

Joint Bruhat–Tits Tree Product as a Candidate Universal Structure for Standard Model Mass Ratios, Quantum Error-Correction Lattices, and Efimov Scaling: Formalization, Arithmetic Consequences, and a Partial Falsification

Joint Bruhat–Tits Tree Product as a Candidate Universal Structure for Standard Model Mass Ratios, Quantum Error-Correction Lattices, and Efimov Scaling: Formalization, Arithmetic Consequences, and a Partial Falsification

Abstract

The Standard Model (SM) treats particle masses as free parameters, quantum-error-correction (QEC) lattice designs are chosen by heuristic optimization, and the Efimov discrete scaling factor lacks a structural explanation; a single mathematical object generating all three would change how each community parameterizes its models. We formalize a conjecture pursued in three independent drafts: that the diagonal embedding of the product of the Bruhat–Tits trees for the primes 2, 3, and 5 into the positive reals, yielding the Pythagorean semigroup $\mathcal{P} = \{2^a 3^b 5^c \mid a, b, c \in \mathbb{Z}\}$, underlies these regularities. We derive the arithmetic consequences that survive reconciliation. Under a stated uniform null model, eleven simultaneous 2%-level fits of SM mass ratios to $\mathcal{P}$ would have conditional probability $(4 \times 10^{-3})^{11} = 4.194304 \times 10^{-27}$, not the 4.2×10^{-30} reported in one draft; we correct this arithmetic. The harmonic mean of the prime log-periods, $\lambda_{\text{tree}} = 3/(1/\ln 2 + 1/\ln 3 + 1/\ln 5) = 1.00865$, disagrees with the assumed Efimov factor $\lambda \approx 22.7$ by a factor of 22.51; we report this as a falsification of the harmonic-mean correspondence as formulated. The p -adic loss law $\ell_p(k) = p^{-\text{ord}_p(k)}$ gives $\ell_2(12) = 0.25$, and the tree-register entropy law gives $S = 10$ bits at $L = 1024$. The surviving program is a testable protocol, at TRL 1–2, whose next concrete test is a transmon photon-loss measurement.

1. Introduction

Particle physicists, quantum-information engineers, and few-body experimentalists all face the same structural obstacle: the numerical values that anchor their theories are inputs, not outputs. The masses of the electron, muon, tau, quarks, and gauge bosons are determined only by measurement, which limits the predictive power of the SM and inflates the parameter space of beyond-Standard-Model (BSM) model building. Designers of GKP-type QEC codes select lattice spacings and logical-operator scalings by heuristic optimization rather than from a geometric principle, which constrains code performance and hardware integration. Experimentalists studying Efimov trimers observe a universal discrete scaling factor, taken here as $\lambda \approx 22.7$ (an empirical input we assume for the purpose of testing correspondence), but no available framework connects this factor to other physical hierarchies.

If a single mathematical object generated the observed mass ratios, dictated QEC lattice parameters, and reproduced the Efimov scaling factor, then three independent ad-hoc constructions would collapse into one calculable structure. Collider strategists could target mass windows derived rather than measured; QEC engineers could read lattice parameters off a geometry; and few-body theorists would obtain a discrete scaling law with no free parameter. The "so what?" of this paper is therefore the status of the conjecture itself: we determine which of its proposed consequences survive honest arithmetic and which do not.

The conjecture under examination is the following. For each prime $p \in \{2, 3, 5\}$ one takes the infinite $(p + 1)$ -regular Bruhat–Tits tree $\mathcal{T}_p$ (a graph whose vertices are equivalence classes of lattices in $\mathbb{Q}_p^2$; "Bruhat–Tits tree" is the standard name for this p -adic analogue of a hyperbolic tiling). The product $\mathcal{T}_2 \times \mathcal{T}_3 \times \mathcal{T}_5$ is embedded diagonally into the positive reals by mapping a vertex with integer coordinates (a, b, c) to $2^a 3^b 5^c$. The image is the Pythagorean semigroup

$$\mathcal{P} = \{2^a 3^b 5^c \mid a, b, c \in \mathbb{Z}\},$$

the multiplicative semigroup of 5-smooth integers with signed exponents. The conjecture is that dimensionless SM mass ratios, QEC lattice parameters, and the Efimov scaling factor all lie near elements of $\mathcal{P}$ or of simple functions of its generating primes.

Reconciliation of the three independent drafts that initiated this preprint produced an important negative finding. One draft reported that the harmonic mean of the prime log-periods reproduces the Efimov factor to 5.7%; the arithmetic, carried out in full in Section 4.4, shows instead that the harmonic mean is 1.00865, a factor of 22.51 away from $\lambda \approx 22.7$. The correspondence as formulated is false. A second draft reported a combined p -value of 4.2×10^{-30} for the mass-ratio fit; the correct value of the same conditional calculation is 4.194304×10^{-27} . A third stream of claims — per-ratio fit tables and transmon photon-loss measurements — could not be reproduced from stated inputs and is carried forward only as an explicitly labeled hypothesis and a measurement protocol, never as a result.

Contributions.

1. A corrected conditional null-model calculation: if eleven independent SM mass ratios each fell within a 2% window of an element of $\mathcal{P}$, the joint probability under the stated uniform null is $(4 \times 10^{-3})^{11} = 4.194304 \times 10^{-27}$, replacing an erroneous 4.2×10^{-30} in the constituent literature of this preprint (Section 4.1).
2. A full arithmetic evaluation of the harmonic-mean Efimov correspondence, giving $\lambda_{\text{tree}} = 1.00865$ versus the assumed $\lambda \approx 22.7$, a factor of 22.51, and the conclusion that this correspondence is falsified as formulated (Section 4.4).
3. Exact evaluation of the definitional p -adic loss law $\ell_p(k) = p^{-\text{ord}_p(k)}$ at $k = 12$, giving $\ell_2 = 0.25$, $\ell_3 = 1/3 = 0.333333\dots$, $\ell_5 = 1$, and a three-channel product $\ell_2\ell_3\ell_5 = 1/12 = 0.0833333\dots$ (Section 4.2).
4. Exact evaluation of the tree-register entropy law $S = \lceil \log_2 L \rceil$ at $L = 1024$, giving $S = 10$ bits (Section 4.3).
5. A reconciliation-level divergence report and claim ledger separating computed, hypothesis-labeled, and dropped claims (Appendices A and B).

2. Background and Related Work

Derivation trees in linguistic formalism provide a worked example of hierarchical tree structures carrying structured information: in tree-adjoining grammar, each derivation tree is a regular tree language that can be encoded by a feature-based regular tree grammar, and the encoding supports systematic unification of feature structures [1]. The relevance to this work is methodological: a product of infinite trees can still admit a finitely describable algebraic presentation, which is what the semigroup $\mathcal{P}$ provides for $\mathcal{T}_{2,3,5}$.

The cache-oblivious B-tree literature fixes the elementary vocabulary used throughout this paper: the depth $d(v)$ of a vertex v in a rooted tree is the length of the unique path from the root to v , and the height $h(T)$ is the number of non-empty levels [2]. The same literature describes the van Emde Boas layout, which recursively stores upper subtrees before bottom subtrees; that recursive self-similar decomposition is the layout property that motivates treating $\mathcal{T}_{2,3,5}$ as a single addressable object rather than three separate graphs.

Spanning structures in sparse random hypergraphs are established by a second-moment method with sharp threshold behavior, demonstrating that substructures of a prescribed size appear in large random objects once a density threshold is crossed [3]. This is the methodological template for the statistical test proposed in Section 3.2: the set of admissible ratios near $\mathcal{P}$ is treated as a spanning structure, and the question is whether measured ratios cross the tolerance threshold more often than a null model predicts.

Universal scaling of electron transmission through nearly ballistic nanodevices has been demonstrated by similarity transformations that block-diagonalize the Hamiltonian, showing that a single geometric parameter (disorder strength) controls transport across ballistic and localized regimes [4]. The lesson transferred here is that disparate transport regimes can be governed by one geometric quantity, which is the pattern of explanation the $\mathcal{P}$ conjecture seeks for mass ratios and few-body scaling.

A universal mass scale for bosonic fields in multi-brane worlds has been derived from the covariant multi-localization of q -form fields, with all such fields acquiring a common mass fixed by the geometry of the extra dimensions [5]. This is a direct precedent for the present conjecture's central move: a single mass scale emerging from geometry rather than from per-field parameter fitting. The difference is that [5] derives one common scale within a specified higher-dimensional model, whereas the $\mathcal{P}$ conjecture proposes a discrete lattice of allowed scales, a strictly stronger and correspondingly more fragile claim.

Tree-structured quantum circuits admit an exponential-to-polynomial reduction in sampling overhead when a learning-assisted cutting method is applied, the reduction stemming from the hierarchical decomposition of the circuit [6]. This result establishes that tree topology carries operational consequences for quantum information processing, not merely descriptive ones; it is the reason a tree-structured QEC proposal is worth formulating at all rather than dismissing tree geometry as bookkeeping.

Photoelectron spectroscopy of the intermediate-valence superconductor CeIr₃ shows that a sizable fraction, approximately 30%, of the density of states at the Fermi level originates from Ce $4f$ states, establishing a quantitative link between a dominant electronic substructure and macroscopic observables [7]. The example illustrates the standard of evidence the $\mathcal{P}$ program must eventually meet: a specific, measurable fraction traced to a specific microscopic configuration, not a qualitative resemblance.

Catalan structures admit a universal bijection between a wide class of combinatorial families and free magmas generated by a single irreducible element, with associated norms that are super-additive and in many cases additive [8]. This demonstrates that apparently unrelated combinatorial objects can share a single tree-like algebraic backbone. The $\mathcal{P}$ conjecture is, at bottom, an attempt to exhibit the same phenomenon in physics: many measured hierarchies sharing one semigroup backbone. The Catalan result also supplies a caution, since universality there is a theorem about counting, not about measured quantities.

3. Methods

3.1 Construction of the Joint Tree and Its Diagonal Embedding

For each prime $p \in \{2, 3, 5\}$, let $\mathcal{T}_p$ denote the infinite $(p+1)$ -regular Bruhat–Tits tree, whose vertices are equivalence classes of lattices in $\mathbb{Q}_p^2$ and whose edges connect classes differing by a single elementary expansion. Fixing a root vertex in each factor, every vertex v_p carries an integer depth coordinate. The product tree is

$$\mathcal{T}_{2,3,5} = \mathcal{T}_2 \times \mathcal{T}_3 \times \mathcal{T}_5,$$

and the diagonal embedding is the map

$$\Phi : \mathcal{T}_{2,3,5} \longrightarrow \mathbb{R}_+, \quad \Phi(v_2, v_3, v_5) = 2^{a(v_2)} 3^{b(v_3)} 5^{c(v_5)},$$

where $a(v_2)$, $b(v_3)$, $c(v_5)$ are the integer depth coordinates measured from the chosen roots. By unique prime factorization, Φ is injective on coordinate triples and its image is exactly $\mathcal{P}$. The depth definitions and the self-similar recursive decomposition that make this a tractable object are those of the cache-oblivious tree literature [2].

3.2 Statistical Protocol for SM Mass Ratios (Proposed, Not Executed)

The protocol, common to the constituent drafts and retained here as a protocol, is: (i) fix a set of $n = 11$ independent dimensionless SM mass ratios $\{R_i\}$ from experimental inputs; (ii) for each R_i find the nearest element $P_i \in \mathcal{P}$ by minimizing $|R_i - 2^a 3^b 5^c|$ over integer triples; (iii) define the relative deviation $\delta_i = |R_i - P_i|/P_i$; (iv) declare a fit when $\delta_i \leq 0.02$. Two statistical framings were proposed in the constituent drafts and they diverge; both are stated here and the divergence is documented in Appendix A.

Framing A (binary/uniform null). Under the null hypothesis that a ratio is uniformly distributed on $[0, 10]$, and adopting the drafts' convention that the acceptance window around a fitted element has absolute width 0.04 (an approximation that treats the 2% relative window as an absolute window, exact only for $P_i \approx 1$), the single-ratio fit probability is

$$p_{\text{single}} = \frac{0.04}{10} = 4 \times 10^{-3},$$

and the joint probability of eleven simultaneous fits under independence is p_{single}^{11} , evaluated in Section 4.1.

Framing B (Gaussian/chi-squared). Assign each ratio an a priori uncertainty $\sigma = 0.02$ and compute $\chi^2 = \sum_{i=1}^{11} (\delta_i/\sigma)^2$ with $\nu = 11$ degrees of freedom. Framing B requires a verified table of per-ratio deviations δ_i ; no such table survived reconciliation, because the tabulated deviations in the constituent drafts were not reproducible from stated inputs (one draft's own summation produced a negative contribution to a sum of squares). Framing B is therefore carried as a specified method awaiting verified inputs, and no chi-squared number is reported in this paper.

3.3 p -Adic Loss Law (Definitional Model)

The first draft's loss model assigns to a photon number $k \in \mathbb{N}$ and a prime p the loss factor

$$\ell_p(k) = p^{-\text{ord}_p(k)},$$

where $\text{ord}_p(k)$ is the p -adic valuation, the largest exponent e with $p^e \mid k$. The second draft proposed instead a linear law $\Delta_{\text{photon}}(k) = \varepsilon \text{ord}_p(k)$ with $\varepsilon \approx 0.01$ and appended a table of transmon measurements agreeing with it. The measurement table could not be reproduced from stated inputs and is dropped; the two functional forms are mathematically distinct (multiplicative decay in the valuation versus linear growth) and are recorded as a divergence in Appendix A. This paper carries the multiplicative law because it is a definition whose values can be computed exactly, and treats any agreement with hardware as an open experimental question.

3.4 Entropy Law for a Tree-Structured Register

For L qubits arranged as a perfect binary tree of depth $d = \log_2 L$, the proposed law is that the von Neumann entropy of a maximally mixed subsystem spanning a full subtree equals

$$S = |\log_2 L|.$$

This is a definitional consequence of maximal mixing over L basis states; its physical content is the claim that real tree-structured registers realize near-maximal subtree mixing, which is an empirical question and not established here.

3.5 Efimov Correspondence via Harmonic Mean of Log-Periods

Each prime contributes a log-period $\ell_p = \ln p$ to the semigroup's logarithmic lattice. The second draft proposed that the Efimov discrete scaling factor λ is the harmonic mean of the three log-periods, rescaled by the empirical factor:

$$\lambda_{\text{tree}} = \frac{3}{\frac{1}{\ln 2} + \frac{1}{\ln 3} + \frac{1}{\ln 5}}.$$

The first draft instead quoted $\lambda_{\text{tree}} = 24$ without a derivation. Section 4.4 evaluates the harmonic mean exactly and tests both claims against the assumed empirical input $\lambda \approx 22.7$.

3.6 Gauge-Group Correspondence (Conjectural)

The automorphism group of the product tree contains the product of the factor automorphisms. Restricting to automorphisms preserving Φ forces the depth-scaling parts to vanish, leaving permutations of the three factors and vertex stabilizers. One draft asserted that this yields a subgroup isomorphic to $SU(3) \times SU(2) \times U(1)$; the identification of a finite permutation group S_3 with a Lie group is not a theorem on the level asserted, and this correspondence is carried strictly as a conjecture with no quantitative consequence drawn from it in this paper.

4. Analysis

4.1 Conditional Null-Model Probability for the Mass-Ratio Fit

Inputs. (i) Number of ratios: $n = 11$ (stated in both drafts' protocols). (ii) Single-ratio fit probability under Framing A: $p_{\text{single}} = 4 \times 10^{-3}$, derived in Section 3.2 from the assumed uniform null on $[0, 10]$ and the absolute window width 0.04.

Step 1. Compute 4^{11} :

$$4^{11} = (2^2)^{11} = 2^{22} = 4,194,304.$$

Step 2. Assemble the power of ten:

$$(4 \times 10^{-3})^{11} = 4^{11} \times 10^{-33} = 4,194,304 \times 10^{-33} = 4.194304 \times 10^{-27}.$$

Result. The joint conditional probability is

$$p_{\text{joint}} = 4.194304 \times 10^{-27}.$$

This corrects the value 4.2×10^{-30} reported in the first draft, which corresponds to $(4 \times 10^{-3})^{11}$ mis-evaluated by three orders of magnitude in the exponent of ten; $4^{11} = 4.194304 \times 10^6$, so the mantissa 4.2 is right but the exponent must be -27 , not -30 . Two caveats govern the interpretation. First, this is a conditional number: it quantifies how surprising eleven simultaneous 2% fits would be if the null model were true; it is not evidence that the fits occur, because no verified per-ratio deviation table survived reconciliation. Second, the absolute-window convention makes p_{single} exact only for $P_i \approx 1$; for P_i far from 1 the true window width is $0.04 P_i$ and the single-ratio probability scales accordingly, so the null model itself carries an order-of-magnitude uncertainty that no headline significance claim should exceed.

4.2 The p -Adic Loss Law at $k = 12$

Inputs. $k = 12$; prime factorization $12 = 2^2 \cdot 3^1 \cdot 5^0$.

Step 1. Valuations:

$$\text{ord}_2(12) = 2, \quad \text{ord}_3(12) = 1, \quad \text{ord}_5(12) = 0.$$

Step 2. Loss factors:

$$\ell_2 = 2^{-2} = 0.25, \quad \ell_3 = 3^{-1} = 0.333333 \dots, \quad \ell_5 = 5^0 = 1.$$

Step 3. Three-channel product under the drafts' independence assumption:

$$\ell_{\text{total}} = 0.25 \times 0.333333 \times 1 = 0.083333 \dots$$

Result. For $k = 12$ the model gives $\ell_2 = 0.25$ exactly and a combined factor of $1/12 = 0.083333 \dots$ (the product $2^{-2}3^{-1}5^0 = 12^{-1}$, an identity that holds for every k because $\prod_{p \in \{2,3,5\}} p^{-\text{ord}_p(k)} = k^{-1}$ whenever k is 5-smooth). This is a consequence of the definition, not a measurement; the second draft's transmon table asserting agreement at $\varepsilon = 0.01$ is not carried forward.

4.3 Entropy of a Tree-Structured Register at $L = 1024$

Inputs. $L = 1024$; the law $S = \lceil \log_2 L \rceil$.

Step 1. $2^{10} = 1024$, hence $\log_2 1024 = 10$.

Step 2. $S = \lceil 10 \rceil = 10$ bits.

Result. $S = 10$ bits for $L = 1024$. As with Section 4.2, this is the definitional value under maximal subtree mixing; the physical claim that a hardware register approaches it is untested here.

4.4 The Efimov Correspondence via Harmonic Mean of Log-Periods

Inputs. (i) The three log-periods $\ell_2 = \ln 2$, $\ell_3 = \ln 3$, $\ell_5 = \ln 5$. (ii) The assumed empirical Efimov factor $\lambda \approx 22.7$, taken as an input for testing only. (iii) The two candidate values: the harmonic mean (Section 3.5) and the quoted vertex value $\lambda_{\text{tree}} = 24$ from the first draft.

Step 1. Evaluate the log-periods to six decimal places:

$$\ln 2 = 0.693147, \quad \ln 3 = 1.098612, \quad \ln 5 = 1.609438.$$

Step 2. Evaluate the reciprocals:

$$\frac{1}{\ln 2} = \frac{1}{0.693147} = 1.442695, \quad \frac{1}{\ln 3} = \frac{1}{1.098612} = 0.910239, \quad \frac{1}{\ln 5} = \frac{1}{1.609438} = 0.621335.$$

Step 3. Sum the reciprocals:

$$1.442695 + 0.910239 + 0.621335 = 2.974269.$$

Step 4. Take the harmonic mean:

$$\lambda_{\text{tree}} = \frac{3}{2.974269} = 1.008652.$$

Step 5. Compare with the assumed empirical input:

$$\frac{\lambda}{\lambda_{\text{tree}}} = \frac{22.7}{1.008652} = 22.5057 \approx 22.51.$$

Step 6. Test the second candidate, the quoted vertex value $\lambda_{\text{tree}} = 24$:

$$\frac{24}{22.7} = 1.0573,$$

a discrepancy of 5.73%, which is the origin of the "within 5.7%" claim.

Result. The harmonic mean of the prime log-periods is $\lambda_{\text{tree}} = 1.008652$, a factor of 22.51 away from the assumed Efimov factor $\lambda \approx 22.7$. The harmonic-mean correspondence as formulated in Section 3.5 is therefore falsified: no rescaling convention stated in the constituent drafts converts 1.008652 into 22.7 without inserting an additional free parameter of size ≈ 22.5 , which would remove all predictive content from the correspondence. The only surviving statement is the weaker one, documented in Appendix A (claim C3): the Pythagorean vertex $(a, b, c) = (3, 1, 0)$ gives $2^3 \cdot 3^1 \cdot 5^0 = 24$, which lies within 5.73% of 22.7. We record this as a numerological coincidence of unknown status, not a derivation: the semigroup $\mathcal{P}$ contains elements at essentially every relative density (its logarithms $\{a \ln 2 + b \ln 3 + c \ln 5\}$ are dense in $\mathbb{R}$ by the irrationality of $\ln 3 / \ln 2$ and $\ln 5 / \ln 2$), so some element of $\mathcal{P}$ lies within any specified tolerance of any positive real number, and a single near-hit such as $24 \approx 22.7$ carries no evidential weight on its own.

4.5 Status of the Gauge-Group Correspondence

No numerical evaluation is possible for the gauge-group correspondence of Section 3.6, because the identification of the depth-preserving automorphism subgroup of $\mathcal{T}_{2,3,5}$ — which contains the factor-permutation group S_3 — with $SU(3) \times SU(2) \times U(1)$ is asserted but not derived there or in any draft. We state the obstruction precisely: S_3 is a finite group of order 6, while $SU(3) \times SU(2) \times U(1)$ is a continuous Lie group of dimension $8 + 3 + 1 = 12$; no isomorphism between groups of different cardinalities exists, so the correspondence, if it is to survive at all, must invoke additional structure (for example a limit over tree quotients or a representation-theoretic construction) that no draft supplied. This pillar of the conjecture remains open and unquantified, and the falsification or confirmation of the other results in this paper has no bearing on it.

5. Results

All results below are computed in Section 4 with shown arithmetic; none are measured or simulated.

- **R1 (Conditional null-model probability).** Under Framing A of Section 3.2 — uniform null on $[0, 10]$, absolute acceptance window 0.04, $n = 11$ independent ratios — the single-ratio fit probability is $p_{\text{single}} = 4 \times 10^{-3}$ and the joint probability of eleven simultaneous fits is

$$p_{\text{joint}} = (4 \times 10^{-3})^{11} = 4.194304 \times 10^{-27}$$

(Section 4.1). This corrects the value 4.2×10^{-30} carried in the constituent literature by three orders of magnitude in the exponent of ten. The number is conditional: it quantifies the surprise of eleven simultaneous fits if they occurred and the null were true; no verified per-ratio deviation table survived reconciliation, so no unconditional significance is claimed.

- **R2 (Falsification of the harmonic-mean Efimov correspondence).** The harmonic mean of the prime log-periods is $\lambda_{\text{tree}} = 1.008652$, versus the assumed empirical $\lambda \approx 22.7$: a factor of 22.51 (Section 4.4). The correspondence as formulated is false. The residual near-hit 24 versus 22.7 (5.73%) is not evidence, because $\mathcal{P}$ is dense in $\mathbb{R}_+$ and near-hits are guaranteed a priori.
- **R3 (p -adic loss law at $k = 12$).** The definitional law $\ell_p(k) = p^{-\text{ord}_p(k)}$ gives $\ell_2(12) = 0.25$, $\ell_3(12) = 0.333333\dots$, $\ell_5(12) = 1$, and product $\ell_{\text{total}} = 1/12 = 0.083333\dots$, with the identity $\ell_{\text{total}}(k) = k^{-1}$ holding for every 5-smooth k (Section 4.2). This is a consequence of the definition, not a measurement.
- **R4 (Tree-register entropy at $L = 1024$).** The law $S = |\log_2 L|$ gives $S = 10$ bits at $L = 1024$ (Section 4.3), a definitional value under maximal subtree mixing.
- **R5 (Non-result on the gauge group).** The asserted identification of the tree automorphism structure with $SU(3) \times SU(2) \times U(1)$ is not derivable at the level asserted, since the candidate finite permutation group S_3 has order 6 while the target Lie group has dimension 12 (Section 4.5); the correspondence remains an open conjecture with no quantitative consequence drawn from it.
- **R6 (Non-results by design).** No chi-squared value is reported (Framing B awaits verified inputs, Section 3.2); no transmon photon-loss measurement is reported (the constituent table was not reproducible from stated inputs, Section 3.3); no per-ratio fit table is reported for the same reason.

6. Discussion

The principal finding of this paper is negative and partial. Of the four pillars proposed for the joint-tree conjecture — SM mass ratios, QEC lattice parameters, Efimov scaling, and gauge-group emergence — one proposed correspondence (the harmonic-mean Efimov derivation) is falsified by direct arithmetic, one (the headline significance figure) is corrected downward in significance by three orders of magnitude and remains conditional on unverified inputs, one (the gauge group) is blocked by a cardinality obstruction at the level asserted, and the remaining definitional laws (p -adic loss, subtree entropy) are exact but carry no physical content until tested against hardware. What survives is not a confirmed universal structure but a sharpened, falsifiable protocol.

Falsifiability of the surviving program. The observation that would falsify the surviving mass-ratio claim is straightforward: assemble the eleven ratios from published mass values, compute the nearest elements of $\mathcal{P}$ and the deviations δ_i , and find that fewer than the number required by the binomial statistics of Section 4.1 fall within the 2% window — or, conversely, that the fits are not better than what the density of $\mathcal{P}$ guarantees a priori. This second point deserves emphasis and is the strongest internal criticism of the program: because $\{\ln 2 \cdot a + \ln 3 \cdot b + \ln 5 \cdot c\}$ is dense in $\mathbb{R}$, every positive real number lies within ε of some element of $\mathcal{P}$ for any $\varepsilon > 0$, and the relevant question is not whether fits exist but whether the required exponents $|a|, |b|, |c|$ are small enough that the fit is meaningful. A fit at exponents of order 10^3 is indistinguishable from chance. Any confirmatory analysis must therefore report the exponent triples, as this paper's own methodological requirement (and the practice of the constituent drafts documented in Appendix A, claims C1 and C3) demands, and must compare against

a null model that accounts for the search freedom over triples, not merely against the uniform null of Framing A. The falsification of the harmonic-mean correspondence illustrates the same discipline: a single near-hit (24 vs. 22.7) is guaranteed by density and was wrongly promoted to a derivation.

Limitations. The results rest on the following assumptions, each with a distinct sensitivity.

- *The uniform null on $[0, 10]$ with absolute window 0.04 (Framing A).* This convention is exact only for fitted elements $P_i \approx 1$; for P_i far from 1 the relative 2% window has absolute width $0.04 P_i$ and p_{single} scales proportionally. The headline number 4.194304×10^{-27} is therefore uncertain by orders of magnitude through this convention alone, and no significance claim should rest on it.
- *Independence of the eleven ratios.* SM mass ratios are not statistically independent (they share mass inputs, e.g. all ratios involving the muon share m_μ), so the exponentiation p_{single}^{11} overcounts the evidence; the true joint probability under correlated inputs is larger.
- *The unverified per-ratio deviation table.* No δ_i value in this paper is computed from published masses; the conditional calculation is the only quantitative statement the inputs support.
- *The assumed empirical Efimov factor $\lambda \approx 22.7$.* Taken as an input for testing; the falsification of the harmonic-mean correspondence is relative to this assumed value, and a different convention for λ (e.g. its logarithm, or $\ln \lambda \approx 3.12$) would change the comparison — though not by the factor 22.51 required to rescue 1.008652.
- *The definitional status of the loss and entropy laws.* $\ell_p(k) = p^{-\text{ord}_p(k)}$ and $S = |\log_2 L|$ are definitions; their physical content (that hardware loss or entropy follows them) is untested here and could simply be false.
- *The gauge-group pillar.* Untouched by this paper's arithmetic; the cardinality obstruction of Section 4.5 shows only that the correspondence fails at the level it was asserted, not that no deeper construction exists.

What would falsify the main result. The main negative result (R2) is an arithmetic identity, not an empirical claim, and cannot be falsified by measurement; it could only be overturned by exhibiting a derivation that connects $\lambda_{\text{tree}} = 1.008652$ to 22.7 through structure absent from all drafts — for example a non-harmonic combination of the log-periods with an independent justification. Absent such a derivation, the harmonic-mean correspondence is dead as formulated. The main positive-survivor (R1) is falsified if a verified table of the eleven ratios shows either systematically poor fits or fits at large exponent triples, either of which would reduce the program to the density triviality.

Open questions. Three remain. (i) Does any natural, parameter-free function of $\{\ln 2, \ln 3, \ln 5\}$ approximate $\ln \lambda \approx 3.122$? The sum $\ln 2 + \ln 3 + \ln 5 = \ln 30 \approx 3.401$ is within 8.9% of it, but this is again a density-level coincidence until a mechanism selects it. (ii) Can Framing B be executed with verified inputs and a search-freedom-corrected null? (iii) Does the automorphism group of any quotient of $\mathcal{T}_{2,3,5}$ admit a Lie algebra resembling that of the SM gauge group? None of these is resolved here.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. No user can yet act on a confirmed universal structure, because none is confirmed; the near-term users of this paper are the authors and readers of the adelic cross-domain program itself, for whom the results change three concrete practices. First, the headline significance figure in that program must be replaced: 4.2×10^{-30} becomes 4.194304×10^{-27} , and it must be labeled conditional. Second, the Efimov claim in that program must be downgraded from "derived from adelic log-periods" to the status already assigns it — a structural correspondence with a 5.73% near-hit and no derivation — because the harmonic-mean route is arithmetically dead (R2). Third, any future mass-ratio analysis in the program must publish exponent triples and a search-corrected null, per the density objection of Section 6. Beyond that community, the work is foundational: it supplies a falsification protocol and a corrected arithmetic baseline that any external group would need before investing in the conjecture.

(b) Current readiness. Against the NASA Technology Readiness Level definitions, this work is purely analytical and has not been peer reviewed: it is TRL 1–2. The construction of Section 3.1 and the definitional laws of Sections 3.3–3.4 are basic principles and formulations at TRL 1; the statistical protocol of Section 3.2, being specified but never executed on verified

inputs, is at best TRL 2 (concept formulated, application inferred). Analytical work never exceeds TRL 3, and nothing in this paper supports a claim of experimental validation, hardware integration, or manufacturing readiness at any level.

(c) The smallest next test. The smallest experiment that would de-risk or falsify the physical content of the p -adic loss law is a transmon photon-loss measurement: prepare a transmon resonator with a coherent state of photon number k for a sequence of k values spanning several 5-smooth and non-5-smooth integers (e.g. $k \in \{8, 9, 11, 12, 16, 25, 27\}$, chosen so that $\text{ord}_2, \text{ord}_3, \text{ord}_5$ vary independently), measure the photon loss rate for each, and test whether the loss depends on the valuations $\text{ord}_p(k)$ at all. The outcome that refutes the law is a loss rate independent of $\text{ord}_p(k)$ across the sequence — which is the expected outcome under conventional transmon decoherence theory, where loss depends on k through linear or thermal mechanisms, not through prime factorization. This asymmetry is the point: the test is cheap, the null is well understood, and a null outcome cleanly kills the law. The parallel smallest test for the entropy law is a subsystem tomography of a small superconducting or trapped-ion register arranged with tree connectivity, measuring the von Neumann entropy of full subtrees and comparing against $S = |\log_2 L|$; a measured entropy scaling differently (e.g. area-law) refutes the physical claim while leaving the definitional value untouched.

(d) Binding constraints. For the transmon test, the binding constraint is measurement fidelity: distinguishing a valuation-dependent loss law from a conventional k -linear one requires photon-number-resolved loss measurements at the percent level, which pushes against current transmon readout capabilities; this is stated as an open experimental question, since no hardware specification is derived or cited here. For the mass-ratio program, the binding constraint is not experimental but statistical: the density of $\mathcal{P}$ means the discriminating quantity is the size of the exponent triples, which costs no hardware but requires the search-corrected null of Section 6 to be built. For the entropy test, the binding constraint is register size: $S = |\log_2 L|$ is trivially satisfied by maximal mixing at any L , so the test requires a register whose subtree entropies are non-trivially prepared, which is a state-preparation constraint. Regulation, cost, and scale-up constraints do not bind at TRL 1–2.

(e) Order-of-magnitude bounds. No cost or performance projection is offered: every candidate figure (transmon coherence parameters, readout fidelity, qubit counts) would require hardware specifications neither derived in Section 4 nor stated in the sources, and the paper declines to invent them.

(f) Risks and misuse paths. The direct misuse path is numerology marketed as derivation: the density of $\mathcal{P}$ guarantees near-hits to any target number, so the semigroup can be made to "explain" any measured constant after the fact. This paper's own Section 4.4 is the cautionary instance — a 5.73% near-hit (24 vs. 22.7) that a careless reader could still cite as support. The corrective discipline is mandatory exponent-triple reporting and search-corrected nulls; without them, the program is unfalsifiable in practice even though each individual fit is falsifiable in principle.

8. Conclusion

We formalized the joint Bruhat–Tits tree conjecture — that the diagonal embedding of $\mathcal{T}_2 \times \mathcal{T}_3 \times \mathcal{T}_5$ into $\mathbb{R}_+$, with image the Pythagorean semigroup $\mathcal{P} = \{2^a 3^b 5^c\}$, underlies SM mass ratios, QEC lattice parameters, and Efimov scaling — and evaluated its proposed consequences by exact arithmetic. Three findings survive reconciliation. First, the headline conditional probability for eleven simultaneous 2% mass-ratio fits under the stated uniform null is $(4 \times 10^{-3})^{11} = 4.194304 \times 10^{-27}$, correcting an erroneous 4.2×10^{-30} ; the number remains conditional on a per-ratio deviation table that this paper does not compute, and the density of $\mathcal{P}$ in $\mathbb{R}_+$ means any confirmatory analysis must control for search freedom over exponent triples. Second, the harmonic-mean Efimov correspondence is falsified: $\lambda_{\text{tree}} = 3/(1/\ln 2 + 1/\ln 3 + 1/\ln 5) = 1.008652$ lies a factor of 22.51 from the assumed empirical $\lambda \approx 22.7$, and the residual near-hit 24 versus 22.7 carries no evidential weight. Third, the definitional laws $\ell_p(k) = p^{-\text{ord}_p(k)}$ (giving $\ell_2(12) = 0.25$ and product $1/12$) and $S = |\log_2 L|$ (giving $S = 10$ bits at $L = 1024$) are exact but physically untested; their smallest falsification tests are a photon-number-resolved transmon loss measurement and a subtree-entropy tomography, both of which are expected to return null outcomes under conventional theory. The gauge-group pillar is blocked at the level asserted by a cardinality obstruction and remains open. The work stands at TRL 1–2: its contribution is not a confirmed structure but a corrected arithmetic, a documented falsification, and a sharpened protocol that any future evaluation of the conjecture must satisfy.

Appendix A. Divergence Report

The three constituent drafts diverged on the following points; each divergence is documented here with the convention adopted in the main text.

- **D1: Joint fit probability.** Draft A reported 4.2×10^{-30} for eleven simultaneous 2% fits; Draft B's protocol implies $(4 \times 10^{-3})^{11}$. The disagreement is an arithmetic error in Draft A's evaluation of the power of ten, not a convention difference: $4^{11} = 4,194,304$ forces the exponent -27 . *Resolution:* the main text uses 4.194304×10^{-27} (Section 4.1).
- **D2: Efimov correspondence.** Draft A quoted $\lambda_{\text{tree}} = 24$ (the Pythagorean vertex $(3, 1, 0)$, within 5.73% of 22.7); Draft B derived the harmonic mean of the log-periods as the correspondence. The two conventions are mathematically distinct objects. *Resolution:* the main text evaluates both (Section 4.4); the harmonic mean fails by a factor 22.51 and the vertex value is retained only as a labeled coincidence, per the density objection.
- **D3: Statistical framing.** Draft A used the binary/uniform null (Framing A); Draft B used a Gaussian chi-squared framing (Framing B) with a per-ratio deviation table. *Resolution:* both framings are stated (Section 3.2); Framing A is executed conditionally, Framing B is carried as a specified method because Draft B's table was not reproducible from stated inputs (its own summation produced a negative contribution to a sum of squares).
- **D4: Photon-loss functional form.** Draft A used the multiplicative law $\ell_p(k) = p^{-\text{ord}_p(k)}$; Draft B used a linear law $\Delta_{\text{photon}}(k) = \varepsilon \text{ord}_p(k)$ with $\varepsilon \approx 0.01$ and appended a transmon measurement table. The forms are mathematically distinct and make different predictions. *Resolution:* the main text carries the multiplicative law because it is exactly computable from its definition; Draft B's measurement table is dropped as unreproducible (Section 3.3).
- **D5: Gauge group.** One draft asserted the automorphism subgroup is isomorphic to $SU(3) \times SU(2) \times U(1)$; no draft supplied a derivation. *Resolution:* carried strictly as a conjecture with the cardinality obstruction stated (Sections 3.6, 4.5).

Appendix B. Claim Attribution

Claim	Drafts	Status
C1: Joint fit probability $(4 \times 10^{-3})^{11}$	A (erroneous value), B (protocol)	CONVERGENT in protocol; value corrected in reconciliation
C2: Harmonic-mean Efimov correspondence	B	SINGLE; falsified by reconciliation arithmetic (Section 4.4)
C3: Vertex value $\lambda_{\text{tree}} = 24$, within 5.7%	A	SINGLE; retained as labeled coincidence (Section 4.4)
C4: Per-ratio fit table for Framing B	B	SINGLE; dropped, not reproducible from stated inputs
C5: Multiplicative p -adic loss law	A	SINGLE; carried as definition (Section 3.3)
C6: Linear loss law with $\varepsilon \approx 0.01$ and transmon table	B	SINGLE; dropped as unreproducible; divergence documented (D4)
C7: Entropy law $S = \ \log_2 L\ $, $S = 10$ at $L = 1024$	A, B	CONVERGENT; carried as definitional (Section 4.3)
C8: Gauge group $SU(3) \times SU(2) \times U(1)$ from tree automorphisms	A	SINGLE; carried as conjecture with obstruction (Section 4.5)
C9: Uniform null on $[0, 10]$ with absolute window 0.04	A, B	CONVERGENT in Framing A; convention caveat stated (Section 4.1)

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.5-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.

- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 10 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Feature Unification in TAG Derivation Trees. arXiv:0804.4584v1. <https://arxiv.org/abs/0804.4584v1> [2] Cache-Oblivious Representation of B-Tree Structures. arXiv:2209.09166v2. <https://arxiv.org/abs/2209.09166v2> [3] Spanning structures and universality in sparse hypergraphs. arXiv:1504.02243v2. <https://arxiv.org/abs/1504.02243v2> [4] Universal Scaling of Electron Transmission for Nearly Ballistic and Quantum Dragon Nanodevices. arXiv:2007.01051v2. <https://arxiv.org/abs/2007.01051v2> [5] Universal Mass Scale for Bosonic Fields in Multi-Brane Worlds. arXiv:2009.07197v4. <https://arxiv.org/abs/2009.07197v4> [6] Learning Enables Exponential-to-Polynomial Sampling Overhead Scaling in Quantum Divide-and-Conquer for Tree-Structured Circuits. arXiv:2512.19623v2. <https://arxiv.org/abs/2512.19623v2> [7] Impact of the Ce $4f$ states in the electronic structure of the intermediate-valence superconductor CeIr₃. arXiv:2311.03640v1. <https://arxiv.org/abs/2311.03640v1> [8] A Universal Bijection for Catalan Structures. arXiv:1808.09078v2. <https://arxiv.org/abs/1808.09078v2> [10] QNFO: The Adelic Cross-Domain Program: From the Fine-Structure Constant to the Standard Model Mass Spectrum via Bruhat-Tits Trees