calibrations) rather than an additional dynamical postulate. This establishes (b).
5. By condition 2 and Theorem E, the fold sector requires a higher-derivative stabiliser. By Proposition E.2, this stabiliser is unique within the minimal derivative class. This establishes (c).
6. By Proposition D and condition 5, the EFT expansion lives inside SMEFT classes. This establishes (d).

Scope and Limitations

What this theorem says. The BCB architecture — records, commitment, TPB, entropy-weighted time, Skyrme-stabilised folds, SMEFT-compatible EFT — is forced within the class of fold-based EFTs with stable records. Any theory in $\mathcal{T}_{fold}$ that satisfies the stated conditions must reproduce the BCB structural elements, possibly under different names or notations.

What this theorem does not say. Theorem U does not claim that BCB is unique among all theories of physics. The restriction to $\mathcal{T}_{fold}$ is stated explicitly. Theories that do not use fold-type solitonic structures, or that abandon microscopic reversibility, or that operate outside the EFT framework, are not addressed.

Why the restriction is not a weakness. The restriction defines the scope of the claim precisely, which is what mathematical theorems require. A claim of unrestricted uniqueness across all conceivable theories would be either false or unfalsifiable. A claim of restricted uniqueness within a well-defined class is testable: anyone who proposes a theory in $\mathcal{T}_{fold}$ that lacks one of the BCB structural elements has either found a counterexample or made an error. That is how mathematics works.

H. The Mountain: Open Problems Required for Full-Strength Uniqueness

Theorem U establishes restricted uniqueness within $\mathcal{T}_{fold}$. Upgrading to full-strength uniqueness — the claim that BCB-type structure is the unique admissible architecture in all of $\mathcal{T}$ — requires three hard results that are not yet proven.

H1. Independent Derivation of the TPB Window

Requirement. Compute or bound TPB_min and TPB_max from substrate constraints without tuning to match the fold spectrum.

Why it matters. Currently, the TPB window is characterised qualitatively (Section C, Proposition C1). For the generation count (H2) to be a genuine prediction, the window must be derived independently — otherwise the argument is circular (the window is adjusted to give three generations, rather than three generations emerging from an independently fixed window).

Formal framework. Using the definitions from Section C — commitment density $\rho\_commit(\tau) := d\,\mathbb{E}[N\_bit] / d\tau$ and conversion invariant $TPB(\tau) := \rho\_commit(\tau)^{-1}$ — the admissibility question is: *for what range of TPB can a substrate generate certifiable and stable records?* Two operational constraints, expressed entirely in the substrate ordering parameter τ, produce a finite window.

Lower bound: TPB_min from certifiability.

A committed bit must be distinguishable, in the record space, from transient substrate fluctuations. Otherwise it is not a record. Formalise this as a hypothesis-testing problem:

Let X(τ) be a substrate observable (or coarse-grained sufficient statistic) used to read out whether commitment occurred. Define two hypotheses:

- H_0: no commitment event occurred in substrate interval $[\tau, \tau + \Delta\tau]$
- H_1: one commitment event occurred in $[\tau, \tau + \Delta\tau]$

A record is certifiable only if there exists a decision rule with error probability below some tolerance δ:

$$P(\text{error}) \le \delta$$

The key constraint: a discrete event cannot be certified with arbitrarily few substrate updates if the readout noise has nonzero variance. Generically, the distinguishability between H_0 and H_1 scales as $(\Delta\mu)^2 / \sigma^2$, where σ^2 shrinks with the number of independent substrate samples in the interval. Below some minimum Δτ, the distributions under H_0 and H_1 overlap too much to separate.

Define $\Delta\tau\_cert(\delta)$ as the minimum substrate progression needed to certify one commitment at confidence $1 - \delta$. Then:

$$TPB\_min \ge \Delta\tau\_cert(\delta)$$

because a bit cannot be committed in fewer substrate ticks than are needed to certify that the commitment occurred. This establishes a strictly positive lower bound whenever the substrate readout carries nonzero noise, killing the "TPB can be arbitrarily small" loophole.

Upper bound: TPB_max from structural maintenance.

Stable composite structures (particles, bound states, measurement apparatus) require enough irreversible commitments per substrate progression to resist perturbations. If commitments are too sparse, the system cannot maintain its structural identity.

Define:

- $\Delta\tau\_pert$: the substrate progression over which uncontrolled perturbations (environmental coupling, internal fluctuations) accumulate enough to push a structure out of its basin of attraction, if not countered by commitments.
- N_min: the minimum number of commitments required within that perturbation window to re-stabilise the configuration (a basin re-anchoring requirement).

Admissibility demands that enough commitments accumulate within the perturbation window:

$$\mathbb{E}[\Delta N\_bit] \approx \rho\_commit \cdot \Delta\tau\_pert = \Delta\tau\_pert / TPB \geq N\_min$$

Rearranging:

$$TPB \leq \Delta\tau\_pert / N\_min =: TPB\_max$$

If TPB exceeds this bound, the commitment density is too low to anchor the structure against perturbation-driven drift out of the stability basin.

The window.

Combining both constraints:

$$\Delta\tau\_cert(\delta) \leq TPB \leq \Delta\tau\_pert / N\_min$$

This is a finite, nonzero admissibility window expressed entirely in τ and commitment counts. No fundamental time is assumed.

Closure programme: eliminating tunability.

The quantities $\Delta\tau\_cert$, $\Delta\tau\_pert$, and N_min currently appear as free parameters. For the window to be a prediction rather than a fit, each must be derived from the same substrate architecture:

(A) *Derive $\Delta\tau\_cert$ from substrate noise statistics.* Choose the readout variable that distinguishes "committed" from "uncommitted." Compute the overlap between the distributions under H_0 and H_1. Solve for the minimum $\Delta\tau$ achieving $P(error) \leq \delta$. This turns TPB_min into a calculation.

(B) *Derive $\Delta\tau\_pert$ from coupling strength and fluctuation spectrum.* Model perturbation accumulation along τ as diffusion in configuration space. Define a basin-exit threshold. Compute the expected exit distance in τ-units to obtain $\Delta\tau\_pert$.

(C) *Derive N_min from the fold stability functional.* Use the coercivity and Skyrme energy bounds from Theorem E. Define the minimal number of discrete anchoring increments required to keep the configuration within the stability basin given the typical perturbation size. This is an energetic margin condition.

Once (A)–(C) are completed, the window is predicted from substrate physics, not fitted to phenomenology. This is the concrete content of the H1 closure task.

Epistemic status. The existence of a finite window is forced by the logic of Theorems A–C; its numerical values are currently underdetermined until tasks (A)–(C) are completed. Until then, TPB bounds function as constrained parameters rather than derived constants. This is an honest intermediate state: the structural necessity is proved, the quantitative prediction awaits calculation.

H2. Spectral Uniqueness → Exactly Three Admissible Bound States

Requirement. Solve the nonlinear fold equation, compute the Sturm–Liouville spectrum, and apply the independently fixed TPB window. Demonstrate that exactly three bound states are admissible.

Why it matters. The three-generation structure of the Standard Model is one of its most striking and unexplained features. If BCB can derive it from the fold spectrum + TPB window, that is a major prediction. If it cannot, or if the number depends on adjustable parameters, the claim of inevitability is significantly weakened.

Path to resolution. This is a numerical computation, not a structural argument. The steps are specified in the closure roadmap of the Tightening Addendum: solve the radial fold equation, compute the fluctuation operator, solve the eigenproblem, filter by TPB admissibility, count. The result is either three or it isn't. There is no room for interpretation.

H3. Integration of the Gauge Group Derivation

Requirement. Demonstrate that the Inevitability Programme's Axioms 1–6 entail the gauge papers' axioms BC1–BC3 + FIM, thereby making the derived gauge group SU(3) × SU(2) × U(1) a downstream consequence of the same inevitability chain. Alternatively, derive the SM gauge group directly within the Axiom 1–6 framework.

Current status. The companion gauge papers (*Minimal Internal Symmetry Theorem*, *Distinguishability Conservation and Gauge Structure*, *Closing the Interfaces*) have derived the SM gauge group from axioms BC1–BC3 + FIM, conditional on three minimal empirical anchors (confinement, two-level internal sector, phase redundancy). The derivation includes: internal dimension bounded to $n \leq 3$ via capacity arguments, each anchor forcing a unique gauge sector, Abelian uniqueness beyond phase redundancy, chirality restriction from BC1, and exclusion of additional gauge factors by entropy minimisation. Detailed derivations are presented in the companion gauge papers. What remains is the axiom-translation task described in Section F.

Why it matters. Theorems F1–F2 of this document establish that gauge symmetry with a compact Lie group is forced. But many compact Lie groups exist. Without an argument that selects the SM group specifically, the Inevitability Programme explains the *form* of gauge symmetry but not its *content*. The gauge papers provide that argument in a closely related framework. Completing the translation would close the gap between restricted and full uniqueness.

Path to resolution. The translation requires showing four implications: (1) Axiom 4 + the record structure from Theorem A/Lemma A1 entails BC1; (2) the structural parsimony implicit in TPB admissibility and EFT truncation entails BC2; (3) Axiom 6 specialised to the gauge sector entails BC3; (4) the abstract fibre-bundle structure of Theorems F1–F2, equipped with the finite distinguishability of Axiom 3, entails FIM. Each implication has a clear scope and does not require new mathematical machinery — it requires careful identification of where the two axiom sets make the same physical demand in different notation. The three empirical anchors must also be shown to follow from (or be compatible with) the Inevitability Programme's axioms plus minimal observational input.

With the mathematical content of the gauge derivation established in the companion papers, the remaining challenge is integration rather than invention. It is a well-defined dictionary task with the mathematical substance already in hand.

Preliminary Axiom Correspondence

To demonstrate that the H3 translation is a genuine mapping and not a rhetorical claim, we prove the first implication as a formal lemma — deriving BC1 from Axiom 4 and Lemma A1 with no additional assumptions about gauge structure, Fisher metric, or internal symmetry — and sketch the remaining three.

Lemma H3.1 (Axiom 4 + Coarse-Graining ⇒ Distinguishability Continuity Equation).

Setup. Let Ω be the substrate microstate space with initial probability measure μ_0. Let $U_\sigma : \Omega \to \Omega$ be an invertible, measure-preserving flow parametrised by the reversible ordering parameter σ (in BCB, $\sigma = \tau$), forming a one-parameter group: $U_{\sigma+\varepsilon} = U_\varepsilon \circ U_\sigma$ and $U_0 = \mathrm{id}$. This is the content of Axiom 4 (microscopic reversibility). Define the evolved measure:

$\mu_\sigma := (U_\sigma)_* \mu_0$, i.e. $\mu_\sigma(A) = \mu_0(U_\sigma^{-1}(A))$ for all measurable $A \subset \Omega$.

Let $\pi : \Omega \to \mathcal{M}$ be the coarse-graining map to the macrostate manifold (Lemma A1), with $\mathcal{M}$ the manifold of physically distinguishable macrostates.

Regularity Condition RC-π (dynamically compatible coarse-graining). The coarse-graining map π is not arbitrary; it is defined operationally by the equivalence relation induced by record indistinguishability under admissible protocols: two microstates are equivalent if and only if they induce the same record macrostate under all admissible readout interactions on the relevant τ-scales. Under this definition, compatibility with U_σ is not assumed but follows from record stability under Axiom 2 — if records are stable under later interactions, then microstates that are record-equivalent at σ remain record-equivalent at $\sigma + \varepsilon$, which is precisely the factorisation condition $\pi \circ U_\sigma = \Phi_\sigma \circ \pi$. Thus Φ_σ is well-defined on the macrostate manifold. Under regularity assumptions (smooth π and smooth U_σ), Φ_σ is a diffeomorphism. The group property of U_σ induces $\Phi_{\sigma+\varepsilon} = \Phi_\varepsilon \circ \Phi_\sigma$ whenever RC-π holds. This is not an arbitrary choice of coarse-graining; it is the specific coarse-graining defined by stable records (Axiom 2), and factorisation is a consequence of that definition rather than an independent assumption.

Define the pushforward macro-measure:

$\nu_\sigma := \pi_* \mu_\sigma$, i.e. $\nu_\sigma(B) = \mu_\sigma(\pi^{-1}(B))$ for all measurable $B \subset \mathcal{M}$.

Step 1 (Macro measure transport). For $B \subset \mathcal{M}$ measurable:

$\nu_\sigma(B) = \mu_\sigma(\pi^{-1}(B)) = \mu_0(U_\sigma^{-1}(\pi^{-1}(B)))$

By RC-π ($\pi \circ U_\sigma = \Phi_\sigma \circ \pi$), we have $U_\sigma^{-1}(\pi^{-1}(B)) = \pi^{-1}(\Phi_\sigma^{-1}(B))$. Therefore:

$\nu_\sigma(B) = \mu_0(\pi^{-1}(\Phi_\sigma^{-1}(B))) = \nu_0(\Phi_\sigma^{-1}(B))$

That is:

$\nu_\sigma = (\Phi_\sigma)_* \nu_0$

The macrostate measure is transported by the induced diffeomorphism. This is derived from micro-level measure preservation (Axiom 4) and the factorisation condition, not assumed.

Step 2 (Density formulation and continuity equation). Assume ν_σ admits a density representation with respect to a reference volume form on $(\mathcal{M}, g)$:

$\nu_\sigma(dm) = \rho_D(m, \sigma)\, dvol_g(m)$

The density representation assumption is satisfied whenever ν_σ is absolutely continuous with respect to the chosen reference volume form; this holds in many smooth settings and can be treated as a regularity condition on the induced flow and the pushforward measure.

Transport of the measure $\nu_\sigma = (\Phi_\sigma)_* \nu_0$ implies:

$(d/d\sigma) \int_{\Phi_\sigma(A)} \rho_D(m, \sigma)\, dvol_g = 0$

for all measurable A. Applying Reynolds' transport theorem on manifolds:

$0 = \int_{\Phi_\sigma(A)} (\partial_\sigma \rho_D + \nabla_a(\rho_D v^a))\, dvol_g$

where v^a is the velocity field generating Φ_σ and ∇ is the Levi–Civita connection of g. Because this holds for all regions A, the integrand vanishes pointwise:

$\partial_\sigma \rho_D + \nabla_a(\rho_D v^a) = 0$. ■

Summary of the logical chain:

1. Axiom 4 ⇒ invertible, measure-preserving microdynamics (U_σ, μ_σ).
2. Lemma A1 ⇒ coarse-graining map π induces macrostate manifold $\mathcal{M}$.
3. RC-π ($\pi \circ U_\sigma = \Phi_\sigma \circ \pi$) ⇒ well-defined induced macroflow Φ_σ.
4. Micro measure preservation + RC-π ⇒ macro measure transport: $\nu_\sigma = (\Phi_\sigma)_* \nu_0$ (Step 1).
5. Macro measure transport + density representation ⇒ continuity equation (Step 2).

The continuity equation is not an added assumption. It is the macro-level expression of microscopic reversibility under coarse-graining. That is BC1.

Corollary H3.2 (BC1 is the macro-level expression of Axiom 4). Under the conditions of Lemma H3.1, the BC1 continuity equation $\partial_\sigma \rho_D + \nabla_a(\rho_D v^a) = 0$ is the coordinate expression of pushforward measure conservation under the induced invertible macroflow. BC1 is therefore the macro-level conservation law corresponding to Axiom 4's invertibility at the substrate level.

What was assumed vs. what was derived.

- *Assumed:* Axiom 4 (invertible microdynamics), Lemma A1 (coarse-graining map exists), RC-π (dynamically compatible coarse-graining, grounded in Axiom 2), regularity of π and U_σ, density representation of ν_σ.
- *Derived:* The induced macroflow Φ_σ, measure transport on $\mathcal{M}$ ($\nu_\sigma = (\Phi_\sigma)_* \nu_0$), and the BC1 continuity equation.
- *Not assumed:* Fisher metric, CP^2, gauge group, internal symmetry, or any structure from the gauge papers.

Proposition H3.3 (Minimal Internal Complexity Under Empirical Equivalence).

Statement. Let $\mathcal{T}$ be the admissible theory class (Axioms 1–6, under the observational equivalence quotient). Consider two candidate internal models $T_1, T_2 \in \mathcal{T}$ that produce identical empirical predictions at accessible scales — i.e. for every admissible preparation/measurement protocol P and every observable outcome set O, the predicted outcome distributions agree: $P_{T_1}(O|P) = P_{T_2}(O|P)$. By the observational equivalence convention, T_1 and T_2 represent the same physical theory; any internal structure distinguishing them is quotiented out as redundancy.

Claim. Among empirically equivalent internal models, the physically admissible representative is the one minimising an internal complexity functional $\mathcal{C}$. Equivalent choices under mild regularity assumptions include:

- *Volume/capacity functional (geometric):* $\mathcal{C}_{vol} := \log \mathrm{Vol}(\mathcal{M}_{int}, g_{int})$, where $(\mathcal{M}_{int}, g_{int})$ is the internal distinguishability manifold.

- *Description-length functional (information-theoretic):* $\mathcal{C}$_MDL := minimum code length of the internal model consistent with predictions.
- *Entropy overhead functional (commitment-theoretic):* $\mathcal{C}$_over := inf over realisations of $\mathbb{E}[\Delta S\_internal]$, the minimal extra entropy/commitment overhead required by the internal sector beyond what is needed to reproduce the same empirical record statistics.

Under empirical equivalence and the observational equivalence quotient, internal structure is selected by minimising $\mathcal{C}$. This recovers the structural role of BC2: among empirically equivalent internal descriptions, redundant internal complexity is excluded via a minimisation principle over admissible internal structure.

Proof sketch. Axiom 3 (finite distinguishability) implies the internal state space cannot carry arbitrarily large resolvable complexity without physical cost: distinguishability capacity is finite in any finite-resource setting. Proposition D (EFT locality under standard conditions) implies that structure below the observational cutoff is integrated out and can only re-enter through finitely many effective operators; any internal degrees of freedom that never affect the accessible operator content are operationally unobservable. Axioms 1–2 (empirical meaning + records) require that physically meaningful differences must correspond to recordable differences in outcome statistics — if two internal structures never change record statistics, their difference cannot be empirically grounded. Under the observational equivalence quotient, those unobservable differences are quotiented out as redundancy, not fundamental content. The remaining permissible freedom is gauge-like redundancy or genuinely physical internal structure. Choosing minimal $\mathcal{C}$ implements the "no gratuitous internal complexity" demand quantitatively. Thus the BC2-style minimisation is not an arbitrary aesthetic preference; it is the operational completion of finite distinguishability + finite observational access. ■

Status: redundancy removal vs. minimality. The observational equivalence quotient eliminates internal structure that is provably unobservable under the admissible protocol class. This removes redundancy. However, BC2 in the gauge papers does more than remove redundancy: it selects, among distinct internal architectures that remain empirically equivalent after redundancy removal, the one minimising internal entropy or distinguishability capacity. These are logically distinct operations — the quotient removes variables with no operational consequences, while the minimality principle selects the simplest representative among remaining empirically equivalent internal geometries. Proposition H3.3 recovers BC2 only under the additional structural principle that fundamental descriptions should not contain unnecessary internal complexity once empirical equivalence is established. This principle is not a new axiom but a completion of Axiom 3 (finite distinguishability) under the observational equivalence quotient: if distinguishability capacity is finite, then among empirically equivalent descriptions, the one consuming less of that finite capacity is preferred. The residual gap is therefore not mathematical but philosophical: BC2 encodes a minimality criterion over admissible internal geometries. The programme is explicit about this step rather than hiding it inside the axioms.

Lemma H3.9 (Gauge Anomalies Violate Axiom 6).

Statement. Let $T \in \mathcal{T}$ satisfy Axioms 1–6, and let G be the compact gauge group derived by Theorems F1–F2. If the gauge sector of T contains a gauge or mixed anomaly, then T violates

Axiom 6 (consistency and stability). Therefore anomaly cancellation (BC3) is a necessary consequence of Axiom 6.

Proof. A quantum gauge theory with compact gauge group G must satisfy Ward identities corresponding to gauge symmetry. These identities enforce current conservation, decoupling of unphysical polarisations, preservation of the positive-norm subspace, and unitarity of the S-matrix.

Suppose a gauge anomaly is present, so that the divergence of the gauge current acquires a quantum correction: $\nabla_\mu J^\mu_a = \mathcal{A}_a \neq 0$. This produces three structural failures:

(1) Failure of gauge invariance at the quantum level. The gauge variation of the effective action becomes $\delta_\alpha \Gamma \neq 0$. The gauge redundancy is no longer a redundancy: unphysical gauge degrees of freedom become dynamical.

(2) Breakdown of Ward identities. Ward identities enforce the cancellation of longitudinal gauge boson contributions. With an anomaly, longitudinal polarisations no longer decouple, gauge fixing cannot remove negative-norm states, and BRST symmetry fails. The physical Hilbert space is no longer a positive-norm quotient space. This directly contradicts the positive-norm requirement of Axiom 6.

(3) Violation of unitarity. Gauge symmetry enforces probability conservation through current conservation. Anomaly-induced current non-conservation leads to violation of the optical theorem, a non-unitary S-matrix, and probability leakage — contradicting the stable vacuum requirement of Axiom 6.

Since an anomalous gauge theory cannot maintain a positive-definite inner product on physical states or a stable vacuum, Axiom 6 requires anomaly cancellation. Therefore BC3 is not an additional assumption but a necessary condition imposed by Axiom 6. ■

Lemma H3.4 (Statistical Distinguishability ⇒ Fisher–Rao Geometry).

Statement. Let $\mathcal{M}_{int}$ denote the manifold of physically distinguishable internal states (the base of the internal bundle in Theorems F1–F2). Assume:

(i) Operational statistical model. For each internal state $\theta \in \mathcal{M}_{int}$ and each admissible measurement protocol P (allowed by Axiom 1), the theory predicts a probability distribution over outcomes x in some measurable outcome space X:

$p(x \mid \theta; P)$.

Two internal states are physically distinguishable only insofar as they yield different predicted outcome distributions for some admissible protocol.

(ii) Coarse-graining as stochastic map. Any admissible processing of outcomes (binning, marginalisation, detector coarse resolution) is represented by a Markov morphism K acting on distributions:

$$p(x) \mapsto (Kp)(y) := \int_X K(y|x)\, p(x)\, dx$$

with $K(\cdot|x)$ a stochastic kernel.

(iii) Monotone distinguishability principle. The intrinsic distinguishability metric on $\mathcal{M}_{int}$ must not increase under admissible coarse-graining of outcome statistics. Equivalently, distinguishability is contractive under Markov morphisms.

Then the Riemannian metric on $\mathcal{M}_{int}$ induced by statistical distinguishability is, up to an overall constant factor, the Fisher information metric:

$$g_{ab}(\theta; P) = \mathbb{E}_{x \sim p(\cdot|\theta;P)} [\, \partial_a \log p(x|\theta;P) \cdot \partial_b \log p(x|\theta;P) \,]$$

Moreover, under the monotonicity requirement, this metric is unique up to scale (Čencov's theorem). This establishes the FIM structure as the canonical internal distinguishability geometry, given only operational distinguishability and admissible coarse-grainings.

Proof sketch (Čencov route). The distributions $\{p(\cdot|\theta; P)\}_\theta$ form a statistical model embedded in the probability simplex (or an appropriate infinite-dimensional analogue). Markov morphisms represent admissible coarse-grainings and induce maps between statistical models. A monotone Riemannian metric is one whose induced statistical distance does not increase under these maps. Čencov's theorem states that the Fisher–Rao metric is the unique (up to positive constant) monotone Riemannian metric on statistical models under Markov morphisms. Therefore the canonical metric measuring distinguishability of internal states through their outcome statistics is Fisher–Rao. ■

Interpretation in Inevitability Programme terms. Axiom 1 (repeatable empirical content) forces outcomes to be treated probabilistically. Axiom 3 (finite distinguishability) restricts to finite-capacity distinguishability structures — finite-dimensional effective internal manifolds at accessible scales. Theorems F1–F2 (bundle + Lie redundancy) give the fibre-bundle structure; this lemma supplies the canonical Riemannian geometry on the base of distinguishable states. Thus "FIM" is not an additional modelling postulate; it is the natural geometry of distinguishability once outcome statistics and admissible coarse-graining are taken seriously.

Scope and caveats.

- *Finite-dimensional effective manifold.* Axiom 3 motivates finite-dimensional effective internal manifolds at accessible scales; a full infinite-dimensional treatment is possible but not needed for the programme's closure logic.
- *Uniqueness is conditional on monotonicity.* The uniqueness claim is strictly "unique among monotone Riemannian metrics under Markov morphisms." That is the relevant physical condition because coarse-graining cannot increase distinguishability.

- *Protocol dependence* is addressed by Proposition H3.5 below.

Proposition H3.5 (Protocol-Independent Fisher Geometry via Information Envelope).

Statement. Let $\mathcal{M}_{int}$ be the manifold of physically distinguishable internal states. Let Π denote the admissible protocol class (preparations + measurements + coarse-grainings allowed under Axiom 1 and finite-resolution constraints of Axiom 3). For each protocol $P \in \Pi$, let $g^{(P)}_{ab}(\theta)$ be the Fisher–Rao metric from Lemma H3.4. Then a protocol-independent internal metric can be defined by either of the following canonical envelopes:

(A) Maximal distinguishability metric (information upper envelope):

$$g^*_{ab}(\theta) := \sup_{P \in \Pi} g^{(P)}_{ab}(\theta)$$

interpreted as the greatest distinguishability accessible to any admissible protocol.

(B) Guaranteed distinguishability metric (information lower envelope):

$$g_{*ab}(\theta) := \inf_{P \in \Pi} g^{(P)}_{ab}(\theta)$$

interpreted as the distinguishability that is robustly accessible across all admissible protocols.

(C) Capacity-weighted metric (canonical average): If Π carries a natural weighting $w(P)$ induced by resource constraints (finite energy, finite resolution), define:

$$\bar{g}_{ab}(\theta) := \int_\Pi g^{(P)}_{ab}(\theta)\, w(dP)$$

Under mild regularity conditions (nonempty Π, boundedness/compactness under resource constraints, measurable dependence of $g^{(P)}$ on P), these constructions are well-defined and yield protocol-independent distinguishability geometry.

Proof sketch. For each P, $g^{(P)}$ is positive semidefinite (standard Fisher property). Under resource constraints, the admissible protocol class Π can be taken as compact (or effectively compact) in an appropriate topology — bounded POVM complexity, bounded coarse-graining resolution, bounded energy. If $P \mapsto g^{(P)}$ is continuous or upper-semicontinuous, the supremum exists (attained) by the extreme value theorem; similarly for the infimum under lower-semicontinuity. Envelope metrics preserve monotonicity under coarse-graining because each $g^{(P)}$ is monotone (Lemma H3.4), and the supremum/infimum over a class of monotone metrics remains monotone. Thus the protocol-independent metric is well-defined and compatible with the Čencov monotonicity requirement. ■

Corollary H3.6 (Canonical internal manifold geometry). The internal Fisher information manifold (FIM) used in the companion gauge papers can be understood as the geometry induced by g^* (maximal distinguishability under admissible protocols) or by $\bar{g}$ (capacity-weighted distinguishability), depending on whether the internal manifold is treated as a "best possible inference geometry" or a "typical finite-resource inference geometry." This resolves the "which protocol?" objection: the FIM structure is protocol-independent at the level required for the gauge derivations.

Proposition H3.7 (Internal Distinguishability Capacity Functional).

Setup. Let ($\mathcal{M}$_int, g) be the internal distinguishability manifold equipped with a protocol-independent Fisher geometry (e.g. $g = g^*$ or $g = \bar{g}$ from Proposition H3.5). Let $\varepsilon > 0$ denote the finite resolution scale on $\mathcal{M}$_int — the smallest Fisher distance for which there exists an admissible protocol achieving a fixed discrimination error threshold δ (e.g. via Chernoff/Fisher bounds), consistent with Axiom 3. This ties the resolution scale to the certifiability logic of Section C: ε is not a free parameter but the internal-manifold analogue of the commitment certifiability threshold.

Definition (ε-capacity via packing number). Define the ε-packing number:

$$N\_pack(\varepsilon; \mathcal{M}\_int, g) := \max\{k : \exists\, \theta_1, \ldots, \theta\_k \in \mathcal{M}\_int \text{ s.t. } d\_g(\theta\_i, \theta\_j) \geq \varepsilon\ \forall\, i \neq j\}$$

This counts the maximum number of mutually distinguishable internal states at resolution ε. The internal distinguishability capacity is:

$$C\_\varepsilon(\mathcal{M}\_int, g) := \log N\_pack(\varepsilon; \mathcal{M}\_int, g)$$

This is the operational measure of "how many distinguishable states does the internal manifold support?" — the natural finite-resolution analogue of entropy on an internal manifold.

Volume bound. If ($\mathcal{M}$_int, g) is compact with uniform curvature bounds, then the packing number is controlled by volume in the standard asymptotic regime (ε sufficiently small):

$$N\_pack(\varepsilon) \approx Vol(\mathcal{M}\_int, g) / Vol(B\_\varepsilon) \sim Vol(\mathcal{M}\_int, g) / \varepsilon\hat{}d \quad (d = \dim \mathcal{M}\_int)$$

Hence:

$$C\_\varepsilon \approx \log Vol(\mathcal{M}\_int, g) - d \log \varepsilon + O(1)$$

This bridges the Fisher geometry to capacity: the volume of the internal manifold in the Fisher metric determines how many distinguishable states it can support at a given resolution.

Linking ε to commitment certifiability. The internal resolution scale ε is not merely analogous to the commitment certifiability threshold Δτ_cert(δ) from Section H1 — the two quantities are structurally linked through statistical discriminability bounds. Let the Chernoff bound for discrimination between internal states θ_1 and θ_2 under optimal protocol P be:

$$P(\text{error}) \leq \exp(-\tfrac{1}{2}\, d\_g(\theta_1, \theta_2)^2 \cdot N\_eff)$$

where N_eff is the effective number of statistically independent substrate samples accumulated over Δτ_cert. Certifiability at error tolerance δ requires:

$\frac{1}{2}\, d\_g(\theta_1, \theta_2)^2 \cdot N\_eff \geq \log(1/\delta)$

Defining $N\_eff \propto \Delta\tau\_cert$ yields $\varepsilon^2 \sim 1/\Delta\tau\_cert$. The two quantities are dual: $\Delta\tau\_cert$ sets resolution in substrate ordering; ε sets geometric resolution in distinguishability space. They are not merely analogous — they are quantitatively linked via statistical discriminability bounds.

Proposition H3.8 (Capacity Monotonicity Under Coarse-Graining).

If K is an admissible coarse-graining (Markov morphism) acting on outcome distributions, then the induced Fisher metric is contractive (Čencov monotonicity from Lemma H3.4). Therefore:

$d\_g(K\theta_1, K\theta_2) \leq d\_g(\theta_1, \theta_2) \Rightarrow N\_pack(\varepsilon)$ does not increase under K

Hence C_ε is non-increasing under admissible coarse-grainings. This ties capacity directly to operational finite resolution: coarser measurements cannot reveal more internal structure.

Connection to the dimension bound in the gauge papers. The companion gauge papers implement a "scissors" criterion using this capacity structure. Empirical anchors define a minimum required capacity C^{min}_ε: Anchor B (two-level internal sector) requires $C\_\varepsilon \geq \log 2$; Anchor A (three-body confinement) imposes a stricter lower bound from antisymmetry constraints; Anchor C (phase redundancy) imposes quotient structure that reduces effective capacity. Candidate internal manifold families (e.g. $CP^{(n-1)}$ acted on by SU(n)) define an available capacity $C^{max}\_\varepsilon(n) := C\_\varepsilon(\mathcal{M}\_int(n), g(n))$. Admissibility demands $C^{max}\_\varepsilon(n) \geq C^{min}\_\varepsilon$. In the companion gauge papers, the admissible internal manifolds $\mathcal{M}\_int = CP^{(n-1)}$ equipped with Fisher/Fubini–Study geometry were analysed quantitatively. The distinguishability capacity functional $C^{max}\_\varepsilon(n)$ exhibits a sharp decrease beyond $n = 3$: the available capacity drops by approximately 23% at the $n = 3 \to 4$ transition and approximately 38% at the $n = 4 \to 5$ transition, while the minimum anchor-imposed capacity remains fixed. This scissors behaviour closes at $n = 3$, yielding the internal dimension bound $n \leq 3$. Detailed calculations appear in the companion gauge trilogy. This is the content of the gauge papers' internal dimension argument, now expressed in terms of an operationally defined capacity functional grounded in the Inevitability Programme's axioms.

Axiom Translation Status.

Translation	Status	Notes
Axioms 1–6 ⇒ BC1	✓ Formal (Lemma H3.1)	Requires RC-π and density regularity
Axioms 1–6 ⇒ BC2	~ Operational (Proposition H3.3)	Depends on observational equivalence quotient and minimal complexity principle
Axioms 1–6 ⇒ BC3	✓ Formal (Lemma H3.9)	Anomaly ⇒ ghost/unitarity violation ⇒ contradicts Axiom 6
Axioms 1–6 ⇒ FIM	✓ Formal (Lemmas H3.4–H3.5, Props H3.7–H3.8)	Protocol envelope + Čencov monotonicity + capacity functional

All four translations are now formally or operationally established: BC1 (Lemma H3.1), BC2 (Proposition H3.3), BC3 (Lemma H3.9), and FIM (Lemma H3.4 + Propositions H3.5–H3.8). The axiom correspondence is structurally closed.

Summary of the Mountain

Hard Result	Status	Difficulty	Impact
H1: Independent TPB window	Formal framework established (certifiability + maintenance bounds); closure requires substrate-specific calculations (A)–(C)	Medium	Eliminates circularity in generation count
H2: Exactly three generations	Numerical computation specified, not yet executed	Medium	Major prediction if successful
H3: SM gauge group integration	Derived in companion gauge papers (BC1–BC3 + FIM); axiom translation to Axioms 1–6 required	Medium (dictionary task)	Closes restricted → full uniqueness gap

---

I. Logical Dependency Map

The full theorem chain, with dependencies made explicit:

```
Axiom 1 (Empirical meaning)
      │
      ▼
Theorem A (Records are forced) ◀── Axiom 2 (Stable records)
      │
      ▼
Lemma A1 (Records require coarse-graining) ◀── Axiom 3 (Finite
distinguishability)
      │
      ▼
Corollary A2 (Coarse-graining + reversibility ⇒ irreversibility) ◀── Axiom 4
(Microscopic reversibility)
      │
      ├──────────────────────────────────────┐
      ▼                                      ▼
Theorem B (Entropy cost per bit)      Theorem C (TPB is forced)
      │                                      │
      │                                      ▼
      │                               Proposition C1 (TPB window motivated)
      │                                      │
      ├──────────────────────────────────────┤
      ▼                                      ▼
Proposition D (EFT locality) ◀── Axiom 5   Theorem E (Skyrme stabiliser
forced)
                              Axiom 6        │
                                             ▼
      │                               Proposition E.2 (Stabiliser unique)
      │                                      │
```

```
    ├─────────────────────────────────────────────┤
    ▼                                              ▼
Theorem F1 (Gauge structure forced)   THEOREM U (Restricted uniqueness)
    │                                              │
    ▼                                              ▼
Theorem F2 (Compact Lie group)        H1, H2 (Full uniqueness programme)
    │
    ▼
┌──────────────────────────────────────────────────────┐
│ GAUGE PAPER TRILOGY (BC1–BC3 + FIM framework)          │
│                                                        │
│ Gauge necessity ◀── BC1 (≈ Axiom 4 + records)          │
│ Internal dim n ≤ 3 ◀── capacity bound                  │
│ Anchor A (confinement) ──▶ SU(3) sector                │
│ Anchor B (two-level)   ──▶ SU(2) sector                │
│ Anchor C (phase)       ──▶ U(1) sector                 │
│ Abelian uniqueness ◀── BC3 (≈ Axiom 6)                 │
│ Chirality restriction ◀── BC1                          │
│                                                        │
│ Result: G_SM = SU(3) × SU(2) × U(1) forced             │
└──────────────────────────────────────────────────────┘
                          │
                          ▼
              H3 (Axiom translation:
              Axioms 1–6 ⟹ BC1–BC3 + FIM)
```

Every arrow is a logical dependency. Every node is either an axiom (stated and defended), a proven theorem (with proof sketch), a motivated proposition (with path to full proof), or an identified integration task (with scope defined). There are no hidden assumptions and no circular dependencies.

The path from "experiments have outcomes" to "the BCB architecture is forced" passes through six axioms, a theorem chain (A–F), and a restricted uniqueness result (U), with two operational propositions. The specific gauge group SU(3) × SU(2) × U(1) is derived in the companion gauge paper trilogy from a closely related axiom framework; completing the axiom translation (H3) would make this derivation a direct consequence of the same inevitability chain. Section J below makes precise what structural uniqueness would follow from closure of H1–H3.

J. From Restricted to Structural Uniqueness

Theorem U establishes restricted uniqueness within the fold-based subclass $\mathcal{T}$_fold ⊂ $\mathcal{T}$. This section makes precise what remains once the three closure tasks H1–H3 are completed.

We separate three levels of uniqueness:

Formal inevitability (Sections A–F). Records, coarse-graining, entropy cost, TPB-like conversion invariant, local EFT structure, and gauge redundancy with compact Lie group are structurally forced by Axioms 1–6 (plus scale separation and cluster decomposition for Proposition D).

Restricted structural uniqueness (Theorem U). Within the class of fold-based EFTs with stable records and finite-derivative truncation, the BCB architecture is forced up to field redefinitions and EFT operator-basis equivalence.

Structural uniqueness within $\mathcal{T}$ (this section). Under closure of H1–H3, the residual freedom reduces to representation choice and coordinate redundancy. No alternative architecture within $\mathcal{T}$ remains.

The goal of this section is to state precisely what is eliminated at each closure step.

J1. What H1 Eliminates (TPB Non-Tunability)

H1 requires independent derivation of:

$$TPB_{min} = \Delta\tau_{cert}(\delta) , \quad TPB_{max} = \Delta\tau_{pert} / N_{min}$$

At present, existence of the window is forced (Proposition C1), but numerical bounds depend on substrate-specific quantities.

Once (A) $\Delta\tau_{cert}$ is derived from substrate noise statistics, (B) $\Delta\tau_{pert}$ from perturbation accumulation under substrate coupling, and (C) N_{min} from fold coercivity margins, then: the TPB window becomes a substrate prediction, the number of admissible fold eigenstates becomes non-adjustable, and no "fit" freedom remains in generation counting.

Formally: H1 eliminates parametric freedom in commitment density. After H1, TPB is no longer a free structural degree of freedom — it is fixed by substrate architecture.

J2. What H2 Eliminates (Generation Ambiguity)

H2 requires solving the fold eigenproblem and applying the independently fixed TPB window.

Let the fluctuation operator spectrum be $\{\lambda_1, \lambda_2, \lambda_3, \lambda_4, \dots\}$. Let admissibility require $\lambda_i \in \mathcal{W}_{TPB}$, where $\mathcal{W}_{TPB}$ is the TPB-derived admissible band. If exactly three eigenvalues fall in the admissible window, then the number of generations is not assumed, not anthropic, not symmetry-imposed — it is spectrally forced.

Formally: H2 eliminates spectral degeneracy freedom. After H2, the number of particle families becomes a theorem rather than an input.

J3. What H3 Eliminates (Gauge Content Ambiguity)

Theorems F1–F2 derive existence of a compact Lie gauge group G. The companion gauge trilogy derives $G = SU(3) \times SU(2) \times U(1)$ from BC1–BC3 + FIM + three empirical anchors.

H3 requires demonstrating: Axioms 1–6 ⟹ BC1–BC3 + FIM.

Sections H3.1–H3.9 show: BC1 follows from Axiom 4 + coarse-graining (Lemma H3.1); FIM follows from statistical distinguishability + monotonicity (Lemma H3.4, Proposition H3.5); BC2 follows from finite distinguishability + the observational equivalence quotient (Proposition H3.3); BC3 follows from Axiom 6 (Lemma H3.9).

Once that translation is completed rigorously, the gauge derivation is no longer external infrastructure.

Formally: H3 eliminates gauge-group freedom within $\mathcal{T}$. After H3, alternative compact Lie groups violate either distinguishability capacity bounds, entropy minimality under empirical equivalence, anomaly cancellation under Axiom 6, or empirical anchors strictly weaker than SM assumptions.

J4. What Remains After H1–H3

Once H1–H3 are completed, remaining freedom falls into four categories:

(1) Field redefinitions. Local invertible transformations of fields $\Psi \mapsto F(\Psi)$ that preserve S-matrix elements. These are equivalence transformations, not distinct theories.

(2) EFT operator basis choice. Operators related by integration by parts, equations of motion, or field redefinitions. These do not change physical predictions.

(3) Coordinate and gauge freedom. Reparametrisation of the substrate ordering parameter τ, the emergent time coordinate t_phys, and gauge potentials under local gauge transformations. These are representational freedoms.

(4) UV completion above cutoff. Proposition D applies at accessible scales. UV completions may differ while flowing to the same EFT. Structural uniqueness is claimed within $\mathcal{T}$ at the level of admissible effective description, not at arbitrarily high-energy completions.

J5. Structural Uniqueness Statement

Theorem SU (Conditional Structural Uniqueness within $\mathcal{T}$).

Hypotheses. Assume:

1. Axioms 1–6 hold.
2. Proposition D holds under empirically satisfied scale-separation and cluster decomposition conditions.
3. **H1:** TPB bounds are derived from substrate microdynamics (TPB window non-tunable).
4. **H2:** The fold eigenvalue problem yields exactly three admissible bound states under the independently derived TPB window.

5. **H3:** Axioms 1–6 entail BC1–BC3 + FIM (via Lemmas H3.1, H3.9, H3.4 and Propositions H3.3, H3.5), and the companion gauge derivation of G = SU(3) × SU(2) × U(1) is therefore internal to $\mathcal{T}$.

Let $\mathcal{T}$_phys ⊂ $\mathcal{T}$ denote theories satisfying these closure conditions.

Conclusion. For any theory T ∈ $\mathcal{T}$_phys, there exists an equivalence transformation E such that E[T] ≡ T_BCB up to:

- (i) Local field redefinitions,
- (ii) EFT operator basis equivalence (integration by parts, EOM redundancy),
- (iii) Gauge redundancy transformations,
- (iv) Coordinate reparametrisations of τ and its monotone calibrations,
- (v) UV completion above the EFT cutoff.

No additional structural freedom remains within $\mathcal{T}$_phys.

Note on logical scope. Theorems A–C, E, F1–F2, and Lemmas H3.1, H3.4, H3.9 follow directly from Axioms 1–6. Proposition D (EFT locality) additionally assumes empirically observed scale separation and cluster decomposition. Structural uniqueness therefore holds within the empirically realised regime where these conditions are satisfied.

Interpretation. This theorem does not claim uniqueness among all conceivable metaphysical models, uniqueness of UV completion, or uniqueness outside the empirically admissible theory space. It claims: within the class of empirically admissible, record-supporting, reversible-substrate, stable-EFT theories satisfying closure conditions H1–H3, the BCB architecture is structurally unique up to representational equivalence.

On the role of Proposition D in structural uniqueness. Theorem SU depends on Proposition D, which assumes empirical scale separation and cluster decomposition at accessible scales. These conditions are not derived from Axioms 1–6 alone — they are additional empirical regularities universally observed in regimes where effective field theory has been tested. The structural inevitability of records, entropy cost, TPB, and gauge redundancy follows from Axioms 1–6 alone. The EFT locality component of structural uniqueness additionally assumes these empirically satisfied regularity conditions. The six-axiom narrative therefore carries an explicit asterisk: six axioms suffice to force the structural backbone; EFT locality requires two additional empirically satisfied conditions.

J6. Why This Is Stronger Than Empirical Adequacy

The Standard Model is empirically adequate but structurally underdetermined: it works, but one can ask why it has three generations, why it has SU(3) × SU(2) × U(1), and why it has the specific operator content it has, without receiving answers from within the framework.

BCB, under H1–H3 closure, would be empirically predictive (Lagrangian paper), structurally inevitable (Inevitability Programme), gauge-content determined (gauge trilogy), spectrally constrained (generation count), thermodynamically grounded (entropy cost), and conversion-calibrated (TPB window).

The difference is categorical: empirical adequacy says "this works." Structural inevitability says "this is the only architecture that works within the admissible class."

J7. Intellectual Honesty Clause

If H1 fails — if the TPB window remains tunable — structural inevitability collapses to parametrised consistency. The architecture is still forced, but the generation count becomes a fit rather than a prediction.

If H2 fails — if the generation count is not spectrally fixed — predictive strength weakens. The framework still constrains but does not determine the number of particle families.

If H3 fails — if the gauge translation cannot be completed — uniqueness remains restricted to Theorem U. The BCB architecture is forced within $\mathcal{T}$_fold, but the specific gauge group requires external input.

The programme's strength lies in its falsifiability. Each closure task is concrete and computationally definable. Failure at any stage does not invalidate the established results — it limits how far inevitability extends.