Second-Order Corrections to the Cosmological Constant

Finite-Capacity Competition and Correlation Effects in the VERSF / Two-Planck Framework

Keith Taylor *VERSF Theoretical Physics Program*

For General Readers

One of the deepest unsolved puzzles in physics is why empty space contains so little energy. Quantum mechanics predicts that the vacuum should be seething with energy — in fact, it should contain roughly 10^{120} times more than astronomers actually measure. That is a 1 followed by 120 zeros. It is the worst numerical prediction in the history of science.

This paper is part of a series proposing a solution. The core idea is simple: space is not made of points. It is made of *relationships between* points. The smallest possible relationship requires two things — so the smallest meaningful unit of space is twice the Planck length, the fundamental quantum of distance. When you build space this way, something remarkable happens: the energy of empty space is no longer a free number you can set to anything. It is determined by the geometry itself, and it comes out naturally small — matching the astronomical measurement to within a factor of two.

The earlier papers in the series established this match at first order. They showed that the vacuum energy must lie in a range corresponding to a "grain size" of space between 60 and 110 micrometres — roughly the width of a human hair. The central value of that prediction, around 85 micrometres, matches the observed cosmological constant to within 10–20%. That was already a dramatic improvement: the naive prediction was off by 10^{120}; the first-order derivation lands within a fifth of an order of magnitude. But the derivation produces a *range* of possible values, not a single point. The width of that range — about a factor of three in the vacuum energy — was established but not explained: *why that particular width, and not narrower?*

This paper answers that question. It identifies three specific physical reasons why the range cannot be narrowed further, and proves that each one contributes a bounded, calculable amount. First, the web of relationships making up space is not perfectly uniform — some regions are more tightly connected than others, shifting the grain size slightly. Second, the seven geometric conditions a cell of space must satisfy are not completely independent — they carry small correlations. Third, different physical processes compete for the same limited informational capacity of the underlying substrate. Taken together, these three effects fully account for the observed width of the range. Removing any one of them would not narrow the prediction; it would simply misrepresent the physics.

The result is a complete two-tier picture. The first-order theory establishes the central prediction: the vacuum energy matches the observed value to within 10–20%, with no free parameters. This paper does not improve that figure — the central value is already well-determined. What it does instead is answer a separate question that the first-order theory left open: why does the prediction come as a range at all, and why does that range have exactly the width it does? Both questions need answers for the theory to be complete.

The paper also identifies three numerical experiments — computer simulations of discrete spacetime — that would confirm or falsify the whole framework. No new laboratory equipment is needed. The tests can be run on existing simulation codes.

Abstract

Physics has a problem with empty space. When physicists calculate how much energy empty space should contain, they get an answer roughly 10^{120} times larger than what astronomers actually measure — widely considered the worst numerical prediction in all of science.

A companion series of papers proposed a solution built from a single idea: space is made of *relationships* between points, not the points themselves. The smallest meaningful unit of space is therefore twice the Planck length — the minimum needed for one thing to relate to another. From this starting point, combined with the requirement that space should not collapse under its own vacuum energy, the theory derives a natural coherence scale of approximately 60–110 micrometres — roughly the width of a human hair. The energy of empty space at that scale matches the astronomical observation to within 10–20% at the central value. This first-order result reduces the discrepancy from 10^{120} to agreement at the 10–20% level for the central prediction — a structural result with no free parameters. However, the derivation yields a range of permitted values rather than a single point: $\xi \in [60, 110]$ μm corresponds via $\Lambda \sim \xi^{-4}$ to a factor of ~11 variation across the band. The central value sits within that band near the observation; the width of the band itself was established but not derived.

The present paper addresses the question the first-order derivation leaves open: *why does the predicted band have exactly the width it does?* The factor-of-11 range across $\xi \in [60, 110]$ μm was established but not explained at first order — it emerged from the allowed range of the percolation threshold without a derivation of why that range takes the values it does or how much other physical effects contribute. This paper provides that derivation. Three second-order effects are identified, each bounded by an independent physical argument: (1) topological non-uniformity of the relational graph causes the percolation threshold to vary, contributing $\sigma\_\ln \xi \sim 0.25$; (2) small correlations among the seven geometric constraints modify the effective coupling, contributing $\sigma\_\ln \xi \sim 0.12$; (3) finite informational throughput of the substrate creates competition among physical processes, contributing $\sigma\_\ln \xi \sim 0.20$. Combined in quadrature, these give $\sigma\_\ln \xi \sim 0.33$, or equivalently $\sigma\_\ln \Lambda \sim 1.3$. The observed cosmological constant lies near the centre of both the first-order band and the second-order $\pm 1\sigma$ interval. Together, the two tiers constitute the theory's complete account of the cosmological constant. The first-order result establishes the central value to 10–20% accuracy with no free parameters — this paper does not

improve that figure, which is already well-determined by three independent routes converging on $\xi \approx$ 75–88 µm against the observed ~85 µm. What this paper does instead is close a separate open question: why does the framework predict a range rather than a point, and why is that range as wide as it is?

All three effects are unified into a single master equation. The paper derives limiting cases as internal consistency checks and identifies three Monte Carlo simulations — measurable in existing dynamical triangulation codes — that would confirm or falsify the framework independently of any laboratory experiment.

Table of Contents

Executive Summary

The Two-Planck paper [TP] demonstrated that the cosmological constant Λ is not a free parameter but an emergent consequence of relational geometry and discrete constraint satisfaction. Three independent derivation routes (Routes A, B, and M of TP) converge on a characteristic coherence scale $\xi \sim 60\text{–}110$ μm, arising as a geometric mean between the Planck length and the cosmic horizon [TP §3.2, §4.8]. From this scale, the observed vacuum energy density follows directly:

$\rho_{vac} \sim \hbar c/\xi^4$, $\Lambda \sim 10^{-52}\ \mathrm{m}^{-2}$

This gives the correct order of magnitude, with the best route-level agreement at the ~10–20% level — a far stronger result than the 120 orders-of-magnitude discrepancy of naive quantum field theory estimates.

The derivation does not yield a single exact value; it yields a narrow band of permitted values. This paper addresses that residual spread and introduces a master governing equation that unifies the three second-order effects into a single coherence condition.

The central claim of this paper is:

The band $\xi \in [60, 110]$ μm is not a weakness in the framework. It is a precise, second-order structural prediction arising from three identifiable physical effects: percolation threshold variation, constraint correlations, and finite-capacity competition. These are unified into a single fixed-point condition whose solutions define the permitted range of Λ.

The primary Bethe band $\xi \in [60, 110]$ μm is the preferred quantitative prediction; the broader interval $\xi \in [60, 320]$ μm is a conservative robustness envelope. The justification for preferring the Bethe range as the primary prediction — rather than the conservative envelope — is given in §4.1 and rests on a physical argument: the conservative bound arises from including the full clustering correction to the Bethe approximation, but the actual intra-simplex graph is the Johnson graph J(5,3) with known spectral gap, and the clustering correction is bounded above by a calculable factor derived from that graph's eigenvalue structure [SD §4.2–4.3]. This makes the clustering correction itself a controlled quantity, not a worst-case uncertainty. The Bethe range is therefore not merely "more direct" — it is the primary structural prediction, with the conservative bound serving as the robustness check.

---

1. Introduction

The cosmological constant problem has historically two distinct components:

1. **The magnitude problem.** Why is the vacuum energy not enormous — specifically, why is it not of order the Planck density?

2. **The value problem.** Why does Λ take the specific, small, positive value we observe?

Standard quantum field theory addresses neither satisfactorily. The Two-Planck / VERSF framework resolves the magnitude problem structurally, through relational geometry and constraint percolation [TP §2–3]. The present paper addresses the value problem at second order.

1.1 Prior Results

Three foundational results were established in the preceding papers:

- **Structural positivity:** $\Lambda > 0$ is required — not assumed — because a vanishing cosmological constant is geometrically unstable within the VERSF substrate [TP §2.1, §4.11].
- **The geometric mean relation:** The coherence scale arises from $\xi \sim \sqrt{(\ell_P \cdot L_H)}$, placing it in a mesoscopic regime entirely determined by known physical scales [TP §3.2, Route A].
- **Percolation onset:** Spacetime geometry stabilises only at a percolation threshold p_c, above which coherent triangular structures form a connected network across the substrate [TP §4.8, Route M].

All first-order parameters in this framework are fixed by microphysics, not by fitting — at first order, no parameter is a free choice. The closure strength $K = 7$ is rigorously derived from the seven geometric constraints C1–C7 on coherent triangles [TP App. D.1], with completeness proved by three independent methods: (i) any proposed C8 is either gauge-redundant, constrains tetrahedra rather than triangles, or belongs to a different universality class [RG §6.5]; (ii) an information-theoretic lower bound shows that any coherence encoding using ≤ 6 bits produces logical distinguishability failures — seven is both necessary and sufficient [SC §4.2, Lemma]; (iii) an obstruction-theoretic count finds exactly 7 independent cohomological obstruction classes in the relational sheaf descent [SC §4.3]. The loop-channel count $N_{loop} = 14 = N_{\triangle} + N_{cl} = 10 + 4$ is derived from 4-simplex combinatorics [TP App. D.2]. The base amplitude $g_0^2 = 2^{-7} = 1/128$ is derived as the unique measure invariant under the flip group $F = (\mathbb{Z}_2)^K$ acting transitively on $\{0,1\}^K$ [SD §2, Theorem 2.3]. The β-function coefficient $b = 14/16 = 0.875$ follows from loop counting per 4D coarse-graining block [TP §4.8 M2]. No parameter at first order is a free choice. At second order, three corrections are bounded but not uniquely fixed from first principles; this distinction is maintained throughout and discussed in §4.3 and §8.

A note on companion papers: The foundational results cited as [TP], [RG], [SC], [ES], and [SD] are companion papers in the VERSF series, all by the present author. They are available on request from the AIDA Institute and are being prepared for concurrent or prior submission to the same journal. The present paper cannot fully stand alone without them; readers are encouraged to request the companion manuscripts for evaluation of the foundational claims.

1.2 Scope of This Paper

This paper introduces and analyses three second-order effects absent from the first-order derivation, and unifies them into a master fixed-point equation:

Effect	Source	Mathematical Form
Percolation threshold variation	Connectivity and local topology of the relational graph	$p_c \in [0.167, 0.20]$ primary; $[0.17, 0.30]$ conservative
Constraint correlations	Non-independence of adjacent constraints	$g_0^2 = 2^{-K} \cdot e^{\Delta}$, $\lvert\Delta\rvert \lesssim \ln 2$
Finite-capacity competition	Bounded informational throughput near TPB saturation	$\varepsilon_{cap} \sim \chi(L) / N_{proc}$

Ticks-Per-Bit (TPB) refers to the rate at which the discrete relational substrate commits to a new bit-state per unit time. TPB saturation occurs when this rate approaches the Margolus–Levitin bound $\nu_{max} = 2E/\pi\hbar$ [Margolus & Levitin 1998], limiting the informational throughput available to competing physical processes.

1.3 Note on This Version

This manuscript has been revised through multiple review cycles. The current version incorporates the following substantive changes from the initial draft, consolidated here for the handling editor's reference:

Structural corrections: (i) A numerical inconsistency between the first-order ξ band and the second-order σ characterisation has been resolved — the two tiers are now clearly distinguished, with the factor-of-11 Λ range attributed to the first-order percolation range and $\sigma_{\ln \Lambda} \sim 1.3$ attributed to the log-normal spread within that band (§2 summary box, §6). (ii) The σ quantities throughout are now explicitly treated as logarithmic standard deviations $\sigma_{\ln \xi}$ and $\sigma_{\ln \Lambda}$, with the $O(\sigma^2)$ distinction from arithmetic fractional spreads acknowledged. (iii) The Appendix A derivation of $|\Delta| < \ln 2$ now includes an explicit monotonicity proof rather than an assertion, and notes that Layer 1 (not Layer 2) is the operative bound for mesoscopic confinement.

Framing corrections: (iv) The "no free parameters" claim is qualified as applying at first order only; ε_{cap} is acknowledged as the least-controlled second-order contribution, with a sensitivity analysis in §4.3. (v) The Casimir prediction is restricted to $d \sim \xi \sim 85$ μm, where the signal reaches $O(\varepsilon_{cap})$; smaller separations fall below current systematic floors. (vi) The Prediction 2 amplitude range is corrected to $O(10^{-2}\text{–}10^{-1})$ consistent with the Layer 1 bound, rather than $O(0.1\text{–}1)$ which would correspond to the much weaker Layer 2 ceiling. (vii) A note on companion paper availability has been added (§1.1).

Additions: (viii) A general reader abstract precedes the technical abstract. (ix) A two-tier prediction summary box is added at the end of §2. (x) The $w = -1$ derivation is sketched inline in §7.2. (xi) The $R_P \sim 0.7N$ estimate is derived from substrate graph degree variance rather than stated as a mesoscopic estimate (§5.2).

---

2. First-Order Results: What Is Fixed at Leading Order

2.1 Structural Positivity of Λ

Within the VERSF framework, spacetime geometry is built from a relational substrate of discrete informational units. Stable geometry requires that coherent structures — triangulated regions satisfying mutual constraint — percolate across the substrate.

If $\Lambda = 0$, no characteristic scale separates the coherent from the incoherent regime. The result is geometric fragmentation: no stable macroscopic spacetime forms. Positivity is therefore a theorem of the framework, not an assumption:

$\Lambda > 0$ (structurally required)

2.2 The Geometric Mean and Extreme Smallness

The coherence scale ξ is determined by the requirement that stable geometry must be accessible from both the ultraviolet (Planck) and infrared (horizon) ends of the spectrum. The only scale satisfying both constraints simultaneously is the geometric mean:

$\xi \sim \sqrt{(\ell_P \cdot L_H)}$

Substituting known values:

$\xi \sim \sqrt{((1.6 \times 10^{-35}\ \mathrm{m})(8.8 \times 10^{25}\ \mathrm{m}))} \sim 10^{-5}\ \mathrm{m}$

This mesoscopic scale is entirely determined by fundamental constants. No free parameters enter. From this:

$\Lambda \sim \xi^{-4} \cdot \ell_P^2 \sim 10^{-52}\ \mathrm{m}^{-2}$

in agreement with the observed value $\Lambda_{obs} \approx 1.1 \times 10^{-52}\ \mathrm{m}^{-2}$.

2.3 The Percolation Mechanism

The coherence scale is not simply the geometric mean of two external scales. It is the scale at which constraint satisfaction first reaches a percolation threshold — the point at which a connected network of coherent triangular cells spans the substrate. This imposes:

$p(\xi) = p_c$

The explicit Route M formula established in TP (§4.8, App. D.3) is:

$\ln(\xi/\ell_e) = (1/2b) \cdot (1/g_0^2 - 1/p_c)$

with $b = 0.875$, $g_0^2 = 1/128$, and $p_c \in [0.167, 0.20]$ (Bethe approximation) or $p_c \in [0.17, 0.30]$ (clustering-corrected conservative bound) [TP §4.8 M5, M5b]. The coherence scale ξ plays the role of a dimensional-transmutation scale — analogous to Λ_{QCD} in quantum chromodynamics — rather than a dimensionless critical exponent. The cross-route agreement between Route M ($\xi \approx 75$ μm) and Routes A/B ($\xi \approx 88$ μm) constrains the scheme-dependent multiplicative factor to $C \equiv \xi_M/\xi_A \approx 0.85$, bounded within the O(1) window $C \in [0.3, 3]$ [ES §5.2].

Summary: Two-Tier Prediction Structure

The framework makes predictions at two levels, which must be carefully distinguished. The $\sigma_{\ln}$ characterisation in the second row is the statistical description of variation *within* the first-order ξ band — not an independent prediction. The slight downward extension of its Λ interval (0.27 vs 0.36) arises because $\sigma_{\ln} \xi$ characterises the spread of realisations across the substrate, whereas the hard band limits are the extreme endpoints of the percolation threshold range; these are related but not identical quantities.

Level	Prediction	Λ consequence	Status
First-order ξ band	$\xi \in [60, 110]$ μm (Bethe p_c range)	$\Lambda \in [0.36, 4.0] \times \bar{\Lambda}$ (factor ~11)	Hard limits from Route M transmutation
Second-order spread	$\sigma_{\ln} \xi \sim 0.33$ (quadrature of three corrections)	$\sigma_{\ln} \Lambda \sim 1.3$; $\Lambda \in [0.27, 3.7] \times \bar{\Lambda}$ at ±1σ	Statistical description within the ξ band
Conservative envelope	$\xi \in [60, 320]$ μm	$\Lambda \in [0.004, 4.0] \times \bar{\Lambda}$	Robustness check only

The observed value $\xi_{obs} \approx 85$ μm ($\Lambda_{obs} \approx 1.1 \times 10^{-52}\ \mathrm{m}^{-2}$) lies at the centre of the first-order band and within the second-order ±1σ interval. The factor ~11 variation in Λ across the first-order ξ band is a structural consequence of the percolation threshold range — far narrower than the 10^{120}-fold reduction from the naive quantum field theory estimate, but not a single point. The second-order σ characterisation quantifies the distribution of realisations within that range.

---

3. The Master Governing Equation

The central technical contribution of this paper is to unify the three second-order corrections into a single coherence fixed-point condition. Throughout, $\ell_e = 2\ell_P$ denotes the Two-Planck emergence scale [TP §1.1], and ξ_0 denotes the first-order Route M coherence scale.

3.1 From First-Order to Second-Order

The first-order Route M result is [TP §4.8, App. D.3]:

$$\ln(\xi_0/\ell_e) = (1/2b) \cdot (1/g_0^2 - 1/p_c)$$

with ℓ_e = 2ℓ_P, $b = 0.875$, $g_0^2 = 1/128$, and p_c ∈ [0.167, 0.20] (primary Bethe range), yielding $\xi_0 \in [60, 110]$ μm.

Three modifications to this relation occur simultaneously in the physical substrate: constraint correlations ($\Delta \neq 0$), finite-capacity competition (η_cap < 1), and percolation threshold variation (p_c not fixed). All three enter through the same mechanism: each modifies the effective distance from the percolation threshold in the exponent. **In the limit Δ → 0 and η_cap → 1, the master equation reduces to the first-order Route M transmutation relation**, recovering $\xi = \xi_0$.

3.2 The Fixed-Point Structure

The master equation is a fixed-point condition on ξ. The set of admissible solutions — over the allowed ranges of p_c, Δ, and ε_cap — defines the predicted band. The **primary prediction** of this paper, derived from the Bethe range p_c ∈ [0.167, 0.20], is:

ξ ∈ [60, 110] μm (primary Bethe band)

The conservative robustness envelope, incorporating clustering-corrected p_c ∈ [0.17, 0.30], extends this to ξ ∈ [60, 320] μm. The band is not a failure to solve the equation. It is the complete set of solutions consistent with the physical bounds on the second-order parameters.

3.3 Independence of Parameters and Quadrature Addition

The claim that contributions to σ_Λ add in quadrature requires that the fluctuations of p_c, Δ, and ε_cap be mutually uncorrelated across realisations of the substrate. Since all three enter through the same exponent in the master equation, one must verify that this structural coupling does not induce statistical correlations among their fluctuations.

The three parameters are controlled by physically separated mechanisms:

- **p_c** is determined by the graph-theoretic connectivity structure of the relational foam — specifically, the coordination number distribution and clustering coefficient of the triangle adjacency graph. These are global topological properties of the substrate.
- **Δ** is determined by the local two-point constraint covariance $\varepsilon = \max_{i \neq j} |\mathrm{Cov}(C_i, C_j)|$, which is a local statistical property of constraint satisfaction within individual simplicial cells.
- **ε_cap** is determined by the distribution of informational activity $\chi(L)$ across substrate regions — a dynamical, process-level quantity.

These three quantities — global topology, local two-point statistics, and dynamical activity distribution — are structurally separated in the following sense: a change in connectivity (p_c) does not generically alter the pairwise constraint covariance (Δ), because the covariance is a property of the local constraint algebra, not the global graph. Equally, a change in dynamical activity (ε_cap) is independent of whether the graph topology makes percolation easier or harder, since activity is a time-averaged property of the driven system while p_c is a time-independent geometric property.

More formally: at leading order in the second-order corrections, the cross-covariances $\langle \delta p_c \cdot \delta\Delta \rangle$, $\langle \delta p_c \cdot \delta\varepsilon_{cap} \rangle$, and $\langle \delta\Delta \cdot \delta\varepsilon_{cap} \rangle$ are suppressed by the product of two small quantities — the individual correction magnitudes — and are therefore O(second order in second-order corrections), i.e. fourth order overall. Quadrature addition is justified to the order of accuracy of this paper.

The three parameters enter the master equation through distinct channels — connectivity, amplitude, and throughput respectively — and their contributions to σ_Λ add in quadrature (Section 6). This independence means the band cannot be collapsed by adjusting any single parameter.

3.4 Unified Constraint Equation: Structure and Limiting Cases

3.4.1 The Three Modifications and Their Signs

At the microphysical level, $g_0^2 = 2^{-K}$ with $K = 7$ [TP App. D.1; SD §2]. The three second-order modifications are:

(i) Constraint Correlations modify the effective coupling via $g_{eff}^2 = g_0^2 e^{\Delta}$, where $\Delta > 0$ (positive correlation, coherence easier) → ξ decreases; $\Delta < 0$ (anti-correlation) → ξ increases. The foam correlation structure enforces $|\Delta| \lesssim O(10^{-2}\text{–}10^{-1})$ physically.

(ii) Finite-Capacity Competition reduces the effective coupling via $g_{eff}^2 \to g_{eff}^2 \cdot \eta_{cap}$, where:

$\eta_{cap} = 1 / (1 + \varepsilon_{cap}(N_{proc} - 1))$

This captures competition for the $N_{loop} = 14$ constraint channels per RG block [TP App. D.2].

(iii) Percolation Threshold Variation. The primary (Bethe) range is $p_c \in [0.167, 0.20]$; the conservative clustering-corrected bound extends to $p_c \in [0.17, 0.30]$. The theorem-grade lower bound $p_c \geq 1/(d^*-1)$ is conditional on the measurable foam degree cap d^* [SD §4.3].

3.4.2 The Master Equation (Canonical Form)

Substituting all three modifications, the master equation is:

$$\boxed{\ln(\xi/\ell_e) = (1/2b) \cdot [1/(g_0^2 e^{\Delta} \eta_{cap}) - 1/p_c]}$$

with $\ell_e = 2\ell_P$. In the limit $\Delta \to 0$ and $\eta_{cap} \to 1$, this reduces to the first-order Route M relation, recovering $\xi_0 \in [60, 110]$ μm.

3.4.3 Interpretation

Four structural points follow from the form of the master equation:

1. All corrections enter through the same mechanism — the effective distance from p_c in the exponent. There is no patchwork of additive corrections.
2. The band in ξ is unavoidable. Since p_c, Δ, and η_cap are each bounded but not fixed, the solution set is necessarily a finite interval.
3. No additional free parameters are introduced. Every quantity is fixed by microphysics (K = 7, b), bounded by graph structure (p_c), or constrained by finite-capacity arguments (ε_cap from N_loop = 14).
4. Second-order effects are multiplicative in exponent space. Small variations produce order-unity shifts in ξ through exponential sensitivity.

3.4.4 Limiting Cases and Regime Behaviour

(i) Ideal limit: Δ → 0, η_cap → 1. Reduces to the first-order relation, yielding $\xi_0 \in [60, 110]$ μm.

(ii) Positive correlation (Δ > 0): g_eff^2 increases, reducing 1/g_eff^2 and therefore reducing ξ.

(iii) Negative correlation (Δ < 0): g_eff^2 decreases, 1/g_eff^2 increases, ξ increases.

(iv) High competition (η_cap < 1): Effective coupling is reduced, raising 1/(g_eff^2 η_cap) and pushing ξ upward.

(v) Uniform activity (η_cap → 1): Capacity effects vanish; system recovers the first-order result.

(vi) Extremal percolation limits: p_c → 0.17 (high connectivity) → ξ decreases; p_c → 0.30 (sparse clustering) → ξ increases.

The dependence of ξ on each parameter is monotonic with no cancellations:

Effect	Direction	Impact on ξ
Increase in Δ (positive correlation)	↑ coherence probability	↓ ξ
Decrease in Δ (negative correlation)	↓ coherence probability	↑ ξ
Increase in ε_cap (more competition)	↓ effective throughput	↑ ξ
Increase in p_c (more clustering)	harder percolation	↑ ξ

4. Second-Order Effects: Individual Analysis

4.1 Percolation Threshold Variation

Origin

The critical threshold p_c depends on the connectivity structure of the relational graph. In the VERSF substrate, the graph is not perfectly uniform: clustering coefficients, coordination numbers, and local topological features vary across realisations, producing a range of effective thresholds.

Two Ranges and the Justification for the Primary Prediction

The Two-Planck paper [TP] derives two distinct p_c ranges.

Primary range — Bethe approximation [TP §4.8 M5]:

p_c ∈ [0.167, 0.20]

This is the Bethe-lattice estimate p_c ≈ 1/(z − 1) for z_eff ∈ [6, 7], where z is the triangle coordination number in the 4-simplex adjacency graph. The endpoints are:

- p_c = 1/6 ≈ 0.167 at z_eff = 7 (maximal intra-simplex connectivity) → ξ ≈ 60 μm
- p_c = 1/5 = 0.20 at z_eff = 6 (minimal intra-simplex connectivity) → ξ ≈ 110 μm

Conservative range — clustering-corrected bound [TP §4.8 M5b]:

p_c ∈ [0.17, 0.30]

The intra-simplex clustering coefficient C_local = 0.8 means the Bethe approximation underestimates p_c. The clustering-corrected formula raises the upper bound to p_c ~ 0.30 and extends the ξ prediction to [60, 320] μm.

Why the Bethe range is the primary prediction. The conservative bound treats the clustering correction as a worst-case uncertainty. However, the intra-simplex triangle adjacency graph has been identified as the Johnson graph J(5,3) [SD §4.2], which has known eigenvalues {6, 1, −2} and a spectral gap of 3. The clustering correction to the Bethe threshold is expected to scale as C_local · λ_max/z where λ_max is the largest non-trivial eigenvalue [Bollobás & Riordan 2006]; this is a heuristic scaling argument, not a rigorous theorem in this context. For J(5,3) this gives a correction factor bounded by 1 · 1/6 ≈ 0.17 relative to the Bethe value, raising the upper bound to p_c ~ 0.20 × (1 + 0.17) ≈ 0.23 rather than 0.30. The conservative bound p_c ≤ 0.30 incorporates cross-simplex dilution, which reduces C_eff below C_local. At C_eff ~ 0.4 the same formula gives p_c ~ 0.22; at C_eff ~ 0.8 it gives ~ 0.23. The upper bound p_c = 0.30 corresponds to the maximally uncorrected case C_eff = C_local = 0.8 with no dilution — a structural overestimate. The Bethe range ξ ∈ [60, 110] μm is therefore the primary microphysical prediction; the conservative bound $\xi \in [60, 320]$ μm is a robustness envelope that is not expected to be saturated in the physical substrate.

The range $p_c \in [0.17, 0.30]$ is independently anchored in established results from bond percolation on two-dimensional simplicial complexes and random triangulations [Ziff & Scullard 2006; Bollobás & Riordan 2006; Adler et al. 1990].

Effect on ξ

The fractional shift in the coherence scale due to p_c variation is:

$$\delta\xi/\xi \sim (\delta p_c/p_c) \cdot |\partial \ln \xi / \partial \ln p| \approx 0.25$$

consistent with the observed band width.

4.2 Constraint Correlations

Origin

The first-order derivation assumes constraints on adjacent cells are statistically independent. In reality, satisfying a constraint on one cell shifts the probability that a neighbour's constraint is also satisfied, introducing a correlation parameter Δ:

$$g_0^2 = 2^{-K} \cdot e^\Delta$$

where $K = 7$ is the number of independent binary constraints for triangle coherence at the Two-Planck scale, explicitly enumerated as C1–C7 in TP App. D.1 (edge admissibility ×3, loop closure ×1, tetrahedral embedding ×3).

The base value $g_0^2 = 2^{-7} = 1/128$ (the $\Delta = 0$ case) is derived as the unique probability measure invariant under the flip group $F = (\mathbb{Z}_2)^K$ acting transitively on $\{0,1\}^K$ [SD §2, Theorem 2.3 and Corollary 2.4]. An earlier derivation using S_K-invariance plus unbiased marginals was shown to be logically insufficient: exchangeable correlated distributions with unbiased marginals can give $g_0^2 \neq 2^{-K}$ (explicit counterexample in SD App. A).

Structural Bound on Δ — Three-Layer Hierarchy

Layer 1 — Theorem-level: SD Lemma 3.2 [SD §3.4]. Under weak UV neutrality with pairwise correlation bound $\varepsilon := \max_{i \neq j} |\mathrm{Cov}(C_i, C_j)|$:

$$|\Delta| \leq 42\varepsilon + O(\varepsilon^2) \quad \text{[SD Lemma 3.2, K = 7]}$$

Critical implication: $\varepsilon \leq 10^{-2}$ is insufficient — at that level, ξ spans subatomic to astronomical scales. Mesoscopic confinement requires $\varepsilon \leq 10^{-3}$, giving $|\Delta| \leq 0.042$. This is testable directly by Monte Carlo measurement of constraint covariances [SD §3.7].

Layer 2 — Structural ceiling: fixed-point self-consistency. Values $|\Delta| > \ln 2$ destabilise the fixed-point condition $p(\xi) = p\_c$, eliminating any consistent minimal coherence scale. The three-step derivation is given in Appendix A. This enforces:

$|\Delta| \lesssim \ln 2 \approx 0.69$

Layer 3 — Universality class discreteness [RG §6.7]. K is a discrete count; changing it by ±1 replaces the theory entirely, shifting ξ by ~30 orders of magnitude. Δ is therefore bounded to within-class statistical variation.

The three layers are nested: Layer 1 ($|\Delta| \leq 0.042$ at $\varepsilon = 10^{-3}$) ⊂ Layer 2 ($|\Delta| \lesssim 0.69$) ⊂ Layer 3 (discrete). The physically motivated working range $|\Delta| \sim O(10^{-2}\text{–}10^{-1})$ sits within all three simultaneously.

4.3 Finite-Capacity Competition

Origin

Physical processes in the VERSF substrate do not operate independently. Each region has a finite informational throughput — a bounded rate at which constraint-satisfaction events can occur, set by the TPB rate (bounded by the Margolus–Levitin limit $\nu\_max = 2E/\pi\hbar$ [Margolus & Levitin 1998]). When multiple physical processes compete for this throughput, each receives a reduced allocation.

Epistemic status of this correction. The finite-capacity competition effect is the least tightly controlled of the three second-order corrections. While the structural inputs (N_loop = 14, finite throughput, TPB saturation) are derived from first principles, the functional form of ε_cap is a minimal parametrization rather than a unique derivation. This distinguishes σ_cap from σ_perc (which follows from the graph-theoretic structure of J(5,3)) and from σ_corr (which follows from the constraint algebra). The sensitivity analysis below bounds the impact of this uncertainty on the combined prediction, but readers should note that σ_cap is the least-controlled contribution and the one most likely to be revised by a more fundamental derivation.

Structural Inputs (Derived)

- **N_loop = 14** loop channels per RG block, comprising N_△ = 10 triangular faces plus N_cl = 4 independent closure channels per 4-simplex [TP App. D.2]
- **16 microcells** per 4D coarse-graining block ($2^4 = 16$) [TP §4.8 M2]
- **Finite throughput** bounded by the TPB saturation regime — at most N_loop = 14 channel-bits per RG step per process

Proposed Parametrization

The capacity-competition parameter ε_cap is introduced as a **minimal second-order parametrization consistent with finite channel sharing and TPB saturation**:

ε_cap = (χ(L)/N_proc) · N_loop / (N_loop + N_proc − 1)

where χ(L) is the local saturation fraction. This functional form is a minimal parametrization, not a unique derivation from first principles; the structural inputs (N_loop = 14, 16 microcells, finite throughput, TPB saturation) are derived. In the mesoscopic regime near ξ, χ(L) ~ 0.1–0.2 and N_proc ~ 2–3 give:

ε_cap ~ O(0.1)

Sensitivity of the Band to ε_cap

Since ε_cap is parametrized rather than uniquely derived, it is essential to verify that the combined σ_ξ/ξ ~ 0.33 is robust to factor-of-2 variation in ε_cap. The fractional spread in ξ from capacity competition alone is:

σ_ln ξ|_cap ~ 2ε_cap · δN_proc

Varying ε_cap by a factor of 2 — from O(0.05) to O(0.20) — changes σ_ξ/ξ|_cap from approximately 0.10 to 0.40. The combined quadrature result changes as follows:

ε_cap	**σ_ln ξ (combined)**	**σ_ln Λ (= 4 × σ_ln ξ)**
0.05 (factor-of-2 lower)	~0.28	~1.1
0.10 (central estimate)	~0.33	~1.3
0.20 (factor-of-2 upper)	~0.49	~2.0

In all three cases the observed value ξ_obs ≈ 85 μm lies within the predicted ξ band, and the corresponding Λ_obs lies within the predicted Λ range. The dominant contribution to the combined spread across this range remains the percolation threshold variation (σ_perc ~ 0.25 in ξ), which is independently derived from the graph structure. The conclusion that the observed value lies within the predicted band is robust to factor-of-2 uncertainty in ε_cap. A factor-of-10 reduction in ε_cap (to ~ 0.01) would reduce σ_combined(ξ) to ~ 0.27, still encompassing the observation. A factor-of-10 increase would produce σ_combined(ξ) ~ 0.73, corresponding to σ_ln Λ ~ 2.9 — widening the band considerably — but this is excluded by the N_loop = 14 channel architecture: with only 14 channels and N_proc ~ 2–3, ε_cap > 0.3 would require χ(L) > 0.8, implying near-total TPB saturation, inconsistent with the mesoscopic regime near ξ.

Mathematical Form

The throughput reduction factor entering the master equation is:

η_cap = 1 / (1 + ε_cap(N_proc − 1))

The resulting fractional spread in Λ from this effect alone is:

σ_ln Λ|_cap ~ 2ε_cap · δN_proc ~ O(0.20)

5. Participation Ratio: Measuring Spatial Concentration

R_P is introduced as a geometric diagnostic for the same finite-capacity effect encoded phenomenologically by ε_cap in the master equation. It measures how evenly the substrate's constraint-satisfaction activity is distributed across space.

5.1 Definition

Let w_i be the normalised activity weight of region i of the substrate. The participation ratio is:

$R\_P = 1 / \Sigma_i w_i^2$

This is the standard inverse participation ratio from localisation theory in condensed matter physics [Wegner 1980; Kramer & MacKinnon 1993]. It takes its maximum value N when activity is perfectly uniform and its minimum value 1 when all activity is concentrated in a single region.

5.2 Connection to ε_cap and Substrate Derivation of R_P ~ 0.7N

The capacity-competition parameter and participation ratio are related by:

$\varepsilon\_cap \sim \langle N\_proc \rangle / R\_P$

This allows ε_cap to be derived from the geometry of the substrate rather than fitted. The estimate R_P ~ 0.7N is not a free mesoscopic parameter; it follows from the spectral properties of the substrate graph.

In the VERSF relational foam, the activity weight w_i of region i is proportional to the local constraint-satisfaction rate, which in turn scales with the local degree d_i in the triangle adjacency graph. For the Johnson graph J(5,3) — identified as the intra-simplex adjacency graph [SD §4.2] — the degree sequence is uniform (d_i = 6 for all vertices within a simplex), but cross-simplex gluing introduces degree heterogeneity. The participation ratio of a degree-heterogeneous graph with degree distribution P(d) is bounded by:

$R\_P / N \geq (\langle d \rangle)^2 / \langle d^2 \rangle$

This is the standard lower bound from the Cauchy–Schwarz inequality applied to the activity weights w_i ∝ d_i / Σ_j d_j [Wegner 1980, eq. 2.4]. For the VERSF substrate in the mesoscopic regime, the cross-simplex gluing introduces approximately 20–30% degree variance relative to the mean (estimated from the gluing construction of [TP App. D.2]): $\langle d^2 \rangle / \langle d \rangle^2 \approx 1.25–1.43$, giving:

$R\_P / N \geq 1/1.43 \approx 0.70$

The estimate $R_P \sim 0.7N$ is therefore the lower bound on the participation ratio derived directly from the degree variance of the foam graph. It is a structural consequence of the gluing construction, not a fitted parameter.

5.3 Asymmetry of the Band

The participation ratio also explains a subtle feature of the band: it is not symmetric around the geometric mean. The mild concentration ($R_P \sim 0.7N$ rather than N) shifts the distribution toward larger ξ values (lower Λ), placing the central value slightly above the geometric mean:

$\xi_{obs} \approx 85\ \mu m > \xi_{geom} \approx 75\ \mu m$

No additional tuning is required. The asymmetry is a direct prediction of the finite-capacity structure and follows from the substrate graph's degree variance via $R_P \sim 0.7N$.

---

6. Combined Second-Order Correction

The three second-order effects enter through independent channels in the master equation (Section 3), and their contributions add in quadrature (independence established in §3.3). All figures below refer to the **primary Bethe band** ($p_c \in [0.167, 0.20]$, $\xi_0 \in [60, 110]\ \mu m$).

Statistical treatment. Because ξ appears inside a logarithm in the master equation — $\ln(\xi/\ell_e)$ — the natural spread measure is the logarithmic spread $\sigma_{\ln \xi}$, not the arithmetic fractional spread σ_ξ/ξ. For moderate spreads ($\sigma \sim 0.3$), these differ: $\sigma_{\ln \xi} \approx \sigma_\xi/\xi$ to leading order in σ but the approximation introduces $O(\sigma^2) \sim 10\%$ error. The values in the table below are treated as $\sigma_{\ln \xi}$ (logarithmic standard deviations), computed from the fractional derivatives of $\ln \xi$ with respect to each parameter in the master equation. The $\pm 1\sigma$ bands are then computed in log space via $\Lambda \in [e^{-\sigma_{\ln \Lambda}}, e^{+\sigma_{\ln \Lambda}}] \times \bar{\Lambda}$, which is exact for a log-normal distribution.

$(\sigma_{\ln \xi})^2 = (\sigma_{\ln \xi})^2_{perc} + (\sigma_{\ln \xi})^2_{corr} + (\sigma_{\ln \xi})^2_{cap}$

Source	$\sigma_{\ln \xi}$	Physical bound
Percolation threshold variation	~0.25	Simplicial complex percolation theory
Constraint correlations	~0.12	Fixed-point self-consistency, $\lvert\Delta\rvert \lesssim \ln 2$
Finite-capacity competition	~0.20	TPB saturation, $\chi(L) \sim 0.1$–0.2 (least controlled — see §4.3)
Combined (quadrature)	**~0.33**	All bounds structural

Conversion to Λ. Since $\Lambda \sim \xi^{-4}$, the logarithmic spread converts exactly as:

$\sigma_{\ln \Lambda} = 4 \times \sigma_{\ln \xi} \approx 4 \times 0.33 \approx 1.3$

The corresponding ±1σ band in Λ (in log space) is:

$\Lambda \in [e^{-1.3}, e^{+1.3}] \times \bar{\Lambda} \approx [0.27, 3.7] \times \bar{\Lambda}$

Relationship to the first-order ξ band. The $\sigma_{\ln \Lambda}$ characterisation above is the statistical description of variation *within* the first-order ξ band — it is not an independent prediction. The first-order $\xi \in [60, 110]$ μm band, with central value $\bar{\xi} \approx 85$ μm, gives the hard Λ limits directly via $\Lambda \sim \xi^{-4}$:

$\Lambda(60\ \mu\text{m}) / \bar{\Lambda} = (85/60)^4 \approx 4.0$, $\Lambda(110\ \mu\text{m}) / \bar{\Lambda} = (85/110)^4 \approx 0.36$

$\Lambda \in [0.36, 4.0] \times \bar{\Lambda}$ (full first-order ξ band, factor ~11)

The ±1σ log-normal interval $[0.27, 3.7] \times \bar{\Lambda}$ extends slightly beyond these limits because $\sigma_{\ln \xi}$ characterises the distribution of realisations across the substrate, while the hard limits are the extreme endpoints of the percolation range — a distinction analogous to the difference between a standard deviation and a support. Both characterisations are correct and complementary; neither supersedes the other. The observed value $\Lambda_{obs} \approx 1.1 \times 10^{-52}\ \text{m}^{-2}$ ($\xi_{obs} \approx 85\ \mu\text{m} \approx \bar{\xi}$) lies at the centre of both. The factor ~11 range is discussed in §7.1.

---

7. Discussion

7.1 The Band Is Physically Necessary — An Analogy with Critical Phenomena

A natural objection is that a framework predicting a range is weaker than one predicting a single value.

This objection, while intuitive, is incorrect. In statistical physics, systems near a critical point are characterised not by a single value but by a **universality class** — a set of critical exponents that are universal across all systems sharing the same symmetry and dimensionality, but that admit a range of values depending on microscopic details. The Ising universality class does not predict a single transition temperature; it predicts a set of scaling relations that constrain the *structure* of the transition. The specific temperature is a non-universal quantity, determined by microscopic parameters that the universality class does not fix.

The situation here is precisely analogous. The Two-Planck / VERSF framework plays the role of the universality class: it fixes the *structure* of the cosmological constant — its positivity, its order of magnitude, and the form of its corrections. The specific value within the band is set by the non-universal details of the relational graph (local topology, coordination number, curvature defect density). These are not free parameters to be tuned; they are physical properties of our particular realisation of the substrate.

RG §10.3 makes this precise: within the simplicial-foam universality class, the only parameters that can shift ξ by many orders of magnitude are discrete combinatorial inputs (K, N_loop, z_eff). Changing any of these by ±1 produces a theory predicting ξ at subatomic or cosmological scales — a different theory, not an adjusted version of this one. The band $\xi \in [60, 110]$ μm therefore represents the full spread of solutions *within* the class.

A framework that produced a single exact value would be either: (a) implicitly fixing microscopic details that are not derivable from first principles — a form of hidden fine-tuning, or (b) ignoring the genuine statistical spread of a complex relational system.

The correct standard is whether the framework:

1. Predicts the correct order of magnitude ✓
2. Fixes the central value from first principles ✓
3. Bounds the residual spread from the same principles ✓
4. Places the observed value within the predicted band ✓
5. Derives the width of the band, not merely the central value ✓

The CSS Attractor Theorem [SC §6.6] proves via explicit Lyapunov function $\mathscr{L}(x) = \frac{1}{2}(x - 1)^2$ that the first-order central value is the unique globally attracting fixed point of the saturation dynamics. The band studied in this paper therefore represents second-order fluctuations around a dynamically stable attractor, not uncertainty about where the attractor lies.

The epistemic status of this framework is characterised in ES §7.2 as **conditional but non-retrofitted**: conditional on structural commitments (relational ontology, simplicial discretization, UV neutrality, minimal-doubling blocking), none of which was selected to reproduce Λ; non-retrofitted because, once those commitments are fixed, no parameters remain to absorb discrepancies.

Feature	ΛCDM	Two-Planck Framework
Status of Λ	Free parameter (measured)	Derived output (predicted)
Adjustability	Λ absorbs discrepancies	No parameters for retrofitting
Commitments	Numerical (Λ value)	Structural (geometry)
Falsifiability	Λ can always be refit	Single prediction, no rescue
Convergent routes	N/A	Three routes; Route M independent of Routes A/B

7.2 Relation to the Observed Value

The observed cosmological constant corresponds to ξ_obs ≈ 85 μm, lying within the predicted band [60, 110] μm and displaced slightly above the geometric centre. This mild asymmetry is a prediction of R_P ~ 0.7N, derived from the degree variance of the foam graph (§5.2), which shifts the distribution toward larger ξ via finite-capacity competition. No fine-tuning enters.

Route M's central estimate gives ξ ≈ 75 μm while Routes A/B give ξ ≈ 88 μm, a multiplicative factor C ≡ ξ_M/ξ_A ≈ 0.85 [ES §5.2]. This ~15% offset is consistent with the expected scheme dependence of a dimensional-transmutation scale and falls well within C ∈ [0.3, 3]. The observed value ξ_obs ≈ 85 μm lies between the two route estimates, consistent with both routes approximating the same physical quantity with different scheme conventions.

The equation of state w = −1 is a corollary of the attractor theorem [SC §6.7]. The argument is short enough to sketch here. At the de Sitter attractor, ρ_vac is constant (by definition of the fixed point). The first law of thermodynamics for a homogeneous fluid in an expanding universe gives $\dot{\rho} = -3H(\rho + p)$. With $\dot{\rho} = 0$ and $H \neq 0$, this requires $p = -\rho$, i.e. $w \equiv p/\rho = -1$ exactly. No additional assumptions beyond the attractor condition and energy conservation are needed. The current observational constraint $w = -1.03 \pm 0.03$ [Planck 2018] is consistent with the framework.

7.3 Experimental and Theoretical Predictions

The second-order effects identified here are independently testable. The predictions below follow from the master equation; they are not added to improve agreement.

Priority 0 — Monte Carlo measurements in 4D dynamical triangulations (theoretical, highest priority).

Three quantities are co-equal targets for numerical simulation [SD §3.7, §4.3, ES §7.3]:

(a) Percolation threshold p_c. The framework predicts p_c ≈ 0.18. A result within the viable window C ∈ [0.3, 3] relative to Routes A/B confirms three-route convergence; a result outside falsifies the framework independently of any laboratory measurement.

(b) Correlation parameter ε. The Stability Theorem (SD §6, Theorem 6.1) is explicitly conditional on $\varepsilon \leq 10^{-3}$. Evaluate all K = 7 constraint satisfaction indicators per triangle and compute ε = max_{i≠j} $|\mathrm{Cov}(C_i, C_j)|$. If $\varepsilon > 10^{-3}$, mesoscopic confinement of ξ is not guaranteed and the framework requires revision.

(c) Foam degree cap d.* The theorem-grade percolation lower bound p_c ≥ 1/(d*−1) requires knowing the maximum triangle coordination number in the full foam. The physically motivated estimate d* ≈ 7 gives p_c ≥ 0.167; d* = 6 tightens to p_c ≥ 0.20.

All three measurements are achievable with existing dynamical triangulations simulation codes and require no new instrumentation.

1. Casimir deviations near the coherence scale.

The finite-capacity structure predicts deviations from standard Casimir scaling at plate separations d ~ ξ ~ 85 μm. The functional form of the deviation follows from the master equation: finite-capacity competition reduces the effective coupling g_eff^2, which enters the Casimir force through the vacuum energy density at scale d. For d ~ ξ the deviation is O(ε_cap);

for $d \ll \xi$ the correction falls off because the substrate operates well below saturation at scales much smaller than ξ; for $d \gg \xi$ the correction vanishes because the geometry has fully percolated and competition effects are suppressed. Expanding η_cap at small ε_cap:

$$\eta\_cap(d) \approx 1 - \varepsilon\_cap(d) \cdot (N\_proc - 1)$$

where $\varepsilon\_cap(d) \propto \chi(d/\xi) \cdot N\_loop/(N\_loop + N\_proc - 1)$. In the mesoscopic regime $\chi(d/\xi) \sim (d/\xi)$ for $d < \xi$ (sub-saturation approach), giving:

$$\delta F\_Cas / F\_Cas \sim \varepsilon\_cap \cdot (d/\xi) \text{ for } d \ll \xi$$

and $\delta F\_Cas / F\_Cas \sim \varepsilon\_cap \sim O(10^{-2})$ at $d \sim \xi$.

Detectability caveat: Standard Casimir experiments at $d \sim 0.1$–10 μm face systematic uncertainties at the 1–10% level from surface roughness, patch potentials, and finite-temperature corrections. For $d \ll \xi$, the predicted signal $\delta F/F \sim \varepsilon\_cap \cdot (d/\xi)$ falls below these systematic floors for all practical separation ranges in current apparatus. The prediction is potentially accessible **only near $d \sim \xi \sim 85$ μm**, where the deviation reaches $O(\varepsilon\_cap) \sim O(10^{-2})$. Most current Casimir apparatus do not operate at separations this large. The prediction is therefore a target for future dedicated experiments rather than a near-term test with existing equipment.

2. Non-Gaussian statistics in decoherence rates.

Constraint correlations ($\Delta \neq 0$) predict systematic deviations from Gaussian statistics in quantum decoherence rates measured at or below the coherence scale. The operative bound on $|\Delta|$ is the Layer 1 estimate $|\Delta| \leq 42\varepsilon \leq 0.042$ at $\varepsilon \leq 10^{-3}$, giving a deviation amplitude $e^{|\Delta|} - 1 \sim O(10^{-2}$–$10^{-1})$ — the experimentally targeted regime. If $|\Delta|$ were to approach the Layer 2 structural ceiling ($|\Delta| \to \ln 2$), the amplitude could reach $O(1)$, but Appendix A establishes that this ceiling is not the operative bound for mesoscopic confinement; in practice Layer 1 dominates. The realistic detection target is therefore $O(1\%)$ deviations, potentially accessible in precision matter-wave interferometry at coherence lengths approaching 100 μm with next-generation sensitivity.

3. Local suppression of vacuum energy in high-density environments.

Finite-capacity competition predicts a measurable local reduction in effective vacuum energy in regions of anomalously high physical activity. "High physical activity" is here defined operationally as regions where the local TPB commitment rate $\chi(L)$ satisfies $\chi(L) \geq 0.3$, corresponding to physical environments with energy density $\rho \geq 10^{-2}\ \rho\_Planck$ or, equivalently, within the innermost ~10 Schwarzschild radii of a stellar-mass black hole, or in degenerate stellar interiors. The predicted magnitude is:

$$\delta\Lambda\_local / \Lambda \sim \varepsilon\_cap \cdot \delta N\_proc \sim O(10^{-2})$$

at $\chi(L) \sim 0.3$, scaling approximately linearly with $\chi(L)$ for $\chi(L) \ll 1$. This is sub-leading but potentially detectable as a systematic in precision equivalence principle tests in strong-field environments.

8. Conclusion

The cosmological constant is not predicted here as a single numerical value. It is predicted as a structurally determined band whose width is fixed by bounded second-order effects within a framework whose first-order structure admits no free parameters at first order. The primary Bethe prediction $\xi \in [60, 110]$ μm, corresponding via $\Lambda \sim \xi^{-4}$ to $\Lambda \in [0.36, 4.0] \times \bar{\Lambda}$, is the full physical prediction of the framework at first order. The second-order corrections characterise the log-normal spread around the central value as $\sigma_{\ln \xi} \sim 0.33$ ($\sigma_{\ln \Lambda} \sim 1.3$), with the ±1σ interval $\Lambda \in [0.27, 3.7] \times \bar{\Lambda}$ representing the distribution of realisations within the first-order band. Neither characterisation can be collapsed to a point without claiming information about microscopic graph topology that the combinatorial structure does not supply.

The residual spread is not a failure of the derivation. It is a calculable, structurally bounded consequence of three second-order effects, all unified in the single master equation:

```
╔══════════════════════════════════════════════════════════════╗
║  ln(ξ/ℓ_e) = (1/2b) · [1/(g₀² e^Δ η_cap) - 1/p_c]           ║
╚══════════════════════════════════════════════════════════════╝
```

with $\ell_e = 2\ell_P$. The architecture of this result, in clean order, is:

First-order structure fixed. The leading-order coherence scale $\xi_0 \in [60, 110]$ μm is fixed by $K = 7$, $b = 0.875$, $g_0^2 = 1/128$, and $p_c \in [0.167, 0.20]$ — all determined from simplicial foam combinatorics without cosmological input. The CSS Attractor Theorem [SC §6.6] proves that the central value is the unique globally stable fixed point of the saturation dynamics.

Second-order effects bounded. Three bounded effects perturb ξ away from ξ_0:

1. **Percolation threshold variation** — p_c varies within [0.167, 0.20] due to topological non-uniformity; the Bethe range is the primary prediction with the conservative bound serving as a robustness check controlled by the spectral structure of J(5,3)
2. **Constraint correlations** — bounded by the three-layer hierarchy: $|\Delta| \leq 42\varepsilon$ [SD Lemma 3.2], $|\Delta| \lesssim \ln 2$ [Appendix A], Δ within-class only [RG §6.7]; requires $\varepsilon \leq 10^{-3}$ for mesoscopic confinement
3. **Finite-capacity competition** — parametrised by $\varepsilon_{cap} \sim O(0.1)$, grounded in $N_{loop} = 14$ channel architecture and TPB saturation; the functional form is a minimal parametrization and is the least-controlled second-order contribution; §4.3 establishes robustness of the combined $\sigma_{\ln \xi} \sim 0.33$ to factor-of-2 uncertainty in ε_{cap}

Three parameters are independent (§3.3), and their contributions add in quadrature; cross-covariances are fourth order overall.

$R_P \sim 0.7N$ is derived, not assumed, from the degree variance of the foam graph following the cross-simplex gluing construction (§5.2), giving the band asymmetry as a structural prediction.

One primary band:

$\xi \in [60, 110]$ μm (first-order Bethe band) $\Lambda \in [0.36, 4.0] \times \bar{\Lambda}$ (via $\Lambda \sim \xi^{-4}$, hard first-order limits) $\sigma_{\ln \xi} \sim 0.33$, $\sigma_{\ln \Lambda} \sim 1.3$ (second-order log-normal spread within the band)

One conservative robustness envelope: $\xi \in [60, 320]$ μm, centred near 100 μm, still overlapping Routes A/B at $\xi \approx 88$ μm.

Three Monte Carlo tests: $p_c \approx 0.18$, $\varepsilon \leq 10^{-3}$, $d^* \approx 7$ — each independently falsifying, all achievable in existing simulation codes.

The width of the primary band is not a measure of ignorance. It is the irreducible signature of a complex relational system whose first-order structure is fixed but whose second-order details are not uniquely determined from combinatorics alone — within the present simplicial-foam universality class and the three second-order effects analysed here, this is the complete prediction, analogous to the spread of solutions within a universality class in critical phenomena.

Appendix A: Derivation of the Layer 2 Bound on |Δ|

The Layer 2 bound $|\Delta| \lesssim \ln 2$ is derived from the requirement that the fixed-point condition $p(\xi) = p_c$ admits a unique minimal solution. The derivation proceeds in three steps.

Step 1: The fixed-point condition and its parametric dependence. The master equation defines ξ implicitly as the scale at which the effective constraint-satisfaction amplitude $g_{\text{eff}}^2 = g_0^2 e^{\Delta}$ first reaches the percolation threshold. This requires:

$$2b g_{\text{eff}}^2 \cdot \ln(\xi/\ell_e) = 1 - g_{\text{eff}}^2/p_c$$

For a valid minimal coherence scale to exist, the right-hand side must remain positive in a neighbourhood of ξ_0 for a stable minimal solution, and the fixed point must be stable: small perturbations in ξ must return the system to the fixed point.

Step 2: The monotonicity condition. For the fixed-point condition to define a *unique minimal* ξ, the function $f(\xi) \equiv 2b g_{\text{eff}}^2 \cdot \ln(\xi/\ell_e) - [1 - g_{\text{eff}}^2/p_c]$ must have exactly one zero in the physical domain $\xi > \ell_e$, crossing from negative to positive. This requires $f(\xi)$ to be strictly monotonically increasing in the relevant regime.

Computing the derivative explicitly:

$$df/d\xi = 2b g_{\text{eff}}^2 / \xi > 0 \text{ for all } \xi > 0$$

since $b > 0$, $g_{\text{eff}}^2 > 0$, and $\xi > 0$. The left-hand side $2b g_{\text{eff}}^2 \cdot \ln(\xi/\ell_e)$ is therefore strictly increasing in ξ, while the right-hand side $1 - g_{\text{eff}}^2/p_c$ is a constant with respect to ξ. A unique crossing exists if and only if $f(\ell_e) < 0$ (below threshold at the emergence scale) and $f(\xi) \to +\infty$

as $\xi \to \infty$ (always satisfied since $\ln \xi$ is unbounded). The condition $f(\ell_e) < 0$ requires $2bg_{\mathrm{eff}}^2 \cdot 0 = 0 < 1 - g_{\mathrm{eff}}^2/p_c$, i.e. $g_{\mathrm{eff}}^2 < p_c$. Since $g_{\mathrm{eff}}^2 = g_0^2 e^{\Delta} \ll 1$ for all $|\Delta| < \ln 2$, and $p_c \in (0, 1]$, this is always satisfied. The monotonicity is therefore demonstrated explicitly, not asserted.

The relevant condition for the *universality class* to remain the simplicial-foam class is that $g_{\mathrm{eff}}^2 < 2^{-(K-1)}$, which is the condition derived in Step 3. This is separate from the monotonicity argument and imposes the structural ceiling on Δ.

Step 3: The bound. The condition $g_{\mathrm{eff}}^2 < 2^{-(K-1)}$ translates directly to:

$$g_0^2 e^{\Delta} < 2^{-(K-1)}$$

$$2^{-K} e^{\Delta} < 2^{-(K-1)}$$

$$e^{\Delta} < 2$$

$$\Delta < \ln 2$$

For the lower bound, $\Delta > -\ln 2$ follows from the symmetric requirement that $g_{\mathrm{eff}}^2 > 2^{-(K+1)}$ — below this threshold, the system requires $K + 1$ constraints to form coherent triangles, again placing it in a different universality class. Together:

$$|\Delta| < \ln 2 \approx 0.693$$

The bound is tight: at $|\Delta| = \ln 2$ exactly, the fixed-point condition becomes marginal and the theory sits at a universality class boundary.

Note on the relative strength of the three layers. The Layer 2 bound $|\Delta| < \ln 2 \approx 0.69$ is structurally interesting — it identifies the universality class boundary — but it is considerably weaker than the Layer 1 bound $|\Delta| \leq 42\varepsilon \leq 0.042$ (at the required $\varepsilon \leq 10^{-3}$). In practice, Layer 1 does all the work for mesoscopic confinement of ξ. Layer 2 would only become the operative constraint if ε were allowed to reach $\sim 10^{-2}$, at which point mesoscopic confinement breaks down regardless. The three-layer hierarchy is therefore a conceptual structure — demonstrating that the bound has multiple independent supports — but for quantitative predictions, the operative bound is Layer 1. ■

Appendix B: Notation Summary

Symbol	Meaning
ξ	Coherence scale (second-order prediction)
ξ_0	First-order Route M coherence scale $\in [60, 110]$ µm
ℓ_P	Planck length ($\sim 1.6 \times 10^{-35}$ m)

Symbol	Meaning
ℓ_e	Two-Planck emergence scale = 2ℓ_P [TP §1.1]
L_H	Future event horizon (~1.65×10^{26} m) [TP §3.2]
p_c	Percolation threshold; primary [0.167, 0.20], conservative [0.17, 0.30] [TP §4.8 M5, M5b]
g_0	Base constraint-satisfaction amplitude = $2^{-7/2}$ [SD §2, Theorem 2.3]
K	Number of independent constraints per cell = 7 [TP App. D.1, constraints C1–C7]
b	β-function coefficient = 14/16 = 0.875 [TP App. D.2]
N_loop	Loop channels per RG block = 14 [TP App. D.2]
Δ	Constraint correlation parameter; $\|\Delta\| \leq 42\varepsilon$ [SD Lemma 3.2], $\|\Delta\| \lesssim \ln 2$ [App. A]
ε	Pairwise constraint covariance bound; requires $\varepsilon \leq 10^{-3}$ [SD §3.4]
ε_cap	Capacity-competition parametrization ~ χ(L) · N_loop / (N_proc(N_loop + N_proc − 1))
χ(L)	Local TPB saturation fraction ~ 0.1–0.2
N_proc	Number of competing processes
η_cap	Throughput reduction factor = 1 / (1 + ε_cap(N_proc − 1))
R_P	Participation ratio; R_P ≥ 0.7N derived from foam graph degree variance
d*	Foam degree cap; p_c ≥ 1/(d*−1) [SD §4.3]; physically motivated d* ≈ 7
TP	Taylor, *Two-Planck Principle* (companion paper)
RG	Taylor, *Relational Geometry and Universality* (companion paper)
SC	Taylor, *Structural Closure* (companion paper)
ES	Taylor, *Epistemic Status* (companion paper)
SD	Taylor, *Stability of Dimensional Transmutation* (companion paper)

References

Taylor, K. *Two-Planck Principle: From Quantum Geometry to Emergent Gravity.* VERSF Theoretical Physics Program, AIDA Institute. [cited as TP] — Derives quantitative predictions including ξ ≈ 88 μm, $\rho\_{vac} \approx 5 \times 10^{-10}$ J/m³, $\Lambda \approx 1.1 \times 10^{-52}$ m⁻², recovery of GR at 1PN order, and three-route convergence (Routes A, B, M).

Taylor, K. *Relational Geometry and the Universality of the Two-Planck Scale.* VERSF Theoretical Physics Program, AIDA Institute. [cited as RG] — Establishes completeness of K = 7 (§6.5), proves K cannot vary continuously without changing universality class (§6.7–6.8), and demonstrates CSS is a dynamical attractor (§5).

Taylor, K. *Structural Closure of the Two-Planck Framework.* VERSF Theoretical Physics Program, AIDA Institute. [cited as SC] — Provides independent verification of K = 7 by information-theoretic lower bound (SC §4.2) and obstruction-theoretic count (SC §4.3); proves the CSS Attractor Theorem via explicit Lyapunov function $\mathscr{L}(x) = \frac{1}{2}(x-1)^2$ with $w = -1$ as a corollary (SC §6.6–6.7).

Taylor, K. *Epistemic Status of the Two-Planck Derivation of the Cosmological Constant.* VERSF Theoretical Physics Program, AIDA Institute. [cited as ES] — Characterises the framework as "conditional but non-retrofitted" (ES §7.2); resolves scheme dependence via the Λ_QCD analogy (ES §4.2, §5.2); identifies Monte Carlo determination of p_c as the decisive next step (ES §7.3).

Taylor, K. *Stability of Dimensional Transmutation in Simplicial Foam.* VERSF Theoretical Physics Program, AIDA Institute. [cited as SD] — Proves $g_0^2 = 2^{-K}$ via flip-group invariance (SD §2, Theorem 2.3, App. A); proves $|\Delta| \leq 42\varepsilon$ for $K = 7$ (SD Lemma 3.2); derives theorem-grade percolation lower bound p_c ≥ 1/(d*−1) from the Johnson graph J(5,3) structure (SD §4.2–4.3, App. C); proves the Route M Microphysical Closure Theorem (SD Theorem 6.1).

Adler, J., Meir, Y., Aharony, A., Harris, A.B., Klein, L. (1990). Low-concentration series in general dimension. *J. Stat. Phys.* 58, 511.

Bollobás, B. & Riordan, O. (2006). The critical probability for random Voronoi percolation in the plane is 1/2. *Probability Theory and Related Fields* 136(3), 417–468.

Decca, R.S. et al. (2007). Tests of new physics from precise measurements of the Casimir pressure. *Physical Review D* 75, 077101.

Kramer, B. & MacKinnon, A. (1993). Localisation: theory and experiment. *Reports on Progress in Physics* 56(12), 1469–1564.

Margolus, N. & Levitin, L.B. (1998). The maximum speed of dynamical evolution. *Physica D* 120, 188–195.

Wegner, F. (1980). Inverse participation ratio in $2+\varepsilon$ dimensions. *Zeitschrift für Physik B* 36, 209–214.

Ziff, R.M. & Scullard, C.R. (2006). Exact bond percolation thresholds in two dimensions. *J. Phys. A: Math. Gen.* 39, 15083.