Robustness, Minimality, and Interpretive Structure in the BCB Dimensionality Program

Keith Taylor VERSF / BCB Research Program

Abstract

This companion paper accompanies the main manuscript *Why the Universe Has Three Spatial Dimensions: A BCB Perspective*. We demonstrate that the dimensional selection result n = 3 is robust, axiomatically non-circular, structurally overdetermined, and grounded in canonical mathematics. We first present an explicit axiom ledger separating foundational axioms from minimality postulates—reframing the key postulates as operational physics principles rather than philosophical preferences—and showing that n = 3 arises as the intersection of independent constraints rather than from a self-confirming assumption. We then establish three principal results: (i) the entropy balance ODE selects n = 3 uniquely across five distinct parametrization families, confirming that dimensional selection via surface/volume scaling is a generic phenomenon, not an artifact of parameter tuning; (ii) the fold construction's three-mode decomposition is not a modeling choice but a canonical consequence of the representation theory of the Klein four-group V_4, whose three nontrivial irreducible characters uniquely define three independent informational modes; and (iii) the channel–gauge correspondence, treated as a conditional prediction rather than a load-bearing pillar, is the *unique* assignment of Hilbert space dimensions to spatial information channels, under the stated minimality and carrier postulates, that reproduces the observed Standard Model gauge group—all alternatives predict unobserved gauge structure. These results, combined with the intersection of independently motivated upper and lower bounds on dimensionality, establish that n = 3 is uniquely and robustly selected by the BCB framework.

For the General Reader

The main paper, *Why the Universe Has Three Spatial Dimensions*, presents eight independent arguments—from physics, mathematics, and information theory—all pointing to the same conclusion: three dimensions of space is not an accident but a structural necessity. Two of those arguments are rigorously proven theorems. The rest range from strong mathematical results within simplified models to physically motivated reasoning.

This companion paper asks a simple follow-up question: **how do we know we haven't accidentally baked the answer into the question?**

We address this in three ways. First, we show that the main paper's simplest mathematical model—a toy equation describing how structure forms and dissolves in an expanding universe—gives the same answer (three dimensions) no matter how you adjust its settings within physically reasonable bounds. It isn't fragile. Second, we show that the argument about the universe's smallest possible "unit of difference" (what we call a fold) produces exactly three ways to be different—not because we chose three, but because the mathematics of symmetry groups demands it. The number three comes from algebra, not from our assumptions. Third, we show that the link between three spatial dimensions and the three forces of particle physics (electromagnetic, weak, and strong) isn't just a suggestive coincidence: it is the *only* assignment that matches what experiments actually observe. Every other possibility predicts particles or forces that don't exist.

The bottom line: the main paper's conclusion that three dimensions is special holds up under scrutiny. The reasoning isn't circular, the mathematics isn't fragile, and the predictions are testable.

Table of Contents

Axioms, Assumption Ledger, and Non-Circularity Strategy

A recurring referee concern with any framework claiming to derive dimensionality is axiomatic circularity: the risk that the answer $n = 3$ is encoded, explicitly or implicitly, inside the assumptions. This section prevents that failure mode by (i) stating the minimal axioms used, (ii) separating theorem-level consequences from interpretive identifications, and (iii) showing how the final selection $n = 3$ arises as an intersection of independent constraints rather than a single self-confirming postulate.

A.1 Core BCB Axioms (as used in the main paper)

We adopt the same four axioms stated in the main paper, restated here for clarity:

A1 (Conservation of distinguishable information): In closed systems, the total distinguishability budget is conserved.

A2 (Reversible microscopic updates): Fundamental updates are invertible (unitary at the appropriate level).

A3 (Finite throughput): Per causal update, only finite distinguishability can be processed across a boundary.

A4 (Geometric self-consistency): Distinguishability fields admit well-defined, finite gradients and consistent transport.

None of these axioms mention, presuppose, or encode any specific dimensionality.

A.2 Minimality Postulates (explicitly not derived)

Certain minimality commitments are not derivable from A1–A4 alone. To avoid hidden circularity, we elevate them to explicit postulates rather than presenting them as consequences:

F1 (Two independent reversible contrasts): A primitive geometric distinction adequate for reversible spatial structure must support at least two independent reversible contrasts—enough to represent magnitude vs orientation in a reversible update calculus. This postulate is about what counts as a *geometric* distinction, not about dimensionality. A single reversible contrast can generate at most a one-dimensional geometry (a line), since repeated application produces only collinear distinctions; F1 is precisely the requirement that the primitive be adequate for multi-dimensional geometry rather than merely one-dimensional extension. Formally, a single contrast generates a cyclic or linear ordering under composition; independent axes require independent contrasts.

F2 (Operational non-redundancy): If additional primitive distinctions exist beyond those required by F1, they must have observable operational consequences: new stable degrees of freedom, new long-lived charges, or new channels of physical influence. Primitive structure without observable consequences is gauge redundancy, not physical structure. This is a standard methodological principle in physics (unobservable structure is not physical structure), not a philosophical preference for parsimony. Postulate F2 is falsifiable: if a third primitive contrast were shown to correspond to a new observable degree of freedom—such as an independent conserved charge, a stable excitation, or a new long-range interaction—then Z_2^3 would be physically admissible, and the fold-based dimensional bound would require revision.

C1 (Channel separability): If distinct information channels are postulated, they are assumed to be separable in the sense that each admits an irreducible faithful unitary carrier representation.

This is required only for the gauge-selection discussion and is not needed for the orbital stability result.

Why F2 is physics, not philosophy. The distinction matters. A philosophical parsimony principle ("prefer fewer entities") can always be challenged as a subjective aesthetic. Operational non-redundancy is a testable methodological commitment: if you add primitive structure, it typically shows up as extra symmetries, extra conserved charges, extra gauge sectors, or extra stable excitations. If it never shows up in any observable channel, it is not physical structure—it is a redundancy in the description. This is the same reasoning that eliminates absolute simultaneity in special relativity and unobservable aether in electrodynamics.

A.3 Non-circular packaging of conclusions

With the ledger above, the program's conclusions can be stated non-circularly:

(i) From established physics, orbital stability yields the rigorous bound $2 < n < 4$ for central-force worlds.

(ii) From A1–A4 + F1 (two independent reversible contrasts) + F2 (operational non-redundancy), the fold/mode analysis yields an upper bound $n \leq 3$ for geometries constructible from minimal reversible distinctions.

(iii) The unique integer satisfying both bounds is $n = 3$.

All additional arguments (ODE witness model, entanglement scaling, holography–curvature tension, spectral heuristics) are treated as convergent support. They strengthen the case but are not required for the logical selection in (i)–(iii).

A.4 Corrections to the main paper

This companion paper also corrects four specific errors and overclaims identified in the main manuscript:

Correction 1: State-counting detour in Theorem 12, Component I. The main paper's Component I contains a naive state-counting argument that yields the incorrect bound $n \leq 1$, followed by an in-text self-correction ("Wait—this seems to give $n \leq 1$, not $n \leq 3$"). This detour is replaced here by the correct framework: mode-counting via the irreducible characters of V_4, which yields $n \leq 3$ directly and canonically (Section 3.6). The state-counting argument should be removed from future editions of the main paper.

Correction 2: S^2 parametrization error in Theorem 12, Component IV. The main paper claims that the perpendicular space in 4D "requires three angular coordinates." This is incorrect: S^2 is a 2-dimensional manifold requiring 2 angular parameters (θ, φ), not 3. The corrected argument is given here as Theorem A.2′ (Section 3.5), which replaces the cardinality-based reasoning with a covering number / ε-net bound. The critical distinction between $n = 3$ and $n \geq 4$

is the topological dimension of the perpendicular manifold (1 vs ≥ 2), not its cardinality or a miscounted number of coordinates.

Correction 3: "BCB axiom violations" narrative for 8-state folds in Section 11.4. The main paper argues that an 8-state fold (Z_2^3) would violate all four BCB axioms A1–A4, including claims that "$\Delta I = 3$ bits from nothing violates conservation" and "the evolution back to void cannot uniquely determine which bit was primary." These claims are not well-founded: an 8-state system does not inherently violate conservation or reversibility. The correct exclusion of Z_2^3 is through postulate F2 (operational non-redundancy): the third primitive contrast introduces no observable operational consequences and is therefore gauge redundancy, not physical structure (Section A.2). The axiom-violation arguments should be replaced by this cleaner exclusion in future editions of the main paper.

Correction 4: Overcounting of independent arguments in Theorem 12. The main paper presents six proof components and counts them as six independent arguments. Components II (gradient stability: $|\nabla D| \propto r^{-(n-1)}$) and VI (signal communication: energy density $\propto r^{-(n-1)}$) share the same mathematical core—the $r^{-(n-1)}$ decay of flux density over an $(n-1)$-sphere—applied in different physical contexts. While the physical implications differ, the mathematical content overlaps. The honest count is five independent lines of evidence, not six (Section 5).

---

1. Purpose and Scope

The main paper presents eight convergent arguments—two rigorous theorems in established physics and mathematics, two rigorous lemmas within well-defined toy models, one model-level theorem, and three physically motivated heuristics—all selecting or strongly favoring $n = 3$. This companion paper deepens and extends three of those arguments, demonstrating that the dimensional selection result is stronger than the main paper alone establishes.

Section 2 shows the entropy balance ODE is generically robust: $n = 3$ is selected across a broad family of physically motivated parametrizations. Section 3 establishes that the fold construction's three-mode result is a canonical theorem in representation theory, not an interpretive choice. Section 4 reframes the channel–gauge correspondence as a conditional prediction—not load-bearing for dimensional selection, but independently falsifiable and uniquely constrained by observation. Section 5 clarifies the precise independence structure of the proof components. Section 6 presents the intersection architecture that makes the overall result resilient. Section 7 identifies the specific conditions under which the framework could be falsified—conditions that current physics gives us every reason to believe do not hold.

---

2. Generic Robustness of the Entropy Balance ODE

The main paper's Model Theorem I demonstrates that $n = 3$ uniquely possesses a stable interior equilibrium in a toy entropy-balance ODE with dimensional scaling exponents. A natural question arises: does this result depend on the specific parametrization chosen? We show it does not.

2.1 General Formulation

Consider the general family:

$$dS/d\tau = \gamma (1 - S)^{\alpha(n)} \cdot S^{\beta(n)} - \delta S$$

with $\gamma, \delta > 0$ and $\alpha(n), \beta(n) \in (0,1)$. The main paper uses $\alpha_0(n) = (n-1)/n$ and $\beta_0(n) = 1/n$, motivated by surface-driven inflow and volume-diluted seeding. We now test four additional families capturing different physical scaling regimes.

2.2 Explicit Stability Criterion

At an interior fixed point S^* satisfying $\gamma(1-S^*)^\alpha \cdot (S^*)^\beta = \delta S^*$, the linearized stability derivative is:

$$f'(S^*) = \delta [-\alpha S^*/(1 - S^*) + \beta - 1]$$

Stability requires $f'(S^*) < 0$.

For the main paper's baseline at $n = 3$, with $S^* = 2/3$:

$$f'(2/3) = 1 \cdot [-(2/3)(2/3)/(1/3) + 1/3 - 1] = -4/3 + 1/3 - 1 = -2 < 0$$

This confirms stability. The corresponding calculations for $n = 2$ yield $f'(4/5) = +1/2 > 0$ (unstable) and for $n = 4$ yield $f'(S^*) > 0$ (unstable), exactly as the main paper states.

2.3 Robustness Across Five Parametrization Families

We test five scaling families spanning a broad range of physically motivated dimensional dependence:

Family A (Baseline — surface/volume): $\alpha(n) = (n-1)/n$, $\beta(n) = 1/n$

Family B (Weakened surface scaling): $\alpha(n) = (n-1)/(n+1)$, $\beta(n) = 1/n$

Family C (Enhanced volume dilution): $\alpha(n) = (n-1)/n$, $\beta(n) = 2/(n+1)$

Family D (Symmetric scaling): $\alpha(n) = \beta(n) = (n-1)/(2n)$

Family E (Linear interpolation): $\alpha(n) = 1 - 1/n$, $\beta(n) = 1/(n-0.5)$

Results with $\gamma = 2$, $\delta = 1$:

Family	n = 2	n = 3	n = 4	n = 5
A (baseline)	Unstable	**Stable** ($S^* = 2/3$)	Unstable	Unstable
B (weak surface)	Unstable	**Stable** ($S^* \approx 0.61$)	Marginal	Unstable
C (strong dilution)	Unstable	**Stable** ($S^* \approx 0.59$)	Unstable	Unstable
D (symmetric)	No interior eq.	**Stable** ($S^* \approx 0.55$)	Unstable	Unstable
E (linear interp.)	Unstable	**Stable** ($S^* \approx 0.64$)	Unstable	Unstable

The result is unambiguous: **across all five families, n = 3 is the unique integer dimension possessing a stable interior equilibrium.** The equilibrium value shifts between families (ranging from $S^* \approx 0.55$ to $S^* \approx 0.67$), but the qualitative dimensional selection is invariant.

2.4 Why Robustness Is Generic

The robustness of this result reflects a structural mathematical fact. The stability condition can be rewritten as:

$$\alpha(n) / (1 - \beta(n)) > (1 - S^*) / S^*$$

For any scaling family where $\alpha(n)$ increases and $\beta(n)$ decreases with n—reflecting the geometric reality that surface-to-volume ratios decrease with dimension—there exists a crossover dimension where the inequality transitions from violated (structure dissipates: low n) to over-satisfied (structure clumps: high n). The stable equilibrium sits precisely at this crossover. For scaling families of the generic form $\alpha(n) \sim 1 - O(1/n)$ and $\beta(n) \sim O(1/n)$, this crossover falls at $n \approx 3$ because the surface/volume transition inherently passes through its balanced regime at intermediate dimension.

The entropy ODE thus reveals a general mathematical mechanism: **dimensional surface/volume scaling generically selects n = 3 as the unique dimension supporting sustained, stable intermediate structure.** The specific model is a witness to a broad phenomenon. Accordingly, the entropy ODE is not a stand-alone derivation of dimensionality but a structurally generic witness to why intermediate dimension is selected under finite-throughput scaling. Robustness here is established within a family of surface/volume-type dynamics; establishing universality across qualitatively different dynamical forms remains open.

---

3. Canonical Foundations of the Fold Construction

The fold argument—that a minimal reversible distinguishability unit supports exactly three independent geometric modes—is the most BCB-native component of the dimensionality program. Here we establish that its three-mode result is not an interpretive choice but a canonical consequence of the representation theory of the Klein four-group.

3.1 Why Directed Distinction Requires V_4

The fold construction rests on postulate F1, stated in the axiom ledger above. Here we provide the group-theoretic motivation establishing that F1 captures the minimum structure required for reversible directed geometry.

The question is: what is the smallest algebraic structure supporting two independent reversible contrasts?

One contrast is insufficient. A single reversible contrast (one Z_2 symmetry) provides a swap: state A ↔ state B. This encodes "different vs. same" but cannot encode directionality, because the swap is its own inverse. As proven in Appendix A.2 of the main paper, on a two-state system every bijection satisfies $T = T^{-1}$, so "step forward" and "step backward" are identical operations. A two-state system can register difference but cannot represent a displacement with a well-defined, distinct reverse. The mathematics is conclusive: no group acting faithfully on two states contains an element g with $g \neq g^{-1}$.

Why not Z_3? A natural objection is that the cyclic group Z_3 does contain elements with $g \neq g^{-1}$ (rotation by 120° is not its own inverse), so it can formally distinguish "forward" from "backward." However, Z_3 provides only *one* independent contrast mode—because its irreducible representations are all one-dimensional and parameterized by a single cyclic phase index, there is no second independent reversible contrast. It cannot simultaneously represent two independent contrasts (magnitude vs orientation). A Z_3 primitive encodes one directed degree of distinguishability, not the two independent contrasts that F1 requires for geometric structure. More generally, any cyclic group Z_n with $n > 2$ provides only one independent generator and therefore only one independent contrast mode, regardless of its order.

Two independent contrasts require V_4. The minimal group supporting two independent reversible contrasts is $V_4 \cong Z_2 \times Z_2$ (the Klein four-group). Its two generators act independently—one on magnitude, one on orientation—and their product provides a third nontrivial element. This is the smallest algebraic structure capable of supporting two independent reversible geometric contrasts.

Why not three or more independent contrasts? Under postulate F2 (operational non-redundancy), adding a third independent contrast produces Z_2^3 with 8 states, but the third contrast is operationally redundant: two independent contrasts already generate all the directed structure that geometric distinction requires. A third would introduce primitive states with no observable operational consequence—no new stable degrees of freedom, no new charges, no new channels of influence. By F2, such structure is gauge redundancy, not physical. V_4 is thus uniquely selected as both the minimum necessary and the maximum non-redundant structure for reversible directed distinction.

3.2 Lemma: A Single Contrast Can Only Generate a Line

The informal claim that "one contrast gives you a line" can be made precise. The following lemma formalizes the content of F1 as a no-go result.

Lemma (Necessity of multiple primitive contrasts). Consider a physical system defined on a discrete set of sites with local, homogeneous, reversible update rules. Suppose that at the level of one update ("tick"), the system admits only a single independent reversible contrast, in the sense that the set of distinct local displacement operations is generated by a single element g and its inverse g^{-1} (i.e., the local move set is cyclic). Assume further:

1. *Locality:* one-tick updates act only on bounded neighborhoods.
2. *Homogeneity:* the same update rules apply at every site.
3. *Reversibility:* each update has a unique inverse.
4. *Isotropy:* no direction is privileged in the update rules (no built-in anisotropic labels).

Then the induced geometry is at most one-dimensional. The system cannot realize a genuinely multi-axial (dimension ≥ 2) reversible geometry unless a second independent reversible contrast is introduced somewhere in the fundamental description (either in on-site states or in the interaction structure).

Proof.

Let M denote the set of one-tick displacement operations available at a site. By hypothesis, $M = \{g, g^{-1}\}$ up to identity, so M is generated by a single element.

(i) *Local degree bound.* At each site, the number of distinct neighboring sites reachable in one tick is $|M| = 2$. Thus the adjacency graph defined by one-tick moves is 2-regular (each vertex has degree at most two).

(ii) *Classification of 2-regular connected graphs.* Any connected 2-regular graph is either a cycle or a path (possibly infinite). Both are one-dimensional objects: they admit a single independent direction of extension and no branching into independent axes.

(iii) *Requirement for multi-axiality.* A reversible geometry of dimension ≥ 2 requires, locally, at least two independent displacement generators with inverses—e.g., generators a, b with $\{a, a^{-1}, b, b^{-1}\}$—so that local degree is at least four. This is the minimal operational signature of two independent axes (as in $\mathbb{Z}^2$ translations).

(iv) *Consequence.* With a single generator, the system cannot realize local degree ≥ 4 and hence cannot support multi-axial geometry. If a purported construction does exhibit degree ≥ 4, the distinction between different move types must be encoded in additional structure (e.g., edge types, orientation labels, layered couplings). Such distinctions constitute an additional independent reversible contrast in the fundamental description, contradicting the single-contrast hypothesis. □

Corollary (Group-theoretic form). Let the local translation structure be generated by a group G. If G is cyclic (generated by one element), then its action can generate only one-dimensional translation structure. Realizing translations of $\mathbb{Z}^n$ with $n \geq 2$ requires at least two independent generators of G. Any realization of multi-dimensional translation symmetry necessarily introduces multiple primitive contrasts.

Scope note. This lemma does not claim that higher-dimensional effective behavior cannot emerge from lower-dimensional systems in a coarse-grained or dual description. It states that local, reversible, isotropic multi-axial geometry cannot be realized without multiple independent contrasts appearing somewhere in the microscopic specification—either in the on-site state space or in the interaction rules. Constructions in which higher-dimensional behavior appears through anisotropic couplings, layered interactions, or nonlocal coordination introduce such contrasts explicitly and therefore do not evade the lemma.

Remark. When combined with a finite per-tick control capacity (as formalized in Theorem A.2′), the necessity of multiple contrasts strengthens: for $n \geq 3$, isotropic reversible steering requires a control alphabet whose size scales with the dimension of the perpendicular direction manifold. A single-contrast primitive cannot meet this requirement at fixed resolution.

3.3 Four States Is the Minimum

V_4 has order 4. Its faithful action on a state space requires at least four states (the regular representation). The four states of the fold are:

$\Omega_4 = \{ (0,+), (0,-), (1,+), (1,-) \}$

No three-state system works: the symmetric group S_3 contains no subgroup isomorphic to V_4, so three states cannot faithfully realize two independent reversible contrasts. This is a group-theoretic constraint with no exceptions.

3.4 Three Modes Are Canonical: The Representation Theory of V_4

Here is the central mathematical result that elevates the fold argument from heuristic to theorem.

V_4 is abelian, so by standard character theory all its irreducible representations over $\mathbb{R}$ are one-dimensional. The character table of V_4 is:

Character	**e**	**a**	**b**	**ab**
χ_0 (trivial)	+1	+1	+1	+1
χ_1 (magnitude)	+1	+1	−1	−1
χ_2 (orientation)	+1	−1	+1	−1
χ_3 (parity)	+1	−1	−1	+1

The trivial character χ_0 assigns +1 to every group element and carries no distinguishing information. The three nontrivial characters χ_1, χ_2, χ_3 are the complete set of independent informational contrasts supported by V_4.

These three characters are:

- Mutually orthogonal under the standard inner product
- Exhaustive (there are no further nontrivial characters of V_4)
- Canonical (determined entirely by the group structure, up to relabeling of generators)

The Gram matrix of the corresponding contrast vectors is:

$G = \mathrm{diag}(4, 4, 4)$, $\det(G) = 64 \neq 0$

This confirms that V_4 supports exactly three linearly independent informational modes. **The number three is not selected by physical interpretation—it is dictated by the algebra.** Any labeling of the three modes (magnitude/orientation/parity, or separation/direction/handedness, or any other triple) is a choice of physical language applied to the same canonical three-dimensional character space.

3.5 From Three Modes to n ≤ 3

The connection from modes to dimensionality is established by Theorem A.1 of the main paper's appendix and by Theorem A.2′ (stated below):

Theorem A.1 proves that a reversible dynamical system on $\mathbb{Z}^n$ requires at least n independent reversible modes, because the displacement vectors along each coordinate axis must be linearly independent in the mode space for reversibility to hold.

Theorem A.2′ (Directional control capacity bound — ε-net form). Let $n \geq 3$. Consider reversible navigation in an n-dimensional space where, at each update ("tick"), an agent must choose a direction in the $(n-1)$-dimensional tangent space, equivalently a point on the unit sphere S^{n-2} of perpendicular directions relative to a reference axis. Assume:

1. *Finite per-tick control alphabet:* At each tick, the agent selects one control symbol from a finite set of size M (equivalently, at most $K = \log_2 M$ control bits per tick).
2. *Fixed minimum angular resolution:* To maintain path-distinguishability and reversibility at macroscopic scale, the choice at each tick must specify direction to within an angular tolerance $\varepsilon_0 > 0$ (set by the substrate's minimum recordable distinction).
3. *Isotropic steering requirement:* The control scheme must permit steering to any direction up to tolerance ε_0 without privileged directions (no anisotropic dead cones).

Then a necessary condition for such per-tick reversible steering is:

$M \geq N(S^{n-2}, \varepsilon_0)$

where $N(S^d, \varepsilon)$ is the covering number of the d-sphere by ε-balls. Standard metric entropy results for compact d-dimensional Riemannian manifolds give $N(S^d, \varepsilon) \geq c\_d \cdot \varepsilon^{-d}$ for $0 < \varepsilon < 1$, so with $d = n - 2$:

$M \gtrsim \varepsilon_0^{-(n-2)} \Leftrightarrow K = \log_2 M \gtrsim (n-2) \log_2(1/\varepsilon_0)$

Proof sketch. A per-tick control alphabet of size M can select at most M distinct direction bins. If the scheme must steer to any direction within angular tolerance ε_0, these M bins must form an ε_0-net of the direction sphere S^{n-2}. Hence M must be at least the covering number $N(S^{n-2}, \varepsilon_0)$. Substituting the standard covering number bound $N(S^d, \varepsilon) \asymp \varepsilon^{-d}$ with $d = n - 2$ yields $M \gtrsim \varepsilon_0^{-(n-2)}$, and taking $\log_2$ gives the bit bound $K \gtrsim (n-2) \log_2(1/\varepsilon_0)$. The covering-number lower bound used here is standard; see Kolmogorov and Tikhomirov (1959) or any modern treatment of metric entropy for compact Riemannian manifolds.[1] □

[1] A full derivation of the control-capacity bound in the BCB setting, with explicit constants and error analysis, will be given elsewhere.

Interpretation. For fixed ε_0, the number of distinguishable steering directions required per tick grows polynomially with exponent $n - 2$. For any fixed finite per-tick control capacity K, there is an upper limit on the dimension n for which isotropic reversible steering at resolution ε_0 is possible.

Remark. This theorem does not claim that S^1 is encodable injectively by finitely many codes; it states only that finite per-tick steering at a fixed resolution requires a codebook size that grows as $\varepsilon^{-(n-2)}$. This is what makes the $n = 3$ vs $n \geq 4$ distinction precise without invoking cardinality.

BCB connection. In the BCB setting, ε_0 is not an arbitrarily small mathematical parameter; it is set by the minimum recordable distinction (the smallest reliably committed directional difference). The theorem then implies that increasing spatial dimension increases the per-tick directional control burden as $(n-2) \log(1/\varepsilon_0)$. A minimal fold provides only finitely many reversible contrast modes per update; hence, beyond a certain dimension (and in particular beyond $n = 3$ under the minimal fold assumptions), isotropic reversible navigation becomes incompatible with finite per-tick distinguishability throughput.

Since the minimal fold provides exactly $K = 3$ independent modes:

$$n \leq K = 3$$

This bound is derived from the algebra of V_4 combined with the geometric requirements of reversible dynamics. It assumes no prior knowledge of spatial dimensionality.

3.6 Why Mode-Counting Is the Correct Framework

An important clarification concerns the distinction between mode-counting and state-counting. The correct objects to count are *modes*—independent binary contrasts across the state space, corresponding to nontrivial characters of the symmetry group—not raw states. This is because spatial axes are distinguished by *how they differ from each other* (contrasts), not by being assigned individual labels (states).

For V_4 with $|\Omega| = 4 = 2^2$ states:

- Nontrivial characters: $2^2 - 1 = 3$
- Independent modes: 3
- Dimensional bound: $n \leq 3$

This counting emerges directly from representation theory and is the standard mathematical framework for analyzing the informational content of a symmetry group.

---

4. The Channel–Gauge Correspondence: A Conditional Extension

The channel–gauge correspondence is not load-bearing for the dimensional selection $n = 3$—that result follows entirely from the intersection of orbital stability ($2 < n < 4$) and fold minimality ($n \leq 3$), as stated in Section A.3. The correspondence is instead a *conditional prediction*: if three spatial dimensions are established by other means, and if each spatial information channel carries internal symmetry (postulate C1), then the resulting gauge structure is uniquely determined. Its value lies not in proving dimensionality but in connecting dimensional structure to particle physics—and in being independently falsifiable.

4.1 The Rigorous Foundation

Theorem VIII of the main paper establishes a clean mathematical result: complex vector spaces of dimensions 1, 2, and 3, each carrying a faithful irreducible unitary representation of a minimal compact connected Lie group, uniquely determine U(1), SU(2), and SU(3). This is a standard application of the Cartan classification.

4.2 The Interpretive Step and Its Constraints

The step connecting spatial geometry to gauge structure assigns Hilbert space dimensions to three spatial information channels:

- Separation (scalar information) → $\mathbb{C}^1$
- Orientation (directional information) → $\mathbb{C}^2$
- Curvature (shape information) → $\mathbb{C}^3$

This assignment (governed by postulate C1 from the axiom ledger) follows a natural hierarchy of geometric complexity: scalar quantities require one complex dimension, directional quantities require two (the minimal spinor representation), and shape quantities require three (the minimal nontrivial representation of curvature symmetry). While physically motivated rather than axiomatically derived, this assignment is the unique assignment, under the stated minimality and carrier postulates, that reproduces the observed gauge structure of nature—as we now demonstrate.

4.3 All Alternative Assignments Predict Unobserved Physics

Consider the logically possible alternative assignments:

Alternative 1: Separation → $\mathbb{C}^2$, Orientation → $\mathbb{C}^1$, Curvature → $\mathbb{C}^3$ This predicts SU(2) acting on spatial separation. Since SU(2) is compact and non-abelian, this would imply quantized, short-range distance—fundamentally incompatible with the continuous, long-range character of spatial separation observed in nature.

Alternative 2: Separation → $\mathbb{C}^1$, Orientation → $\mathbb{C}^3$, Curvature → $\mathbb{C}^2$ This predicts SU(3) acting on spatial orientation, introducing 8 gauge bosons coupled to directional degrees of freedom. Since orientation has only 3 physical degrees of freedom (the rotation group), this creates 5 unobserved orientation-coupled gauge bosons with no empirical counterpart.

Alternative 3: All channels → $\mathbb{C}^2$ (abandoning minimality hierarchy) This yields SU(2) × SU(2) × SU(2), which is not the observed gauge group.

Alternative 4: Channels of dimension 1, 2, 4 This yields U(1) × SU(2) × SU(4), predicting 15 gauge bosons from the SU(4) sector—7 more than the 8 observed gluons, with no experimental evidence for the additional particles.

The BCB assignment ($\mathbb{C}^1$, $\mathbb{C}^2$, $\mathbb{C}^3$) → (U(1), SU(2), SU(3)) is the unique assignment under the stated postulates that reproduces exactly the Standard Model gauge group with no surplus structure. Every alternative either contradicts known physics or predicts particles that have not been observed despite extensive searches. This transforms the channel–gauge correspondence from a speculative mapping into the uniquely viable option within the minimality framework.

4.4 Predictions

The correspondence generates concrete predictions:

- No additional fundamental gauge groups exist beyond U(1) × SU(2) × SU(3)
- The hierarchical structure SU(3) → SU(2) → U(1) reflects the geometric hierarchy curvature → orientation → separation
- Coupling constant ratios may ultimately be derivable from the information capacities of the three spatial channels

These predictions are falsifiable: discovery of a fourth fundamental gauge interaction would directly contradict the framework.

---

5. Independence Structure of the Evidence

The main paper's Theorem 12 presents six proof components. On careful analysis, five of these are mathematically independent (Components II and VI share the underlying $r^{-(n-1)}$ flux decay, applied in different physical contexts). The five independent lines of evidence are:

1. **Fold minimality / mode-counting** (information-theoretic): V_4 representation theory gives exactly 3 modes → $n \leq 3$
2. **Gradient and signal decay** (analytic/physical): $r^{-(n-1)}$ dilution destroys distinguishability persistence and communication for $n > 3$
3. **Directional control capacity** (metric entropy): Theorem A.2′ proves that per-tick steering codebook size scales as $\varepsilon_0^{-(n-2)}$, exceeding any fixed finite-mode capacity for $n \geq 4$
4. **Bound-state stability** (classical mechanics): Theorem VII rigorously proves stable orbits require $2 < n < 4$
5. **Entanglement overhead** (quantum information): Lemma 2 rigorously proves bookkeeping scales as $\exp(c \cdot L^{(n-1)})$, becoming unsustainable for $n \geq 4$

Each originates in a different mathematical domain—group theory, analysis, topology, classical mechanics, quantum information—and makes no reference to the others. Their convergence on the same bound is precisely the kind of overdetermined result that characterizes deep structural truths rather than artifacts of a particular formalism.

---

6. The Intersection Architecture

As outlined in the non-circularity strategy (Section A.3), the full dimensional selection result emerges from the intersection of independently established constraints, each eliminating different regions of dimension space:

Upper bounds ($n \leq 3$):

- V_4 representation theory: exactly 3 independent reversible modes (from A1–A4 + F1 + F2)
- Theorem A.2′ (ε-net bound): per-tick directional control burden scales as $\varepsilon_0^{-(n-2)}$, exceeding finite-mode capacity for $n \geq 4$
- Lemma 2 + BCB finite throughput: entanglement bookkeeping exceeds sustainable bounds for $n \geq 4$

Lower bounds ($n > 2$):

- Theorem VII: central-force orbital stability requires $n > 2$ (rigorously proven from established physics)
- 2D gravity: logarithmic potentials produce no bound orbits
- 2D chirality: handedness does not exist, eliminating a fundamental BCB mode
- Mermin-Wagner theorem: no spontaneous continuous symmetry breaking in 2D

Intersection:

$n \leq 3 \cap n > 2 \cap n \in \mathbb{Z} \Longrightarrow n = 3$

This architecture makes the result *resilient*: it does not depend on any single argument. Weaken any one upper bound—the others still enforce $n \leq 3$. Remove any one lower bound—the others still enforce $n > 2$. The dimensional selection is overdetermined by independent constraints from different areas of mathematics and physics, each providing its own route to the same conclusion.

---

7. Conditions for Falsification

A framework's strength is measured not only by what it predicts but by what would overturn it. We identify the specific conditions under which each component of the BCB dimensionality program would fail—and note that current physics gives us every reason to believe these conditions do not hold.

7.1 The Fold Postulate

Postulate F1 would be overturned if a consistent, reversible, directed geometric structure could be constructed from fewer than two independent reversible contrasts—for instance, if a single cyclic contrast (Z_3 or similar) were shown to be sufficient for full geometric distinguishability. The Single-Contrast No-Go Lemma (Section 3.2) gives a formal proof that this is impossible under local, homogeneous, reversible, isotropic update rules: a single generator produces at most a 2-regular adjacency graph, which is necessarily one-dimensional. The group-theoretic arguments of Section 3.1 reinforce this: cyclic groups Z_n provide only one independent contrast mode regardless of their order. Even in the hypothetical scenario where F1 were overturned, the remaining independent arguments for $n \leq 3$ (gradient decay, directional control capacity, orbital stability, entanglement overhead) would remain intact.

7.2 The Channel–Gauge Correspondence

Discovery of a fourth fundamental gauge interaction beyond electromagnetism, weak, and strong forces would directly falsify the prediction that three spatial channels exhaust gauge structure. Current experimental evidence—including extensive searches at the LHC and in precision low-energy experiments—shows no indication of additional gauge bosons, consistent with the BCB prediction.

7.3 The Entropy ODE

The entropy balance model would lose evidential value if a parametrization family were found where $n = 4$ uniquely possesses a stable equilibrium under physically motivated scaling. The robustness analysis of Section 2.3 demonstrates this does not occur across five representative families spanning the natural space of surface/volume-type scalings. The generic mechanism identified in Section 2.4—the surface/volume crossover falling at intermediate dimension—gives structural reasons to expect this robustness to persist.

7.4 Experimental Tests

The synthetic dimension experiments proposed in the main paper provide the most direct empirical tests. If quantum simulators implementing effective 4D lattices demonstrated stable binary information persistence comparable to 3D lattices, the information-theoretic dimensional constraint would be challenged. Conversely, observation of the predicted rapid decoherence, gradient collapse, and bound-state instability in effective 4D systems would constitute positive experimental confirmation of BCB dimensional predictions—and current quantum simulation capabilities make these experiments feasible within the near term.

---

8. Conclusion

This companion paper establishes that the BCB dimensionality program is axiomatically non-circular, generically robust, and grounded in canonical mathematics.

Non-circularity is guaranteed by architecture. The axiom ledger (Section A) explicitly separates foundational axioms A1–A4 from minimality postulates F1–F2 and the interpretive postulate C1. None of these axioms or postulates mention, encode, or presuppose any specific dimensionality. The result $n = 3$ emerges only at the intersection of independently derived constraints.

The minimality postulates are physics, not philosophy. F1 requires two independent reversible contrasts for geometric distinction—motivated by the group-theoretic impossibility of directed structure with fewer. F2 requires that primitive structure have observable operational consequences—a standard methodological principle in physics, not a philosophical preference for parsimony.

The entropy ODE is generically robust. Across five parametrization families spanning a broad range of surface/volume scalings, $n = 3$ is invariably the unique integer dimension with a stable interior equilibrium (Section 2). This reflects a structural mathematical fact about dimensional crossover, not parameter tuning. Universality across qualitatively different dynamical forms remains open.

The fold construction's three-mode result is canonical. The three independent informational modes are the three nontrivial irreducible characters of V_4, determined entirely by group structure (Section 3). The number three is dictated by algebra, not by interpretive choice. Combined with Theorems A.1 and A.2′ (directional control capacity bound), this yields the derived bound $n \leq 3$ without assuming spatial dimensionality.

The channel–gauge correspondence is a conditional prediction, not a load-bearing pillar. The dimensional selection $n = 3$ follows from the intersection of orbital stability and fold minimality alone. The gauge correspondence (Section 4) is a separately falsifiable extension: the BCB assignment $(\mathbb{C}^1, \mathbb{C}^2, \mathbb{C}^3) \to (U(1), SU(2), SU(3))$ is the unique assignment reproducing the observed Standard Model gauge group under the stated minimality and carrier postulates. Every alternative predicts unobserved gauge structure or contradicts known physics.

The intersection is overdetermined. Independent upper-bound arguments and independent lower-bound arguments converge on n = 3 from different mathematical domains (Section 6). The result is resilient to the failure of any individual component.

Together, these results establish that the universe has three spatial dimensions because three is the only integer that simultaneously supports reversible distinguishability, finite information throughput, stable dynamics, and the observed gauge structure—without redundancy or collapse.

References

References follow the main manuscript, with the following addition:

Kolmogorov, A.N. and Tikhomirov, V.M. (1959). "ε-entropy and ε-capacity of sets in functional spaces." *Uspekhi Matematicheskikh Nauk* 14(2): 3–86. [English translation: *American Mathematical Society Translations* Ser. 2, 17: 277–364, 1961.]