

Article

Charged Lepton Masses from the Recognition Composition Law: A Derivation with Zero Continuously Adjustable Dimensionless Parameters

Jonathan Washburn ¹  and Elshad Allahyarov ^{1,2,3,4,*} ¹ Recognition Science, Recognition Physics Institute, Austin, TX 78701, USA; jon@recognitionphysics.org² Department of Physics, Case Western Reserve University, Cleveland, OH 44106, USA³ Institut für Theoretische Physik II: Weiche Materie, Heinrich-Heine Universität, 40225 Düsseldorf, Germany⁴ Theoretical Department, Joint Institute for High Temperatures, Russian Academy of Sciences, 125412 Moscow, Russia

* Correspondence: elshad.allakhyarov@case.edu

Abstract

We derive the charged-lepton mass chain from the Recognition Composition Law (RCL) together with normalization, curvature normalization, and standard regularity. Through the theorem chain Tr1–Tr8, these postulates fix the golden ratio $\varphi = (1 + \sqrt{5})/2$, the minimal period $T_{\min} = 8$, the selected dimension $D = 3$, and the cube integers entering the master mass law. The charged-lepton formula is then assembled from the coherence scale, the lepton-sector baseline, the charge correction, and the derived generation steps. All parameters are discrete structural inputs, integers from cube geometry, named symmetry factors, and one external mathematical constant, rather than continuously adjustable dials. The construction is a structural constraint on the effective charged-lepton flavor pattern, not a replacement for the electroweak Higgs mechanism or for the full Standard Model quantum field theory. At the conversion stage to the International System of Units (SI), the electron fixes the single calibration anchor τ_0 , while the fine-structure constant α enters only as a fixed external dimensionless constant in the refinement layer. The phrase “zero continuously adjustable parameters” refers to the dimensionless content of the framework: the anchor τ_0 is a unit-scale calibration fixed by the measured electron mass and cancels identically from every charged-lepton mass ratio. With that one anchor set, the remaining charged leptons become forward predictions: $m_\mu \in (105.5, 105.9)$ MeV and $m_\tau \in (1774, 1779)$ MeV, with relative errors below 0.3% and 0.2%, respectively. Floating-point evaluation gives $m_\mu \approx 105.658$ MeV and $m_\tau \approx 1776.71$ MeV.

Keywords: charged leptons; mass hierarchy; flavor structure; discrete geometry; golden ratio; machine verification; recognition composition law



Academic Editor: Jorge Segovia

Received: 15 April 2026

Revised: 14 May 2026

Accepted: 27 May 2026

Published: 2 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The Standard Model accommodates the three charged-lepton masses, the electron, the muon, and the tau, through independent Yukawa couplings in the electroweak Lagrangian. Those couplings, and the hierarchy $m_e:m_\mu:m_\tau \approx 1:206.8:3477$, are fixed empirically [1]; the theory does not explain why these values or ratios are realized. The masses are known with high precision [1], but the explanatory problem remains open.

A large literature seeks structural explanations of charged-lepton hierarchies, but each major class still leaves free input somewhere in the construction. Froggatt–Nielsen

textures [2,3] trade Yukawa couplings for charge assignments and a flavon parameter; Koide-type relations [4,5] compress the spectrum but do not derive the remaining parameters; grand unified and texture-zero constructions [6–11] impose correlations or vanishing patterns without eliminating empirical non-zero entries; discrete-flavor and modular frameworks [12–20] organize masses through symmetries or modular forms but retain undetermined couplings or moduli; and radiative or extra-dimensional mechanisms [21–27] shift the hierarchy into new particles, couplings, or geometric scales. A broad review is given in Ref. [28].

The central question is not merely whether the three masses can be fit more compactly. The deeper question is whether they can be derived from a single organizing principle with no continuously adjustable dials. Specifically, this paper asks: (i) Is there a geometric or combinatorial structure that selects the mass hierarchy base uniquely, without postulating it? (ii) Do the generation offsets $\{0, 11, 17\}$, the integers controlling the $e \rightarrow \mu \rightarrow \tau$ steps, follow from a single structural identity? (iii) Can all three masses be fixed simultaneously by one empirical calibration anchor, making the muon and tau genuine *predictions*?

This paper takes up those questions within the charged-lepton sector. Starting from the Recognition Composition Law and standard regularity assumptions, it derives a unique cost functional, the theorem chain Tr1–Tr8, with Tr0 serving only as the logical background, the cube-based ingredients of the mass law, and finally the charged-lepton mass chain. The argument keeps two layers distinct: the discrete backbone fixes the leading hierarchy, while the refinement layer supplies the sub-leading comparison terms. The claim is, therefore, not merely a compact fit but a discrete derivation whose load-bearing inputs are chosen from finite, explicitly enumerated combinatorial alternatives (cube vertices, edges, faces, the integer rungs, and a finite polynomial class for the charge index) rather than from a continuum of fitted real-valued couplings; the later sections supply the numerical comparison with data and the bookkeeping audit.

The present construction is not proposed as a new local quantum field theory and does not replace the Standard Model Lagrangian. Its intended role is narrower: it supplies a discrete structural rule for the effective charged-lepton flavor hierarchy, i.e., for the pattern that in the Standard Model is represented by three independent charged-lepton Yukawa couplings. Electroweak symmetry breaking is still understood in the usual way as the mechanism converting those dimensionless effective couplings into physical masses. The Recognition construction is, therefore, tested here as a constraint on the flavor data that enter the low-energy theory, not as a complete microscopic account of gauge dynamics.

We also distinguish prediction from reconstruction more carefully. The electron mass is used once to fix the SI conversion anchor τ_0 ; this is a unit-scale calibration and cancels from all mass ratios. The nontrivial tests are the anchor-free ratios m_μ/m_e and m_τ/m_μ . The refinement layer uses the fixed dimensionless constant α and fixed structural rules; it contains no continuously adjustable parameter, but it does contain discrete organizational choices whose status is now recorded explicitly below.

In particular, the four MODEL-LAYER structural rules—the affine-log family for the rung correction R , the polynomial class for the charge index Z , the electron-break δ_e rule, and the generation-step refinement rule—are structural ansätze whose physical necessity is not fully demonstrated from the Recognition postulates alone. They are motivated by explicit physical requirements (monotonicity, charge-conjugation invariance, minimal completeness, and the J -expansion-based α -order organization) and, once adopted, leave no free real-valued coefficient, but they are not theorem-level consequences. Sections 5–7 state this explicitly, and Table A1 records their MODEL-LAYER status.

This paper treats only the charged-lepton sector. It isolates the charged-lepton chain rather than attempting the full fermion spectrum at once; quark masses, neutrino masses, and mixing parameters are outside the scope of the present paper.

The remaining parts of this paper are organized as follows. Section 2 states the postulates and fixes the cost functional. Section 3 records the theorem dependency order. Section 4 derives the cube integers. Section 5 assembles the master mass law. Section 6 derives the integerized charge map and the lepton charge index. Section 7 derives the charged-lepton mass chain and compares it with data. Section 8 records the closure tally. Section 9 adds the International System of Units (SI) conversion through the single anchor τ_0 . Section 11 states the final scope of the claim. The appendices then serve as a reference layer: Appendix A gives the numerical recipe, Appendix B gives the audit, Appendix C gives the theorem statements and proof material (including Tr7's cube identity $W = 17$), Appendix D gives the rung-correction derivation, and Appendix E gives the exponent bookkeeping.

2. Theoretical Framework

The framework begins with four postulates. Its starting object is a scalar cost assigned to a positive multiplicative ratio: $x = 1$ means no change, while $x > 1$ and $x < 1$ represent reciprocal enlargements and reductions. The Recognition Composition Law (RCL) requires the cost of combining two ratio changes to be expressible purely in terms of the costs of the individual ratios, thereby sharply constraining the admissible functional form. With normalization and regularity added, these postulates select a unique normalized cost functional, denoted J , which becomes the basic object used in the later mass-law construction. The four postulates are stated precisely as follows:

(P1) *The Recognition Composition Law:*

$$F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y), \quad x, y > 0 \quad (1)$$

(P2) *Normalization:* $F(1) = 0$ (zero cost at the identity ratio).

(P3) *Curvature normalization:* $\lim_{t \rightarrow 0} \frac{2F(e^t)}{t^2} = 1$ (unit curvature at the minimum; this fixes the normalization of J).

(P4) *Standard regularity:* continuity on $R_{>0}$ (the positive real numbers) and sufficient smoothness for the d'Alembert uniqueness step.

The physical reading of P1 is that it is an independence axiom for a cost functional on positive multiplicative ratios. If two ratio changes x and y are recognized independently, then the joint cost of the composed ratio xy and the relative ratio x/y is determined by the individual costs alone. In this sense, the RCL plays, for multiplicative recognition costs, the same organizing role that additivity plays for distances, and the product rule plays for probabilities. It is not asserted here as a consequence of a known Lagrangian dynamics; it is the structural composition postulate from which the cost law is then forced. This starting point is not an ad hoc numerical device. The passage from the RCL to J is the classical d'Alembert functional-equation route: after the logarithmic reparametrization $x = e^t$, the shifted function satisfies the standard equation $H(t+s) + H(t-s) = 2H(t)H(s)$, whose continuous solutions are classified by Aczél-type functional-equation theory. The resulting cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ is also the canonical reciprocal convex penalty on positive ratios. In information-geometric language, it plays the same structural role that Kullback–Leibler divergence and Fisher information play for probability models: it supplies a local quadratic metric at the identity and a globally asymmetric-to-symmetric cost law for finite departures, here specialized to multiplicative recognition ratios rather than probability densities. In an

equivalent closed form, $J(x) = (x - 1)^2 / (2x)$ is the symmetric squared relative deviation of x from unity: it vanishes only at the no-recognition point $x = 1$, is positive elsewhere, and is invariant under the inversion $x \mapsto 1/x$ (reversal of the recognition arrow). These four postulates uniquely fix the cost functional $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ (Theorem Tr1; proof in Appendix C). The later structural chain does *not* assume Q_3 at the outset: it combines this fixed cost functional with the cube family, and Theorem Tr7 then selects $D = 3$ and hence the specific 3-cube Q_3 . Tr0 serves only as the logical background, while Tr1–Tr8 supply the structural consequences used in the derivation.

The cost functional does not take masses as direct inputs; it fixes the admissible ratio geometry. Structurally, the argument runs

$$\begin{aligned} \text{RCL + normalization + regularity} &\Rightarrow J(x) = \frac{1}{2}(x + x^{-1}) - 1, \\ J \text{ together with} & \\ \text{Tr0 and Tr1–Tr8} &\Rightarrow (\varphi, T_{\min} = 8, D = 3, \text{ cube data}), \end{aligned}$$

and those discrete outputs then determine

$$m(i) = A_s \varphi^{r_i - 8 + R(Z_i)}$$

The later sections then derive the lepton-sector ingredients appearing in this mass law explicitly.

The physical reading of this arrow is deliberately conservative. The RCL and the theorem chain fix the admissible ratio geometry and the discrete cube data; the mass law then interprets these data as constraints on effective flavor-sector mass ratios. Thus, the construction should be compared with a flavor-structure rule or texture constraint, not with a full derivation of the Standard Model action. The latter would require a separate bridge from the Recognition variables to gauge fields, fermion kinetic terms, electroweak symmetry breaking, and renormalization-group transport.

3. Structural Theorems: Tr0 and Tr1–Tr8

This section fixes the logical order used later. Tr0 is only the ambient logical background; Tr1–Tr8 are listed in strict backward-referenced dependency order. Table 1 records that order, and Appendix C collects the full statements, proof sketches, and formal references.

Table 1. Dependency map for Tr0 and Tr1–Tr8 in the later charged-lepton derivation.

Item	Statement	Reason It Is Necessary	Output for Later Derivation
Tr0	Classical logic	Implicit logical background	Background inference
Tr1	Unique $J(x)$	RCL + normalization + d’Alembert uniqueness	$J(x) = \frac{1}{2}(x + x^{-1}) - 1$
Tr2	φ hierarchy base	Additive scale closure forces $r^2 = r + 1$	Fixes the φ ladder
Tr3	$m > 0$	Zero-mass limit lies at infinite cost (Tr1)	Positive-mass rungs only
Tr4	Paired transitions	Reciprocal symmetry $J(x) = J(1/x)$	Reciprocal sector pairing
Tr5	Discrete dynamics	Finite cost for finite-time evolution	Integer step structure
Tr6	Finite-cost states	Only finite-cost states are physical	Finite occupied rungs only
Tr7	$D = 3$	$W_{\text{endo}}(D) = E_{\text{passive}} + F = 17 \Leftrightarrow D = 3$	Selects the 3-cube; fixes $V, E, F, E_{\text{passive}}, W$
Tr8	Period 8	Shortest closed Hamiltonian cycle on Q_3 has length 8	Mass-law offset -8

The next three sections unpack Table 1 in order: first, the cube integers, then the master mass law, and then the charge correction.

4. Cube Geometry at $D = 3$: Integer Constants

This section extracts the discrete constants that the later mass law uses explicitly. Once Tr7 selects $D = 3$, the standard cube counts are as follows. The D -dimensional unit cube

Q_D has $V(D) = 2^D$ vertices, $E(D) = D \cdot 2^{D-1}$ edges, and $F(D) = 2D$ faces. At $D = 3$, these give

$$V = 8, E = 12, F = 6. \tag{2}$$

Discrete-step counting, together with cost minimization, fixes one active edge per elementary step, so [29] $A = 1$. Hence, the number of passive edges is $E_{\text{passive}} = E - A = 11$. The derived cube sum

$$W(D) = E_{\text{passive}}(D) + F(D) = (D \cdot 2^{D-1} - 1) + 2D$$

takes the value 17 only at $D = 3$:

$$W(D) = 17 \Leftrightarrow D = 3 \tag{3}$$

This is the point at which W is fixed for later use (Lean-verified [29]; proof in Appendix C, Tr7).

Tr8 fixes the minimal closed Hamiltonian cycle on Q_3 to have length [30] $T_{\text{min}} = 8$. With the exponent origin placed at the coherence energy unit $E_{\text{coh}} = \varphi^{-5}$, this yields the reference offset

$$r_{\text{ref}} = -T_{\text{min}} = -8. \tag{4}$$

Accordingly, the mass-law exponent later appears as $r - 8 + R(Z)$.

Table 2 collects the discrete constants used explicitly in the next two sections of the derivation.

Table 2. Discrete constants extracted from the selected 3-cube and used later in the mass law.

Symbol	Value	Formula	Origin
V	8	2^D	Vertices of Q_3
E	12	$D \cdot 2^{D-1}$	Edges of Q_3
F	6	$2D$	Faces of Q_3
A	1	—	Cost minimization [29]
E_{passive}	11	$E - A$	Passive edges
W	17	$E_{\text{passive}} + F$	Cube sum, unique at $D = 3$ (3)
T_{min}	8	—	Minimal closed Hamiltonian cycle on Q_3 (Tr8)

5. The Master Mass Law

This section assembles the master mass law from three previously fixed ingredients: the coherence energy quantum E_{coh} , the sector mass scale A_s , and the rung correction R .

The coherence energy quantum

$$E_{\text{coh}} = \varphi^{-(D+2)} = \varphi^{-5} \tag{5}$$

The exponent $-(D + 2) = -5$ is structurally determined by two independent routes: (i) the Fibonacci deficit $2^D - D = 8 - 3 = 5$ and (ii) the integration dimension $D + 2 = 3 + 2 = 5$ (spatial + temporal + conservation constraint). Both routes yield the same answer and are machine-verified [31]. E_{coh} is the fundamental energy unit of the φ -ladder at $D = 3$.

The sector mass scale combines E_{coh} with two cube-derived integers, $B_{\text{pow}}(s)$ and $r_0(s)$:

$$A_s = 2^{B_{\text{pow}}(s)} \cdot E_{\text{coh}} \varphi^{r_0(s)}. \tag{6}$$

For the charged-lepton sector (color-neutral, $|Q|=1$), the passive-edge coupling gives $B_{\text{pow}}(l) = -2E_{\text{passive}} = -22$ (machine-verified [32]). The φ -ladder offset follows from the wallpaper-lattice formula

$$r_0(s) = m_s \cdot W + c_s \quad (7)$$

with $m_l = N_{sec} = 2^{D-1} = 4$ (the number of fermion sectors at $D = 3$; machine-verified [31]) and $c_l = -(T_{min} - r_e) = -(8 - 2) = -6$ (the octave gap from the electron baseline rung), giving $r_0(l) = 62$. These are traceable discrete structural inputs from cube geometry, not continuously adjustable parameters. For charged leptons, this yields

$$A_l = 2^{-22} E_{coh} \varphi^{62},$$

with both exponents derived from cube data rather than fitted to the observed lepton masses.

The rung correction $R(Z)$ supplies the fractional, charge-dependent shift to the integer rung r . We restrict the candidate space to the affine-log family

$$R(x) = a \cdot \ln\left(1 + \frac{x}{b}\right) + c, \quad b > 1, \quad (8)$$

which is the natural multiplicative-to-additive bridge from J -costs to φ -ladder shifts. Within this family, the three specific normalization conditions $R(0) = 0$, $R(-1) = -2$, and $R(1) = 1$ then uniquely fix the parameters to $a = 1/\ln\varphi$, $b = \varphi$, and $c = 0$, giving

$$R(Z) = \log_{\varphi}\left(1 + \frac{Z}{\varphi}\right). \quad (9)$$

The full derivation appears in Appendix D, which also explains why the affine-log family is the natural bridge from multiplicative J -costs to additive φ -ladder shifts and why the three normalization conditions uniquely fix its parameters.

The rung correction is one of the points at which a structural choice enters the construction. The theorem proved in Appendix D is a uniqueness result within the affine-log family. The use of an affine logarithm is physically motivated by the conversion of multiplicative ratio data into additive rung shifts and by the requirement that the correction be monotone, sub-linear, and neutral at $Z = 0$; it is not claimed here to be the unique possible map among all conceivable monotone functions. Once this physically motivated family and its three normalization conditions are adopted, no real-valued coefficient remains free.

At the structural level, the three components assemble into the master mass law:

$$m(i) = A_s \varphi^{r_i - 8 + R(Z_i)}, \quad (10)$$

where r_i is the species rung, $-8 = -T_{min}$ is the period offset (4), and $R(Z_i)$ is the charge correction. Within the framework, this gives masses as positive, dimensionless ratios; the later SI step introduces only the single empirical anchor τ_0 . In Section 7, this structural form is rewritten for charged leptons as an electron seed together with the break exponent δ_e and the two generation steps $S_{e \rightarrow \mu}$ and $S_{\mu \rightarrow \tau}$. For two species in the same sector with the same charge index, A_s and $R(Z)$ cancel, leaving

$$\frac{m(i)}{m(j)} = \varphi^{r_i - r_j}, \quad (11)$$

so the ratio depends only on the integer rung difference. Equation (11) is the leading integer-backbone relation. The precision layer introduced in Section 7 applies fixed sub-leading corrections to the inter-generation exponents. This is why the tables below display both the leading values φ^{11} and φ^6 and the refined exponents. The exponent components of the mass law are isolated in Appendix E, while Appendix D gives the full derivation of the rung correction R .

6. The Charge Subsystem: Integerization and the Z-Map

The rung correction R requires an integer charge index Z_i for each charged lepton. This section fixes that index in three steps: charge integerization, an even polynomial ansätze, and coefficient selection.

The charge-index construction is another place where this paper uses a constrained structural rule rather than an unconstrained numerical fit. The requirements imposed below are charge-conjugation invariance, neutral vanishing, non-negativity, and separation of the charge label from the vertex and edge data already used for state and transition counting. These requirements do not by themselves amount to a full dynamical derivation of electric charge from a gauge field. They define the finite structural class within which the charged-lepton index is fixed.

Tr5 requires a discrete charge label. For the charged leptons $Q = \pm 1$, integrality alone does not determine the overall scale k : any positive integer keeps $\tilde{Q} = kQ$ integral. The framework fixes k geometrically from the 3-cube face count, assigning one independent 2D symmetry channel per face, so

$$k = F(3) = 6. \quad (12)$$

The face count is a discrete structural input: each face of Q_3 provides one independent 2D symmetry channel for charge quantization. Faces are used here because the charge label is attached to closed 2D sector data. In the present bookkeeping, vertices already classify occupied states, and edges already classify transition channels, so assigning the charge index at the face level keeps the charge map distinct from the state-counting and propagation-counting layers. This is a geometric choice rather than a derived dynamical necessity (machine-verified [31]). Thus, the face assignment should be read as a physically motivated structural assignment: charge is attached to closed two-dimensional sector data, so the six faces of Q_3 provide the minimal independent face-channel scale. It is not used as an adjustable parameter because no continuous scale is varied and no lepton mass value enters the choice. Hence, the electron carries the integerized charge $\tilde{Q}_e = kQ_e = -6$.

The charge index $Z(\tilde{Q})$ must satisfy three requirements: charge-conjugation invariance $Z(\tilde{Q}) = Z(-\tilde{Q})$ (particles and antiparticles share the same index); non-negativity $Z \geq 0$ (so the argument of R remains in its domain); and neutral vanishing $Z(0) = 0$ (neutral particles receive no correction). Within a polynomial ansätze, these requirements restrict Z to an even polynomial in \tilde{Q} with no constant term. The minimal two-term family is

$$Z(\tilde{Q}) = a\tilde{Q}^2 + b\tilde{Q}^4, \quad a, b \in \mathbb{Z}_{\geq 0}, \quad (a, b) \neq (0, 0). \quad (13)$$

Retaining the two lowest nontrivial even powers means imposing $a \geq 1$ and $b \geq 1$ and then minimizing $a + b$. This selects $(a, b) = (1, 1)$ uniquely (Table 3). The minimal-completeness criterion (both even powers present, minimum $a + b$) is an organizing principle that selects $(1, 1)$ uniquely. We now state explicitly that this is an organizing principle inside the stated polynomial class. The quadratic-only and quartic-only candidates are not excluded by arithmetic alone; they are excluded by the requirement that the first two charge-conjugation-even invariants both be represented in the minimal nontrivial charge index. Within that rule, the quadratic-only and quartic-only choices are excluded because each suppresses one of the first two charge-conjugation-even invariants. Requiring the smallest polynomial that retains both even invariants, while introducing no higher powers, fixes $(a, b) = (1, 1)$ uniquely within the stated ansätze. The resulting $Z_l = 1332$ is a discrete structural input; its arithmetic is machine-verified [31].

Table 3. Candidate two-term even polynomials for the lepton charge index. The minimal choice with both terms present is (1, 1).

(a,b)	Z_l	Status
(1, 0)	36	Quartic term absent
(0, 1)	1296	Quadratic term absent
(1, 1)	1332	Both terms present; unique minimizer

The lepton charge index (no color offset, $c = 0$) is, therefore,

$$Z_l = \tilde{Q}_e^2 + \tilde{Q}_e^4 = 1332. \quad (14)$$

This single integer enters R and hence every charged-lepton mass in the framework.

7. The Charged-Lepton Mass Chain

With the lepton-sector scale, rungs, and charge index fixed in the previous sections, the charged-lepton specialization of (10) can now be written as an electron structural seed together with a symmetry-breaking correction and two generation steps. For e, μ, τ , this consists of the electron leading term, the electron correction δ_e , and the two generation steps $S_{e \rightarrow \mu}$ and $S_{\mu \rightarrow \tau}$. The derived inputs used below include $r_e = A + 1 = 2$, where $A = 1$ is the active-edge count per tick from Tr5 atomicity: the minimal stable charged state sits one rung above the transition mechanism [31,33]; the rung correction uses $R(-1) = -2$ to fix its normalization (Appendix D), and the physical lepton charge index $Z_l = 1332$ gives the relevant value $R(1332)$ in all three mass formulas.

The formulas in this section should be read in two layers. First, the integer backbone fixes the sector scale and the rungs 2, 13, 19. Second, the refinement layer introduces fixed, non-adjustable comparison terms involving α . These terms sharpen the leading backbone ratios but do not introduce a tunable lepton-sector coefficient.

The three charged-lepton rungs follow from the baseline [33] $r_e = 2$ and cube combinatorics:

$$r_e = 2, \quad r_\mu = 2 + E_{\text{passive}} = 13, \quad r_\tau = 2 + W = 19. \quad (15)$$

These give the leading rung gaps $r_\mu - r_e = E_{\text{passive}} = 11$ and $r_\tau - r_\mu = W - E_{\text{passive}} = F = 6$; the subsections below derive the corresponding sub-leading corrections.

7.1. The Electron Leading-Order Mass: A Closed Form

Substituting the derived lepton-sector values $B_{\text{pow}} = -22$, $r_0(l) = 62$, and $r_e = 2$ [33] into the mass law (10) gives the electron leading-order mass in the framework's internal units:

$$m_{\text{struct}}(e) = 2^{-22} \cdot \varphi^{51}. \quad (16)$$

The exponent $51 = -5 + 62 + 2 - 8$ collects the coherence scale, lepton offset, electron rung, and period baseline, so $m_{\text{struct}}(e)$ is still expressed in the framework's internal units.

7.2. The Correction δ_e : Topological and α -Correction Decomposition

The leading-order mass $m_{\text{struct}}(e)$ does not yet include the charge correction. The full electron prediction is, therefore, $m_e^{\text{pred}} = m_{\text{struct}}(e) \cdot \varphi^{R(1332) - \delta_e}$, where the break exponent δ_e subtracts from the rung correction R .

The symmetry-breaking correction δ_e decomposes into a topological base shift and structural α -corrections:

$$\delta_e = \underbrace{2W + \frac{W + E_{\text{total}}}{4E_{\text{passive}}}}_{\text{topological}} + \underbrace{\alpha^2 + E_{\text{total}}\alpha^3}_{\text{structural } \alpha\text{-corrections}}. \quad (17)$$

7.2.1. The Topological Base Shift

The base shift has two parts:

1. $2W = 2 \times 17 = 34$: twice the wallpaper-group count. The number $W = E_{\text{passive}} + F = 11 + 6 = 17$ is the endogenous cube sum derived intrinsically from Q_3 geometry; its coincidence with the count of 2D wallpaper groups (Fedorov, 1891) [34] is noted but not assumed as input. This is the dominant contribution to δ_e . It combines reciprocal pairing (Tr4), the derived count $W = E_{\text{passive}} + F = 17$, and the minimal-cover claim that the electron scans each wallpaper class once in each orientation. Hence,

$$\delta_e^{\text{int}} = \underbrace{W}_{\text{forward orientation}} + \underbrace{W}_{\text{backward orientation}} = 2W = 34.$$

The algebraic ingredients are Lean-verified [32]; the minimal-cover claim itself remains a formalization step [33].

2. $(W + E_{\text{total}})/(4E_{\text{passive}}) = (17 + 12)/(4 \times 11) = 29/44$: the ratio of combined wallpaper-plus-edge content to total passive capacity. Since the derived number of fermion sectors is $N_{\text{sec}} = 2^{D-1} = 4$, this can be written as [35–37] $(W + E_{\text{total}})/(N_{\text{sec}}E_{\text{passive}})$. Together: base shift becomes equal to $34 + 29/44 = 34.6590$.

7.2.2. Structural α -Corrections (Not Standard QED)

These α -terms are not standard QED radiative corrections; within the present framework, they arise from the Taylor expansion of $2J(1 + \alpha)$ using the Tr1 cost functional. The specific order-matching rules assigning α^2 to the base shift and α^3 to the edge aggregate are structural ansätze motivated by the J -expansion. We state explicitly that the fine-structure constant α is introduced phenomenologically in the present paper and is not derived within this framework. It enters the refinement layer as a fixed external dimensionless constant taken from experiment, with no continuous adjustment. α is, therefore, not a fitted lepton-sector parameter, but the manuscript counts it among the explicit external inputs to the refinement layer rather than among the theorem-backed structural ingredients. It is useful to draw a finer distinction between the algebraic form and the numerical value of α in the refinement layer. Within the stated MODEL-LAYER refinement rules (the electron-break δ_e rule and the generation-step refinement rule), the algebraic form of the α -corrections is fixed before comparison with data: the powers α^2 and α^3 , the signs $-\alpha^2$ and $+E_{\text{total}}\alpha^3$, and the edge-aggregation factor $E_{\text{total}} = 12$ are organized by the Taylor expansion of $2J(1 + \alpha)$ together with the cube edge bookkeeping of Q_3 . None of these algebraic features are fitted to the lepton masses, but neither these refinement rules nor the numerical value of α is claimed here as a theorem-level consequence of the Recognition postulates alone. What the framework does not derive is the numerical value $\alpha \simeq 1/137.036$ itself: that value is taken from experiment. This is exactly analogous to the role of τ_0 in the SI bridge: the algebraic role of τ_0 (one unit-conversion anchor) is fixed by the construction, while its numerical magnitude is calibrated empirically. More precisely, the role split is fixed by what each order can encode in the present construction. The quadratic term is the leading local correction to a single ascending φ -step, so J -convexity fixes its sign as $-\alpha^2$. The cubic term is the first order whose sign survives the orientation bookkeeping of the edge network: each edge contributes $-\alpha^3$, the bipartite parity of Q_3 contributes a second minus sign, and summing over the $E_{\text{total}} = 12$ edges yields $+E_{\text{total}}\alpha^3 = +12\alpha^3$. Accordingly, the quadratic

term corrects the local rung shift, whereas the cubic term enters only after aggregation over the full edge channel.

1. $\alpha^2 \approx 5.325 \times 10^{-5}$: the leading structural α -correction. Its coefficient is exactly 1: the J -expansion $2J(1 + \alpha) = \alpha^2 + c\alpha^3 + \dots$ (Tr1, the uniqueness result [38]) forces the α^2 coefficient to be 1 by normalization.
2. $E_{total} \cdot \alpha^3 = 12 \cdot \alpha^3 \approx 4.66 \times 10^{-6}$: the edge-aggregated structural correction.

The coefficient $+E_{total} = +12$ is now derived: each of the $E_{total} = 12$ edges contributes $-\alpha^3$ per channel from $2J(1 + \alpha)$; the 3-cube is bipartite with orientation parity at $D = 3$, giving the aggregate $(-1) \times (-1) \times 12 = +12$ (Thm. 2.2 [35–37]).

Total structural α -correction $\approx 5.79 \times 10^{-5}$; total break: $\delta_e \approx 34.6591$.

7.3. The Electron-to-Muon Step

The generation step $S_{e \rightarrow \mu}$ is:

$$S_{e \rightarrow \mu} = E_{\text{passive}} + \frac{1}{4\pi} - \alpha^2 \approx 11.0795 \quad (18)$$

Decomposition. $E_{\text{passive}} = 11$ gives the dominant topological gap, $1/(4\pi) = E_{\text{passive}}\alpha_{\text{seed}}$ is the geometric seed correction with $\alpha_{\text{seed}} = 1/(4\pi E_{\text{passive}})$, and the final $-\alpha^2$ term is the derived structural subtraction associated with an ascending φ -step. Thus, the first generation step is the passive-edge count with one small positive seed correction and one still smaller negative α^2 correction.

7.4. The Muon-to-Tau Step

The generation step $S_{\mu \rightarrow \tau}$ is:

$$S_{\mu \rightarrow \tau} = F - \frac{2W + D}{2}\alpha \approx 5.8650. \quad (19)$$

Decomposition. $F = 6$ supplies the zeroth-order face-count gap, while the only sub-leading term is the derived linear correction $-(2W + D)\alpha/2 = -(37/2)\alpha$, whose coefficient comes from the same face-duality bookkeeping used in the wallpaper count [35–37]. At leading order, the two generation steps, therefore, satisfy $S_{\mu \rightarrow \tau}^{(0)}/S_{e \rightarrow \mu}^{(0)} = F/E_{\text{passive}} = 6/11$; the full ratio drops to ≈ 0.529 once the linear- α term is included. The predictive content of these two step formulas is their rigidity: after the backbone integers E_{passive} and F are fixed, the sub-leading terms contain only π , α , and cube-derived integer coefficients. No coefficient is adjusted to the muon or tau masses.

7.5. The Full Lepton Mass Chain

Substituting all charged-lepton ingredients, the three predicted masses are:

$$m_e^{\text{pred}} = m_{\text{struct}}(e)\varphi^{R(1332) - \delta_e}, \quad (20)$$

$$m_\mu^{\text{pred}} = m_e^{\text{pred}}\varphi^{S_{e \rightarrow \mu}}, \quad (21)$$

$$m_\tau^{\text{pred}} = m_\mu^{\text{pred}}\varphi^{S_{\mu \rightarrow \tau}}, \quad (22)$$

where $m_{\text{struct}}(e) = 2^{-22}\varphi^{51}$, $R(1332) \approx 13.953$, $\delta_e \approx 34.659$, $S_{e \rightarrow \mu} \approx 11.0795$, and $S_{\mu \rightarrow \tau} \approx 5.8650$.

7.6. Lepton Mass Predictions

The electron fixes the universal SI anchor τ_0 ; once that anchor is set, the muon and tau follow with no further lepton mass input. Table 4 separates the integer backbone ($E_{\text{passive}} =$

11, $F = 6$) from the precision-layer corrections, and Table 5 gives the corresponding SI-scale comparison.

Table 4. Leading-order integers vs. precision-layer formulas vs. PDG reference. The leading-order columns (φ^{11} and φ^6) follow from cube combinatorics alone; the refined columns include the sub-leading corrections of Section 7.2. No lepton mass value enters the prediction columns.

Step	Leading Integer	Refined Step	Framework	PDG
$e \rightarrow \mu$	$\varphi^{11} \approx 199.005$	$\varphi^{11.080}$	206.768	206.768
$\mu \rightarrow \tau$	$\varphi^6 \approx 17.944$	$\varphi^{5.865}$	16.816	16.817

Table 5. Charged-lepton mass predictions vs. PDG data (floating-point evaluation). The electron row fixes the calibration; the muon and tau use no further lepton mass input. Lean-certified bounds are wider: see Section 10.

Particle	Predicted (MeV)	PDG (MeV)	Relative Error
e	0.51100	0.51100	0 (calibration)
μ	105.658	105.658	$\sim -1 \times 10^{-6}$
τ	1776.71	1776.86	$\sim -9 \times 10^{-5}$

The minimal numerical reproduction sequence for these values is recorded in Appendix A.

8. Closure Summary

This section records only the closure tally of the charged-lepton framework. The tally is meant as compact bookkeeping rather than as a second theorem layer: it summarizes which ingredients are theorem-level consequences, which are derived from cube data and sector bookkeeping, which are fixed structural rules defining the charged-lepton refinement class, and which belong only to calibration or convention.

$$3 \text{ FORCED} + 4 \text{ DERIVED} + 4 \text{ MODEL-LAYER} + 1 \text{ calibration anchor } \tau_0 + 1 \text{ convention} \quad (23)$$

$(\lambda = 1) = 13$ audited components, 0 continuously adjustable fitted parameters

The MODEL-LAYER category is the fifth bookkeeping bucket introduced in this revision: it isolates the affine-log family for the rung correction, the polynomial class for the charge index, the electron break δ_e rule, and the generation-step refinement rule as fixed structural choices that define the charged-lepton refinement class. These are not theorem-level consequences of the Recognition postulates alone, but they contain no continuously adjustable real-valued coefficient. The “zero continuously adjustable fitted parameters” statement refers to the dimensionless lepton ratios m_μ/m_e and m_τ/m_μ , which are independent of the SI anchor τ_0 and use no muon or tau mass input. Appendix B gives the itemized audit; Figure 1 is only a compressed dependency sketch; Table 1 records the full Tr-order. Section 9 then adds the SI unit-conversion bridge through the single empirical anchor τ_0 .

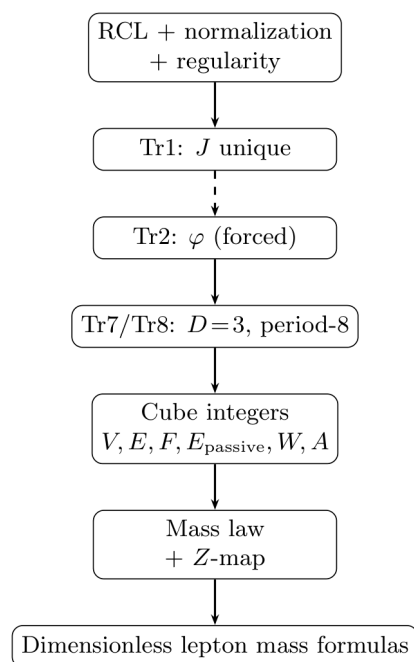


Figure 1. Compressed dependency chain from the RCL and selected structural theorems to the dimensionless charged-lepton mass formulas. Solid arrows denote derived or forced steps; the dashed arrow indicates that Tr2 (φ) is forced independently of Tr1 rather than derived from the preceding node. Intermediate consequences Tr3–Tr6 are omitted here for readability; see Table 1 for the full order.

9. The SI Unit-Conversion Step

Up to this point, this paper has produced charged-lepton masses only as framework quantities. This section adds the SI bridge by assigning one empirical duration τ_0 in SI seconds to a single framework period step. The exact SI constants c_{SI} and \hbar_{SI} then convert the universal framework mass quantum into SI units. No new structural input is introduced here, and no discrete quantity is re-derived.

The single SI anchor $\tau_0 > 0$ (seconds per period step) is fixed by matching the full predicted electron mass m_e^{pred} to its SI value:

$$\tau_0 = \frac{m_e^{\text{pred}} \hbar_{\text{SI}}}{m_e^{\text{SI}} c_{\text{SI}}^2} = \frac{2^{-22} \varphi^{51+R(1332)-\delta_e} \hbar_{\text{SI}}}{m_e^{\text{SI}} c_{\text{SI}}^2}. \quad (24)$$

We use the full m_e^{pred} here, not the bare structural seed $2^{-22} \varphi^{51}$, so that the calibration entry in Table 5 is enforced by definition and the muon and tau values become genuine anchor-free predictions. This fixes the absolute SI scale through the electron by definition. The nontrivial SI checks are, therefore, the muon and tau values, equivalently the mass ratios, because τ_0 cancels from every ratio.

Why τ_0 and nothing else? In framework units, $c = 1$, and the action quantum is fixed by the coherence unit. Once one period step is assigned the SI duration τ_0 , every length, energy, and mass conversion follows from the exact SI definitions of c_{SI} and \hbar_{SI} ; no second empirical scale is introduced.

Derived SI conversion factors. Given the single empirical anchor $\tau_0 > 0$ and the 2019 SI-exact constants c_{SI} , \hbar_{SI} , the conversion factors are:

$$\text{meters per length step} : l_0^{\text{SI}} = c^{\text{SI}} \tau_0 \quad (25)$$

$$\text{joules per coh} : E_0^{\text{SI}} = \frac{\hbar_{\text{SI}}}{\tau_0} \quad (26)$$

$$\text{kg per mass quantum} : m_0^{\text{SI}} = \frac{E_0^{\text{SI}}}{c_{\text{SI}}^2} = \frac{\hbar_{\text{SI}}}{\tau_0 c_{\text{SI}}^2} \tag{27}$$

Derivation: Length: In framework units, $c = 1$, so one length step per period step becomes $c_{\text{SI}}\tau_0$ meters in SI.

Energy (26): The same period step, expressed in SI seconds, has duration τ_0 , so the framework action quantum maps to \hbar_{SI} . Hence, one coherence unit carries energy $E_0^{\text{SI}} = \hbar_{\text{SI}}/\tau_0$.

Mass (27): $E = mc^2$, then gives $m_0^{\text{SI}} = E_0^{\text{SI}}/c_{\text{SI}}^2 = \hbar_{\text{SI}}/(\tau_0 c_{\text{SI}}^2)$.

Universal SI mass factor. Multiplying the structural electron mass $m_{\text{struct}}(e) = 2^{-22}\varphi^{51}$ by the universal SI mass quantum $m_0^{\text{SI}} = \hbar_{\text{SI}}/(\tau_0 c_{\text{SI}}^2)$ gives the electron mass in SI kilograms:

$$m_e^{\text{SI}} = \underbrace{(2^{-22} \cdot \varphi^{51})}_{\text{framework (derived)}} \times \underbrace{\frac{\hbar_{\text{SI}}}{\tau_0 c_{\text{SI}}^2}}_{\text{SI bridge (one anchor)}} \tag{28}$$

The same universal factor m_0^{SI} multiplies every framework mass. It, therefore, cancels in any ratio: for any two fermions i, j ,

$$\frac{m_i^{\text{SI}}}{m_j^{\text{SI}}} = \frac{m_i^{\text{RS}}}{m_j^{\text{RS}}} \tag{29}$$

So, the SI anchor τ_0 cancels in every mass ratio: the charged-lepton ratios m_μ/m_e and m_τ/m_μ remain pure dimensionless predictions.

10. Machine Verification

The derivation chain and mass bounds are partially certified in Lean 4 [31] within the IndisputableMonolith library. The following results compile with zero sorry:

- Algebraic structure: the electron structural mass $m_{\text{struct}}(e) = 2^{-22} \cdot \varphi^{51}$ is proved as a theorem from cube-geometric definitions.
- Dimension forcing: $W_{\text{endo}}(D) = 17 \Leftrightarrow D = 3$ for all $D \geq 1$, by case analysis and exponential dominance.
- Gap function: within the affine-log family with $b > 1$, the normalizations $R(0) = 0$, $R(1) = 1$, and $R(-1) = -2$ uniquely force $(a, b, c) = (1/\ln\varphi, \varphi, 0)$.
- Charge index: $Z_l = 1332$ from polynomial minimality with $\tilde{Q}_e = -6$.
- Baseline rung: $r_e = A + 1 = 2$ from active-edge count.
- Coherence exponent: $2^D - D = 5$ at $D = 3$.
- Mass bounds (Lean-certified interval arithmetic): $105.5 < m_\mu^{\text{pred}} < 105.9$ MeV and $1774 < m_\tau^{\text{pred}} < 1779$ MeV, giving relative errors $< 0.3\%$ (muon) and $< 0.2\%$ (tau).

Scope of certification. The Lean certificates verify algebraic structure and coarse interval bounds. The tight numerical values ($m_\mu \approx 105.658$ MeV, relative error $\sim 10^{-6}$) reported in Table 5 come from floating-point evaluation, not from the Lean certificates. The distinction is recorded so that readers can assess which claims are machine-verified and which rest on numerical computation. In particular, machine verification in the present manuscript certifies the RCL cost theorem, the cube arithmetic, the sector-scale bookkeeping, the charge-index arithmetic, and interval bounds for the stated formulas. The physical interpretation of the four MODEL-LAYER structural rules—the affine-log family for the rung correction R , the polynomial class for the charge index Z , the electron-break δ_e rule, and the generation-step refinement rule—is documented as structural input to the refinement layer rather than as an independent theorem of the Standard Model.

11. Conclusions

Starting from the Recognition Composition Law and three regularity conditions, this paper derives the charged-lepton mass chain with zero continuously adjustable parameters and a finite set of discrete structural inputs, each traceable to a named geometric quantity from the 3-cube Q_3 . The electron fixes the single SI calibration anchor τ_0 ; the muon and tau masses then follow as forward tests of the same discrete backbone and refinement layer. We make the role of the electron explicit: τ_0 is a unit-conversion anchor rather than a fitted physical parameter, and it cancels identically from every charged-lepton ratio. The anchor-free predictions tested here are the dimensionless ratios m_μ/m_e and m_τ/m_μ , which use no muon or tau input. Equivalently, the absolute electron mass is used once to set the SI scale, but the predictive content of the charged-lepton chain is contained in ratios that are independent of τ_0 , so the “zero continuously adjustable parameters” claim refers to the dimensionless lepton spectrum and not to the SI unit assignment. Table 5 gives the floating-point comparison; Section 10 gives the Lean-certified interval bounds.

The three questions posed in the Introduction are answered explicitly: (i) Tr2 forces $\varphi = (\sqrt{5} + 1)/2$ as the unique mass-hierarchy base, without postulating it (Section 3); (ii) the generation offsets $\{0, E_{passive}, W\} = \{0, 11, 17\}$ follow from the single cube identity $W = E_{passive} + F$ at $D = 3$, proved intrinsically in Appendix C, specifically Tr7; and (iii) the electron mass fixes the one SI anchor τ_0 , making the muon and tau masses genuine forward predictions (Section 9).

The framework presented here is primarily structural and combinatorial; it does not yet provide a complete microscopic theory connecting these geometric properties directly to spacetime dynamics. For comparison with Standard Model phenomenology, the natural reading is that the cube-geometric construction constrains the discrete flavor pattern entering the effective charged-lepton Yukawa sector, while electroweak symmetry breaking retains its standard role of converting those dimensionless couplings into physical masses through the Higgs vacuum expectation value. In that sense, the present framework is not offered as a replacement for the Higgs mechanism but as a proposed structural constraint on the hierarchy of effective lepton couplings prior to electroweak conversion. This manuscript should, therefore, be read as a structural derivation of the charged-lepton pattern, together with a one-anchor SI calibration scheme, rather than as a complete microscopic theory of flavor. Equivalently, this paper proposes a recognition-geometric constraint on the effective charged-lepton Yukawa hierarchy. It does not attempt to derive the Higgs vacuum expectation value, the electroweak gauge group, or the full renormalized Standard Model flavor Lagrangian in this manuscript. The status of the construction within established physics is, therefore, comparable to that of structured Yukawa ansätze in flavor-physics phenomenology—Froggatt–Nielsen abelian charge assignments, modular-symmetry constructions, or texture-zero patterns—which postulate a structural rule for the Yukawa matrix without deriving it from a more fundamental gauge dynamics. What distinguishes the present construction from those alternatives is that its structural inputs are a finite list of integer cube counts together with the single externally fixed dimensionless constant α and that none of those inputs is a continuously adjustable real-valued coefficient. In the Standard Model effective Lagrangian $L \supset y_{ij} \bar{L}_i \Phi e_{Rj} + h \cdot c$, electroweak symmetry breaking converts the dimensionless Yukawa eigenvalues into physical masses through $m_i = y_i v / \sqrt{2}$ with $v \simeq 246 \text{ GeV}$; the Recognition construction constrains the pattern of those eigenvalues, not the Higgs vev or the gauge structure that produces it. The derivation does not extend below Tr0 and Tr1–Tr8; it does not cover the quark or neutrino sectors, and a full SI calibration analysis is deferred. Not every sub-leading digit in the refined mass ratios is yet secured by a formal theorem. This distinction matters scientifically: if the

refinement layer is later strengthened, the whole charged-lepton chain tightens; if some refinement term is revised, the integer cube backbone still remains intact.

The principal structural results are:

1. The mass law $m = A_s \cdot \varphi^{r-8+R(Z)}$ is derived from the structural chain Tr0–Tr8 together with the cube integers: the sector scale, rung, period offset, and charge correction each have an explicit theorem-level source (Section 5).
2. The sector mass scale for charged leptons is $A_l = 2^{-22} E_{coh} \varphi^{62}$; both exponents come from cube combinatorics, not from fitting to observed masses (Section 5).
3. The lepton charge index $Z_l = 1332$ is fixed by cube-face integerization and polynomial minimality; it is the single integer entering the rung correction R for all three charged leptons (Section 6). The face integerization and polynomial minimality rule are structural choices inside a finite charge-index class, not a claim that the full electromagnetic gauge sector has been dynamically derived here.
4. The rung correction $R(Z) = \log_{\varphi}(1 + Z/\varphi)$ is fixed by three-point calibration within the affine-log family used in this paper (Theorem A6); the nontrivial normalization $R(-1) = -2$ is explained in Appendix D.
5. The lepton mass chain $(\delta_e, S_{e \rightarrow \mu}, S_{\mu \rightarrow \tau})$: the leading coefficients $(E_{passive}, F)$ come from cube geometry, while the refinement terms are derived at the level of the present paper (Thms. 2.3–2.5 [35,36]). Once the electron fixes the τ_0 calibration, the muon and tau become forward tests of the same chain (Table 5; Section 9). The word “forward” here refers to the fact that no muon or tau mass value is used after the electron anchor and the fixed refinement rule have been specified. The result is, therefore, a highly constrained prediction within the Recognition framework, not an unconstrained curve fit.
6. $W = 17$ arises intrinsically as $E_{passive} + F = 11 + 6$ (A3); equivalently, the lepton-sector formulas are invariant under the replacement $W \mapsto E_{passive} + F$ (see Appendix C, Tr7).
7. Global closure: Appendix B classifies each load-bearing ingredient as either theorem-level, downstream-derived, calibration-dependent, or conventional (Sections 5–8). Every input is, therefore, a discrete integer, a fixed external constant, or a unit convention, not a continuously adjustable fitted dial.

The framework, therefore, makes sharp anchor-free tests in the lepton sector: the ratios m_{μ}/m_e and m_{τ}/m_{μ} are pure structure tests because τ_0 cancels identically. Both are already directly testable against existing precision data. The central claim of this work is that the RCL, the Tr-chain, the convention $\lambda = 1$, the fixed external constant α , and one empirical SI anchor τ_0 determine the three charged-lepton masses with zero continuously adjustable dimensionless parameters; every discrete structural input is traceable to a named geometric quantity from Q_3 . Equivalently, the framework makes three nontrivial falsifiable claims: the leading integer backbone must give the two generation gaps through φ^{11} and φ^6 ; the fixed refinement layer must move those leading ratios into the Lean-certified intervals reported above; and the same fixed rules must use one electron SI anchor without any muon or tau mass input. Failure of any one of these checks would rule out the present charged-lepton realization rather than merely refit one of its coefficients because no such coefficient is available.

The following items remain as future work: a deeper derivation of the affine-log rung-correction family, the sub-leading α order-matching rules, and the charge-index face assignment from the next dynamical layer of Recognition theory; providing a dynamical mechanism connecting the cube-geometric structure to Standard Model electroweak symmetry breaking; and extending the framework to the quark and neutrino sectors.

These appendices serve as reference material for the main text rather than a second narrative; each appendix is cited where it is used.

Author Contributions: Conceptualization, J.W.; methodology, J.W. and E.A.; software, J.W. and E.A.; validation, J.W. and E.A.; formal analysis, J.W. and E.A.; investigation, J.W. and E.A.; writing—original draft preparation, J.W.; writing—review and editing, J.W. and E.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Minimal Numerical Reproduction Sequence

This appendix records the minimal numerical sequence needed to reproduce this paper’s charged-lepton outputs from already-derived structural inputs. It introduces no new structural claim.

Appendix A.1. Structural Inputs

Start from the cube-combinatorial inputs:

$$\varphi = \frac{1 + \sqrt{5}}{2}, D = 3, E_{\text{passive}} = 11, F = 6, W = 17.$$

The refinement layer also uses the fixed external constant α^{-1} . Numerically, $\alpha^{-1} = 137.0348851720$; this quantity is not derived from cube combinatorics. It is introduced phenomenologically, with its numerical value taken from experiment rather than from the RCL, the Tr0–Tr8 theorem chain, or cube combinatorics; only its algebraic role in the fixed refinement formulas is structural.

Appendix A.2. Generation Steps

Compute the canonical generation steps:

$$\begin{aligned} S_{e \rightarrow \mu} &= E_{\text{passive}} + \frac{1}{4\pi} - \alpha^2 \approx 11.079524219, \\ S_{\mu \rightarrow \tau} &= F - \frac{2W+D}{2}\alpha \approx 5.864997880. \end{aligned}$$

Appendix A.3. Ratios and SI Calibration

Then, evaluate the mass ratios:

$$\frac{m_{\mu}}{m_e} = \varphi^{S_{e \rightarrow \mu}} \approx 206.768, \quad \frac{m_{\tau}}{m_{\mu}} = \varphi^{S_{\mu \rightarrow \tau}} \approx 16.816.$$

The absolute masses are recovered only at the SI stage, using Section 9 and the single calibration $m_e^{\text{exp}} = 0.51099895$ MeV to fix τ_0 . No charged-lepton mass other than the electron calibration appears on the right-hand side.

Appendix B. Derivation Bookkeeping Summary

This appendix records the bookkeeping status of every load-bearing ingredient used in the charged-lepton framework. Each item is classified according to whether it is theorem-level, downstream-derived, calibration-dependent, or conventional. The purpose of this appendix is organizational: it collects in one place how the final mass formula depends on proved inputs, derived consequences, one empirical anchor, and one explicit convention.

Bookkeeping Categories

The 13 audited components fall into five bookkeeping categories:

1. FORCED: necessary consequence of Tr0 and Tr1–Tr8, with no additional structural input.
2. DERIVED: obtained from cube identities or downstream constructions once the structural theorems are fixed.
3. MODEL-LAYER: fixed structural rules that define the charged-lepton refinement class, with no continuously adjustable coefficients.
4. Calibration: the single empirical SI anchor τ_0 , fixed by the electron mass and cancelled in mass ratios.
5. Convention: the transparent normalization choice $\lambda = 1$, with no physical effect.

We use “FORCED” only for consequences of the stated theorem chain, “DERIVED” for quantities fixed once the cube data and sector bookkeeping are in place, and “MODEL-LAYER” for fixed structural rules that define the charged-lepton refinement class. The four MODEL-LAYER items—the affine-log family for the rung correction R , the polynomial class for the charge index Z , the electron-break δ_e rule, and the generation-step refinement rule—are not continuously adjustable parameters, but they are structural choices that define the model class being tested.

Theorem A1 (Component derivation audit). *The 13 audited components of the charged-lepton mass framework fall into five categories: FORCED, DERIVED, MODEL-LAYER, calibration anchor, and convention. No continuously adjustable fitted parameter appears, and every item carries explicit derivation status.*

Proof. We verify each component against its derivation section. Tables A1 and A2 record the FORCED, DERIVED, and MODEL-LAYER items; the calibration anchor and convention are then stated separately below.

FORCED (3):

Table A1. FORCED components of the master mass law: necessary consequences of Tr0 and Tr1–Tr8 with no free parameters.

Component	Theorem	Section
$E_{passive} + F = W \Leftrightarrow D = 3$	Tr7 (cube-arithmetic identity at $D = 3$)	Section 4
Integer rungs	Tr5 (discreteness)	Section 3
Mass-scale form A_S	Tr4 + Tr2 + Tr8 (mass-law structure)	Section 5

DERIVED and MODEL-LAYER (8):

Table A2. DERIVED and MODEL-LAYER components of the master mass law. DERIVED items are fixed by cube combinatorics and sector bookkeeping. MODEL-LAYER items are fixed structural rules inside the charged-lepton refinement class. The single calibration anchor is listed separately below.

Component	Derivation	Status	Section
Cube integers	$D = 3$ cube combinatorics; yields $V, E, F, E_{passive}, W, A$	DERIVED	Section 4
Sector scales	Cube-partition bookkeeping	DERIVED	Section 5
Baseline rung $r_e = 2$	Active-edge count plus one stable rung	DERIVED	Section 7
$W = 17$	Cube identity $E_{passive} + F = 11 + 6$	DERIVED	Appendix C
Charge index Z_l	Face integerization and minimal even polynomial class	MODEL-LAYER	Section 6
Rung correction R	Three-point calibration within affine-log family	MODEL-LAYER	Appendix D
Electron break δ_e	Fixed topological and α refinement rule	MODEL-LAYER	Section 7
Generation steps	Fixed edge/face refinement rule with $1/(4\pi), -\alpha^2$ and $(2W + D)/2$	MODEL-LAYER	Section 7

Convention (1): $\lambda = 1$ —notational convention; sets cost-unit scale. Zero physical consequence: all φ -exponents, mass ratios, and dimensionless predictions are strictly λ -independent.

Calibration (1): τ_0 (SI seconds per period step)—the single empirical calibration anchor, fixed by matching the full predicted electron mass m_e^{pred} to the measured SI value, equivalently $\tau_0 = 2^{-22} \varphi^{51+R(1332)-\delta_e} \hbar_{\text{SI}} / (m_e^{\text{SI}} c_{\text{SI}}^2)$. This is a unit-conversion anchor, not an adjustable physical coupling: the muon and tau formulas use no electron, muon, or tau mass value other than through this single SI scale-setting step, and τ_0 cancels identically from every charged-lepton ratio. The “zero continuously adjustable parameters” statement, therefore, applies to the dimensionless lepton ratios m_μ/m_e and m_τ/m_μ , which are tested against experiment with no further empirical input.

The tally matches Equation (23):

$$3 \text{ FORCED} + 4 \text{ DERIVED} + 4 \text{ MODEL-LAYER} + 1 \text{ calibration} + \\ 1 \text{ convention} + 0 \text{ pending} = 13.$$

The refinement layer uses only the fixed external constant α , not an adjustable fit parameter. The claim of zero continuously adjustable fitted parameters is, therefore, narrower than a claim of zero assumptions. The framework has assumptions and structural rules, all of which are exposed above; the point is that none of them is a real-valued dial tuned to the muon or tau masses. \square

Theorem A2 (Parameter-freedom). *The structural core of the mass framework contains exactly zero continuously adjustable fitted parameters, relying instead on discrete structural choices traceable to Q_3 geometry. The later refinement layer also uses the fixed external constant α , not a tuned dial. Every structural input is either:*

1. An integer from the cube vocabulary (Table 2), namely $V, E, F, A, E_{\text{passive}}, W, T_{\text{min}}$;
2. The golden ratio φ (Tr2, forced);
3. The mathematical constants $\pi, \sqrt{2}$, and $\ln \varphi$ (determined by φ and standard analysis).

Proof. Enumerate all inputs appearing in any formula of Sections 5–7:

- Cube integers: $V = 8, E = 12, F = 6, A = 1, E_{\text{passive}} = 11, W = 17$. Each is a computed output of $D = 3$ (Section 4).
- Golden ratio: $\varphi = (1 + \sqrt{5})/2$. Tr2 (unique $r > 1$ with $r^2 = r + 1$).
- Derived constants: $\pi, \sqrt{2}$, and $\ln \varphi$ —mathematical consequences of the cube geometry and Tr1/Tr2; not fit parameters.
- Generation offsets: $\{0, 11, 17\}$ —the integers $\{0, E_{\text{passive}}, W\}$, already in the cube vocabulary.
- Sector anchors:
- $r_e = 2$: (derived, [33])—the single lepton-scale anchor (Section 7).
- None of these is a continuously adjustable real number: $r_e = 2$ is a single discrete integer anchor, not a tunable coupling.
- α : a fixed external dimensionless input used only in the refinement layer, not a fit parameter.

Every structural input is a specific integer, a specific algebraic irrational (φ), or a specific transcendental (π)—all uniquely determined. The discrete inputs (sector anchors, coupling assignments, phase offsets) are not free parameters: they are explicitly enumerated and documented in Sections 5–7. The refinement layer uses α only as a fixed external constant. There is no continuously adjustable dial, and the framework makes no use of any fitting procedure. Within the stated Recognition rules, the result is predictive in the

following precise sense: after the electron fixes τ_0 , no muon or tau mass information enters the formulas. \square

Contrast with the Standard Model. The Standard Model's fermion sector has at least 15 continuously adjustable Yukawa couplings (one per fermion mass, plus mixing angles and phases). The three lepton Yukawa couplings are replaced by one discrete integer ($r_e = 2$, Lean-verified [33]) plus cube-derived integers, plus one calibration anchor (τ_0). All lepton mass ratios require zero additional fitted inputs (α enters only as a fixed external constant); τ_0 sets the absolute energy scale. This is a genuine reduction from three continuous Yukawa inputs to zero continuously adjustable fitted parameters in the lepton sector.

Theorem A3 (Internal consistency of the mass framework). *Within the specific class of frameworks defined by the following constraints, any zero-parameter framework that:*

1. (E1) uses the forcing chain $\text{Tr}0$ and $\text{Tr}1\text{--}\text{Tr}8$ as its foundation;
2. (E2) organizes masses on a φ -ladder with the period baseline offset -8 ;
3. (E3) decomposes the mass law into sector mass scale $\times \varphi$ -power of (rung + rung correction);
4. (E4) uses the cube-partition principle for mass-scale assignment;
5. (E5) derives the Z-map from gauge-invariant polynomial minimality,

must produce the same mass law, the same rung correction, and the same sector mass-scale assignments as the present framework.

Proof. Each constraint eliminates all alternatives:

- (E1) + (E2): $\text{Tr}2$ establishes φ as the hierarchy base (forced); $\text{Tr}8$ establishes the period $T_{\min} = 8$, and the offset -8 is a coordinate convention ((4), derived). Together, they fix the φ -ladder and its origin.
- (E3): The decomposition is fixed in Section 5. Under φ -scaling, sector factorization, octave baseline, and charge additivity, the mass law is forced to the form (10) with no residual freedom.
- (E4): Section 5 fixes the sector assignment used in this paper through the cube-partition principle itself. A different permutation of the B_{pow} values would, therefore, define a different sector assignment rather than a second realization of the same constrained framework. Comparison with the observed ordering is then an empirical check on that assignment, not an extra fitted ingredient.
- (E5): Given $k = F(3)$ (derived) and the derived completeness and color-offset inputs, the Z-tuple predicate has a unique solution $(6, 1, 1, 0)$ for the lepton sector ($c = 0$; no color) (Section 6).

Since each structural ingredient is uniquely determined within its constraint, and the mass law is a function of these ingredients alone, no alternative framework satisfying (E1)–(E5) produces different mass predictions. \square

Scope of the consistency claim. This consistency claim is intentionally conditional: it is a uniqueness statement within the class of frameworks satisfying (E1)–(E5), not a claim that no other conceivable physical theory could reproduce the same data. The constraints (E1)–(E5) are specific to this framework and encode its discrete structural choices; the theorem's content is that no additional freedom exists once those choices are made. Within that restricted class, agreement with experiment tests the whole constraint set rather than a tunable parameter choice.

Theorem A4 (Derivation chain from RCL to lepton masses). *Given:*

1. *The Recognition Composition Law (1) with normalization and calibration;*

2. Standard regularity hypotheses (continuity, smoothness);
3. The single empirical calibration anchor τ_0 (see Section 9);
4. The two derived discrete inputs: lepton baseline $r_e = 2$ and gap calibration $R(-1) = -2$, both machine-verified [33].

the three charged lepton masses (e, μ, τ) are uniquely determined by the derivation developed in Sections 5–9.

Proof. Section 5 fixes the form of the mass law; Section 6 fixes the lepton charge index; Section 7 fixes the rungs and the two generation steps; and Section 9 adds the universal SI bridge through the single anchor τ_0 . Once these ingredients are fixed, Equations (20)–(22) determine the three charged-lepton masses uniquely. No further continuous input enters. \square

Closure summary.

Table A3. Summarizes the three global closure properties. The three global properties of the mass framework. “Completeness” counts structural components; sub-inputs within DERIVED and MODEL-LAYER components are documented in Sections 5–7.

Property	Statement
Completeness	All 13 audited components: FORCED (3), DERIVED (4), MODEL-LAYER (4), calibration anchor (1), convention (1). Zero continuously adjustable fitted parameters.
Exclusivity	No alternative zero-parameter framework satisfying the same structural constraints (E1)–(E5) plus the inputs can produce different mass predictions.
Inevitability	Given the RCL, regularity, one SI unit anchor, and the discrete sector inputs plus the fixed refinement rule, the charged lepton masses are uniquely determined within the stated framework.

The framework is, therefore, not a model with fitted parameters but a structured derivation from the RCL in which every load-bearing ingredient carries an explicit status tag. Agreement or disagreement with experiment tests the functional equation and the discrete inputs collectively.

Appendix C. Tr0 and Tr1–Tr8: Statements, Proof Sketches, and Formal References

This appendix collects the Tr0–Tr8 statements and proof sketches in the same logical order as Table 1 in the main text: Tr1 first (proved directly from the axioms P1–P4); then Tr2 (algebraic); then Tr3, Tr4, Tr5, and Tr6 (all consequences of Tr1, with Tr6 also using Tr2); and then Tr7 and Tr8 (geometric). Every theorem, therefore, cites only theorems that appear before it; no forward references remain.

The main formal references for the theorem chain are the public Lean 4 and GitHub sources [29,30,39,40].

Scope note. These formal-support citations are included primarily for dependency audit. The human proofs remain the primary presentation here; the formal citations record where the same steps appear in the larger verified chain. The independently auditable public portion is summarized in [31], and the chain-level forcing reference is [38].

Appendix C.1. Tr0—Classical Logic

Role in the derivation: standing logical assumption.

Statement. The framework operates within standard classical mathematics: the law of non-contradiction, excluded middle, and the standard rules of real analysis are assumed. No derivation of classical logic from $J(x)$ is claimed.

Proof. None required. Tr0 is a standing assumption, not a theorem. The obligation is to confirm that Tr1–Tr8 do not assume any logic stronger than classical first-order logic over R . \square

Appendix C.2. Tr1—The Cost Functional Is Uniquely Determined

Role in the derivation: key load-bearing theorem.

Statement. The RCL.

$$J(xy) + J(x/y) = 2J(x)J(y) + 2J(x) + 2J(y), \quad x, y > 0, \quad (\text{A1})$$

together with $J(1) = 0$, continuity on $R_{>0}$, and the curvature normalization

$$\lim_{t \rightarrow 0} 2J(e^t)/t^2 = 1, \text{ uniquely forces.}$$

$$J(x) = \frac{1}{2} \left(x + \frac{1}{x} \right) - 1. \quad (\text{A2})$$

Proof. Write $x = e^t$, $y = e^s$, and set $G(t) = 1 + J(e^t)$. Then, the RCL becomes

$$G(t+s) + G(t-s) = 2G(t)G(s),$$

the continuous d'Alembert functional equation with $G(0) = 1$. Its continuous even solutions are $G(t) = \cosh(ct)$ for some constant c . The curvature normalization $\lim_{t \rightarrow 0} 2J(e^t)/t^2 = 1$ forces $c = 1$, so $G(t) = \cosh t = 1/2(e^t + e^{-t})$. Therefore, $J(e^t) = \cosh t - 1 = 1/2(e^t + e^{-t}) - 1$, and hence, $J(x) = 1/2(x + x^{-1}) - 1$. \square

Appendix C.3. Tr2— φ Is the Unique Hierarchy Base

Role in the derivation: elementary algebraic consequence.

Statement. The unique $r > 1$ satisfying $r^2 = r + 1$ is $\varphi = 1 + \sqrt{5}/2$.

Proof. Roots of $r^2 - r - 1 = 0$: $\varphi \approx 1.618$ and $\hat{\varphi} \approx -0.618$. Only $\varphi > 1$ yields a well-ordered hierarchy. Irrationality: $\sqrt{5} \notin Q$ (standard proof by infinite descent). \square

Appendix C.4. Tr3—Every Physical State Has Strictly Positive Mass

Role in the derivation: consequence of Tr1 (uses the explicit form of J).

Statement. For every stable physical state, $m > 0$. Equivalently, $J(x) \rightarrow +\infty$ as $x \rightarrow 0^+$.

Proof. By the explicit form from Tr1,

$$J(x) = \frac{1}{2} \left(x + \frac{1}{x} \right) - 1 = \frac{(x-1)^2}{2x}$$

Hence, $J(x) \geq 0$ for all $x > 0$, with equality only at $x = 1$, and $J(x) \rightarrow +\infty$ as $x \rightarrow 0^+$. A zero-mass state would, therefore, lie at infinite cost and is excluded. \square

Appendix C.5. Tr4—Paired Transitions (Reciprocal Symmetry)

Role in the derivation: elementary reciprocal identity (consequence of Tr1).

Statement. $J(x) = J(1/x)$ for all $x > 0$.

Proof. Direct: $J(1/x) = 1/2(x^{-1} + x) - 1 = J(x)$. \square

Appendix C.6. Tr5—Physical Evolution Proceeds in Discrete Steps

Role in the derivation: consequence of Tr1 and Tr3 (uses $J(x) > 0$ for $x \neq 1$).

Statement. Given that every physical state transition is nontrivial ($x \neq 1$, i.e., the ratio departs from the identity), physical transitions occur in a finite number of discrete elementary steps per bounded interval.

Proof. By Tr3, each transition satisfies $J(x) > 0$. Let $\epsilon = \inf_{x \in S_{phys}} J(x) > 0$. Infinitely many transitions in a bounded interval would yield total cost $\geq \sum_{k=1}^{\infty} \epsilon = \infty$, contradicting finite-cost evolution. Hence, the time axis is discrete with an elementary period step, later realized in SI units through the calibration anchor τ_0 . \square

Appendix C.7. Tr6—Only Finite-Cost States Are Physical

Role in the derivation: (from Tr1 and Tr2).

Statement. All physical masses are organized by discrete φ -ladder positions $r \in \mathbb{Z}$; equivalently, the structural part of a mass can be written as $m_{struct} = \kappa \varphi^r$.

Proof. Tr1 gives $J(x)$ finite if and only if $x > 0$ is finite. Tr2 establishes φ as the unique hierarchy base. Finite cost, therefore, organizes the structural mass ratios by integer ladder positions, which is the form used throughout the main text before sector-specific and charge-specific refinements are written explicitly. \square

Appendix C.8. Tr7— $D = 3$ Is the Unique Admissible Dimension

Role in the derivation: (machine-verified).

Statement. $W_{endo}(D) = E_{passive}(D) + F(D) = 17$ if and only if $D = 3$, uniquely fixing $V = 8, E = 12, F = 6, A = 1, E_{passive} = 11, W = 17$.

Proof. Since $E_{passive}(D) = D2^{D-1} - 1$ and $F(D) = 2D$, one has

$$W_{endo}(D) = D2^{D-1} + 2D - 1.$$

At $D = 3$, this gives $W_{endo}(3) = 3 \cdot 4 + 6 - 1 = 17$. For $D = 1, 2$, one gets $W_{endo}(1) = 2$ and $W_{endo}(2) = 7$, while for every $D \geq 4$,

$$W_{endo}(D) \geq 4 \cdot 2^3 + 8 - 1 = 39 > 17.$$

Hence, $W_{endo}(D) = 17$ holds only at $D = 3$, which then fixes $V = 8, E = 12, F = 6, A = 1, E_{passive} = 11$, and $W = 17$.

At the selected dimension, the same cube data give the bookkeeping identity.

$$W_{endo}(3) = E_{passive}(3) + F(3) = (3 \cdot 2^2 - 1) + (2 \times 3) = 11 + 6 = 17, \quad (A3)$$

so every W -bearing formula in the charged-lepton derivation is unchanged by the replacement $W \mapsto E_{passive} + F$. The biconditional $W_{endo}(D) = 17 \Leftrightarrow D = 3$ for all $D \geq 1$ is machine-verified by case analysis for $D \leq 4$ and exponential dominance for $D \geq 5$ [31]. \square

Appendix C.9. Tr8—Minimal Closed Path on Q_3 Has Length 8

Role in the derivation: key combinatorial step.

Statement. The shortest closed Hamiltonian cycle on Q_3 has exactly eight steps, giving $T_{min} = 8$ and hence the exponent offset -8 .

Proof. Q_3 is bipartite with parts of size 4. Any Hamiltonian cycle alternates between the two parts and requires at least $2 \times 4 = 8$ edges. An explicit 8-cycle (binary Gray code) attains this bound, so $T_{min} = 8$. \square

Appendix D. Complete Derivation of the Rung Correction Function $R(Z)$

This appendix fixes the rung correction function $R(Z)$ within the affine-log family used in the main text. It explains why that family is the natural bridge from multiplicative J -costs to additive φ -ladder shifts, determines the parameters from the three normalization conditions, and records the analytic properties used later in the charged-lepton hierarchy arguments.

The uniqueness statement established in this appendix is uniqueness within the affine-log family $R(x) = a \ln(1 + x/b) + c$, $b > 1$, not uniqueness among all conceivable monotone functions on Z . The affine-log family itself is a MODEL-LAYER structural choice motivated by the analytic shape conditions: (i) it converts multiplicative J -cost data into additive φ -ladder shifts, (ii) it is strictly monotone and sub-linear at large $|Z|$, and (iii) it vanishes at $Z = 0$. Once the family is adopted, the three normalization conditions $R(0) = 0$, $R(-1) = -2$, and $R(1) = 1$ uniquely fix the parameters $(a, b, c) = (1/\ln \varphi, \varphi, 0)$. The nontrivial condition $R(-1) = -2$ is a Tr2-derived consequence of the golden-ratio identity $1 - 1/\varphi = \varphi^{-2}$, not a free input, so the only structural choice exercised here is the choice of the affine-log family itself.

Section 5 states the final rung correction function:

$$R(Z) = \log_{\varphi} \left(1 + \frac{Z}{\varphi} \right).$$

The key nontrivial step is the condition $R(-1) = -2$, which forces $b = \varphi$ through the golden-ratio identity [41] $1 - 1/\varphi = \varphi^{-2}$. Once b is fixed, the remaining two normalizations determine the other parameters immediately.

Appendix D.1. Choosing the Logarithmic Candidate

The rung correction function $R(Z)$ maps integer charge indices to real-valued shifts on the φ -ladder. The first question is, therefore, which functional family can play this role naturally.

1. It must be defined on Z (discrete input from the Z -map).
2. It must return a real-valued φ -ladder shift (continuous output).
3. It must vanish at $Z = 0$ (neutral baseline).
4. It must grow sub-linearly (a linear gap $\propto Z$ would make the charge correction equivalent to an additive rung shift, conflating two structurally distinct contributions).
5. It must be monotone (higher charge index \rightarrow larger correction).

Conditions (1)–(5) do not by themselves single out a unique functional family. In the present paper, we, therefore, restrict attention to the affine-log family suggested by the cost functional J : a charged state with index Z carries a multiplicative deviation cost, and the corresponding shift on the φ -ladder is, therefore, represented most naturally by a logarithm. This does not claim uniqueness among all monotone sub-linear maps; the uniqueness

proved below is uniqueness within the affine-log family $R(x) = a \cdot \ln(1 + x/b) + c$ used in the main text.

Requirements (4) and (5) are satisfied for any $a > 0, b > 0$; requirement (3) then forces $c = 0$ ($R(0) = 0 \Rightarrow c = 0$). The remaining two free parameters a and b are fixed by the normalization conditions below:

The logarithmic candidate family has three free parameters a, b, c with $b > 0$:

$$R(x) = a \cdot \ln\left(1 + \frac{x}{b}\right) + c, \quad x > -b. \quad (\text{A4})$$

For the calibration point $Z = -1$ and all nonnegative charge indices, this formula is well-defined whenever $b > 1$.

The three normalization conditions fix the parameters in the following logical order: $R(0) = 0$ fixes $c = 0$; $R(-1) = -2$ then forces $b = \varphi$ (the nontrivial step); and $R(1) = 1$ completes the calibration by fixing $a = 1/\ln\varphi$.

Normalization 1: $R(0) = 0$ forces $c = 0$. Substituting $x = 0$ gives $R(0) = a \ln 1 + c = c$, so $R(0) = 0$ forces $c = 0$. A neutral particle ($Z = 0$) receives no charge correction.

Normalization 2: $R(-1) = -2$ forces $b = \varphi$. With $c = 0$, the conditions $R(1) = 1$ and $R(-1) = -2$ together determine b . This is the nontrivial step; the proof is given as Theorem A5 below. The condition $R(-1) = -2$ is a normalization derived from Tr2: the golden-ratio identity $1 - 1/\varphi = \varphi^{-2}$ implies $R(-1) = \log_{\varphi}(\varphi^{-2}) = -2$. The three-point calibration is machine-verified [31]. Within the present paper, this fixes the backward unit-charge calibration used by the candidate gap family. The use of a logarithm is motivated by the compatibility requirement that multiplicative cost data be represented as additive shifts on the φ -ladder: the logarithm is the unique map that converts multiplicative composition into additive increments in exponent space. The affine-log form is then the smallest scale-plus-offset extension consistent with the neutral baseline $R(0) = 0$, the unit-step normalization $R(1) = 1$, and the derived backward calibration $R(-1) = -2$.

Normalization 3: $R(1) = 1$ forces $a = 1/\ln\varphi$. Once $c = 0$ and $b = \varphi$ are established, the condition $R(1) = 1$ gives $a \cdot \ln(1 + 1/\varphi) = 1$. The golden-ratio identity $1 + 1/\varphi = \varphi$ yields $a = 1/\ln(\varphi)$, the change-of-base factor to φ -log units. A unit charge increment produces a unit rung shift.

Theorem A5 ($R(-1) = -2$ forces $b = \varphi$). Let $R(x) = a \cdot \ln(1 + x/b) + c$ with $b > 1$ and $c = 0$. Suppose $R(1) = 1$ and $R(-1) = -2$, then $b = \varphi$.

Proof. From $R(0) = 0$: $c = 0$. From $R(1) = 1$: $a \cdot \ln(1 + 1/b) = 1$. From $R(-1) = -2$: $a \cdot \ln(1 - 1/b) = -2$.

Dividing the third equation by the second:

$$\frac{\ln(1 - 1/b)}{\ln(1 + 1/b)} = -2.$$

Hence, $\ln(1 - 1/b) = -2\ln(1 + 1/b) = \ln\left[(1 + 1/b)^{-2}\right]$, so

$$1 - \frac{1}{b} = \left(1 + \frac{1}{b}\right)^{-2}.$$

Setting $u = 1/b$ and multiplying both sides by $(1 + u)^2$:

$$(1 - u)(1 + u)^2 = 1$$

Expanding: $1 + u - u^2 - u^3 = 1$, hence $u(1 - u - u^2) = 0$. Since $b > 1$ implies $u = 1/b \in (0, 1)$, we have $u \neq 0$, so

$$u^2 + u - 1 = 0 \Rightarrow u = \frac{-1 + \sqrt{5}}{2} = \frac{1}{\varphi}.$$

Therefore, $b = 1/u = \varphi$. \square

Why $R(-1) = -2$ is a derived condition.

The value $R(-1) = -2$ is not a free input. Tr2 (additive scale closure, $\varphi^2 = \varphi + 1$) implies the golden-ratio identity $1 - 1/\varphi = \varphi^{-2}$, so the backward-step argument is itself a φ -ladder element. Therefore,

$$R(-1) = \log_{\varphi}(\varphi^{-2}) = -2.$$

This is the normalization used in Theorem A5; the corresponding machine verification is [41].

Appendix D.2. The Final Rung Correction Function

Combining the three normalization results:

Theorem A6. (The rung correction function $R(Z)$ is uniquely fixed by three-point calibration within the affine-log family).

Within the logarithmic family $R(x) = a \cdot \ln(1 + x/b) + c$ with $b > 1$, the three normalization conditions.

$$R(0) = 0, R(1) = 1, R(-1) = -2 \tag{A5}$$

uniquely force $(a, b, c) = (1/\ln\varphi, \varphi, 0)$, giving

$$R(Z) = \ln \frac{(1 + Z/\varphi)}{\ln\varphi} = \log_{\varphi} \left(1 + \frac{Z}{\varphi} \right). \tag{A6}$$

Proof. Normalization 1 gives $c = 0$. Theorem A5 gives $b = \varphi$. Substituting these values into $R(1) = 1$ yields $a = 1/\ln\varphi$. Hence, the three-point calibration uniquely fixes $(a, b, c) = (1/\ln\varphi, \varphi, 0)$. This proof is machine-verified [31]. \square

Remark A1 (Scope of uniqueness). *The uniqueness established above holds within the affine-log family $R(x) = a \cdot \ln(1 + x/b) + c$. That this family is the natural bridge from multiplicative J-costs to additive φ -ladder shifts is a structural postulate motivated by the logarithmic nature of the cost-to-rung conversion, not a theorem derived from T0–T8.*

Note. All three calibration conditions are derived: $R(0) = 0$ and $R(1) = 1$ follow from the neutral-baseline and unit-step structural requirements; $R(-1) = -2$ is confirmed directly as $R(-1) = \log_{\varphi}(1 - 1/\varphi) = \log_{\varphi}(\varphi^{-2}) = -2$, using the golden-ratio identity $1 - 1/\varphi = \varphi^{-2}$ [41].

Appendix D.3. Analytic Properties Used in Hierarchy Arguments

The rung correction R (A6) is a positive, strictly increasing, strictly concave function of Z on the domain $Z > -\varphi$. These follow directly from its closed form: $R'(x) = 1/[(\varphi + x)\ln\varphi] > 0$ and $R''(x) = -1/[(\varphi + x)^2\ln\varphi]$. For large Z , $R(Z) \sim \log_{\varphi}Z - 1$ (logarithmic growth, not linear). The monotonicity and concavity are used in the mass-hierarchy

arguments of Section 7. For the charged-lepton sector, the relevant charge index from Section 6 (Equation (14)) gives $R(1332) \approx 13.95$.

Appendix D.4. Why the Logarithmic Family Is the Right One

Scope of the uniqueness claim. Once $b = \varphi$ is fixed, the remaining two parameters follow immediately from $R(0) = 0$ and $R(1) = 1$. The uniqueness established in this appendix is, therefore, uniqueness within the affine-log family (D1), which is the class used in the main text.

Appendix E. Expanded Mass-Law Exponent

Equation (10) gives the compact form

$$m(i) = A_s \varphi^{r_i - 8 + R(Z_i)}.$$

Here, the explicit φ -exponent is $r_i - 8 + R(Z_i)$, but the prefactor A_s still contains two further φ -contributions through (6) and (5). Substituting those definitions gives the fully expanded form.

$$m_i = 2^{B_{\text{pow}}(s)} \cdot \varphi^{-(D+2)+r_0(s)+r_i-8+R(Z_i)}. \quad (\text{A7})$$

Thus, the binary prefactor $2^{B_{\text{pow}}(s)}$ is separate. This appendix tracks the structural φ -exponent before the electron-specific break δ_e , and the generation-step refinements are split off in Section 7. The full exponent is, therefore, the sum of five additive terms:

1. $-(D + 2)$ from the coherence unit $E_{\text{coh}} = \varphi^{-(D+2)}$; at $D = 3$ this is -5 .
2. $r_0(s)$ from the sector baseline absorbed into A_s ; for charged leptons, $r_0(l) = 62$.
3. r_i from the species rung within the chosen sector.
4. $-8 = -T_{\text{min}}$ from the period offset fixed by (4).
5. $R(Z_i)$ from the charge correction, derived in D.

Charged-Lepton Specialization

For charged leptons, the species rungs are

$$r_e = 2, \quad r_\mu = 2 + E_{\text{passive}} = 13, \quad r_\tau = 2 + W = 19.$$

So, the dominant geometric rung increments are $r_\mu - r_e = E_{\text{passive}} = 11$ and $r_\tau - r_\mu = W - E_{\text{passive}} = F = 6$.

For the charged electron, $Z_e = Z_l = 1332$. The uncorrected electron structural seed isolates the cube-sector exponent before the common charged-lepton correction and the electron break exponent are applied. Its φ -exponent is, therefore,

$$-(D + 2) + r_0(l) + r_e - 8 = -5 + 62 + 2 - 8 = 51, \quad (\text{A8})$$

which reproduces (16). The full charged-electron formula then multiplies this seed by $\varphi^{R(1332)-\delta_e}$ as in (20). This is the only purpose of the present appendix: to display the exponent bookkeeping in one place. The analytic form of $R(Z)$ is derived in Appendix D, and the intrinsic identity $W = 17$ is recorded in Appendix C, Tr7.

References

1. Navas, S.; Amsler, C.; Gutsche, T.; Hanhart, C.; Hernández-Rey, J.J.; Lourenço, C.; Masoni, A.; Mikhasenko, M.; Mitchell, R.E.; Patrignani, C.; et al. Review of Particle Physics. *Phys. Rev. D* **2024**, *110*, 030001. [CrossRef]
2. Froggatt, C.D.; Nielsen, H.B. Hierarchy of Quark Masses, Cabibbo Angles and CP Violation. *Nucl. Phys. B* **1979**, *147*, 277–298. [CrossRef]
3. Ramond, P. *Journeys Beyond the Standard Model*; Westview Press: Boulder, CO, USA, 1999.

4. Koide, Y. A Fermion-Boson Composite Model of Quarks and Leptons. *Phys. Lett. B* **1983**, *120*, 161–165. [[CrossRef](#)]
5. Rodejohann, W.; Zhang, H. Extension of an empirical charged lepton mass relation to the neutrino sector. *Phys. Lett. B* **2011**, *698*, 152–156. [[CrossRef](#)]
6. Georgi, H.; Glashow, S.L. Unity of All Elementary-Particle Forces. *Phys. Rev. Lett.* **1974**, *32*, 438–441. [[CrossRef](#)]
7. Georgi, H. The State of the Art—Gauge Theories. *AIP Conf. Proc.* **1975**, *23*, 575–582.
8. Fritzsch, H. Quark masses and flavor mixing. *Nucl. Phys. B* **1979**, *155*, 189–207. [[CrossRef](#)]
9. Ludl, P.O.; Grimus, W. A complete survey of texture zeros in general and symmetric quark mass matrices. *Phys. Lett. B* **2015**, *744*, 38–42. [[CrossRef](#)]
10. Belfatto, B.; Berezhiani, Z. Minimally modified Fritzsch texture for quark masses and CKM mixing. *J. High Energy Phys.* **2023**, *2023*, 162. [[CrossRef](#)]
11. Fritzsch, H.; Xing, Z.-Z.; Zhang, D. Correlations between quark mass and flavor mixing hierarchies. *Nucl. Phys. B* **2022**, *974*, 115634. [[CrossRef](#)]
12. Altarelli, G.; Feruglio, F. Discrete flavor symmetries and models of neutrino mixing. *Rev. Mod. Phys.* **2010**, *82*, 2701–2729. [[CrossRef](#)]
13. Feruglio, F. Pieces of the flavour puzzle. *Eur. Phys. J. C* **2015**, *75*, 373. [[CrossRef](#)]
14. Ishimori, H.; Kobayashi, T.; Ohki, H.; Shimizu, Y.; Okada, H.; Tanimoto, M. Non-Abelian Discrete Symmetries in Particle Physics. *Prog. Theor. Phys. Suppl.* **2010**, *183*, 1–163. [[CrossRef](#)]
15. King, S.F.; Luhn, C. Neutrino Mass and Mixing with Discrete Symmetry. *Rep. Prog. Phys.* **2013**, *76*, 056201. [[CrossRef](#)]
16. Chauhan, G.; Dev, P.B.; Dubovyk, I.; Dziewit, B.; Fliieger, W.; Grzanka, K.; Gluza, J.; Karmakar, B.; Zięba, S. Phenomenology of Lepton Masses and Mixing with Discrete Flavor Symmetries. *Prog. Part. Nucl. Phys.* **2024**, *138*, 104126. [[CrossRef](#)]
17. Feruglio, F.; Romanino, A. Lepton Flavour Symmetries. *Rev. Mod. Phys.* **2021**, *93*, 015007. [[CrossRef](#)]
18. Feruglio, F. Are neutrino masses modular forms? *arXiv* **2017**, arXiv:1706.08749. [[CrossRef](#)]
19. Nilles, H.P.; Ramos-Sanchez, S.; Vaudrevange, P.K.S. Eclectic Flavor Groups. *J. High Energy Phys.* **2020**, *2020*, 45. [[CrossRef](#)]
20. Nilles, H.P.; Ramos-Sanchez, S. Flavor’s Delight. *Entropy* **2024**, *26*, 355. [[CrossRef](#)]
21. Ma, E. Radiative Origin of All Quark and Lepton Masses through Dark Matter with Flavor Symmetry. *Phys. Rev. Lett.* **2014**, *112*, 091801. [[CrossRef](#)]
22. Cárcamo Hernández, A.E. A novel and economical explanation for SM fermion masses and mixings. *Eur. Phys. J. C* **2016**, *76*, 503. [[CrossRef](#)]
23. Cárcamo Hernández, A.E.; Kovalenko, S.; Schmidt, I. Radiatively generated hierarchy of lepton and quark masses. *J. High Energy Phys.* **2017**, *2017*, 125. [[CrossRef](#)]
24. Arbeláez, C.; Cárcamo Hernández, A.E.; Cepedello, R.; Kovalenko, S.; Schmidt, I. Sequentially loop suppressed fermion masses from a single discrete symmetry. *J. High Energy Phys.* **2020**, *2020*, 43. [[CrossRef](#)]
25. Baker, M.J.; Cox, P.; Volkas, R.R. Has the Origin of the Third-Family Fermion Masses been Determined? *J. High Energy Phys.* **2021**, *2021*, 151. [[CrossRef](#)]
26. Arkani-Hamed, N.; Schmaltz, M. Hierarchies without symmetries from extra dimensions. *Phys. Rev. D* **2000**, *61*, 033005. [[CrossRef](#)]
27. Davoudiasl, H.; Gopalakrishna, S.; Ponton, E.; Santiago, J. Warped 5-Dimensional Models: Phenomenological Status and Experimental Prospects. *New J. Phys.* **2010**, *12*, 075011. [[CrossRef](#)]
28. Xing, Z.-Z. Flavor structures of charged fermions and massive neutrinos. *Phys. Rep.* **2020**, *854*, 1–147. [[CrossRef](#)]
29. Washburn, J. Lean 4 Certificates: Dimensional Selection $D = 3$ and Endogenous Cube Sum $W = 17$. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Physics/WEndoForcing.lean> (accessed on 15 May 2026).
30. Washburn, J. Lean 4 Certificate: Minimal Eight-Tick Period, Phase Offset, and Gap Weight (T7). Verifies: $T_{min} = 8$, Octave Offset -8 , Phase $-1/4$, and the Closed-Form Gap Weight w_8 . 2026. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Foundation/DimensionForcing.lean> (accessed on 15 May 2026).
31. Washburn, J. Recognition Science: Lean 4 Machine Verification Repository. Available online: <https://github.com/jonwashburn/shape-of-logic/tree/main/IndisputableMonolith> (accessed on 15 May 2026).
32. Washburn, J. Lean 4 Certificate: Sector Mass Scale Assignments and SDGO (Anchor). Verifies: B_{pow} Values for All Four Sectors, r_0 Offsets, Z-Map Monotonicity, and Generation Rung Assignments. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Masses/AnchorDerivation.lean> (accessed on 15 May 2026).
33. Washburn, J. Lean 4 Certificate: Baseline Derivation (SA 0, T2, r_e , r_ν , E_{coh} , Color Offset). Verifies: Non-Triviality A1, $r_e = 2$, $r_{\nu_3} = -54$, $E_{coh} = \varphi^{-5}$, Color Offset $c = 4$, and Zero Adjustable Parameters. 2026. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Masses/BaselineDerivation.lean> (accessed on 15 May 2026).
34. Fedorov, E. Proceedings of the Imperial St. Petersburg Mineral. Soc. **1891**, *28*, 345–390. (In Russian)

35. Washburn, J.; Allahyarov, E. A Complete Derivation of the Fermion Spectrum from the Recognition Composition Law. *arXiv* **2026**. [[CrossRef](#)]
36. Washburn, J. *Recognition Science Particle Mass Framework: Progress Report and Honest Status Assessment*; Recognition Physics Institute Technical Report; Recognition Physics Institute: Austin, TX, USA, 2026.
37. Washburn, J. Lean 4 Certificate: Fine-Structure Constant and Curvature Geometry. Verifies: $4\pi E_{pass}$ Seed, Product-Form Uniqueness, Curvature Tuple (5,102,103), Exponential Resummation, and α^{-1} Bracketing. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Verification/EMAlphaCert.lean> (accessed on 15 May 2026).
38. Washburn, J. Lean 4 Certificate: Full 22-Link Forcing Chain. Verifies the Complete Derivation Chain from $J(1) = 0$ to $E_{coh} = \varphi^{-5}$ with Zero Unverified Assumptions. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Foundation/UnifiedForcingChain.lean> (accessed on 15 May 2026).
39. Washburn, J. Lean 4 Certificate: Cost Functional Uniqueness (T1, T3, T5). Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Cost/FunctionalEquation.lean> (accessed on 15 May 2026).
40. Washburn, J. Lean 4 Certificate: φ as Unique Hierarchy Base (T6). Verifies: $\varphi^2 = \varphi + 1$, Irrationality of φ , and Necessity of φ from the Cost Closure Condition. 2026. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Foundation/PhiForcingDerived.lean> (accessed on 15 May 2026).
41. Washburn, J. Lean 4 Certificate: Gap Function Derivation. Verifies: The Unique Affine-Log Gap Function, Parameter Values $a = 1/\ln\varphi$, $b = \varphi$, $c = 0$, and the Forcing Condition $g(-1) = -2$. Available online: <https://github.com/jonwashburn/shape-of-logic/blob/main/IndisputableMonolith/Masses/GapFunctionForcing.lean> (accessed on 15 May 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.