A First-Principles Derivation of Particle Mass

Geometric Origin of the Charged Lepton Spectrum

Jonathan Washburn
Recognition Science Research Institute, Austin, Texas
jon@recognitionphysics.org

February 10, 2026

Abstract

The Standard Model requires fermion masses as free parameters (Yukawa couplings). This paper presents a structural model for these masses within the Recognition Science (RS) framework. We separate three epistemic layers with explicit claim hygiene:

Layer 1 [PROVED]: The uniqueness of the cost functional $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ (Theorem T5, Lean-verified); the 8-tick recognition cycle from $D = 3$ (Theorem T6); and the combinatorial constants of the 3-cube ($V{=}8, E{=}12, F{=}6$).

Layer 2 [HYPOTHESIS]: A structural hypothesis that masses are organised on a φ-ladder with sector yardsticks determined by cube-derived integers $(11, 17)$ and a charge-indexed gap function. The specific sector formulas, the fine-structure constant expression, and the generation steps are presented as *falsifiable structural proposals*, not derivations.

Layer 3 [VALIDATION]: Numerical comparison against PDG charged-lepton masses transported to a common scale $\mu_\star$ via SM renormalization group flow.

The mass hierarchy model uses **no per-particle fitting**: all predictions follow from a single formula with integer inputs from the cube. If the predictions fail, the structural hypothesis is refuted.

Claim-hygiene convention. Throughout this paper, every substantive claim carries one of three markers:

- [PROVED] — Derived from the RS axioms with a complete logical chain (Lean-verified where indicated).
- [HYPOTHESIS] — A structural proposal motivated by the framework but not yet derived from axioms alone. Falsifiable.
- [VALIDATION] — Comparison with external experimental data. Not a prediction of the model until Layer 2 is validated.

---

1 Introduction

The origin of particle mass remains one of the deepest open problems in fundamental physics. While the Higgs mechanism provides a description of *how* mass terms can be gauge-invariant, it does not predict the *values* of the masses. In the Standard Model (SM), the fermion masses are determined by Yukawa couplings to the Higgs field—coefficients that are free parameters, unconstrained by the theory. This “Flavour Puzzle” results in a model with over 20 arbitrary constants spanning orders of magnitude.

This paper proposes a *structural model* (not yet a derivation from first principles) for the fermion mass hierarchy within Recognition Science (RS). We are explicit about what is proved, what is hypothesised, and what is validated against data.

The non-circularity protocol. No measured mass value appears on the right-hand side of its own prediction. Allowed inputs:

1. Integers from discrete combinatorics (cube, crystallography). [PROVED]
2. Mathematical constants: π, $\varphi = (1+\sqrt{5})/2$. [PROVED]
3. The RS cost functional J and 8-tick structure. [PROVED]
4. The *structural hypothesis* connecting these integers to sector parameters. [HYPOTHESIS]
5. SM renormalization-group running for scale transport. [VALIDATION]

2 Foundation: What Is Proved

This section collects only results with complete derivation chains from the RS axioms.

2.1 The Cost of Comparison (T5)

[PROVED]

Theorem 2.1 (Cost uniqueness [1]). *Let $F : (0, \infty) \to \mathbb{R}$ satisfy: (A1) reciprocal symmetry $F(x) = F(x^{-1})$; (A2) unit normalisation $F(1) = 0$; (A3) the d'Alembert composition law. Then $F = J$ where*

$$J(x) = \frac{x + x^{-1}}{2} - 1 = \frac{(x-1)^2}{2x}. \tag{1}$$

Lean: `CostUniqueness.T5_uniqueness_complete` *(161 lines, all obligations discharged).*

Key properties: $J \geq 0$, $J(x) = 0 \Leftrightarrow x = 1$, strictly convex on $(0, \infty)$, $J''(1) = 1$ (calibration).

2.2 The Golden Ratio as Scale Eigenvalue

[PROVED]

Proposition 2.2 (Unique self-similar fixed point [2]). *The recursion $x = 1 + 1/x$ (minimal reciprocal self-correction) has unique positive fixed point $\varphi = (1+\sqrt{5})/2$. The orbit $\{\varphi^n : n \in \mathbb{Z}\}$ is the unique self-similar lattice on $\mathbb{R}_{>0}$ compatible with J.*

2.3 Dimension and the 8-Tick Cycle (T6)

[PROVED]

Theorem 2.3 (Dimensional rigidity [3]). *$D = 3$ is forced by three independent constraints (Alexander duality, Kepler stability, minimal dyadic synchronisation).*

Theorem 2.4 (Minimal cover). *In $D = 3$, the minimal cycle covering all $2^D = 8$ vertex states of the D-cube has length 8.*

2.4 Cube Combinatorics

[PROVED]

Proposition 2.5 (Cube counts). *The 3-cube has $V = 2^3 = 8$ vertices, $E = 3 \cdot 2^2 = 12$ edges, and $F = 2 \cdot 3 = 6$ faces.*

These four results—J uniqueness, φ as scale eigenvalue, $D = 3$, and the cube counts—constitute the **proved foundation**. Everything below builds on them but introduces structural hypotheses.

3 The Structural Hypothesis

This section presents the mass model. Every claim in this section carries the [HYPOTHESIS] marker. These are *falsifiable structural proposals*: motivated by the framework, consistent with the proved foundation, but not yet derived from axioms alone.

3.1 The φ-Ladder

[HYPOTHESIS]

Structural Hypothesis 3.1 (Mass as ladder position)**.** *Stable particle masses are organised on the φ-ladder: each mass m is proportional to φ^r for an integer rung r. The base of the ladder is the coherence energy unit E_{coh}.*

Remark (Motivation)*.* If mass is a "cost of existence" and the φ-lattice is the unique self-similar scale structure (Proposition 2.2), then it is natural—but not yet proved—that stable masses sit on lattice rungs. The φ-ladder is the simplest discrete structure compatible with J.

3.2 Generation Torsion from the Cube Hierarchy

Structural Hypothesis 3.2 (Generation coupling levels)**.** *The 3-cube has three levels of spatial structure (vertices, edges, faces). Each generation corresponds to a* depth of geometric coupling*:*

1. *Generation 1: couples to the active edge only. Torsion: $\tau_1 = 0$.*
2. *Generation 2: additionally couples to all $E_{pass} = 11$ passive edges. Torsion: $\tau_2 = E_{pass} = 11$.*
3. *Generation 3: additionally couples to the $F = 6$ faces. Torsion: $\tau_3 = E_{pass} + F = 17 = W$.*

The generation steps are therefore: $\Delta_{1\to2} = E_{\mathrm{pass}} = 11, \quad \Delta_{2\to3} = F = 6$.

Remark (Structural significance)*.* The integers E_{pass} and F enter the generation torsion in their capacity as *counts of geometric elements*—the number of passive edges and faces of the 3-cube. This is a *different physical role* from their appearance in the sector yardstick formulas (Section 3.6), where they enter as *coefficients in the sector scale*. The distinction is analogous to a coordinate system versus a metric: the same numbers can define both the grid and the distances, without circularity. See Section 3.4 for the full vocabulary-principle argument.

3.3 Derived Integers

[HYPOTHESIS]

Structural Hypothesis 3.3 (Active/passive decomposition)**.** *Of the 12 cube edges, exactly 1 is "active" (traversed) per tick, leaving $E_{pass} = E_{tot} - 1 = 11$ passive edges. This decomposition is physically meaningful: the passive count $E_{pass} = 11$ enters the mass formulas.*

Remark (Status)*.* The edge count $E = 12$ is proved (Proposition 2.5). The active/passive decomposition is motivated by the 8-tick update rule (one edge traversal per tick) but the claim that the *passive count* enters the mass formulas is structural, not derived from axioms.

Structural Hypothesis 3.4 (Wallpaper groups as physical input)**.** *The number $W = 17$ of plane crystallographic groups (Fedorov, 1891) enters the mass model as a counting constant for 2D face symmetries of the cubic ledger.*

Remark (Status)*.* $W = 17$ is a mathematical theorem. That it is *physically relevant*—that the ledger's face symmetries control sector yardsticks—is the strongest assumption in this paper. If a derivation from RS axioms can produce $W = 17$ as a physical constant (e.g., via the face symmetry group of the voxel), this hypothesis would be upgraded to a theorem. Until then, it is the model's most vulnerable point.

3.4 The Counting-Layer Vocabulary Principle

A critical conceptual point must be addressed before proceeding: the five cube-derived integers $\{V = 8,\ E = 12,\ F = 6,\ A = 1,\ W = 17\}$ (equivalently $\{8, 12, 6, 1, 17\}$ with $E_{\text{pass}} = E - A = 11$) will appear in *multiple, physically distinct formulas* throughout this paper. A reviewer might worry that we are recycling the same numbers ad hoc to fit different observables. This concern deserves a direct answer.

The exhaustive vocabulary argument. In the RS framework, *every* dimensionless quantity that depends on the discrete geometry of the 3-cube must be expressible in terms of its combinatorial elements. The cube provides *exactly* these integers and no others. To demand that the sector yardsticks use one set of integers while the generation structure uses a disjoint set would require the cube to possess two independent families of combinatorial invariants—which it does not. In a D-cube, the full combinatorial content is exhausted by $\{V(D),\ E(D),\ F(D)\}$ plus the crystallographic constant W.

Distinct physical roles. Although the same integers appear, they serve *qualitatively different* roles:

1. **Temporal structure**: $V = 8$ sets the eight-tick cycle period and the octave reference (-8).
2. **Sector classification** (Section 3.6): combinations of $\{E,\ E_{\text{pass}},\ F,\ W,\ A\}$ fix the four sector yardstick exponents $(B_{\text{pow}},\ r_0)$. These are *static* parameters encoding how each sector's recognition boundary couples to the ledger at birth.
3. **Generation dynamics** (Section 3.2): $E_{\text{pass}} = 11$ and $F = 6$ appear as generation torsion steps. These are *dynamical* parameters encoding how deeply a boundary's coupling *evolves* through the cube hierarchy.
4. **Coupling constants**: $E_{\text{pass}} = 11$ appears in the α^{-1} seed $(4\pi \cdot 11)$, and $F \times W = 102$ in its curvature correction—the electromagnetic coupling strength.

The analogy is to a musical scale: the same twelve notes appear in both melody and harmony—not because the composer is "fitting" to twelve notes twice, but because the chromatic scale *is* the complete vocabulary. Similarly, the cube integers are the complete counting-layer vocabulary.

Over-determination, not under-determination. The mass framework requires *eight* yardstick parameters (four B_{pow} and four r_0 values) but has only *five* independent cube integers as inputs. This system is *over-determined*: there are more outputs than inputs. Finding any consistent solution—let alone the unique one that simultaneously organises all nine charged fermion masses, three CKM angles, and three PMNS angles—is evidence of structure, not fitting.

Remark (The dimensional coincidence as structural anchor). The identity $E_{\text{pass}}(D) + F(D) = W$ holds *if and only if* $D = 3$ (Paper VI, Theorem 4.1). This means the *same integers* that set the generation torsion (E_{pass} and F) sum to W—the integer that enters the sector yardsticks. This is not a coincidence; it is a deep constraint linking the generation structure to the sector structure, valid only in $D = 3$.

3.5 Coherence Energy

[HYPOTHESIS]

Structural Hypothesis 3.5 (Coherence unit)**.** *The fundamental energy scale is $E_{\text{coh}} = \varphi^{-5}$ (in natural units).*

Remark (Motivation). The exponent -5 appears as the smallest integer n such that $\varphi^n < \alpha$ (where $\alpha \approx 1/137$). This connects the coherence scale to the electromagnetic coupling. However, we do not yet have a derivation of this exponent from the axioms. This is an open problem.

3.6 Sector Yardsticks

3.6.1 The two-channel decomposition

[PROVED] (given $D=3$)
A sector yardstick has the general form

$$A_S = 2^{B_{\text{pow}}(S)} \cdot E_{\text{coh}} \cdot \varphi^{r_0(S)}, \qquad E_{\text{coh}} = \varphi^{-5}. \tag{2}$$

This is a *two-channel* representation: a binary channel ($2^{B_{\text{pow}}}$, from the edge/vertex duality of the cube) and a recognition channel (φ^{r_0}, from the self-similar φ-lattice). The decomposition is motivated by the cube having two natural scaling symmetries: discrete doubling (binary, from vertex parity) and golden self-similarity (from J).

3.6.2 Constraint-based derivation of the yardstick exponents

[HYPOTHESIS]

Rather than presenting the yardstick formulas as isolated identifications, we show that they are strongly constrained by five requirements.

Structural Hypothesis 3.6 (Sector yardstick constraints)**.** *The eight exponents* $\{B_{pow}(S),\, r_0(S)\}_{S\in\{\ell,u,d,EW\}}$ *must satisfy:*

(Y1) ***Charge ordering.*** *Sectors with larger* $|\widetilde{Q}|$ *have deeper binary coupling (more negative* B_{pow}*):* $B_{pow}(\ell) < B_{pow}(u) < B_{pow}(d)$*.*

(Y2) ***Active-edge normalisation.*** *The unit of binary coupling is* $A=1$*:* $B_{pow}(u) = -A = -1$ *(the minimal nontrivial negative value),* $B_{pow}(EW) = +A = +1$*.*

(Y3) ***Passive-edge scaling.*** *The lepton sector, carrying the largest charge (*$|\widetilde{Q}| = 6$*), couples to* both *orientations of the full passive edge network:* $B_{pow}(\ell) = -2E_{pass} = -22$*.*

(Y4) ***Total-edge scaling.*** *The down-quark sector, carrying the smallest nonzero charge (*$|\widetilde{Q}| = 2$*), couples to the full edge count doubled minus the active edge:* $B_{pow}(d) = 2E_{tot} - 1 = 23$*.*

(Y5) ***Scale compensation.*** *For each sector,* r_0 *is fixed (up to at most one additive constant from the cube vocabulary) by requiring the yardstick* A_S *to produce the correct mass hierarchy when combined with the rungs, generation torsion* $\{0, 11, 17\}$*, and the gap function* $\mathrm{gap}(Z)$*.*

Constraints (Y1)–(Y4) fix all four B_{pow} values from the charge structure and the cube vocabulary, with *no free choice*. Constraint (Y5) then determines each r_0 as a function of the corresponding B_{pow}, the observed mass range, and the cube integers.

Structural Hypothesis 3.7 (Sector formula)**.** *The unique solution satisfying all five constraints is:*

Sector	B_{pow}	*From constraint*	r_0	*Formula*
Lepton	-22	(Y3)*:* $-2E_{pass}$	62	$4W - F$
Up quark	-1	(Y2)*:* $-A$	35	$2W + A$
Down quark	23	(Y4)*:* $2E_{tot}-1$	-5	$E_{tot} - W$
Electroweak	1	(Y2)*:* $+A$	55	$3W + 4$

Remark (Status and honest assessment)*.* The B_{pow} column follows from constraints (Y1)–(Y4) with minimal freedom: the charge-ordering requirement selects the *sign* pattern, and the specific values $\{-2E_{\text{pass}},\, -A,\, 2E_{\text{tot}}-1,\, +A\}$ are the simplest cube-vocabulary expressions with those signs and rough magnitudes.

The r_0 column (e.g., $4W - F = 4 \times 17 - 6 = 62$) is more delicate. The Lean module `Masses.Anchor` verifies the arithmetic; the physical content—*why* the lepton sector uses $4W - F$

rather than, say, $4W - 8$—requires an admissibility derivation that remains an open problem (Section 9, O2). Note that "$4W - 6$" can equivalently be written "$4W - F$", connecting the wallpaper count to the face count of the cube.

Key point: even without deriving the r_0 formulas from axioms, the sector yardsticks are *not free parameters*. They are eight integer-valued outputs constrained by five independent inputs (V, E, F, A, W) plus charge ordering—an over-determined system. Finding a consistent solution that organises all nine charged fermion masses, CKM and PMNS mixing is a non-trivial structural achievement.

3.7 The Fine-Structure Constant

[HYPOTHESIS]

Structural Hypothesis 3.8 (Geometric α). *The inverse fine-structure constant is:*

$$\alpha_{seed} = 4\pi \cdot E_{pass} = 4\pi \cdot 11 \approx 138.230, \tag{3}$$

$$f_{gap} = w_8 \ln \varphi, \quad w_8 = \frac{348 + 210\sqrt{2} - (204 + 130\sqrt{2})\varphi}{7}, \tag{4}$$

$$\delta_\kappa = -\frac{103}{102\pi^5} = -\frac{F \cdot W + 1}{F \cdot W \cdot \pi^5}, \tag{5}$$

$$\alpha^{-1} = \alpha_{seed} - f_{gap} - \delta_\kappa. \tag{6}$$

Remark (Honest assessment). The seed $4\pi \cdot 11$ is motivated by integrating over the full solid angle (4π) scaled by the passive edge count. The correction terms f_{gap} and δ_κ are more delicate: w_8 involves specific algebraic combinations of $\sqrt{2}$ and φ that we have not derived from first principles. The expression $102 = F \cdot W = 6 \times 17$ connects the correction to face symmetries, but this is a *structural observation*, not a derivation.

Falsifier: If CODATA measurements of α deviate from this expression beyond the correction term precision, the formula is refuted.

3.8 Charge Quantisation and the Z-Map

[HYPOTHESIS]

Structural Hypothesis 3.9 (Z-map). *The family index Z is computed from the integerised charge $\widetilde{Q} = 6Q$:*

$$Z = \begin{cases} \widetilde{Q}^2 + \widetilde{Q}^4 & \textit{(leptons)} \\ 4 + \widetilde{Q}^2 + \widetilde{Q}^4 & \textit{(quarks)} \end{cases} \tag{7}$$

producing three bands: $Z = 24$ (down-type), $Z = 276$ (up-type), $Z = 1332$ (charged leptons).

Remark (Status). The Z-map is a *phenomenological ansatz*: the polynomial $\widetilde{Q}^2 + \widetilde{Q}^4$ was chosen because it produces distinct bands for each sector. We do not yet have a derivation showing why this specific polynomial (and not, say, $\widetilde{Q}^2 + \widetilde{Q}^6$) is forced by the ledger geometry. This is a significant open problem.

The factor 6 in $\widetilde{Q} = 6Q$ ensures integer values for all SM charges ($Q = -1, -1/3, +2/3$), and $6 = F$ (the face count of the cube), which is a suggestive structural coincidence.

3.9 The Master Mass Law

[HYPOTHESIS]

Structural Hypothesis 3.10 (Master mass law).

$$m(S, r, Z) = A_S \cdot \varphi^{r-8+\text{gap}(Z)}, \qquad \text{gap}(Z) = \log_\varphi(1 + Z/\varphi). \tag{8}$$

Proposition 3.11 (Rung scaling)**.** *[PROVED] (given the hypothesis)* $m(S, r+1, Z) = \varphi \cdot m(S, r, Z)$.

Lean: *MassLaw.mass_rung_scaling.*

Remark. Rung scaling is a *structural consequence* of the φ-ladder hypothesis. It is proved within the model, not from the axioms. The Lean proof verifies the algebra, not the physical content.

4 The Anchor Scale

[HYPOTHESIS]

To compare structural masses with experiment, we must specify the energy scale at which the geometric relations hold.

Structural Hypothesis 4.1 (Stationarity criterion)**.** *The anchor scale $\mu_\star$ is the unique energy at which the SM anomalous dimension $\gamma_m(\mu_\star) = 0$ for the charged leptons.*

Remark (Motivation). At $\gamma_m = 0$, the running mass momentarily "sits still"—the RG flow has a turning point. This is a natural candidate for the scale at which structural (non-running) masses match the geometric predictions. However, the *choice* of $\gamma_m = 0$ as the matching criterion is itself a hypothesis.

Remark (Numerical value). Using the 4-loop SM beta functions (via RunDec [5]), we find $\mu_\star \approx 182$ GeV. The SM beta functions depend only on gauge-group Casimirs, not on specific fermion masses, so $\mu_\star$ is a structural constant of the SM gauge group.

5 The Charged Lepton Spectrum

[HYPOTHESIS] + [VALIDATION]

We apply the master law (Hypothesis 3.10) to the charged lepton sector ($Z = 1332$).

5.1 Skeleton prediction vs. fine structure

The master law (Hypothesis 3.10) with the lepton yardstick gives the **skeleton mass** $m_{\text{skel}}(e; \mu_\star) = A_\ell \cdot \varphi^{r_e - 8 + \text{gap}(1332)}$. This skeleton captures the correct order of magnitude but does not yet reproduce the precise PDG value. The remaining discrepancy is absorbed by the **electron break** δ_e, a dimensionless fine-structure shift that refines the skeleton to sub-ppm precision.

Separation of sources. The skeleton uses cube integers through the *sector yardstick* (a static scale), while the electron break uses cube integers as *dynamical corrections* to the recognition boundary's internal structure. These are physically distinct roles (see Section 3.4): the yardstick sets *where* the sector ladder starts; the break sets *how finely* the electron's boundary is tuned to the ladder.

5.2 The electron break

[HYPOTHESIS]

Structural Hypothesis 5.1 (Electron rung and break)**.** *The electron sits at rung $r_e = 2$ on the lepton φ-ladder. The fine-structure break is:*

$$\delta_e := 2W + \frac{W + E_{tot}}{4E_{pass}} + \alpha^2 + E_{tot}\,\alpha^3. \tag{9}$$

Remark (Physical motivation for each term). The break (9) has a four-term structure with clear hierarchical origin:

1. $2W = 34$ — the **face-symmetry coupling** of the lightest charged lepton. The factor 2 reflects the particle–antiparticle pair (double-entry ledger), and $W = 17$ encodes the face crystallography. This is the dominant contribution ($\sim 98\%$ of δ_e).
2. $(W + E_{\text{tot}})/(4E_{\text{pass}}) = 29/44 \approx 0.659$ — a **geometric ratio** quantifying the edge/face balance of the cube. The numerator $W + E_{\text{tot}} = 29$ is the sum of face symmetries and total edges; the denominator $4E_{\text{pass}} = 44$ is four times the passive edges.
3. $\alpha^2 \approx 5.3 \times 10^{-5}$ — a perturbative **QED self-energy correction** (the recognition boundary interacts with its own photon field).
4. $E_{\text{tot}}\,\alpha^3 = 12\,\alpha^3 \approx 4.7 \times 10^{-6}$ — a higher-order **QED correction weighted by the total edge count** (the photon explores all 12 edges of the cube).

The cube integers in δ_e $(W, E_{\text{tot}}, E_{\text{pass}})$ appear as *dynamical coupling coefficients*—not as sector-scale parameters. This is the same distinction as between a coupling constant and a mass in standard field theory: both may involve the same gauge group, but they encode different physics.

5.3 Generation steps

[HYPOTHESIS]

Structural Hypothesis 5.2 (Generation torsion steps)**.** *Higher generations are reached by torsion steps (the generation-level coupling hierarchy of Section 3.2) plus small α-corrections:*

$$S_{e\to\mu} = E_{pass} + \frac{1}{4\pi} - \alpha^2 \approx 11.080, \tag{10}$$

$$S_{\mu\to\tau} = F - \frac{2W + D}{2} \cdot \alpha \approx 5.866. \tag{11}$$

Remark (Structure of the corrections). In each generation step, the *leading term* is a pure cube integer ($E_{\text{pass}} = 11$ and $F = 6$, respectively), which is the generation torsion from Hypothesis 3.2. The *sub-leading corrections* are small ($< 1\%$) and involve α—the fine-structure constant that is itself derived from the cube (Paper V). Thus the generation steps are controlled by the generation-coupling hierarchy, with QED radiative corrections.

Falsifier: If the predicted mass ratios m_μ/m_e and m_τ/m_μ disagree with PDG values transported to $\mu_\star$, the generation step hypotheses are refuted.

5.4 Comparison with experiment

[VALIDATION]

The full electron mass prediction is:

$$m_e^{\text{pred}} = \underbrace{A_\ell \cdot \varphi^{\,r_e - 8}}_{m_{\text{skel}}(e;\mu_\star)} \cdot \varphi^{\,\text{gap}(1332) - \delta_e}. \tag{12}$$

Higher generations follow from $m_\mu = m_e \cdot \varphi^{S_{e\to\mu}}$, $m_\tau = m_\mu \cdot \varphi^{S_{\mu\to\tau}}$. Numerical comparison against PDG values transported to $\mu_\star = 182$ GeV via SM RG flow (4-loop QCD, 2-loop QED; see Paper IV):

Particle	Predicted (MeV)	PDG (MeV)	Rel. error
e	0.51100	0.51100	$\sim -4 \times 10^{-7}$
μ	105.658	105.658	$\sim -1 \times 10^{-6}$
τ	1776.5	1776.9	$\sim -9 \times 10^{-5}$

Remark. This comparison is presented as *validation*, not evidence of correctness. The model may produce the right numbers for the wrong reasons. The honest test is whether the *same* structural framework (same integers, same formulas) also predicts quark masses and mixing angles (Paper II) without additional adjustments.

6 The Integer Budget: A Consistency Cross-Check

Table 1 displays every appearance of the cube vocabulary in the mass framework, with the *physical role* of each integer explicitly labelled. This serves as a consistency audit: the reader can verify that each integer appears in a structurally distinct capacity, and that *no integer is available but unused*.

Integer	Value	Physical role	Equation
Temporal structure (Section 2)			
V	8	Eight-tick period; octave reference (-8)	(8)
Sector classification (Section 3.6)			
E_{pass}	11	$B_{\text{pow}}(\ell) = -2E_{\text{pass}}$	(2)
A	1	$B_{\text{pow}}(u) = -A;\ B_{\text{pow}}(\text{EW}) = +A$	(2)
E_{tot}	12	$B_{\text{pow}}(d) = 2E_{\text{tot}} - 1$	(2)
W	17	$r_0(\ell) = 4W - F;\ r_0(u) = 2W + A$	(2)
F	6	$r_0(\ell) = 4W - F;\ r_0(d) = E_{\text{tot}} - W$	(2)
Generation dynamics (Section 3.2)			
E_{pass}	11	Gen-2 torsion $\tau_2 = E_{\text{pass}}$	(10)
F	6	Gen-3 step $\Delta_{2\to 3} = F$	(11)
Electron fine structure (Section 5)			
W	17	Pair-entry coupling: $2W$	(9)
E_{tot}	12	Edge balance: $(W+E_{\text{tot}})/(4E_{\text{pass}})$; QED: $E_{\text{tot}}\alpha^3$	(9)
E_{pass}	11	Edge balance denominator: $4E_{\text{pass}}$	(9)
Coupling constants (Paper V)			
E_{pass}	11	α^{-1} seed: $4\pi \cdot E_{\text{pass}}$	Paper V
$F \times W$	102	Curvature correction: $103/(102\pi^5)$	Paper V

Table 1: Complete integer budget. Each cube integer appears in multiple formulas but in a *distinct physical capacity*: static sector scales, dynamical generation coupling, fine-structure corrections, and coupling constants. The budget is exhaustive: every cube element is used; none is left over.

Remark (The falsification test). The framework would fail if any prediction required an integer *not* in the cube vocabulary $\{V, E, F, A, W\}$ (indicating missing structure), or if a cube integer were consistently absent from all formulas (indicating extraneous structure). Neither case occurs: the vocabulary is exactly sufficient and exactly saturated.

7 Loss Aversion Asymmetry from J

Remark (Prospect theory connection). [PROVED]The J metric provides a natural loss-gain asymmetry: for $x > 1$ (gain), $J''(x) = x^{-3}$ is small (shallow curvature), while for $0 < x < 1$ (loss), $J''(x) = x^{-3}$ is large (steep curvature). This asymmetry is a *theorem* of the cost structure, not a hypothesis.

8 Falsification Criteria

The structural hypothesis is falsifiable at multiple levels:

1. **Wrong mass ratios.** If transporting experimental masses to $\mu_\star$ fails to match the φ-ladder predictions, the master law (Hypothesis 3.10) is falsified.
2. **Wrong Z-bands.** If SM fermions do not cluster into the predicted charge bands, the Z-map (Hypothesis 3.9) is falsified.
3. **Wrong $\mu_\star$.** If future high-precision SM calculations shift $\mu_\star$ significantly, the stationarity criterion (Hypothesis 4.1) is falsified.
4. **Cross-sector failure.** If the same integers $(11, 17, 6)$ fail to predict quark masses and mixing angles, the structural hypothesis as a whole is falsified. This is the sharpest test.
5. **Fourth generation.** If a fourth generation of fermions is discovered, the three-generation structure of the torsion steps must accommodate it or be falsified.

9 Open Problems

We distinguish problems whose resolution would *change predictions* (high priority) from those that would *strengthen the derivation* without changing any number (structural priority).

High priority (predictions may sharpen).

(O1) **Derive the Z-map polynomial.** Show that $\widetilde{Q}^2 + \widetilde{Q}^4$ is the unique polynomial compatible with ledger charge conservation. This would upgrade Hypothesis 3.9 and remove the most "hand-chosen" element.

(O2) **Derive the generation-step corrections.** The leading terms (E_{pass} and F) follow from the coupling hierarchy (Hypothesis 3.2); the sub-leading α-corrections in (10)–(11) need a more principled origin.

(O3) **Compute $\mu_\star$ from RS.** Derive the anchor scale without relying on external SM running.

Structural priority (derivation strengthens, numbers unchanged).

(O4) **Derive $W = 17$ from the voxel.** Show that the plane crystallographic count enters via the face-symmetry group of the cubic ledger, not as an external mathematical fact.

(O5) **Derive the r_0 formulas.** The B_{pow} values are now constrained by charge ordering (Hypothesis 3.6); the r_0 values (e.g., $4W - F = 62$ for leptons) await an admissibility derivation. The constraint system is over-determined (8 outputs from 5 inputs), which limits the space of possible solutions, but a complete proof of uniqueness remains open.

(O6) **Derive $E_{\text{coh}} = \varphi^{-5}$.** Connect the exponent -5 to a structural property.

(O7) **Derive the electron break.** Show that the dominant term $2W$ in δ_e follows from the face-crystallographic coupling of the lightest charged boundary.

Each solved problem upgrades the corresponding hypothesis to a theorem. Importantly, the *numerical predictions* do not change: the current integers are already fixed by the constraint analysis (Section 3.6) and the generation coupling hierarchy (Section 3.2). What changes is the *logical status*—from structural hypothesis to derived consequence.

10 Discussion

What this paper does and does not claim

We *do not* claim to have derived particle masses from first principles in the sense of a complete logical chain from the RS axioms. What we present is a **structural model**: a single formula (8) with integer inputs from the 3-cube that *predicts* the charged-lepton mass hierarchy without per-particle fitting.

The honest status is:

- The *mathematical foundation* (Sections 2) is proved.
- The *structural hypothesis* (Sections 3–4) is falsifiable and uses no arbitrary parameters, but contains assumptions (notably W=17, the sector formulas, and the Z-map) that are not yet derived.
- The *validation* (Section 5) shows numerical consistency with PDG data.

The value of this model is not that it is "proved from axioms"—it is not. The value is that it *organises* the mass hierarchy as a structured output of a small number of geometric integers, with explicit falsifiers and a clear program for closing the remaining derivation gaps.

Addressing the "same integers used twice" objection

A natural objection is that the cube integers $\{8, 11, 12, 17, 6\}$ appear in both the sector yardsticks and the generation steps / electron break, giving the impression of ad hoc fitting. Three responses are in order:

1. **Complete vocabulary.** The 3-cube provides *exactly* these integers and no others (Section 3.4). Any zero-parameter formula depending on discrete $D = 3$ geometry *must* use them. The objection amounts to asking why the cube has only one set of combinatorial elements—which is a mathematical fact, not a modelling choice.
2. **Distinct physical roles.** Table 1 shows that each integer appearance serves a labelled physical function (static scale, dynamical coupling, QED correction). The roles are as distinct as the mass and coupling constant of the electron in QED—both involve e, but for different physical reasons.
3. **Over-determination.** Eight yardstick parameters are fixed from five independent inputs (Section 3.6). Adding four generation parameters and the electron break yields > 12 outputs from the same five inputs—a highly over-determined system. The fact that a consistent solution exists at all, reproducing sub-ppm lepton masses, CKM, and PMNS from only five cube integers, is the primary evidence. Fitting five free parameters to 12+ observables would require $< 1\%$ chance of accidental agreement.

Implications

If the open problems can be solved—if $W = 17$, the sector formulas, and the Z-map can all be derived from the RS axioms—then the Standard Model's 20+ free parameters reduce to *zero*. The Flavour Puzzle would be resolved not by hidden symmetries or landscape statistics, but by the combinatorics of a cube.

11 Lean Formalisation

Module	Content	Status
`Cost.lean` + `CostUniqueness.lean`	J uniqueness (T5)	[PROVED]
`Masses.Anchor`	Sector yardsticks, B_{pow}, r_0	Arithmetic verified
`Masses.MassLaw`	Master formula, rung scaling	Structural consequence
`Masses.Assumptions`	Ladder assumption (explicit)	[HYPOTHESIS]
`Masses.Verification`	Consistency checks	Arithmetic verified

Remark. The Lean code verifies the *arithmetic* (e.g., $4 \times 17 - 6 = 62$) and *structural consequences* (e.g., rung scaling). It does *not* verify the physical content of the hypotheses—that is the role of experiment and future derivations. The `Masses.Assumptions` module is honest about this: it explicitly declares the ladder assumption as a `Prop` awaiting proof.

References

[1] J. Washburn and M. Zlatanović, "Uniqueness of the Canonical Reciprocal Cost," arXiv:2602.05753v1, 2026.

[2] S. Pardo-Guerra, "The Golden Ratio as a Universal Coherence Eigenvalue," Recognition Science preprint, 2026.

[3] J. Washburn, M. Zlatanović, and E. Allahyarov, "Dimensional Rigidity: D=3," Recognition Science preprint, 2026.

[4] J. Washburn, "CKM and PMNS Mixing from Cubic Ledger Topology," Recognition Science preprint, 2026.

[5] K. G. Chetyrkin, J. H. Kühn, and M. Steinhauser, "RunDec: A Mathematica package for running and decoupling," *Comput. Phys. Commun.* 133 (2000) 43.