

Article

A Structural Origin of the Charged-Lepton Hierarchy

Bin Li 

Research Department, Silicon Minds Inc., Clarksville, MD 21029, USA; binli.siliconminds@gmail.com

Abstract

The charged-lepton masses are free Yukawa-sector parameters in the Standard Model, whereas their measured pole-mass ratios display a highly structured hierarchy and satisfy the Koide relation to notable accuracy. This paper develops a conditional mathematical-physics proposal in which these dimensionless regularities arise from a charge-neutral parent carrier-defect architecture before effective Higgs–Yukawa read-out. The assumptions of the construction are stated explicitly as structural postulates and are separated from their derived consequences. The central rule assigns equal primitive weight to admissible internal sectors that are indistinguishable at the level where they first become exposed; protected sectors are removed before counting, and later refinements are conditional on previously selected sectors. Under this rule, the Koide relation follows as an equal-power theorem between the democratic parent component and the orthogonal branch-splitting component of the charged-lepton root-amplitude state. A minimal endpoint construction then yields a rapidly stabilizing charged tower for the electron–muon ratio. Because deeper charged terms are too small to remove the remaining residual, the framework assigns that residual to the continuation-dual neutral branch. The resulting neutral overlap gives a leading solar-angle target of 33.21 degrees and closes the electron–muon ratio at the present experimental precision; the Koide relation then fixes the corresponding tau ratios. The construction does not replace the Standard Model but is proposed as a selection rule for the boundary values of effective charged-lepton Yukawa parameters, with pole masses used because the claimed invariant is attached to completed asymptotic particle read-out. Running parameters, the absolute mass scale, and the full Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix remain outside the present derivation. The proposal has explicit failure conditions: improved measurements can exclude the predicted tau ratios or solar-angle target, and the claimed conditional uniqueness fails if a different counting scheme satisfies the same postulates while producing different endpoint weights.

Keywords: charged-lepton hierarchy; Koide relation; lepton mass ratios; discrete symmetry; equal-weight counting; neutral-parent structure; endpoint expansion; PMNS mixing



Academic Editor: Jorge Segovia

Received: 16 June 2026

Revised: 14 July 2026

Accepted: 20 July 2026

Published: 21 July 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 Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The charged-lepton masses are among the simplest precisely measured quantities in particle physics and among the least structurally explained. They obey the strong hierarchy

$$m_e \ll m_\mu \ll m_\tau,$$

but in the Standard Model their values are fixed by three independent Yukawa parameters rather than derived from a deeper organizing principle. The present paper asks whether the dimensionless charged-lepton hierarchy can instead arise as a carrier-readable invariant of a common charge-neutral particle architecture.

The work develops a quantitative specialization of the carrier-resolution framework introduced in Ref. [1]. In that framework, particle species are interpreted as different carrier-readable manifestations of a common loop-detectable codimension-two identity defect. A neutral parent resolves asymmetrically into a lepton-facing Z_2 -Lorentz branch and a positive branch supporting nested Z_3 closure, while higher Z_n structures act as conditional refinements rather than as additional ordinary particle generations. The present paper does not repeat the full ontology and topological motivation of that work. It asks a narrower quantitative question: whether the same neutral-parent architecture and closure-refinement logic can determine the charged-lepton mass ratios.

The relationship between the two papers is structural rather than merely terminological. Ref. [1] introduced the carrier, codimension-two identity support, neutral-parent resolution, protected internal components, carrier-facing exposure, nested Z_n refinement, saturation of a same-branch correction sequence, and neutral completion of a residual closure seam. The present manuscript applies these principles to a different observable. Instead of magnetic response in a completed hadronic branch, it studies the small charged endpoint of the lepton-facing branch. The detailed counting coefficients derived below are specific to the charged-lepton problem, but the organizing logic is inherited from the broader carrier-resolution architecture rather than introduced solely to reproduce the lepton spectrum.

The charged-lepton masses also exhibit the remarkable Koide relation [2–4],

$$\frac{m_e + m_\mu + m_\tau}{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2} = \frac{2}{3}.$$

Because this relation is naturally expressed in square-root mass variables, it is commonly represented as an angle condition in a three-dimensional root-amplitude space. In the present framework, that geometry is not taken as fundamental. It is interpreted as a continuous representation of a more primitive structural equality: equal normalized power in the democratic parent component and in the complementary branch-splitting component. Thus the Koide cone represents

$$P_{\text{dem}} = P_{\text{split}},$$

rather than introducing an independently selected continuous mass angle.

The relevant source structure is a charge-neutral parent, denoted P_0 . The motivation is not that every observed process is literally the decay of a single hidden particle. Rather, positive, negative, and neutral particle channels repeatedly occur in correlated conservation and closure patterns. Neutron beta decay,

$$n \rightarrow p + e^- + \bar{\nu}_e,$$

is a familiar example: the charged positive, charged negative, and neutral channels form a coordinated closure pattern rather than three unrelated primitive outputs. This does not prove the neutral-parent interpretation, but it motivates treating particle identities as different read-outs of a common neutral organization.

The first persistent resolution used here is

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+).$$

The arrow denotes structural resolution or carrier read-out, not temporal evolution, particle decay, subgroup inclusion, or a Lagrangian transition amplitude. The Z_2^- side is the lepton-

facing branch. After carrier embedding it supports a charged transverse read-out and a neutral continuation read-out,

$$Z_2^- \longrightarrow \ell_\perp^- \oplus \nu_\parallel.$$

The $Z_2^+ \rightarrow Z_3^+$ side is the positive closure branch. After read-out, it is compatible with a hadron-facing threefold closure structure. The asymmetry is essential: a perfectly mirrored polarized resolution would tend to recombine into the unresolved neutral closure, whereas persistent distinguishable matter requires complementary but non-identical read-out branches.

The charged-lepton hierarchy is therefore not treated as a property of the negative branch in isolation. The electron is interpreted as a small charged endpoint exposure of the lepton-facing branch, but the count that selects this endpoint is measured against the full neutral-parent closure architecture. Its first non-protected exposure occurs when the completed positive-side closure is embedded into its first higher environment,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The successive structures

$$((Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4) \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow Z_7 \hookrightarrow \dots$$

are nested conditional refinements of an already selected endpoint. They are not new gauge groups, new colors, or fourth, fifth, and higher ordinary generations.

The counting rule used in the paper is level-wise structural indifference. Before metric, energetic, or carrier-dynamical distinctions have emerged, admissible sectors that are indistinguishable by the available branch, monodromy, protection, and closure data have no structural basis for unequal primitive weights. Protected, identified, or cancelling sectors are first removed or quotiented, and later substructure is counted conditionally inside the sector in which it appears. Such later refinement does not retroactively alter the primitive weight assigned at an earlier exposure level.

This equal-weight rule should not be interpreted as a fitted probability distribution or as a claim that all post-read-out physical states have equal dynamical weight. It is a premetric admissibility principle. A more complete carrier theory should explain how this primitive structural indifference is preserved or transformed in the effective read-out. The present manuscript investigates the consequences that follow if the charged-lepton branch inherits this natural premetric weighting.

Throughout the paper, m_e , m_μ , and m_τ denote the physical charged-lepton rest masses [5–7]. The internal construction first fixes relative root amplitudes. Their squares are then interpreted as carrier-localized rest-energy coefficients,

$$m_i = M_\ell r_i^2,$$

where the common charged-lepton scale M_ℓ cancels from all ratios. Appendix A states the minimal assumptions under which these root-amplitude ratios may be compared with physical inertial-mass ratios.

The internal space invoked here is not an ordinary microscopic region of spacetime. It is a premetric admissibility structure whose primitive data are branches, sectors, exclusions, identifications, monodromies, protected components, and closure relations. Vectors, projections, and angles are used only as representations of those structural data, just as discrete graphs may be represented by matrices and finite groups by linear representations. The Koide angle is therefore not introduced as a fundamental physical angle. It is the Euclidean representation of an equal-sector-power invariant.

The guiding hierarchy is

$$\begin{aligned} P_0 &\longrightarrow \text{premetric admissibility and closure} \\ &\longrightarrow \text{carrier read-out} \\ &\longrightarrow \text{dimensionless mass invariants.} \end{aligned}$$

The carrier supplies Lorentz-readable propagation, localization, and the common physical scale. The premetric architecture supplies the dimensionless branch relations that may survive the read-out.

The first quantitative consequence is the Koide relation itself. The charged-lepton root vector is decomposed into a democratic parent component and an orthogonal branch-splitting component. Equal primitive power in those two unresolved sectors gives

$$P_{\text{dem}} = P_{\text{split}},$$

which is algebraically equivalent to the Koide relation. In this interpretation, Koide is not imposed as an empirical constraint after the mass ratios have been chosen. It is the first carrier-facing consequence of the equal-sector-power rule.

The second and principal quantitative result is the charged endpoint expansion. The small electron mass is interpreted as one surviving endpoint unit in the first exposed non-protected chamber environment, with successive corrections generated by conditional higher-interface refinements. The resulting tower is

$$\left(\frac{m_e}{m_\mu}\right)_{\text{ch}} = \frac{1}{3^2(4! - 1)} \left[1 + \sum_{k=5}^{\infty} \prod_{n=5}^k \frac{2/n}{3(n! - 1)} \right].$$

The factor $4! - 1$ is the first ordered exposure count after removal of the protected closure chamber, while the factor 3^2 records the leading charged-family and compensational-family projection. The higher factors

$$q_n = \frac{2/n}{3(n! - 1)}$$

are conditional refinements inside an already selected endpoint. Consequently, they enter as nested products rather than as freely adjustable additive corrections.

This structure parallels the closure-interface logic already used in Ref. [1]: a rule-generated same-branch sequence is extended until it saturates, and a residual that cannot be removed by deeper terms is treated as evidence of a distinct neutral-parent completion seam. The detailed coefficients are different because the observables and exposure spaces are different, but the diagnostic principle is the same.

Numerically, the charged tower stabilizes rapidly and gives

$$\left(\frac{m_\mu}{m_e}\right)_{\text{ch}} = 206.768291043.$$

This lies within approximately 1.8σ of the observed inverse electron–muon mass ratio. More importantly, the later charged terms are many orders of magnitude too small to remove the remaining central difference. The residual is therefore not naturally interpreted as an omitted deeper charged-interface correction.

Neutral-parent closure supplies a continuation-dual compensational read-out associated with the neutrino sector. The compensational overlap is not imported from the measured Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix. It is obtained from a

democratic three-direction neutral exposure with one protected chamber removed from the Z_5 two-slot interface:

$$\sin^2 \theta_{12}^{\text{str}} = \frac{1}{3} - \frac{1}{3\binom{5}{2}} = \frac{3}{10},$$

and hence

$$\frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30}.$$

The neutral completion of the charged endpoint is therefore $q_5 q_6 (7/30)$, rather than a term fitted from solar-neutrino data. The same calculation gives the independent leading target

$$\sin^2 \theta_{12}^{\text{str}} = 0.30000, \quad \theta_{12}^{\text{str}} = 33.21^\circ.$$

This value can be compared with the empirical solar-mixing region described by the PMNS framework [5,8–10], but it is not used as an input to the charged-lepton calculation.

The completed endpoint prediction is

$$\left(\frac{m_\mu}{m_e} \right)_{\text{pred}} = 206.768282689,$$

which differs from the observed central value by approximately -1.9×10^{-8} , or about -0.004σ at the quoted present experimental precision. Once this endpoint is fixed, the Koide sector-power condition determines the large- τ branch:

$$\left(\frac{m_\tau}{m_e} \right)_{\text{pred}} = 3477.441636, \quad \left(\frac{m_\tau}{m_\mu} \right)_{\text{pred}} = 16.8180612.$$

Thus the construction is sequential: charged endpoint counting plus neutral compensation determines m_μ/m_e , and the Koide sector-power relation then fixes the tau branch. No continuous parameter is adjusted to the measured electron–muon ratio, solar angle, or tau mass.

The proposal is deliberately conditional. The structural assumptions are stated explicitly in Section 3; the subsequent formulas are derived within that architecture. The paper does not claim to provide a local microscopic Lagrangian for the premetric domain, since ordinary spacetime locality and continuous field evolution are not assumed to be fundamental there. Nor does it derive the absolute charged-lepton scale, the full Standard Model Yukawa sector, or the complete PMNS matrix. Its narrower claim is that several otherwise separate charged-lepton and neutral-overlap regularities may be carrier-facing manifestations of one neutral-parent admissibility structure.

The paper is organized as follows. Section 2 introduces the carrier-defect architecture and its terminology. Section 3 separates the structural assumptions from the derived claims. Section 4 formulates level-wise equal-weight counting and conditional refinement. Section 5 derives the Koide relation as an equal-sector-power theorem. Section 6 develops the charged endpoint expansion from the electron-zero boundary through the higher-interface tower. Section 7 examines the minimality and conditional uniqueness of the factors $4! - 1$, 3^2 , q_n , and $7/30$. Section 8 evaluates the stabilized charged tower. Section 9 derives the neutral-overlap completion and the solar-angle target. Section 10 explains the relation to the Standard Model Yukawa description, pole masses, running parameters, and conventional flavor models. Section 11 discusses the broader particle architecture and the limits of the interpretation. Section 12 summarizes the independent predictions and explicit failure conditions. The final sections state the scope of the framework and present the conclusions.

Appendix A discusses the conversion from internal root amplitudes to physical mass ratios. Appendix B gives additional structural motivation for the Z_2 -Lorentz branch, the positive $Z_2 \rightarrow Z_3$ closure branch, and the higher nested embedding hierarchy. Appendix C presents a minimal Lorentz-read-out model. Appendix D develops the detailed neutral-compensation construction underlying the main-text derivation of the factor $7/30$.

The overall logical structure of the derivation is summarized in Figure 1.




Figure 1. Compact roadmap of the argument. The charged-lepton hierarchy is developed as a specialization of the carrier-resolution framework of Ref. [1]. Once the structural assumptions are stated, no continuous coefficient is adjusted to the measured charged-lepton ratios or the leading solar-angle target. The charged same-branch tower is continued until it saturates, after which the remaining closure seam is assigned to the neutral continuation sector.

2. Carrier-Defect Architecture and Basic Terminology

This section defines the structural language used in the paper. The goal is not to replace the Standard Model as an effective field theory. The Standard Model successfully describes already read-out particles, fields, charges, and interactions, but it does not explain why the charged-lepton mass ratios take their observed values. The present paper asks whether those ratios may reflect a pre-read-out internal structure whose effects become physical only after carrier embedding.

This distinction is motivated by standard physics itself. In quantum field theory the vacuum is not an empty mathematical background; it carries field structure, fluctuations, phases, symmetry breaking, and renormalized parameters [11–13]. The physical masses and couplings measured in experiment are not bare labels. They are read-out quantities after embedding into the vacuum structure of the theory. The present framework asks whether the embedding of particle-like defects into this carrier can leave calculable traces in dimensionless mass ratios.

2.1. Carrier

The term *carrier* denotes the physically effective vacuum structure through which a localized internal configuration becomes a particle-like read-out. It should not be understood as an ordinary material ether, nor as a substance placed inside spacetime. In the present usage, the carrier is the spacetime vacuum viewed in its active physical role: the support of local fields, phases, defects, renormalized parameters, and observable particle manifestations.

The logical order is

internal configuration $\longrightarrow$ carrier embedding/read-out $\longrightarrow$ effective particle.

The internal configuration supplies structural data such as branches, sectors, exclusions, identifications, monodromy, and closure relations. The carrier supplies the physical read-out environment: Lorentz-readable propagation, effective field representation, localization, rest-energy scale, and long-field stress-energy behavior.

This terminology is useful because the way a particle-like object is embedded into the vacuum carrier is an under-explored structural question. The Standard Model describes the resulting carrier-facing fields and interactions, but it does not by itself determine whether different particles may originate as different read-outs of a common pre-read-out structure.

2.2. Defect, Read-Out, and Pregeometric Internal Space

A *defect* is a localized stable carrier manifestation of a nontrivial internal structure. A particle is therefore not treated here as a structureless point inserted by hand, but as the carrier-facing read-out of a localized defect whose internal identity can be continued, closed, and recognized across admissible stages.

A *read-out* is the process by which an admissible internal configuration becomes a physical manifestation describable by the usual effective language of particle physics: fields, masses, charges, spinorial behavior, mixing matrices, and stress-energy. The present paper does not claim that charged-lepton masses exist as ordinary physical masses before carrier embedding. Instead, the internal counting structure fixes relative root amplitudes. After read-out, the squared amplitudes become relative rest-energy coefficients.

The associated *internal defect space* is called *pregeometric* or *premetric* in a limited sense. Before carrier embedding, the defect-internal structure has not yet been incorporated into the spacetime-vacuum read-out that gives ordinary metric, energetic, and field-theoretic descriptions. It should therefore not be assumed to carry an intrinsic distance, volume, angle, or physical energy scale. Its primitive information is combinatorial and topological: branch labels, sector relations, protected channels, excluded channels, identifications, monodromy, and closure constraints.

This is why the Z_n notation used below should not be read as a replacement for Standard Model gauge groups at the observable level. Before Lorentz-covariant read-out, the defect pocket is not yet a differentiable internal manifold carrying the Standard Model groups as its primary description. The finite Z_n structures describe pregeometric internal continuation and closure. Continuous gauge and spinorial descriptions become

physically appropriate only after carrier read-out [14–17]. Appendix B gives the structural argument for why the primitive Z_2 branch opposition is read out as spinorial behavior for free fermions, while the positive $Z_2 \rightarrow Z_3$ closure and its higher embeddings belong to the positive/hadron-facing side. Appendix C gives a minimal toy Lorentz-read-out model illustrating this distinction.

2.3. Neutral Parent and Asymmetric Branch Architecture

The central structural object is the *charge-neutral parent*, denoted P_0 . It is not introduced as a new observed particle, nor as a literal object that simultaneously decays into all possible channels. It is an unresolved admissibility structure whose different carrier read-outs may appear as distinct matter-particle channels.

The motivation for introducing such a neutral parent is severalfold. First, observed particle processes do not treat charge as an isolated label; charge appears in closed combinations. Standard Model interactions preserve total electric charge, and physical reaction channels are organized so that charge, spin, lepton-family labels, baryon number, and other conserved or approximately conserved quantum numbers are consistently balanced [5,11,13]. This suggests that particle identities are not arbitrary independent labels, but mutually constrained read-outs of a larger closure structure.

Second, it would be structurally unnatural to assume that all matter particles exist as unrelated primitive species while the total charge content of allowed physical processes nevertheless closes exactly to neutral combinations. The empirical success of charge conservation and the repeated appearance of compensating particle channels suggest a more economical possibility: charged and neutral particles may be different resolved aspects of an underlying charge-neutral organization.

Third, familiar weak processes already point in this direction. Neutron beta decay,

$$n \rightarrow p + e^- + \bar{\nu}_e,$$

does not produce an isolated charged particle. It produces a correlated charged-positive, charged-negative, and neutral compensating structure. This does not prove the existence of a neutral parent in the present sense, but it motivates using a charge-neutral parent as a theory of origin: a structural source from which positive, negative, and neutral read-outs emerge together under conservation and closure constraints.

The first resolved architecture of the neutral parent is therefore written schematically as

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+).$$

This resolution is intentionally asymmetric. If the two polarized sides were simple mirror copies, they would tend to recombine or annihilate at the unresolved level and would not generate stable distinguishable matter read-outs. A persistent matter architecture therefore requires complementary but non-identical branches: a lepton-facing Z_2^- branch and a positive closure branch $Z_2^+ \rightarrow Z_3^+$.

The Z_2^- side is the lepton-facing branch-resolution side. It supports a charged transverse read-out and a neutral continuation read-out,

$$Z_2^- \longrightarrow \ell_{\perp}^- \oplus \nu_{\parallel}.$$

Here $\ell_{\perp}^-$ denotes the charged-lepton read-out direction, while $\nu_{\parallel}$ denotes the neutral compensating continuation direction. These symbols are structural labels, not claims that the electron, muon, tau, and neutrinos are already present as independent particles before read-out.

The $Z_2^+ \rightarrow Z_3^+$ side is the positive closure branch. It is the counterpart of the charged lepton-facing branch and supplies the first completed positive-side closure. In later carrier-facing language this side is compatible with hadron-facing threefold closure, but it should not be identified with the full $SU(3)$ color gauge group at the pregeometric level. It records a threefold closure condition inside the defect architecture. Appendix B explains why this positive closure begins with Z_3 , whereas the free lepton-facing branch remains Z_2 -primitive before spinorial carrier read-out.

Thus the charged-lepton hierarchy is not treated as a property of the negative Z_2^- branch in isolation. It is a projection of the full neutral-parent resolution

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+).$$

This point is essential for the endpoint calculation: the charged endpoint is a leakage of the lepton-facing branch, but the count that selects it is controlled by the full neutral-parent closure architecture.

2.4. Positive-Side Nested Embedding Hierarchy

Once the positive-side closure

$$Z_2^+ \rightarrow Z_3^+$$

is completed, it is exposed through nested carrier-readability environments,

$$((Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4) \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow Z_7 \hookrightarrow \dots$$

These higher Z_n structures are not additional particle generations, not new colors, and not new independent gauge sectors. They are nested embedding environments of the already completed positive-side closure.

This distinction explains why the first endpoint leakage count begins with

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The Z_2 level separates complementary read-out directions of the neutral parent, but it is not yet a completed positive closure. The first completed positive-side support is the embedded $Z_2^+ \rightarrow Z_3^+$ closure. The first non-protected exposure chamber therefore appears only when this completed positive closure is embedded into a Z_4 environment. The detailed motivation for this assignment is given in Appendix B.

In the charged-lepton calculation, the electron endpoint is not a purely leptonic count. It is a negative-side boundary leakage measured against the full neutral-parent closure architecture. The leading factor $4! - 1$ arises from the first exposed ordered chamber space after the protected closure channel has been removed. The later $Z_4 \hookrightarrow Z_5$, $Z_5 \hookrightarrow Z_6$, and deeper embeddings refine the already selected endpoint conditionally rather than opening independent new endpoints. This is why the charged endpoint expansion has a nested product structure rather than a sum of unrelated corrections.

2.5. Root Amplitude and Physical Mass Read-Out

The charged-lepton hierarchy is represented using the root vector

$$\mathbf{r}_\ell = (\sqrt{m_e}, \sqrt{m_\mu}, \sqrt{m_\tau}).$$

In the present interpretation, the square-root variables are internal root amplitudes. Their squared carrier read-out gives physical mass weights:

$$m_i = M_\ell r_i^2,$$

where M_ℓ is a common charged-lepton carrier scale. Since this scale cancels in ratios, the paper focuses on dimensionless quantities such as m_e/m_μ and m_τ/m_e .

This representation is not chosen merely for convenience. The Koide relation is naturally expressed in square-root mass variables, suggesting that the relevant invariant is not a direct linear mass sum, but a sector-power relation among root amplitudes [2–4]. The detailed physical-mass interpretation is given in Appendix A.

2.6. Neutral Compensation

A *neutral compensational read-out* is a neutral continuation channel required to complete the closure of a charged or asymmetric branch resolution. In the charged-lepton setting, this role is associated with the neutrino sector.

This concept is introduced because the charged endpoint expansion nearly, but not completely, closes the observed electron–muon ratio. The remaining discrepancy is too large to be supplied by deeper charged-layer terms, yet it has the correct size to be interpreted as a neutral compensation effect. Thus the neutrino sector is not treated as an unrelated addition to the charged-lepton hierarchy. It appears as the neutral continuation channel needed when the charged endpoint structure does not close by itself.

The present paper does not claim to derive the full PMNS matrix. It does, however, derive the particular solar-sector compensational overlap needed to close the charged endpoint residual. The observed solar mixing angle is therefore used only as an external consistency comparison, not as an input to the charged-lepton hierarchy [5,8–10]. Appendix B explains why this neutral continuation is naturally paired with the charged Z_2^- lepton-facing branch rather than with the positive Z_3 -closure side.

2.7. Scope of the Terminology

The terminology introduced here organizes the structural calculation; it does not replace the Standard Model as the effective description of carrier-facing fields and interactions. The quantitative claims concern dimensionless charged-lepton mass ratios and the associated neutral overlap. They do not constitute a complete formal theory of premetric admissibility, carrier read-out, the absolute mass scale, or neutrino mixing.

The calculation uses the architecture

$$P_0 \rightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+), \quad ((Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4) \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow \dots,$$

together with

$$Z_2^- \rightarrow \ell_\perp^- \oplus \nu_\parallel.$$

Within these declared data, the endpoint expansion is the charged-lepton projection of the broader neutral-parent closure structure rather than an isolated numerical rule. Figure 2 summarizes the architecture used below.




Figure 2. Compact schematic of the neutral-parent carrier-defect architecture. The unresolved neutral parent P_0 resolves asymmetrically into a lepton-facing Z_2^- branch and an embedded positive closure branch $Z_2^+ \rightarrow Z_3^+$. The lepton-facing branch supports charged and neutral read-outs, while the positive closure branch is exposed through nested embedding environments $Z_4, Z_5, Z_6, \dots$. The electron endpoint is a charged leakage, but its count is controlled by the full neutral-parent architecture.

3. Structural Postulates and the Status of the Claims

Because the present proposal departs from conventional bottom-up approaches to flavor physics, it is important to distinguish clearly between structural assumptions and derived consequences. The framework developed here is therefore formulated conditionally. It does not assume that the charged-lepton hierarchy has already been derived from a more general mathematical theory of premetric admissibility. Instead, it introduces a small set of explicit structural postulates that specify how admissible sectors are organized before carrier read-out and then examines what quantitative consequences follow.

This formulation should not be interpreted as an incomplete effective field theory awaiting a more microscopic Lagrangian description. The premetric domain considered here precedes ordinary spacetime geometry, local dynamics, and continuous field evolution; consequently, the appropriate foundational completion is expected to be a mathematical theory of admissibility, protected sectors, closure, and carrier read-out, rather than another local dynamical model defined on an already existing spacetime manifold.

Within this framework, the role of the postulates is analogous to that of axioms in mathematics: they define the admissible structural setting. The subsequent endpoint expansion, neutral-compensation algebra, and the charged-lepton mass-ratio predictions are then derived consequences of those postulates rather than independently adjusted numerical rules. This distinction is essential for assessing both the predictive content and the limitations of the proposal.

Axiom 1 (Carrier-readable defect). *A persistent particle channel is the Lorentz-readable manifestation of a localized internal defect. The internal structure fixes dimensionless relative amplitudes, whereas carrier embedding supplies the common rest-energy scale and the ordinary field-theoretic representation.*

Axiom 2 (Neutral-parent closure). *The charged-negative, positive-closure, and neutral-continuation channels are complementary resolutions of a common charge-neutral admissibility structure P_0 . No charged branch is counted as an isolated primitive object.*

Axiom 3 (Level-wise indifference). *At a fixed exposure level, admissible sectors that are indistinguishable by the available branch, monodromy, protection, and closure data carry equal primitive weight. Unequal weights require an already available structural distinction and may not be inserted merely to improve numerical agreement.*

Axiom 4 (Protected-sector quotient). *Identity, closure, or otherwise protected chambers are removed or quotiented before exposed endpoint channels are counted. A protected chamber may constrain closure without contributing a local endpoint leakage unit.*

Axiom 5 (Conditional refinement). *Once a parent sector has received its primitive weight, a later embedding refines that sector conditionally. Deeper substructure multiplies the previous weight and may not retroactively reweight the already selected parent exposure.*

Axiom 6 (Quadratic mass read-out). *For charged leptons, the carrier-localized rest-mass coefficient is quadratic in the corresponding internal root amplitude, $m_i = M_\ell r_i^2$, with a common charged-lepton scale M_ℓ . Hence the present construction determines mass ratios but not the absolute scale.*

These postulates have different logical roles. The first, second, and sixth specify the physical dictionary. The third through fifth specify the counting calculus. The branch architecture

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+), \quad (Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4 \hookrightarrow Z_5 \hookrightarrow \dots$$

is the minimal architecture investigated in this paper. The words “minimal” and “unique” are used only relative to these declared data: no claim is made that abstract mathematics forbids every imaginable flavor model. Rather, Section 7 shows that the numerical weights used here are fixed once the postulates, the stated branch architecture, and the requirement of one first charged endpoint are held fixed.

For clarity, Table 1 records the division between inputs, structural assumptions, and calculated consequences.

Table 1. Logical status of the ingredients and results. Experimental values are not used to tune the integer factors in the construction.

Item	Status	Role
Standard Model charged-lepton sector	Background theory	Supplies the effective Higgs–Yukawa language and measured pole masses
Neutral-parent branch architecture	Structural postulate	Specifies the admissible charged, positive-closure, and neutral channels
Equal primitive weight and protected quotient	Structural postulates	Define the premetric counting rule
Quadratic root-amplitude read-out	Structural postulate	Converts amplitude ratios to mass ratios
Koide relation	Derived theorem	Follows from equal power of democratic and splitting sectors
$3^2(4! - 1)$ and q_n	Derived within the stated architecture	Count the first endpoint and its conditional refinements
$\sin^2 \theta_{12}^{\text{str}} = 3/10$	Derived leading neutral overlap	Not taken from oscillation data
m_μ/m_e and tau ratios	Predictions	Compared with experiment only after the count is completed
Absolute scale and full PMNS matrix	Not derived	Explicitly outside the present scope

The benefit of this formulation is that the proposal can fail cleanly. A derivation of the postulates from a deeper theory would strengthen it; a contradiction among the

postulates, a non-unique count satisfying the same constraints, or an excluded prediction would invalidate it in its present form. The next section now states the counting calculus implied by the level-wise postulates.

4. Equal-Weight Counting in Premetric Defect-Internal Spaces

The present paper uses one structural principle: equal primitive weight among indistinguishable admissible sectors at the level where they first become exposed. The principle is not specific to leptons. It applies to any premetric internal state space associated with a carrier-readable defect. Its role is to specify how primitive weights are assigned before metric, dynamical, energetic, or carrier-facing distinctions have appeared.

The internal spaces considered here are pregeometric in a limited but important sense. They are not ordinary spacetime regions equipped with distances, angles, volumes, or metric weights. Before carrier-facing read-out, the available data are combinatorial and topological: admissible branches, identified sectors, protected channels, excluded identity elements, connections, monodromy, and closure relations. In such a setting, there is no intrinsic reason to regard one admissible sector as larger, closer, stronger, or more important than another unless some additional structure distinguishes it.

This is the same kind of distinction that appears throughout topology. Metric size, material composition, and geometric embedding are not the primary data; connectivity, holes, boundaries, deformation class, and obstruction structure are. A premetric defect-internal space is used in the same spirit: it distinguishes admissible sectors only by the branching, identification, and closure data available before metric read-out [18–22].

This motivates the following rule.

Equal-weight defect-internal counting principle. *In a premetric internal defect space, admissible sectors that are indistinguishable by the available branch, connection, monodromy, and closure data are assigned equal primitive weight at the level at which they first become exposed. Protected, excluded, or identified sectors are removed or quotiented before the count is performed. If later substructure appears inside an already exposed sector, it is counted conditionally within that sector and does not retroactively change the primitive weight assigned at the earlier level.*

The level-wise qualification is essential. A defect-internal structure is not treated as a completed tree of final outcomes given all at once. It is exposed in stages. At a given stage, only the currently exposed admissible sectors exist as distinguishable alternatives. Later subbranches are refinements inside already formed sectors; they do not retroactively change the primitive weights assigned at the earlier level.

Formally, if X indistinguishable sectors are exposed at a level, each receives primitive weight

$$w_i = \frac{1}{X}.$$

If sector i later exposes Y_i admissible subsectors, those subsectors are counted conditionally:

$$w_{j|i} = \frac{1}{Y_i}, \quad w_{ij} = w_i w_{j|i} = \frac{1}{X} \frac{1}{Y_i}.$$

Thus the number Y_i of later subsectors does not alter the earlier weight w_i . The later refinement is conditional on the prior formation of sector i .

This staged rule is analogous to familiar constructions in conditional probability, filtrations, branching processes, cell complexes, and simplicial refinements. In conditional probability one writes

$$P(i, j) = P(i)P(j|i),$$

rather than assigning the parent event a weight proportional to the number of later refinements it may contain. Likewise, in a filtration or cellular construction, coarse structure is fixed before finer structure is attached. The later refinement does not rewrite the primitive count at the earlier level.

The principle is deliberately conservative. It does not claim that all physical states have equal probability after metric read-out. It applies only at the premetric internal counting stage, before the carrier supplies scales, stiffnesses, propagation dynamics, effective couplings, and field-theoretic response. Once those structures are present, unequal physical responses may arise. But before they are present, assigning unequal primitive weights to sectors that are identical by the available internal data would introduce an arbitrary bias.

A continuous representation may be used without changing the primitive status of the invariant. This is standard in mathematics. A graph is discrete, but it may be represented by an adjacency matrix, a Laplacian, a spectrum, or an embedding. A simplicial complex is combinatorial, but it may be drawn geometrically. A finite group is algebraic, but it may be studied through linear representations. In each case, continuous or linear language is a representation of a more primitive structural object, not a replacement for it.

The same convention is used here. Root-amplitude vectors and geometric language provide efficient representations of normalized sector counts. They do not make spacetime geometry fundamental at the premetric stage. The counting-native statement is primary; the continuous representation is used only when it faithfully encodes the same invariant. In this sense, the geometric language used below is a calculational and visualization tool, while the underlying principle remains level-wise equal-weight counting in defect-internal space.

5. Koide Relation as an Equal-Weight Theorem

We first apply the equal-weight defect-internal counting principle to the charged-lepton root sector. The neutral parent is not assumed to contain an already actual Koide mass geometry. Before carrier read-out, it contains admissible internal potentialities, not three physical charged-lepton masses. The Koide relation is therefore not postulated as a pre-existing geometric mass pattern. It is interpreted as the carrier-facing read-out of an equal sector-power count in an unresolved internal root space [2–4].

Define the charged-lepton root vector

$$\mathbf{r}_\ell = (\sqrt{m_e}, \sqrt{m_\mu}, \sqrt{m_\tau})$$

and its total root power

$$P_{\text{tot}} = |\mathbf{r}_\ell|^2 = m_e + m_\mu + m_\tau.$$

Here m_e, m_μ, m_τ denote the physical charged-lepton rest masses. The square-root variables are interpreted as internal root amplitudes whose squared carrier read-out gives localized rest-energy weights:

$$m_i = M_\ell r_i^2,$$

with a common charged-lepton carrier scale M_ℓ . This common scale cancels from all mass ratios. The stress-energy interpretation of this read-out is discussed in Appendix A.

The vector $\mathbf{r}_\ell$ belongs to an internal root-amplitude representation of the charged-lepton read-out sector. Its components label charged-lepton branch amplitudes, not directions in ordinary spacetime. The democratic component is the equal-weight unresolved parent direction. In root-vector representation it is

$$\mathbf{r}_{\text{dem}} = \frac{\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}}{3} (1, 1, 1).$$

Its sector power is

$$P_{\text{dem}} = |\mathbf{r}_{\text{dem}}|^2 = \frac{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2}{3}.$$

The branch-splitting sector is the complementary non-democratic sector, with power

$$P_{\text{split}} = P_{\text{tot}} - P_{\text{dem}}.$$

Before carrier-facing actualization, the democratic and branch-splitting sectors are complementary potential sectors of the same normalized internal root state. Neither sector has yet acquired an independent physical scale, metric weight, or dynamical preference. Therefore, by the equal-weight principle, the two unresolved sectors carry equal primitive root power:

$$P_{\text{dem}} = P_{\text{split}}.$$

This is the counting-native form of the Koide relation. The familiar 45° representation is only the Euclidean visualization of this sector-power equality after the discrete counting invariant has been represented in root-vector coordinates.

Theorem 1 (Koide relation from equal sector-power counting). *If the normalized charged-lepton root state is selected by equal primitive internal weight between the democratic parent sector and the branch-splitting sector,*

$$P_{\text{dem}} = P_{\text{split}},$$

then the charged-lepton masses satisfy

$$\frac{m_e + m_\mu + m_\tau}{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2} = \frac{2}{3}.$$

Proof. Since

$$P_{\text{tot}} = P_{\text{dem}} + P_{\text{split}},$$

the equal-sector condition gives

$$P_{\text{tot}} = 2P_{\text{dem}}.$$

Substituting the expressions above,

$$m_e + m_\mu + m_\tau = 2 \frac{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2}{3}.$$

Rearranging yields

$$\frac{m_e + m_\mu + m_\tau}{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2} = \frac{2}{3}.$$

□

The same theorem explains why the Koide relation can be written as a geometric cone without making ordinary spacetime geometry fundamental. Introduce the normalized root vector

$$\hat{\mathbf{r}}_\ell = \frac{\mathbf{r}_\ell}{|\mathbf{r}_\ell|}$$

and the democratic unit direction

$$\hat{\mathbf{d}} = \frac{1}{\sqrt{3}}(1, 1, 1).$$

Then

$$P_{\text{dem}} = P_{\text{tot}}(\hat{\mathbf{r}}_\ell \cdot \hat{\mathbf{d}})^2.$$

The equality $P_{\text{dem}} = P_{\text{split}}$ is therefore equivalent to

$$(\hat{\mathbf{r}}_\ell \cdot \hat{\mathbf{d}})^2 = \frac{1}{2}, \quad \hat{\mathbf{r}}_\ell \cdot \hat{\mathbf{d}} = \frac{1}{\sqrt{2}}.$$

Thus the Koide cone is not an independent postulate. It is a faithful continuous representation of the discrete two-sector equality

$$P_{\text{dem}} : P_{\text{split}} = 1 : 1.$$

The axes of this representation label admissible internal branch components, not directions in already-given physical space.

It is important not to reverse the logic. The parent is not assumed to contain actual charged-lepton masses or an actual Koide geometry before read-out. Rather, it contains admissible internal potentialities. Since the democratic and branch-splitting sectors are potential sectors of the same unresolved normalized root state, they have equal primitive status before carrier-facing actualization. Their equal-weight resolution appears, after read-out, as the Koide relation and, in a convenient Euclidean representation, as the Koide cone.

The remaining task is therefore not to postulate the charged-lepton mass ratios independently. Once the equal-weight principle fixes the sector-power relation, the hierarchy is reduced to determining the small endpoint displacement inside the same internal counting architecture. The following section applies the same level-wise counting principle to the first exposed chamber structure that opens the electron endpoint.

The equal-weight interpretation of the Koide relation can be visualized as the root-space decomposition shown in Figure 3.




Figure 3. Root-space representation of the Koide sector-power condition. The charged-lepton root vector $\mathbf{r}_\ell$ decomposes into a democratic parent component $\mathbf{r}_{\text{dem}}$ and an orthogonal branch-splitting component $\mathbf{r}_{\text{split}}$. The Koide relation is equivalent to the equal-power condition $P_{\text{dem}} = P_{\text{split}}$, represented geometrically by a 45° root-space decomposition. The geometry is used only as a representation of the underlying sector-power count.

6. From the Electron-Zero Boundary to the Endpoint Expansion

The Koide sector-power relation fixes one relation among the three charged-lepton masses. To determine the hierarchy, one must still identify the small endpoint displacement that makes the electron mass nonzero. This section develops that endpoint step. We begin at the electron-zero boundary, then turn on the small electron endpoint and derive the proposed endpoint expansion for m_e/m_μ .

The endpoint calculation should be read within the neutral-parent architecture introduced in Section 2 and clarified in Appendices B and C. The electron endpoint is a charged

leakage of the lepton-facing Z_2^- branch, but it is not selected inside that branch alone. It is measured against the full asymmetric neutral-parent resolution

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+),$$

and, in particular, against the first exposure environment of the completed positive-side closure,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

Thus the chamber counts below refer to nested embedding environments of the positive closure branch, not to independent new particle generations, new colors, or new gauge sectors.

6.1. Electron-Zero Koide Boundary

The smallness of the electron suggests that the first endpoint analysis should be formulated at the electron-zero boundary,

$$m_e = 0.$$

Let

$$b = \sqrt{\frac{m_\mu}{m_\tau}}.$$

Then the Koide relation becomes

$$\frac{1+b^2}{(1+b)^2} = \frac{2}{3}.$$

Solving gives

$$b^2 - 4b + 1 = 0,$$

so that the physical branch is

$$b = 2 - \sqrt{3}.$$

Therefore the electron-zero Koide boundary gives

$$\boxed{\sqrt{\frac{m_\mu}{m_\tau}} = 2 - \sqrt{3} = \tan \frac{\pi}{12}}.$$

The trigonometric form is only a compact representation of the same boundary value. The underlying claim is not that a literal spacetime angle is fundamental. Rather, the electron-zero boundary is the root-vector representation of a three-branch boundary of the two-sector Koide invariant.

The electron-zero solution is not exact phenomenologically, but it shows that the dominant μ - τ hierarchy is already constrained by the Koide sector-power relation. The small nonzero electron can therefore be treated as an endpoint displacement rather than as a freely fitted independent mass.

6.2. Exact Koide Displacement from a Nonzero Electron Endpoint

Define the two root ratios

$$a = \sqrt{\frac{m_e}{m_\mu}}, \quad b = \sqrt{\frac{m_\mu}{m_\tau}}.$$

Then the charged-lepton root vector may be written, up to an overall scale, as

$$\mathbf{r}_\ell \sim (ab, b, 1).$$

The Koide relation becomes

$$\frac{a^2b^2 + b^2 + 1}{(ab + b + 1)^2} = \frac{2}{3}.$$

Equivalently,

$$(a^2 - 4a + 1)b^2 - 4(a + 1)b + 1 = 0.$$

Solving for the physical branch gives

$$b(a) = \frac{2 + 2a - \sqrt{3a^2 + 12a + 3}}{a^2 - 4a + 1}.$$

At $a = 0$,

$$b(0) = 2 - \sqrt{3},$$

as required.

Thus once the electron endpoint parameter a is known, the Koide sector-power relation fixes the remaining μ - τ ratio:

$$\frac{m_\mu}{m_\tau} = b(a)^2.$$

The problem of deriving the charged-lepton hierarchy is therefore reduced to deriving one small endpoint parameter,

$$a^2 = \frac{m_e}{m_\mu}.$$

6.3. Leading Endpoint Leakage Count

We now apply the same equal-weight counting principle to the small nonzero electron endpoint. The first exposed non-protected chamber is not a free-standing Z_4 object. It is the first embedding environment of the completed positive-side closure,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The ordered four-slot exposure environment contains

$$4!$$

chambers. One chamber is protected by the identity/closure condition and does not contribute to endpoint leakage. Therefore the number of available non-protected leakage chambers is

$$4! - 1 = 23.$$

The endpoint leakage must also be projected through the charged-family and compensational-family overlap structure. At leading order this gives a three-by-three family-pair denominator,

$$3^2.$$

Therefore the leading number of admissible endpoint leakage channels is

$$N_e^{(0)} = 3^2(4! - 1) = 207.$$

The leading electron endpoint is then one equal-weight leakage unit:

$$\boxed{\frac{m_e}{m_\mu} \simeq \frac{1}{3^2(4! - 1)} = \frac{1}{207}.$$

Numerically,

$$\frac{1}{207} = 0.0048309179\dots,$$

which already identifies the correct scale of the observed pole-mass ratio,

$$\frac{m_e}{m_\mu} \simeq 0.0048363317.$$

The interpretation is that the electron endpoint is not selected by a free carrier-Lagrangian parameter and not by the lepton-facing Z_2^- branch in isolation. It is the first equal-weight leakage unit made available when the completed positive-side closure ($Z_2^+ \rightarrow Z_3^+$) is embedded into the Z_4 exposure environment, after the protected closure chamber has been removed from the count.

6.4. Nested Higher-Interface Corrections

The leading endpoint leakage can be refined by higher embedding environments. After the first exposure

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4,$$

the same completed positive-side closure is conditionally embedded into successively higher environments,

$$Z_4 \hookrightarrow Z_5, \quad Z_5 \hookrightarrow Z_6, \quad Z_6 \hookrightarrow Z_7, \quad \dots$$

These steps should not be read as adding independent new layers or opening independent new electron endpoints. They are nested refinements of an endpoint already selected at the first exposure embedding.

Therefore the charged endpoint tower is written as

$$\frac{m_e}{m_\mu} = \frac{1}{3^2(4! - 1)} [1 + q_5 + q_5 q_6 + q_5 q_6 q_7 + \dots].$$

The product form expresses conditional refinement: q_5 refines the endpoint inside the Z_5 embedding environment, q_6 refines the already Z_5 -refined endpoint inside the Z_6 environment, and so on.

Applying the same structural counting rule to the deeper embedding environments gives

$$q_n = \frac{2/n}{3(n! - 1)}, \quad n \geq 5.$$

Here $n! - 1$ is the non-protected ordered chamber count at the Z_n embedding level, the factor 3 is the residual single-family projection after the endpoint channel has already been selected, and $2/n$ is the two-sided spinorial residue distributed through the n -slot embedding environment.

Thus

$$q_5 = \frac{2/5}{3(5! - 1)} = \frac{2/5}{357} = 0.001120448\dots,$$

and

$$q_6 = \frac{2/6}{3(6! - 1)} = \frac{1}{6471} = 0.000154536\dots$$

The all-order endpoint formula is therefore

$$\boxed{\frac{m_e}{m_\mu} = \frac{1}{3^2(4! - 1)} \left[1 + \sum_{k=5}^{\infty} \prod_{n=5}^k \frac{2/n}{3(n! - 1)} \right]}.$$

Keeping only the Z_5 correction gives

$$\frac{m_e}{m_\mu} \approx \frac{1}{207} \left[1 + \frac{2/5}{3(5! - 1)} \right].$$

Keeping the Z_6 nested correction gives

$$\frac{m_e}{m_\mu} \approx \frac{1}{207} \left[1 + \frac{2/5}{3(5! - 1)} + \frac{2/5}{3(5! - 1)} \frac{2/6}{3(6! - 1)} \right].$$

Section 8 evaluates this endpoint expansion numerically and shows that the charged tower stabilizes by the Z_6 embedding level. The remaining central residual will then be interpreted as a neutral-parent closure obstruction requiring neutrino-sector compensation.

7. Minimality and Uniqueness of the Counting Construction

The endpoint formula is numerically striking only if its ingredients are fixed before comparison with the data. This section therefore states the precise sense in which the construction is minimal and examines the most obvious alternatives.

7.1. Why the First Exposure Is Z_4

The Z_2^- branch distinguishes the charged transverse and neutral continuation directions, but it does not provide the completed positive closure against which a stable charged endpoint can be measured. The positive branch first closes at $Z_2^+ \rightarrow Z_3^+$. An endpoint is an exposure of that completed closure rather than another copy of it; hence the first external ordered environment has four slots. Starting at Z_3 would identify closure with exposure, while starting at Z_5 or higher would skip the first admissible environment and violate level-wise selection.

Proposition 1 (First-exposure count). *Under neutral-parent closure, protected-sector quotient, and first-level selection, the number of exposed ordered chambers at the leading endpoint level is $4! - 1$.*

Proof. The first environment external to the completed threefold positive closure has four ordered slots and therefore $4!$ ordered chambers. Exactly one chamber is the identity/closure continuation of the already completed parent and is protected. The quotient therefore leaves $4! - 1$ exposed chambers. Retaining the identity chamber violates the protected-sector postulate; deleting more than one chamber requires an additional structural distinction not present at this level. $\square$

7.2. Why the Leading Family Factor Is 3^2

The endpoint is not a count of three observed charged leptons inserted by hand. It is a projection between a three-direction charged-family read-out and a three-direction compensational-family organization. Before a specific endpoint is selected, both indices remain unresolved. Equal primitive weight therefore applies to the ordered pair, producing 3×3 possibilities. A single factor of three would erase one of the two unresolved projections; a factor 3^3 would add an independent family index that is absent from the architecture.

Proposition 2 (Leading endpoint denominator). *With one endpoint leakage unit, one four-slot first exposure, one protected closure chamber, and two unresolved three-direction projections, the leading endpoint weight is uniquely*

$$w_e^{(0)} = \frac{1}{3^2(4! - 1)} = \frac{1}{207}.$$

7.3. Why Higher Terms Multiply Rather than Add Independently

After the Z_4 endpoint has been selected, the Z_5 environment does not open a second primitive electron endpoint. It resolves the already selected endpoint. By conditional refinement, the joint weight is the parent weight times the conditional weight. The same argument repeats at each later stage. Consequently, the admissible series has the form

$$w_e^{(0)}(1 + q_5 + q_5q_6 + q_5q_6q_7 + \cdots),$$

not $w_e^{(0)}(1 + q_5 + q_6 + q_7 + \cdots)$. The latter would treat later refinements as independent first exposures and would retroactively change the primitive endpoint measure.

At level n , the ordered non-protected chamber count is $n! - 1$. The charged endpoint has already been chosen, so only one residual three-direction projection remains. The continuation residue has two opposed orientations distributed over n slots, giving $2/n$. These three independent pieces fix

$$q_n = \frac{2/n}{3(n! - 1)}, \quad n \geq 5.$$

Replacing $2/n$, 3 , or $n! - 1$ by another factor requires, respectively, a different orientation content, a different residual family projection, or a different protected ordered chamber space. Such a replacement is a different model, not an alternative implementation of the stated one.

7.4. Why the Neutral Factor Is $7/30$

The neutral sector begins with three indistinguishable continuation directions. The leading solar overlap is therefore democratic, $1/3$, before protected closure is imposed. At the first neutral correction, two continuation slots are selected from five. There are $\binom{5}{2} = 10$ such pairs, one of which is protected by closure. Because the overlap belongs to one of the three neutral directions, the protected displacement is $1/[3\binom{5}{2}] = 1/30$. Hence

$$\sin^2 \theta_{12}^{\text{str}} = \frac{1}{3} - \frac{1}{30} = \frac{3}{10}.$$

The endpoint uses the closure-aligned complement and then projects back through one neutral direction:

$$\Omega_v = \frac{1 - 3/10}{3} = \frac{7}{30}.$$

No oscillation datum enters these fractions. Using $3/10$ directly in the endpoint would select the transverse overlap rather than its closure-aligned complement; omitting the final division by three would fail to project the neutral correction onto the charged endpoint.

7.5. Scope of the Uniqueness Claim

The propositions above establish conditional uniqueness, not uniqueness among all conceivable theories of flavor. The result can be summarized as follows:

Theorem 2 (Conditional uniqueness of the endpoint tower). *Fix the neutral-parent branch architecture, first-exposure rule, level-wise indifference, one protected identity chamber at each ordered embedding level, two opposed continuation orientations, and conditional refinement. Then the leading endpoint denominator, higher refinement factors, and leading neutral overlap are fixed as*

$$3^2(4! - 1), \quad q_n = \frac{2/n}{3(n! - 1)}, \quad \Omega_v = \frac{7}{30}.$$

Any different numerical rule must alter at least one declared structural assumption.

This theorem is the appropriate answer to the concern that many unrelated integer formulas might approximate the same mass ratio. The empirical agreement alone does not establish the architecture; the scientific claim is that a small, explicit set of structural assumptions generates several linked quantities and can be tested jointly.

8. Numerical Evaluation

We now evaluate the endpoint expansion derived in Section 6. The leading endpoint estimate is

$$\left(\frac{m_e}{m_\mu}\right)_{Z_4} = \frac{1}{3^2(4! - 1)} = 0.004830917874 \dots$$

Including the first two nested corrections gives

$$\left(\frac{m_e}{m_\mu}\right)_{Z_5} = \frac{1}{207} \left[1 + \frac{2/5}{3(5! - 1)} \right] = 0.004836330668 \dots,$$

and

$$\left(\frac{m_e}{m_\mu}\right)_{Z_6} = \frac{1}{207} \left[1 + \frac{2/5}{3(5! - 1)} + \frac{2/5}{3(5! - 1)} \frac{2/6}{3(6! - 1)} \right] = 0.004836331504 \dots$$

For comparison, using the 2022 Committee on Data (CODATA) recommended values [6,7,23]:

$$m_e = 0.51099895069(16) \text{ MeV}, \quad m_\mu = 105.6583755(23) \text{ MeV},$$

one obtains

$$\left(\frac{m_e}{m_\mu}\right)_{\text{obs}} = 0.004836331699 \dots,$$

or equivalently

$$\left(\frac{m_\mu}{m_e}\right)_{\text{obs}} = 206.7682827(46).$$

The uncertainty in the inverse ratio is about 4.6×10^{-6} and is dominated by the muon-mass uncertainty.

The endpoint leakage estimates for m_e/m_μ are summarized in Table 2, and the corresponding inverse ratios are shown in Table 3. The leading $1/207$ estimate already identifies the correct scale. The Z_5 correction moves the result close to the observed value, and the Z_6 correction stabilizes the charged endpoint tower. The remaining difference is small and experimentally meaningful at the present precision, but it is only at the level of about two standard deviations.

Table 2. Endpoint leakage estimates for m_e/m_μ .

Approximation	Formula	Value
Leading Z_4 leakage	$\frac{1}{3^2(4! - 1)}$	0.004830917874
Including Z_5	$\frac{1}{207} \left[1 + \frac{2/5}{3(5! - 1)} \right]$	0.004836330668
Including Z_6	$\frac{1}{207} \left[1 + \frac{2/5}{3(5! - 1)} + \frac{2/5}{3(5! - 1)} \frac{2/6}{3(6! - 1)} \right]$	0.004836331504
Observed	m_e/m_μ	0.004836331699

Table 3. Comparison of the endpoint prediction for m_μ/m_e with the observed mass ratio.

Approximation	Predicted m_μ/m_e	Difference from Observation	Pull
Leading Z_4 leakage	207.000000000	$+2.31717 \times 10^{-1}$	$\sim 5.0 \times 10^4 \sigma$
Including Z_5	206.768326805	$+4.410 \times 10^{-5}$	$\sim 9.6 \sigma$
Including Z_6	206.768291043	$+8.335 \times 10^{-6}$	$\sim 1.8 \sigma$

The large pull for the leading approximation should not be misread as a large absolute error; it reflects the high precision of the electron–muon mass ratio. After the Z_6 correction, the absolute residual in the inverse ratio is

$$\Delta_{\text{abs}} = \left(\frac{m_\mu}{m_e} \right)_{Z_6} - \left(\frac{m_\mu}{m_e} \right)_{\text{obs}} = 8.33 \times 10^{-6},$$

which is about 1.8σ relative to the present uncertainty. Equivalently, in the direct ratio,

$$\left(\frac{m_e}{m_\mu} \right)_{Z_6} = 0.004836331504 \dots,$$

while

$$\left(\frac{m_e}{m_\mu} \right)_{\text{obs}} = 0.004836331699 \dots$$

The reason the charged tower stabilizes so quickly is that the endpoint rule is nested. The correction factors are

$$q_n = \frac{2/n}{3(n! - 1)}, \quad n \geq 5.$$

Explicitly,

$$\begin{aligned} q_5 &= 1.120448179 \times 10^{-3}, & q_6 &= 1.545356205 \times 10^{-4}, \\ q_7 &= 1.890019751 \times 10^{-5}, & q_8 &= 2.066850203 \times 10^{-6}. \end{aligned}$$

But the successive contributions to the endpoint series are products:

$$\begin{aligned} q_5 &= 1.120448179 \times 10^{-3}, \\ q_5 q_6 &= 1.731491546 \times 10^{-7}, \\ q_5 q_6 q_7 &= 3.272553220 \times 10^{-12}, \\ q_5 q_6 q_7 q_8 &= 6.763877286 \times 10^{-18}. \end{aligned}$$

Thus the Z_7 and Z_8 terms do not appreciably change the displayed prediction. The all-order nested value is already stabilized, to the precision relevant here, at the Z_6 level:

$$\left(\frac{m_e}{m_\mu} \right)_{\text{tower}} = 0.004836331504 \dots,$$

or

$$\left(\frac{m_\mu}{m_e} \right)_{\text{tower}} = 206.768291043 \dots$$

Next insert the resulting endpoint value

$$a = \sqrt{m_e/m_\mu}$$

into the exact Koide displacement formula

$$b(a) = \frac{2 + 2a - \sqrt{3a^2 + 12a + 3}}{a^2 - 4a + 1}.$$

Then

$$\frac{m_\mu}{m_\tau} = b(a)^2.$$

The resulting μ – τ ratios are shown in Table 4. The uncertainty in the observed ratio is dominated by the present experimental uncertainty in the tau mass.

Table 4. Muon–tau ratio from the endpoint-corrected Koide constraint. Here Δ_{obs} denotes prediction minus observation.

Case	m_μ/m_τ	Δ_{obs}
Leading endpoint $m_e/m_\mu = 1/207$	0.05946571732	$+4.53 \times 10^{-6}$
Charged tower through Z_6	0.05945988605	-1.31×10^{-6}
Observed pole-mass ratio	0.05946119(30)	—

The important point is structural. Once the endpoint displacement is specified, the Koide sector-power relation fixes the remaining charged-lepton root direction. The nested charged tower captures the electron endpoint to high absolute accuracy and stabilizes rapidly. The Z_7 and Z_8 terms show that the remaining central electron–muon discrepancy is not removed by extending the same charged sequence. This is the central diagnostic observation of the paper. It suggests that the remaining gap should be interpreted not as another charged endpoint correction, but as a residual neutral-parent closure obstruction. The next section interprets this obstruction as the part supplied by the neutrino compensational sector.

9. Residual Closure and the Neutrino Compensation Sector

The nested endpoint-leakage rule gives a highly accurate charged-sector prediction for the electron–muon mass ratio. More importantly, it shows where the charged calculation stabilizes. Once the Z_5 and Z_6 nested corrections are included, the subsequent Z_7 and Z_8 terms are far too small to remove the remaining central-value discrepancy. The residual gap is therefore not naturally explained as another correction of the same charged endpoint tower.

In the neutral-parent framework, this is structurally significant. The charged endpoint is not selected in isolation. A resolved weak branch contains a charged transverse lepton and a neutral compensational continuation,

$$Z_2^- \rightarrow \ell_\perp^- \oplus \nu_\parallel.$$

The charged endpoint tower computes the charged-side leakage before full neutral-parent closure. The remaining residue is interpreted here as a small closure obstruction left when one attempts to close the neutral parent using only the charged endpoint structure. The neutrino sector is therefore not introduced as an unrelated empirical addition; it is the neutral continuation branch required to patch this closure seam. The essential neutral-overlap derivation is included here so that the main argument is logically complete. Appendix D retains the fuller geometric interpretation and consistency checks.

9.1. Neutral Exposure and Protected Two-Slot Displacement

Before PMNS read-out, the neutral continuation sector is represented by three indistinguishable compensational directions. Level-wise indifference assigns the primitive solar-sector exposure

$$w_{\text{dem}}^{(v)} = \frac{1}{3}.$$

At the Z_5 neutral interface, the relevant overlap uses two slots. Equal weighting gives $\binom{5}{2} = 10$ unordered pairs. One pair is closure protected and does not contribute to transverse mixing. Projected through one of three neutral directions, its displacement is

$$\delta_{\text{prot}} = \frac{1}{3\binom{5}{2}} = \frac{1}{30}.$$

The leading transverse neutral overlap is consequently

$$\sin^2 \theta_{12}^{\text{str}} = \frac{1}{3} - \frac{1}{30} = \frac{3}{10}.$$

This is an output of the chamber rule, not an insertion of the measured solar angle.

9.2. Closure-Aligned Complement and Charged-Endpoint Projection

The charged residual is patched by the closure-aligned part of the neutral continuation. Its complement is $1 - 3/10 = 7/10$. Projection through one of the three neutral directions gives

$$\Omega_v = \frac{7/10}{3} = \frac{7}{30}.$$

Thus the same construction yields both the solar target and the endpoint correction; they are not independently adjustable.

Let

$$x_{\text{ch}} = \left(\frac{m_e}{m_\mu} \right)_{\text{ch}}$$

denote the charged endpoint prediction stabilized through Z_6 :

$$x_{\text{ch}} = \frac{1}{207} [1 + q_5 + q_5 q_6],$$

where

$$q_5 = \frac{2/5}{3(5! - 1)}, \quad q_6 = \frac{2/6}{3(6! - 1)}.$$

Numerically,

$$q_5 = 1.120448179 \times 10^{-3}, \quad q_6 = 1.545356205 \times 10^{-4},$$

so that

$$q_5 q_6 = 1.731491546 \times 10^{-7}.$$

The charged tower gives

$$x_{\text{ch}} = 0.0048363315040, \quad x_{\text{ch}}^{-1} = 206.768291043.$$

By contrast, using the CODATA charged-lepton values gives

$$x_{\text{obs}}^{-1} = \left(\frac{m_\mu}{m_e} \right)_{\text{obs}} = 206.768282708 \pm 0.0000046,$$

or, in concise notation,

$$\left(\frac{m_\mu}{m_e}\right)_{\text{obs}} = 206.7682827(46).$$

Equivalently,

$$x_{\text{obs}} = 0.00483633169896.$$

Thus the relative residual left by the charged tower is

$$\Delta_{\text{res}} = \frac{x_{\text{obs}}}{x_{\text{ch}}} - 1 = 4.03109 \times 10^{-8}.$$

In the inverse ratio this same residual is

$$x_{\text{ch}}^{-1} - x_{\text{obs}}^{-1} = 8.3350 \times 10^{-6}.$$

Relative to the experimental uncertainty of about 4.6×10^{-6} , this is approximately 1.8σ . The charged tower is therefore already a highly accurate central-value prediction, but it leaves a stable residual whose size is not removed by the deeper charged Z_7 and Z_8 terms.

The residual has the scale expected from a secondary projection of the already selected Z_5 - Z_6 interface. The neutral correction is therefore written as

$$\Delta_\nu = q_5 q_6 \Omega_\nu,$$

where $q_5 q_6$ fixes the interface scale and Ω_ν is the neutral-overlap fraction. Appendix D derives

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}, \quad \Omega_\nu = \frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30}.$$

Thus the compensational correction is

$$\Delta_\nu = q_5 q_6 \frac{7}{30}.$$

Numerically,

$$\Delta_\nu = (1.731491546 \times 10^{-7}) \frac{7}{30} = 4.04015 \times 10^{-8}.$$

This matches the central scale required to close the charged endpoint residual. The associated solar-sector value is

$$\theta_{12}^{\text{str}} = 33.21^\circ,$$

which lies close to the empirical solar-neutrino mixing region [5,10]. The observed solar angle is not used as an input; it is used only as an external comparison.

The neutral-compensated endpoint prediction is therefore

$$x_{\text{full}} = x_{\text{ch}}(1 + \Delta_\nu) = x_{\text{ch}} \left[1 + q_5 q_6 \frac{7}{30} \right].$$

Equivalently,

$$x_{\text{full}} = \frac{1}{207} [1 + q_5 + q_5 q_6] \left[1 + q_5 q_6 \frac{7}{30} \right].$$

This gives

$$x_{\text{full}} = 0.00483633169940, \quad x_{\text{full}}^{-1} = 206.768282689.$$

Hence

$$x_{\text{full}}^{-1} - x_{\text{obs}}^{-1} = -1.87 \times 10^{-8},$$

which is about -0.004σ relative to the present experimental uncertainty. The electron–muon endpoint hierarchy is therefore closed to present experimental precision without inserting the observed solar PMNS angle as a fitted input.

The same relation can be read in reverse as a central-value consistency check. If the charged endpoint residual is used to infer the solar projection required by the correction, then

$$\frac{\cos^2 \theta_{12}}{3} = \frac{\Delta_{\text{res}}}{q_5 q_6},$$

and therefore

$$\sin^2 \theta_{12}^{\text{inv}} = 1 - 3 \frac{\Delta_{\text{res}}}{q_5 q_6} = 0.30157.$$

This central value is close to the structural prediction

$$\sin^2 \theta_{12}^{\text{str}} = 0.30000.$$

Because the present uncertainty in m_μ/m_e is much larger than the final neutral-corrected residual, this inverse reading should not be interpreted as a precision extraction of the solar angle. Its role is more modest: it shows that the central residual left by the charged endpoint tower has the size expected from the first neutral compensational overlap count.

Once m_μ/m_e is fixed, the Koide sector-power theorem fixes the large- τ branch. Set

$$Y = \frac{m_\mu}{m_e} = x_{\text{full}}^{-1}, \quad T = \frac{m_\tau}{m_e}.$$

The Koide relation becomes

$$\frac{1 + Y + T}{(1 + \sqrt{Y} + \sqrt{T})^2} = \frac{2}{3}.$$

Choosing the physical large- τ branch gives

$$\sqrt{T} = 2(1 + \sqrt{Y}) + \sqrt{3Y + 12\sqrt{Y} + 3}.$$

Therefore

$$\left(\frac{m_\tau}{m_e}\right)_{\text{pred}} = 3477.441636, \quad \left(\frac{m_\tau}{m_\mu}\right)_{\text{pred}} = 16.8180612.$$

For comparison, using $m_\tau \simeq 1776.93(9)$ MeV gives approximately

$$\left(\frac{m_\tau}{m_e}\right)_{\text{obs}} = 3477.37(18), \quad \left(\frac{m_\tau}{m_\mu}\right)_{\text{obs}} = 16.81769(85),$$

with uncertainties dominated by the tau mass. The tau-sector comparison is therefore much less precise than the electron–muon test, but it is compatible at the present tau-mass precision scale.

The complete two-step hierarchy prediction is summarized in Table 5.

This agreement should not be interpreted as a complete derivation of the PMNS matrix or of neutrino dynamics. The result derives only the particular solar-sector neutral compensational overlap required by the charged endpoint residual. A complete PMNS theory would still need to derive the atmospheric angle, reactor angle, charge-conjugation–parity (CP) phase, mass ordering, and full neutrino mass structure.

Together with the Koide sector-power theorem, the corrected endpoint prediction gives a compact two-step account of the charged-lepton hierarchy:

$$\text{endpoint count} + \text{neutral compensation} \Rightarrow m_\mu/m_e,$$

$$\text{Koide sector-power equality} \Rightarrow m_\tau/m_e.$$

In this two-step structure, the tau mass is no longer a freely adjustable quantity once the electron–muon endpoint ratio has been fixed. Future higher-precision measurements of the tau mass therefore provide a direct test of the proposed endpoint-plus-neutral-compensation mechanism.

The logic of the endpoint tower and the subsequent neutral-compensation correction is summarized schematically in Figure 4.




Figure 4. Nested endpoint expansion and residual neutral compensation. The leading electron endpoint is opened at the first exposure environment $(Z_2^+ \rightarrow Z_3^+) \leftrightarrow Z_4$. Higher charged corrections are conditional refinements $q_5, q_5q_6, q_5q_6q_7, \dots$, which rapidly stabilize by the Z_6 level. The remaining charged-sector residual is then interpreted as a neutral-parent closure obstruction supplied by the neutral compensational channel. In the present counting interpretation, the first neutral overlap gives the structural factor $7/30$, so the correction is $\Delta_v = q_5q_6(7/30)$.

Table 5. Complete charged-lepton hierarchy prediction. The endpoint-counting tower, corrected by the neutral-compensational factor $q_5 q_6 (7/30)$, predicts m_μ/m_e . The factor $7/30$ is the structural replacement for $\cos^2 \theta_{12}/3$, corresponding to the solar-overlap prediction $\sin^2 \theta_{12} = 3/10$. The Koide sector-power theorem then fixes the large- τ branch. The electron–muon comparison is limited by the present uncertainty in the muon–electron mass ratio, while the tau-sector comparison is limited mainly by the experimental uncertainty in the tau mass.

Ratio	Structural Prediction	Observed Value	Residual (Pull)
m_μ/m_e	206.768282689	206.768282708(46)	-1.9×10^{-8} (-0.004σ)
m_τ/m_e	3477.441636	3477.37(18)	$+0.076$ ($\sim 0.43\sigma$)
m_τ/m_μ	16.8180612	16.81769(85)	$+3.69 \times 10^{-4}$ ($\sim 0.43\sigma$)

10. Interface with the Standard Model and Renormalized Masses

The proposal is intended as a selection layer beneath, not a replacement for, the Standard Model. It is therefore necessary to state how the structural quantities used above are related to conventional particle physics.

10.1. Yukawa Couplings as Effective Read-Out Parameters

In the Standard Model, after electroweak symmetry breaking,

$$m_i(\mu) = \frac{v(\mu)}{\sqrt{2}} y_i(\mu),$$

where $y_i(\mu)$ are renormalized charged-lepton Yukawa couplings and $v(\mu)$ is the renormalized Higgs expectation value. The Standard Model correctly describes how these parameters enter amplitudes and how they run, but it does not determine their three boundary values from an internal flavor principle [11–13].

The present framework addresses that missing selection question. Its root amplitudes do not bypass the Higgs mechanism. Rather, after carrier read-out they are mapped to dimensionless ratios of the effective Yukawa boundary values:

$$\frac{y_i}{y_j} = \frac{m_i}{m_j} = \frac{r_i^2}{r_j^2}$$

at the specified physical comparison point. The common factor v and the common charged-lepton carrier scale cancel. In this sense, the model is analogous in logical placement to flavor models that seek to explain Yukawa textures, while differing in the structural mechanism used to select them.

10.2. Why Pole Masses Are Compared

Renormalized masses run with the scale and scheme; a numerical relation among running parameters must therefore specify both. The invariant claimed here is attached to a completed, isolated particle read-out. Accordingly, the natural first comparison is with pole masses, defined by the position of the propagator pole and associated with the asymptotic one-particle state. Radiative corrections are not denied: they are part of the map from renormalized Lagrangian parameters to the completed pole observable.

This choice is a physical hypothesis and a testable one. The paper does not claim that the same bare numerical relation holds for running modified minimal-subtraction ($\overline{\text{MS}}$) masses at every scale. A complete dynamical theory would need a matching condition

$$\mathbf{r}_{\text{str}} \longrightarrow \{y_i(\mu_0)\} \longrightarrow \{m_i^{\text{pole}}\},$$

including electroweak and quantum-electrodynamical (QED) threshold corrections. The present work identifies the first and last members of this chain at the level of ratios; deriving the full matching map remains open.

10.3. Relation to Established Flavor Approaches

Froggatt–Nielsen models generate hierarchical Yukawa entries through charges and powers of a small symmetry-breaking parameter [24]. Non-Abelian discrete-family models organize multiplets and mixing textures through finite-group representations [25–27]. Family-gauge approaches to Koide-type relations seek radiative protection or dynamical stabilization of the mass formula [28]. Those programs begin with conventional fields, representations, and effective operators.

The present construction differs in three ways. First, its finite Z_n labels describe pre-read-out continuation and closure rather than observable gauge groups. Second, the small hierarchy is obtained from protected chamber counts rather than a fitted expansion parameter. Third, the charged and neutral quantities are linked by one parent closure architecture: the solar overlap is not an independent mixing texture appended to the charged-lepton sector.

These differences do not establish superiority. They define a distinct, more primitive hypothesis that must earn credibility through linked predictions and through eventual embedding into an effective Lagrangian. Conversely, the successful Standard Model treatment of interactions is not evidence against the possibility that its Yukawa boundary data are selected by a deeper admissibility structure.

10.4. What a Complete Interface Would Require

A full implementation must provide: (i) an explicit Lorentz-covariant field or defect realization; (ii) a matching calculation to a renormalized Yukawa matrix; (iii) proof that radiative corrections preserve or predictably deform the structural root direction; (iv) a derivation of the absolute charged-lepton scale; and (v) extension to neutrino masses and the complete PMNS matrix. The toy model in Appendix C addresses only the first item at an illustrative level. These limitations are substantial, but they are separate from the narrower conditional result established here: the stated structural postulates select definite dimensionless ratios.

11. Discussion

The preceding sections give the main charged-lepton result. The Koide relation is interpreted as an equal-sector-power theorem, the small electron endpoint is interpreted as a first exposed leakage of the neutral-parent architecture, and the remaining charged-sector residual is closed by a neutral compensational overlap. The central conceptual lesson is that the charged-lepton hierarchy should not be modeled as an isolated lepton-only pattern. It appears to depend on the coupled structure of the lepton-facing branch, the positive closure branch, and the neutral compensational branch of a common parent.

11.1. Particle Structure as a Coupled Whole

A key point of the present work is that particle structure should be modeled as a coupled whole rather than as a collection of isolated islands. In the Standard Model,

electrons, neutrinos, quarks, nucleons, and gauge interactions are described by distinct fields and representations, but observed processes are organized by charge closure, weak transitions, spinorial read-out, neutral compensation, and confined positive-side closure. This suggests that the deeper structural problem should be formulated at the level of the whole particle architecture.

The charged-lepton hierarchy illustrates this point directly. In the present proposal, the electron endpoint is a leakage of the lepton-facing Z_2^- branch, but the count that selects it is controlled by the full asymmetric neutral-parent resolution

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+).$$

Thus the positive-side closure branch participates in the charged-lepton mass hierarchy, even though it is not itself a charged lepton. The relevant endpoint count begins with the first exposure embedding of the positive closure,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4,$$

and the higher corrections arise from nested embeddings of the same positive-side structure.

The neutral sector also participates. The charged endpoint tower stabilizes by the Z_6 embedding level, and deeper charged terms are too small to close the observed electron–muon residual. The residual is therefore interpreted as the part of the neutral-parent closure that belongs to the neutral continuation branch,

$$Z_2^- \rightarrow \ell_\perp^- \oplus \nu_\parallel.$$

In this sense, the lepton mass ratios are not controlled by the charged leptons alone. They are controlled by how the charged branch, the positive closure branch, and the neutral compensational branch fit together inside one carrier-defect architecture.

The neutron beta-decay pattern provides a useful physical analogy:

$$n \rightarrow p + e^- + \bar{\nu}_e.$$

This process does not display an isolated electron, proton, or neutrino as an independent structural island. It displays a correlated positive, negative, and neutral closure pattern. The present framework takes this as motivation for using a charge-neutral parent as a theory of origin: the observed particle channels are treated as different carrier read-outs of a common neutral organization.

At a more primitive level, the neutral parent may be viewed as a carrier-readable manifestation of an archetype loop: a primitive closure structure whose different embeddings generate lepton-facing, positive/hadron-facing, and neutral-compensating read-outs. The full development of this archetype-loop particle architecture is beyond the scope of the present paper. The narrower claim made here is that the charged-lepton hierarchy already shows evidence of whole-architecture control: Koide sector-power balance, endpoint leakage, positive-side nested embedding, and neutral compensation are different projections of one neutral-parent closure structure.

11.2. Implications for PMNS Mixing

The neutrino sector enters the present paper through a residual closure diagnostic. The charged endpoint tower gives a highly accurate bare charged-sector prediction, but after the Z_6 embedding level the remaining discrepancy is too large to be closed by deeper charged-tower terms. The proposed interpretation is that this residual is not another

charged endpoint correction. It is the part of the neutral-parent closure belonging to the neutral compensational read-out.

This gives a natural link to PMNS mixing. In the present framework, the PMNS matrix is not an unrelated neutrino-sector object added after the charged-lepton hierarchy is fixed. It is interpreted as an overlap structure of the neutral compensational channel. The charged endpoint calculation points to the need for a neutrino-like neutral continuation, while the first neutral compensational overlap supplies the closure scale required by the residual.

The solar-sector overlap is not inserted from the observed PMNS matrix, but structurally derived. The primitive neutral-continuation sector begins from a democratic three-direction compensational exposure. The first neutral displacement is then identified with one protected closure chamber in the Z_5 two-slot interface. This gives

$$\sin^2 \theta_{12}^{\text{str}} = \frac{1}{3} - \frac{1}{3\binom{5}{2}} = \frac{3}{10},$$

and therefore

$$\frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30}.$$

The neutral correction entering the endpoint residual is consequently

$$\Delta_\nu = q_5 q_6 \frac{7}{30}.$$

Thus the solar-sector factor is a structural prediction of the neutral-compensation count, not an empirical parameter fitted from neutrino data.

The resulting value,

$$\theta_{12}^{\text{str}} = 33.21^\circ,$$

lies close to the empirical solar mixing region described by the PMNS framework [5,8–10]. This comparison is important, but it should be read in the correct direction. The observed solar angle is not used to tune the charged-lepton hierarchy. Rather, the charged-lepton residual predicts a solar compensational overlap that is independently close to the observed solar-neutrino mixing scale.

This should not be read as a full derivation of the PMNS matrix. The present paper derives only the particular solar-sector overlap that enters the charged-lepton residual. A complete PMNS theory would need to derive the full compensation-overlap basis, including the atmospheric angle, reactor angle, CP phase, and mass-ordering dependence. The present result supplies one nontrivial overlap component and shows why the neutral sector must participate in completing the charged-lepton hierarchy.

11.3. Relation to Previous Approaches to the Charged-Lepton Hierarchy

The charged-lepton hierarchy and the Koide relation have been studied from several viewpoints. Some approaches treat Koide's relation as an empirical mass formula or as a geometric regularity in square-root mass space [2–4,29–31]. Others seek to protect or explain it through radiative mechanisms, family gauge symmetry, flavor symmetries, discrete flavor groups, mass-matrix textures, or Froggatt–Nielsen-type hierarchical charges [24–28,32]. These approaches have produced many suggestive structures, but the charged-lepton masses usually still depend on fitted parameters, chosen charges, assumed textures, or model-specific symmetry-breaking patterns.

The present proposal differs in its starting point. It does not begin with a Yukawa texture, mass matrix, or continuous flavor symmetry. It begins with a premetric internal counting rule: sectors indistinguishable at the level at which they first become exposed

carry equal primitive weight. Koide's relation is then interpreted as the carrier-facing root-vector representation of

$$P_{\text{dem}} = P_{\text{split}}.$$

In this view, the Koide cone is not a fundamental geometric postulate and not a fitted mass relation. It is the continuous representation of a sector-power balance.

The proposal also differs by treating the charged-lepton hierarchy as a whole-architecture problem. The electron endpoint is selected by the positive-side nested embedding, while the remaining charged residual is closed by a neutral compensational sector. Thus the charged-lepton masses are not generated entirely inside an isolated charged-lepton sector. The hierarchy is a projection of the coupled neutral-parent architecture.

The treatment of the electron is correspondingly distinctive. Instead of treating m_e as a small fitted Yukawa coupling or suppressed texture entry, the present proposal interprets it as one endpoint leakage unit selected by the first exposed non-protected chamber of

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The leading factor $3^2(4! - 1) = 207$ and the higher nested corrections are proposed as chamber-counting consequences of the neutral-parent carrier-defect architecture, not as independent numerical adjustments.

The treatment of the neutrino sector is also distinctive. The neutral correction is not introduced as a freely adjustable mixing parameter. The solar-sector overlap entering the residual is derived as

$$\frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30},$$

equivalent to the prediction

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}.$$

Thus the charged-lepton hierarchy and one PMNS-related overlap are linked by the same residual-closure logic, rather than treated as unrelated phenomena.

Finally, once the corrected electron–muon ratio is obtained, the Koide sector-power theorem fixes the large- τ branch. Thus m_τ is not an independent fitted parameter in the proposed mechanism. Future higher-precision tau-mass measurements can therefore directly test the combined endpoint-plus-neutral-compensation structure.

11.4. A Solar-Angle Prediction from Neutral Compensation

The neutral-compensation mechanism also gives a scoped prediction for the solar neutrino mixing angle. This point is important because the solar angle is not used as an input in the charged-lepton endpoint calculation. The charged endpoint tower first fixes the interface scale $q_5 q_6$. The neutral-overlap count then supplies the compensational fraction

$$\Omega_\nu = \frac{7}{30}.$$

In the normalized 1-2 neutral subspace, this fraction is interpreted as

$$\Omega_\nu = \frac{\cos^2 \theta_{12}^{\text{str}}}{3}.$$

Therefore

$$\frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30},$$

so that

$$\cos^2 \theta_{12}^{\text{str}} = \frac{7}{10}, \quad \sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}.$$

Equivalently,

$$\sin^2 \theta_{12}^{\text{str}} = 0.30000$$

and

$$\theta_{12}^{\text{str}} = \arcsin \sqrt{\frac{3}{10}} = 33.21^\circ.$$

This should be read as a leading structural target, not as a fitted parameter. Current global analyses of neutrino oscillation data give a solar angle close to this value. For example, NuFIT 6.0 reports approximately

$$\sin^2 \theta_{12} = 0.307\text{--}0.308_{-0.011}^{+0.012}, \quad \theta_{12} = 33.68_{-0.70}^{+0.73}^\circ,$$

depending on the data variant and mass ordering [10]. The structural prediction

$$\sin^2 \theta_{12}^{\text{str}} = 0.30000, \quad \theta_{12}^{\text{str}} = 33.21^\circ$$

therefore lies within the present empirical solar-mixing region.

The significance of this comparison is not that the present experimental data already force the value $3/10$. They do not. The present uncertainty in the solar angle is still large enough that nearby values remain allowed. Rather, the point is that the same neutral-overlap count which closes the central charged-lepton endpoint residual also predicts a specific solar-sector target. In this sense the solar angle becomes an independent future test of the neutral-compensation interpretation.

There are two possible outcomes. If future solar-neutrino and reactor data converge toward

$$\sin^2 \theta_{12} = 0.30000,$$

then the leading neutral-overlap count would be strongly supported. If future data converge to a nearby but distinct value, the present framework would require a calculable higher-order correction to the leading neutral-overlap fixed point. Such corrections could arise from the full PMNS embedding, from carrier read-out effects, or from higher neutral overlap chambers not included in the leading Z_5 two-slot count.

Thus the conservative prediction of the present paper is

$$\sin^2 \theta_{12} = 0.30000$$

as the leading neutral-compensation target, with possible subleading corrections to be derived only in a more complete theory of the full PMNS matrix. This gives the proposal a clear falsifiable direction: future precision measurements of the solar mixing angle should determine whether the neutral-overlap value $3/10$ is the final value or the leading term of a slightly corrected expansion.

11.5. Status of the Proposal

The present framework is a structural derivation theory for dimensionless particle invariants. It is not a local microscopic field theory of the premetric domain, because spacetime geometry, Lagrangian dynamics, and continuous field evolution are not assumed to be fundamental there. Instead, it specifies a pre-read-out admissibility structure and asks which dimensionless relations survive carrier read-out.

Equal primitive weighting is natural at that level: sectors that are indistinguishable by the available branch, monodromy, protection, and closure data have no structural basis for unequal weights. The principle is therefore not a numerical fit, although a fuller carrier theory should explain how this structural indifference is transmitted into effective masses and mixings.

Within the stated architecture, the same rule yields the Koide sector-power equality, the small electron endpoint, the rapidly stabilizing charged tower, the neutral completion factor $7/30$, the leading target

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}, \quad \theta_{12}^{\text{str}} = 33.21^\circ,$$

and the downstream tau ratios. The remaining open task is to embed the endpoint and neutral-overlap rules in a general mathematics of admissible chamber exposure and carrier read-out. The present paper does not derive the other PMNS parameters or the complete Standard Model interaction structure.

The claim is therefore limited but testable: the charged-lepton hierarchy may be a carrier-readable invariant of a neutral-parent admissibility structure. A persistent experimental exclusion of the solar target or tau ratios would falsify the leading construction, while a nearby discrepancy would require a calculable higher-order neutral or carrier-read-out correction.

12. Predictions, Cross-Checks, and Falsifiability

The framework should be judged by linked consequences rather than by the accuracy of one isolated integer expression. Table 6 summarizes the quantities fixed by the present leading construction.

Table 6. Outputs of the present construction and their testing status.

Quantity	Structural Result	Test
Charged tower	m_μ/m_e 206.768291043	= Diagnoses the size and sign of the neutral residual
Neutral-compensated ratio	m_μ/m_e 206.768282689	= Direct pole-mass comparison
Solar neutral overlap	$\sin^2 \theta_{12}^{\text{str}} = 0.30000$	Future global oscillation fits
Tau–electron ratio	$m_\tau/m_e = 3477.441636$	Improved tau-mass measurements
Tau–muon ratio	$m_\tau/m_\mu = 16.8180612$	Improved tau-mass measurements
Higher charged corrections	Factorially suppressed nested products	Any larger same-sequence correction would contradict the tower

The first line is a cross-check rather than the final result: the charged tower must stop at a residual of the scale later supplied by neutral closure. If its deeper terms could be rearranged to fit the data, the neutral prediction would lose content. The factorial suppression forbids that rearrangement within the stated model.

The following outcomes would falsify or force revision of the present leading construction:

- (F1) A future pole-mass determination excludes either predicted tau ratio after the electron–muon ratio is fixed.
- (F2) Global oscillation data exclude $\sin^2 \theta_{12} = 3/10$ at a level that cannot be attributed to a derived higher-order neutral term.

- (F3) A formally consistent implementation of the same postulates yields a different leading endpoint denominator, a different q_n , or a different neutral overlap. This would disprove the claimed conditional uniqueness.
- (F4) A microscopic Standard Model matching calculation shows that no carrier-readable defect realization can produce the required pole ratios without introducing equivalent continuous fitting parameters.
- (F5) The framework fails to extend coherently to the remaining PMNS observables or produces mutually inconsistent sector assignments.

The last two tests are theoretical rather than experimental, but they are no less important. The current numerical agreement is evidence worth investigating, not proof of the architecture. The strongest future test is overconstraint: the same protected-chamber rules should determine new quantities without adding a new numerical prescription for each sector.

13. Scope and Limitations

The present calculation is limited to dimensionless charged-lepton mass ratios and the leading neutral overlap that closes the residual endpoint seam. It does not determine the absolute charged-lepton scale, the full PMNS matrix, or a complete formal theory of premetric admissibility and carrier read-out.

Three foundational tasks remain. First, the electron-zero boundary and the first endpoint count should be derived from a general topology of the charged-lepton branch. Second, the factors

$$3^2(4! - 1), \quad q_n = \frac{2/n}{3(n! - 1)}$$

should emerge as theorems of a general chamber-exposure calculus rather than only as the minimal rules fixed within the present architecture. Third, the factor $7/30$ should appear as a special case of a broader neutral-compensation algebra.

The neutrino result is correspondingly restricted. The paper derives the leading solar-sector target

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}, \quad \frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30},$$

but not the atmospheric angle, reactor angle, CP phase, mass ordering, or neutrino mass spectrum.

The use of pole masses is appropriate to the completed asymptotic read-out considered here, but a fuller theory must connect the structural ratios to renormalized Yukawa parameters and their scale dependence. Likewise, the appendices provide a finite closure architecture and a toy Lorentz-readable dictionary, not a derivation of the complete Standard Model weak, color, confinement, or hadronic sectors.

14. Conclusions

This paper has developed a conditional structural account of the charged-lepton hierarchy. In the proposed premetric admissibility space, equal primitive power in the democratic and branch-splitting sectors gives

$$P_{\text{dem}} = P_{\text{split}},$$

whose carrier-facing root-amplitude representation is the Koide relation. The Koide cone is therefore interpreted as a representation of a discrete sector-power equality rather than as an independently chosen continuous mass geometry.

The small electron endpoint is selected by the first non-protected exposure of the positive-side closure and refined conditionally through higher interfaces:

$$\left(\frac{m_e}{m_\mu}\right)_{\text{ch}} = \frac{1}{3^2(4! - 1)} \left[1 + \sum_{k=5}^{\infty} \prod_{n=5}^k \frac{2/n}{3(n! - 1)} \right].$$

The tower stabilizes rapidly, giving

$$\left(\frac{m_\mu}{m_e}\right)_{\text{ch}} = 206.768291043.$$

Because deeper charged terms are too small to remove the residual, the remaining seam is assigned to the neutral continuation branch.

The leading neutral count gives

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}, \quad \frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30},$$

and hence the completed endpoint prediction

$$\left(\frac{m_\mu}{m_e}\right)_{\text{pred}} = 206.768282689.$$

The Koide sector-power relation then fixes

$$\left(\frac{m_\tau}{m_e}\right)_{\text{pred}} = 3477.441636, \quad \left(\frac{m_\tau}{m_\mu}\right)_{\text{pred}} = 16.8180612.$$

Thus endpoint counting plus neutral completion determines the electron–muon hierarchy, while the Koide condition determines the tau branch.

These quantities are parameter-free targets within the declared architecture, not inputs used to tune its integer factors. Improved tau and solar-mixing measurements can therefore test the leading construction directly. A distinct counting scheme satisfying the same postulates but yielding different endpoint weights would likewise defeat the conditional uniqueness claim.

The framework is not presented as a microscopic Lagrangian theory of the premetric domain. Its appropriate completion is a mathematical theory of admissible objects, protected sectors, composition and closure, together with a carrier read-out map to effective masses, mixings, and interactions. At its present stage, the result is more limited: one neutral-parent architecture organizes the Koide relation, the electron endpoint, charged-tower saturation, a solar-angle target, and the tau branch. This linked structure is the principal evidence that the charged-lepton hierarchy may be a carrier-readable invariant rather than an isolated numerical coincidence.

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 author.

Conflicts of Interest: The author is employed by Silicon Minds Inc. The company had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A. From Internal Root Amplitudes to Physical Rest Mass

This appendix clarifies the sense in which the mass ratios derived in the main text are ratios of physical charged-lepton rest masses rather than ratios of abstract internal labels. The distinction is important. The equal-weight counting argument is formulated in a premetric internal defect space, whereas physical mass is measured as rest energy, inertial response, and gravitational source. The purpose of this appendix is not to derive the complete microscopic carrier dynamics. Rather, it is to state the minimal read-out assumptions under which the internal root-amplitude ratios computed in the paper become ratios of physical rest masses.

Appendix A.1. Three Levels of Description

We distinguish three levels.

First, the internal counting theory assigns a normalized root-amplitude direction

$$\hat{\mathbf{r}}_\ell = \frac{(r_e, r_\mu, r_\tau)}{\sqrt{r_e^2 + r_\mu^2 + r_\tau^2}}.$$

The main text argues that this direction is fixed by equal-weight defect-internal counting. This level is premetric and scale-free.

Second, a carrier read-out map embeds the internal defect state into a localized physical excitation. This step supplies a common charged-lepton energy scale and converts internal root amplitudes into localized rest-energy densities.

Third, at distances large compared with the defect core, the localized excitation is described by an effective one-particle stress-energy tensor. The integrated rest energy is the physical rest mass [33].

The present paper derives only the first level and uses the minimal structure of the second and third levels to justify comparison with observed charged-lepton mass ratios [5–7,23]. The absolute charged-lepton scale and the detailed carrier Lagrangian are left for future work.

Appendix A.2. Minimal Read-Out Assumptions

Let r_i be the internal root amplitude for charged-lepton channel $i = e, \mu, \tau$. The minimal read-out assumptions are:

1. **Common carrier channel.** The three charged leptons are read out through the same charged-lepton carrier sector. Therefore a common carrier scale M_ℓ multiplies the three channels.
2. **Quadratic energy read-out.** The localized rest-energy density is quadratic in the internal amplitude. Equivalently, the primitive amplitude is r_i , while the physical energy weight is r_i^2 . This is consistent with the standard quadratic dependence of field energies and quantum probabilities on amplitudes [11,13].
3. **Normalized localization.** For each channel, the carrier produces a localized rest-frame energy profile $\rho_i(\mathbf{x})$ normalized by

$$\int d^3x \rho_i(\mathbf{x}) = 1.$$

Differences in short-distance profile shape are allowed, but they do not alter the integrated rest-energy coefficient at the level of mass ratios considered here.

4. **Long-field stress-energy equivalence.** At distances large compared with the defect core, the localized carrier excitation is represented by the standard effective

stress-energy tensor of a massive particle with rest mass equal to its integrated rest energy [33].

Under these assumptions, the rest-frame energy density of channel i has the schematic form

$$T_i^{00}(\mathbf{x}) = M_\ell r_i^2 \rho_i(\mathbf{x}) + \text{short-distance derivative/core terms.}$$

After integration over the localized core, derivative and profile details renormalize the common carrier scale or the normalized profile, while the channel-dependent coefficient is r_i^2 . Thus

$$m_i = \int d^3x T_i^{00}(\mathbf{x}) = M_\ell r_i^2.$$

Consequently,

$$\frac{m_i}{m_j} = \frac{r_i^2}{r_j^2}.$$

This is the sense in which the internal root-amplitude ratios are identified with physical charged-lepton mass ratios. The paper does not declare an internal number to be mass by definition. Rather, it assumes a carrier read-out in which the internal amplitude controls the coefficient of the localized rest-energy density.

Appendix A.3. Why the Read-Out Is Quadratic

The quadratic form

$$m_i \propto r_i^2$$

is not an arbitrary choice. It is the standard way in which an amplitude or field mode contributes to a positive energy functional. In ordinary field theory, energies and probabilities are typically quadratic in amplitudes: a normal mode with amplitude A contributes energy proportional to $|A|^2$, and a quantum probability is quadratic in a wave amplitude [11,13]. Likewise, if r_i is the internal root amplitude of a defect channel, the corresponding positive carrier-energy weight is naturally r_i^2 .

This is also why the Koide relation is naturally written in terms of square roots of masses. The square-root variables are treated as the primitive internal amplitudes, while the physical masses are their quadratic carrier read-out:

$$r_i = \sqrt{m_i / M_\ell}.$$

The main text therefore works in root space not for algebraic convenience alone, but because the internal variables are amplitudes whose squared norms supply physical rest-energy weights.

Appendix A.4. Relation to Stress-Energy and the Long-Field Limit

Physical mass is not only an internal energy label. In relativistic physics, rest mass is the invariant coefficient of a localized stress-energy source. For a localized excitation with worldline $z^\mu(\tau)$, the long-field effective stress-energy tensor takes the standard point-particle form [33]

$$T_{i,\text{eff}}^{\mu\nu}(x) = m_i \int u^\mu u^\nu \frac{\delta^{(4)}(x - z_i(\tau))}{\sqrt{-g}} d\tau,$$

where u^μ is the four-velocity. In a local rest-frame or asymptotically stationary long-field description, the integrated rest energy may be written schematically as

$$m_i = \int_\Sigma T_i^{\mu\nu} n_\mu \xi_\nu d\Sigma,$$

where ξ^μ denotes the relevant rest-frame time-translation direction and n^μ is the unit normal to the spatial slice Σ . Therefore, if the carrier read-out gives

$$\int d^3x T_i^{00}(\mathbf{x}) = M_\ell r_i^2,$$

then the same quantity $M_\ell r_i^2$ is the inertial rest energy and the long-field gravitational source coefficient. In the weak-field limit, it is the quantity that contributes to the source term for the metric through the effective stress-energy tensor. Thus the mass ratios computed from r_i^2 are ratios of the same masses that appear in the long-wavelength gravitational and inertial descriptions.

Appendix A.5. Why Only Ratios Are Claimed

The counting argument fixes the normalized direction

$$(r_e, r_\mu, r_\tau)$$

in internal root space. It does not fix the common carrier scale M_ℓ . This is why the present paper predicts ratios such as

$$\frac{m_e}{m_\mu} = \frac{r_e^2}{r_\mu^2}, \quad \frac{m_\mu}{m_\tau} = \frac{r_\mu^2}{r_\tau^2},$$

rather than the absolute masses in MeV.

This limitation is also consistent with the structure of the Koide relation [2–4]. Koide is homogeneous under a common rescaling

$$m_i \mapsto \lambda m_i.$$

It therefore constrains the direction of the charged-lepton mass vector, not its overall scale. The present theory explains this direction as an internal counting invariant and leaves the absolute scale as a separate carrier-read-out problem.

Appendix A.6. Comparison with Standard Particle-Physics Language

In the Standard Model, charged-lepton masses are written schematically as

$$m_i = \frac{v}{\sqrt{2}} y_i,$$

where v is the electroweak scale and y_i are dimensionless Yukawa couplings [5,12,13]. The Standard Model does not explain the hierarchy among the three y_i . The present framework may be viewed, at the level of ratios, as replacing the unexplained hierarchy of Yukawa coefficients by a hierarchy of internal carrier-read-out weights:

$$y_i \propto r_i^2.$$

The common scale remains external to the counting argument, while the relative weights are fixed by the defect-internal invariant.

Thus the comparison with measured charged-lepton masses is meaningful in the same sense that a theory of Yukawa ratios would be meaningful: it targets the dimensionless hierarchy among physical rest masses, not the absolute electroweak or carrier scale [5–7,23].

Appendix A.7. What Is Not Claimed Here

Several difficult questions are not solved in this appendix.

First, we do not derive the full microscopic carrier Lagrangian. We only state the minimal form of the read-out needed for the internal root-amplitude ratios to become physical rest-mass ratios.

Second, we do not derive the common charged-lepton scale M_ℓ . The present paper is a theory of the dimensionless hierarchy, not of the absolute MeV scale.

Third, we do not claim that all short-distance core profiles are identical. The claim is weaker: after normalization and common charged-lepton carrier read-out, the channel-dependent integrated rest-energy coefficient is r_i^2 .

Fourth, we do not replace the effective field-theoretic description of charged leptons. At long wavelength the result must reduce to the standard description of a localized particle with mass m_i . The present proposal concerns the internal origin of the ratios among these masses.

With these limitations stated, the bridge from internal counting to physical mass is defensible: the equal-weight argument fixes the relative root amplitudes; the carrier read-out maps squared amplitudes to localized stress-energy; the integrated stress-energy is the physical rest mass; and the same mass appears in the long-field effective stress-energy tensor.

Appendix B. Why Z_2 Is the Primitive Free-Fermion Read-Out and Why $Z_3 \rightarrow Z_4 \rightarrow Z_5 \rightarrow \dots$ Belongs to the Positive Closure Side

This appendix explains the structural role of the discrete monodromy symbols used in the main text. The goal is not to replace the Standard Model groups $SU(2)$ and $SU(3)$. Rather, the point is to clarify which finite pregeometric closure structures are being represented before carrier read-out, and why they appear as familiar continuous structures only after the defect becomes a Lorentz-readable physical particle.

Appendix B.1. Finite Monodromy Before Carrier Read-Out

In the present framework, a particle is treated as a localized carrier-supported defect. Before carrier read-out, the defect pocket is not assumed to be an ordinary differentiable internal manifold equipped with a metric, spin connection, gauge bundle, or continuum group action. Its primitive data are more limited: continuation, closure, branch opposition, exclusion, and admissible return. This use of finite branching and monodromy language is in the same broad mathematical spirit as standard uses of covering spaces, monodromy, and topological defect classification, although the present carrier-defect interpretation is specific to this paper [15,18,21].

For this reason the notation

$$Z_2, \quad Z_3, \quad Z_4, \quad \dots$$

is used to describe finite internal monodromy and closure structure. It does not mean that the Standard Model gauge groups are being replaced by finite cyclic groups at the observable field-theory level. The finite symbols describe pre-read-out closure rules inside the defect pocket. The usual continuous structures, including spinorial $SU(2)$ and color $SU(3)$, become appropriate only after carrier embedding supplies a Lorentz-readable and gauge-readable effective description [11,13,16].

Thus the intended logical order is

finite internal closure $\longrightarrow$ carrier embedding $\longrightarrow$ continuous effective field representation.

The finite structure is not the observed gauge theory. It is the pregeometric closure pattern whose carrier read-out may be represented by the observed gauge and spinorial language.

Appendix B.2. Z_2 as the Primitive Free-Fermion Branch Structure

The first nontrivial internal distinction of a neutral parent is a two-sided branch opposition. This is denoted by Z_2 . It records the minimal fact that an unresolved neutral configuration can be resolved into complementary read-out directions:

$$P_0 \longrightarrow Z_2^- \oplus Z_2^+.$$

At this stage Z_2 is not yet the continuous spin group. It is the primitive branch opposition: two complementary orientations of the same unresolved parent.

After carrier read-out, however, a free fermion must be represented spinorially. The familiar spinorial fact is that a 2π rotation does not return a spinor to itself, but changes its sign, while a 4π rotation restores it:

$$2\pi : \psi \mapsto -\psi, \quad 4\pi : \psi \mapsto \psi.$$

This two-valued return behavior is the standard double-cover structure of spinorial representations [11,14]. It is precisely the type of two-valued return structure that a primitive Z_2 branch opposition can support after Lorentz-readable carrier embedding.

The carrier-facing continuous representation of this two-valued structure is the spin lift

$$\text{Spin}(3) \simeq SU(2),$$

rather than the rotation group $SO(3)$ itself [11,14,17]. Thus the present interpretation is

$$Z_2 \text{ internal branch opposition} \xrightarrow{\text{carrier read-out}} SU(2) \text{ spinorial free-fermion representation.}$$

The Z_2 structure is therefore not a substitute for $SU(2)$. It is the pregeometric primitive that makes a spinorial $SU(2)$ read-out natural once the defect becomes a Lorentz-readable free fermion.

This also explains why the charged-lepton side is fundamentally Z_2 -like in the present paper. A charged lepton is a free fermion-facing read-out of the neutral parent. Its primitive internal branch structure is two-sided, and after carrier embedding this two-sided structure appears in the usual spinorial language of free fermions.

Appendix B.3. Lepton-Facing and Neutral-Compensating Z_2 Read-Outs

The Z_2^- side of the neutral parent supports the lepton-facing read-out. In the charged-lepton application it is useful to resolve this side schematically as

$$Z_2^- \longrightarrow \ell_{\perp}^- \oplus \nu_{\parallel}.$$

Here $\ell_{\perp}^-$ denotes the charged transverse read-out and $\nu_{\parallel}$ denotes the neutral continuation read-out. The symbols $\perp$ and $\parallel$ are structural labels: they indicate different carrier embeddings of the same Z_2 -origin branch, not literal directions in ordinary spacetime.

This distinction is important. The electron, muon, tau, and neutrino channels are not assumed to be unrelated primitive particles before read-out. They are different carrier-facing manifestations of an underlying two-sided lepton branch. The charged branch supplies the charged-lepton endpoint hierarchy; the neutral continuation branch supplies the compensational channel needed to close the small residual left by the charged endpoint tower. The later appearance of the solar-sector angle should therefore be understood as a

structural neutral-overlap prediction used to close the charged endpoint residual, not as an empirical PMNS input or as a full derivation of neutrino mixing [5,8,10].

Appendix B.4. Z_3 as the First Completed Positive-Side Closure

The positive side of the neutral parent is not merely another free Z_2 fermion branch. It continues toward the first completed positive-side closure:

$$Z_2^+ \longrightarrow Z_3^+.$$

The Z_3^+ structure is interpreted as the primitive threefold closure condition of the positive branch. In carrier-facing language, this side is compatible with hadron-facing threefold closure. This is why it is natural to associate the positive side with the confined or hadron-facing sector.

Again, Z_3 should not be identified with the full Standard Model $SU(3)$ color gauge group. The relation is instead

$$Z_3 \text{ pregeometric threefold closure} \xrightarrow{\text{carrier read-out}} SU(3) \text{ color-compatible effective description.}$$

The finite Z_3 records the minimal closure requirement: incomplete threefold sectors are not carrier-readable as isolated asymptotic positive-side objects, while completed threefold closure can support a hadron-facing read-out. The continuous $SU(3)$ gauge description belongs to the effective field-theory layer after carrier read-out [5,12,13,16].

Thus the neutral-parent architecture used in this paper is asymmetric:

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+).$$

The negative side supports free-fermion lepton-facing read-outs; the positive side supports the first completed closure branch. Both sides belong to the same charge-neutral parent.

The notation $Z_2^+ \rightarrow Z_3^+$ should be read as an internal embedding relation, not as the addition of an independent Z_3 layer on top of a completed Z_2 particle. The positive branch is a nested closure structure: the Z_2^+ branch orientation is completed only through its embedding into the threefold Z_3^+ closure. In this sense, the positive-side primitive relevant for later exposure is not a standalone Z_2^+ , but the embedded closure object

$$(Z_2^+ \rightarrow Z_3^+).$$

Appendix B.5. Nested Embedding into $Z_4, Z_5, Z_6, \dots$

The higher structure

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4 \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow Z_7 \hookrightarrow \dots$$

is attached to the positive closure side. It is not a sequence in which new layers are simply added externally to previous layers. Rather, it is a nested embedding hierarchy. The already formed positive closure $(Z_2^+ \rightarrow Z_3^+)$ is first embedded into a Z_4 exposure environment; that embedded structure may then be embedded into a Z_5 environment, and so on.

Thus $Z_4, Z_5, Z_6, \dots$ should not be interpreted as additional particle generations, new colors, or new independent gauge sectors. They are carrier-readability environments in which the same positive-side closure structure is exposed, tested, and refined at successively higher interface order. This distinction is important because observed color and electroweak gauge structures are already described by the Standard Model effective field theory; the present Z_n hierarchy is instead a pre-read-out closure hierarchy proposed to underlie selected dimensionless mass ratios [5,11,13].

The reason the first exposure begins with the embedding into Z_4 is structural. The Z_2 level provides branch opposition, but it does not by itself give a completed positive closure. The first completed positive support is the embedded closure

$$(Z_2^+ \rightarrow Z_3^+).$$

Only after this threefold closure exists can it be placed into a larger exposure environment. The first such environment is Z_4 :

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The Z_4 slot is therefore not a fourth color. It is the first carrier-readable exposure environment of the completed positive-side closure. Likewise, the later steps

$$Z_4 \hookrightarrow Z_5, \quad Z_5 \hookrightarrow Z_6, \quad Z_6 \hookrightarrow Z_7, \quad \dots$$

are nested embedding refinements of the same positive-side closure architecture.

This explains the logic of the endpoint expansion in the main text. The electron endpoint is a negative-side charged leakage, but it is not selected inside the negative Z_2^- branch alone. It is measured against the full neutral-parent architecture. Since the positive side first becomes complete as the embedded closure $(Z_2^+ \rightarrow Z_3^+)$, the first non-protected exposure chamber available to the charged endpoint is the ordered chamber space of the embedding

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

Therefore the factor

$$4! - 1$$

appearing in the leading endpoint count is not arbitrary. It is the ordered four-slot exposure chamber of the first embedding environment of the completed positive-side closure, with the protected identity/closure chamber removed.

Appendix B.6. Why the Higher Endpoint Corrections Are Nested

Once the electron endpoint has been opened through the first embedding environment

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4,$$

later environments do not create independent new electron endpoints. They refine the already selected endpoint leakage conditionally. This is why the charged endpoint expansion has the nested form

$$\frac{m_e}{m_\mu} = \frac{1}{3^2(4! - 1)} [1 + q_5 + q_5q_6 + q_5q_6q_7 + \dots].$$

The factors $q_5, q_6, q_7, \dots$ are not independent additive parameters. They are conditional refinement weights associated with the successive nested embedding environments

$$Z_4 \hookrightarrow Z_5, \quad Z_5 \hookrightarrow Z_6, \quad Z_6 \hookrightarrow Z_7, \quad \dots$$

Equivalently, the hierarchy should be read as

$$((Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4) \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow Z_7 \hookrightarrow \dots,$$

rather than as a list of independent layers added side by side.

The rule

$$q_n = \frac{2/n}{3(n! - 1)}, \quad n \geq 5,$$

therefore has the following interpretation. The factor $n! - 1$ counts the non-protected ordered chamber possibilities at the Z_n embedding level; the factor 3 is the residual single-family projection after the endpoint has already been selected; and $2/n$ is the two-sided spinorial residue distributed through the n -slot embedding environment.

The rapid suppression of these terms is also structurally meaningful. The deeper environments are not equally strong new openings. They are higher conditional refinements of an endpoint already selected at the first exposure embedding. This is why the charged tower stabilizes by the Z_6 level, and why the remaining discrepancy cannot plausibly be closed by simply continuing the charged tower to $Z_7, Z_8, \dots$

Appendix B.7. Summary of the Read-Out Hierarchy

The structural hierarchy used in the paper can now be summarized as

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+),$$

with

$$Z_2^- \longrightarrow \ell_{\perp}^- \oplus \nu_{\parallel},$$

and the positive-side nested embedding hierarchy

$$((Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4) \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow Z_7 \hookrightarrow \dots$$

The Z_2^- side gives the free-fermion lepton-facing primitive whose carrier read-out is spinorial. The $Z_2^+ \rightarrow Z_3^+$ side gives the positive closure primitive whose carrier-facing completion is compatible with the hadron side. The higher Z_n structures are not layers added as independent sectors; they are nested embedding environments of the same completed positive closure.

In summary, this appendix explains why the endpoint expansion in the main text is not an isolated numerical ansatz. It is the charged-lepton projection of a neutral-parent architecture in which the free-fermion Z_2 branch, the positive embedded $Z_2 \rightarrow Z_3$ closure, and the higher

$$(Z_2 \rightarrow Z_3) \hookrightarrow Z_4 \hookrightarrow Z_5 \hookrightarrow \dots$$

embedding hierarchy play distinct but connected roles. A full development of this carrier-defect particle architecture, including its complete relation to the Standard Model gauge and confinement sectors, is beyond the scope of the present paper and is deferred to future work.

Appendix C. A Toy Lorentz Read-Out Model for the Z_2 Lepton Branch and the Embedded $Z_2 \rightarrow Z_3$ Positive Branch

This appendix gives a minimal toy model showing how the finite internal structures used in the paper may become Lorentz-readable after carrier embedding. The purpose is not to derive the full Standard Model representation theory. Rather, the purpose is to demonstrate structural compatibility between two ingredients used in the main text:

$$Z_2^- \quad \text{for the free lepton-facing branch,}$$

and

$$Z_2^+ \rightarrow Z_3^+ \quad \text{for the embedded positive closure branch.}$$

The first naturally admits a spinorial free-fermion read-out. The second naturally admits a hadron-facing closure read-out in which isolated positive-side fragments are not carrier-readable as free asymptotic particles. The toy-model language should be read as a structural read-out dictionary, not as a replacement for the Standard Model gauge description [5,11,13,16].

Appendix C.1. Lorentz Read-Out of the Z_2^- Lepton Branch

Consider a minimal internal two-state branch variable

$$s \in \{+, -\}, \quad s \sim -s \quad \text{only after a completed two-step return.}$$

This is a toy representation of the Z_2 internal opposition. Before carrier read-out, s is not yet an ordinary spin degree of freedom in spacetime. It is only a two-valued continuation label inside the defect-internal structure.

A Lorentz-readable carrier representation must assign this two-valued internal opposition to a field-like object whose state is not returned to itself by a single 2π rotation, but is returned by a 4π rotation. The minimal carrier-facing rule is therefore

$$R(2\pi) : \psi \mapsto -\psi, \quad R(4\pi) : \psi \mapsto \psi.$$

This is the standard two-valued behavior of spinorial representations, associated with the double-cover relation between spin groups and rotation/Lorentz groups [11,14,17]. In a spatial rest-frame representation, it is described by the spin lift

$$\text{Spin}(3) \simeq SU(2),$$

and in a Lorentz-covariant setting by the corresponding spinorial representation of $\text{Spin}(3, 1)$ [11,13,14]. Thus the toy read-out map is

$$Z_2^- \text{ internal branch opposition} \xrightarrow{\text{carrier read-out}} \text{spinorial free-fermion behavior.}$$

This does not mean that Z_2 is the Standard Model weak group. The finite Z_2 structure is the internal primitive. The continuous spinorial representation appears only after the defect is embedded into a Lorentz-readable carrier. In this sense, Z_2^- is the correct primitive for a free lepton-facing branch: it supplies the two-valued return structure whose carrier-facing representation is spinorial.

The lepton-facing branch used in the main text is therefore written schematically as

$$Z_2^- \longrightarrow \ell_{\perp}^- \oplus \nu_{\parallel}.$$

The charged read-out $\ell_{\perp}^-$ and the neutral continuation read-out $\nu_{\parallel}$ are two carrier-facing manifestations of the same Z_2^- -origin branch. Their distinction is not that one is spinorial and the other is not. Rather, both inherit the primitive two-valued branch structure, but they are embedded differently into the carrier: one as a charged transverse read-out and the other as a neutral continuation read-out. The latter appearance of the solar-sector angle should therefore be understood as a structural neutral-overlap prediction used to close the charged endpoint residual, not as an empirical PMNS input or as a full derivation of neutrino mixing [5,8,10].

A compact toy representation is

$$\Psi_{Z_2^-} = \begin{pmatrix} \psi_{\ell} \\ \psi_{\nu} \end{pmatrix},$$

where ψ_ℓ denotes the charged carrier read-out component and ψ_ν denotes the neutral continuation component. Under a spinorial 2π carrier rotation,

$$\Psi_{Z_2^-} \mapsto -\Psi_{Z_2^-},$$

while under a 4π rotation,

$$\Psi_{Z_2^-} \mapsto \Psi_{Z_2^-}.$$

This is the minimal Lorentz read-out of the Z_2^- lepton-facing branch.

Appendix C.2. Lorentz Read-Out of the Embedded $Z_2^+ \rightarrow Z_3^+$ Positive Branch

The positive side is different. It is not a second free Z_2 spinorial branch. Its Z_2^+ polarity becomes carrier-readable only through completion into the threefold positive closure

$$Z_2^+ \rightarrow Z_3^+.$$

The toy model should therefore not assign an isolated free spinor to Z_2^+ alone. Instead, the carrier-facing positive object is an embedded closure object:

$$\Phi_+ = \Phi_{(Z_2^+ \rightarrow Z_3^+)}.$$

A minimal toy closure rule is the following. Let

$$c_j, \quad j = 1, 2, 3,$$

denote three internal positive-side closure components. They are not assumed to be ordinary observed particles before carrier read-out. They are internal closure positions. A carrier-readable positive-side object must satisfy the threefold closure condition

$$c_1 c_2 c_3 \sim 1,$$

or, equivalently at the level of internal phases,

$$\omega_1 \omega_2 \omega_3 = 1.$$

In the simplest cyclic toy case, one may take

$$\omega_j \in Z_3, \quad \omega = e^{2\pi i/3}, \quad \omega^3 = 1.$$

The point is not that hadrons are literally Z_3 phase products. The point is that the positive branch becomes carrier-readable only after a threefold internal closure has been completed.

This gives the qualitative confinement-like rule:

$$c_j \quad \text{not carrier-readable as an isolated asymptotic object,}$$

but

$$c_1 c_2 c_3 \quad \text{carrier-readable as a closed positive-side object.}$$

After carrier read-out, such a threefold closure can be represented by a color-compatible effective description. Thus the relation is

$$Z_3^+ \text{ pregeometric threefold closure} \xrightarrow{\text{carrier read-out}} \text{hadron-facing threefold effective structure.}$$

This is only a structural analogy at the present level. The finite Z_3^+ closure should not be identified with the full $SU(3)$ color gauge theory. The latter belongs to the effective field-theoretic carrier-facing layer, while Z_3^+ denotes the internal closure primitive. The standard color description and the associated confinement problem belong to quantum chromodynamics (QCD) rather than to the finite toy closure rule itself [5,12,13,16].

Appendix C.3. Nested Lorentz-Readable Embeddings of the Positive Closure

The positive closure is not followed by a sequence of externally added independent layers. Rather, the completed positive closure is nested into larger carrier-readability environments:

$$((Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4) \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow Z_7 \hookrightarrow \dots$$

In the toy model, each embedding environment supplies additional ordered chambers in which the same completed positive closure may be exposed and refined. This use of ordered chambers and nested refinement is a structural counting assumption of the present framework, not a standard QCD or electroweak construction.

At the first exposure level,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4,$$

the ordered chamber space contains $4!$ possible chamber orderings. One chamber is protected by the identity/closure condition and is not available for endpoint leakage. Thus the first non-protected exposure count is

$$4! - 1.$$

This is the structural origin of the leading endpoint factor used in the main text.

At the next embedding level,

$$Z_4 \hookrightarrow Z_5,$$

the previously selected endpoint is not reopened as a new independent endpoint. It is refined conditionally inside the Z_5 environment. Likewise, the Z_6 and deeper levels refine already selected structure rather than adding independent channels. This is why the endpoint expansion takes the nested product form

$$1 + q_5 + q_5 q_6 + q_5 q_6 q_7 + \dots,$$

rather than a simple additive sum of unrelated chamber corrections.

The Lorentz read-out interpretation is therefore different on the two sides of the neutral parent:

$$Z_2^- \longrightarrow \text{free spinorial lepton-facing read-out,}$$

whereas

$$Z_2^+ \rightarrow Z_3^+ \longrightarrow \text{closed positive-side hadron-facing read-out.}$$

The first supports free asymptotic fermion behavior. The second supports a closure condition whose incomplete components are not carrier-readable as isolated asymptotic positive-side particles.

Appendix C.4. A Minimal Read-Out Table

The toy read-out dictionary is summarized in Table A1. The table should not be read as a replacement for the Standard Model field content. Its role is only to show why the

structural assignments used in the charged-lepton endpoint calculation are natural inside the neutral-parent carrier-defect architecture.

Table A1. Minimal toy read-out dictionary for the neutral-parent architecture. The entries describe how the finite pre-read-out structures used in the endpoint calculation may be represented after carrier embedding. The table is not a replacement for the Standard Model field content; it is a structural dictionary connecting internal closure data to the effective read-out roles used in the paper.

Internal Structure	Primitive Role	Toy Carrier Read-Out	Physical Interpretation
Z_2^-	Two-valued branch op- position	$\psi \mapsto -\psi$ under a 2π car- rier rotation	Free spinorial lepton-facing read-out
$Z_2^- \rightarrow \ell_\perp^-$	Charged transverse em- bedding	ψ_ℓ	Charged-lepton channel
$Z_2^- \rightarrow \nu_\parallel$	Neutral continuation em- bedding	ψ_ν	Neutrino compensational channel
$Z_2^+ \rightarrow Z_3^+$	Positive threefold clo- sure	$c_1 c_2 c_3 \sim 1$	Hadron-facing closed posi- tive sector
$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4$	First exposure environ- ment	$4! - 1$ non-protected chambers	Leading electron endpoint leakage count
$Z_4 \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow \dots$	Nested embedding re- finement	$q_5, q_5 q_6, \dots$	Conditional endpoint correc- tions

Table A1 should be interpreted only as a toy read-out dictionary. It records the structural roles needed for the endpoint calculation: the lepton-facing side is Z_2 -primitive and spinorial after carrier read-out; the positive side is an embedded $Z_2 \rightarrow Z_3$ closure; and the higher Z_n structures are nested exposure environments that refine the already selected endpoint rather than introducing independent new particle sectors.

Appendix C.5. Scope of the Toy Model

This appendix is deliberately limited. It does not derive the full Lorentz group, the Standard Model weak interaction, the $SU(3)$ color gauge theory, confinement dynamics, or the PMNS matrix. It only shows that the finite internal structures used in the main text have plausible Lorentz-readable carrier manifestations:

Z_2^- as a spinorial free-fermion primitive,

and

$Z_2^+ \rightarrow Z_3^+$ as an embedded positive closure primitive.

The full development of this carrier-defect particle architecture is beyond the scope of the present paper and is deferred to future work. For the present argument, the toy model is sufficient to justify the structural distinction needed in the endpoint calculation: the lepton-facing side is Z_2 -primitive and spinorial after carrier read-out, while the positive side is an embedded $Z_2 \rightarrow Z_3$ closure whose higher Z_n structures are nested exposure environments.

Appendix D. Neutral-Overlap Closure and the Solar Compensational Factor

This appendix explains the neutral-compensation mechanism used in Section 9. The purpose is not to derive the full PMNS matrix. The narrower purpose is to explain why the charged endpoint residual is naturally patched by a neutral overlap factor

$$\frac{7}{30'}$$

and why this factor can be read as the solar-sector compensational overlap associated with the neutrino continuation channel.

The main point is that the charged endpoint tower computes only the charged transverse leakage. It nearly closes the electron–muon ratio, but after the Z_5 and Z_6 refinements the remaining discrepancy is too large to be supplied by deeper charged-tower terms. The residual is therefore interpreted as a small unclosed seam in the full neutral-parent architecture. The neutrino sector enters as the neutral continuation channel needed to close this seam.

Appendix D.1. The Closure Problem Left by the Charged Tower

The charged endpoint tower computes the leakage associated with the charged transverse branch,

$$Z_2^- \rightarrow \ell_\perp^-.$$

This charged leakage is not selected in isolation. It is measured against the full neutral-parent resolution

$$P_0 \longrightarrow Z_2^- \oplus (Z_2^+ \rightarrow Z_3^+),$$

and its leading endpoint count begins at the first positive-side exposure environment,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The higher charged corrections then refine this endpoint through nested positive-side embeddings,

$$Z_4 \hookrightarrow Z_5 \hookrightarrow Z_6 \hookrightarrow \cdots.$$

The important observation is that the charged tower stabilizes rapidly. After the Z_5 and Z_6 refinements are included, the next charged terms are too small to remove the remaining discrepancy with the observed electron–muon mass ratio. Thus the residual is not naturally another ordinary charged endpoint correction. Instead, it indicates that the charged branch has nearly closed its own endpoint leakage, but has not yet closed the full neutral-parent architecture.

In this sense, the residual should be read as a small unclosed seam in the parent structure:

$$\text{charged tower residual} = \text{unclosed neutral-parent balance.}$$

The role of the neutral compensational sector is to patch this seam. The neutrino sector is therefore not added as an unrelated empirical ingredient. It is interpreted as the neutral continuation branch needed to complete the same asymmetric parent closure:

$$Z_2^- \rightarrow \ell_\perp^- \oplus \nu_\parallel.$$

Appendix D.2. Neutral Mixing Structure Before PMNS Read-Out

To make the neutral compensation mechanism precise, one must distinguish the pre-read-out neutral overlap structure from the full carrier-facing PMNS matrix. Before carrier read-out, the present theory does not assume that physical neutrino flavor states and physical neutrino mass eigenstates already exist as ordinary particles. It assumes only a neutral continuation sector whose internal directions may later be read out in PMNS-like form.

We denote the pre-read-out neutral compensation space by

$$\mathcal{H}_\nu^{\text{comp}} = \text{span}\{c_1, c_2, c_3\}.$$

The three directions c_1, c_2, c_3 are primitive compensational-family directions. They are not directions in ordinary space, and they are not yet physical neutrino mass eigenstates. They label the three neutral continuation possibilities paired with the charged-family closure problem.

A second set of directions is needed to describe closure read-out. Let

$$\mathcal{H}_\nu^{\text{cl}} = \text{span}\{n_1, n_2, n_3\}$$

denote the neutral closure basis. For the purpose of defining overlap weights, both sets of directions are taken to be normalized in the pre-read-out overlap representation. The inner products below therefore measure sector-power overlaps rather than ordinary spatial angles.

After charged-family read-out, the three compensational directions may be labeled by the charged-lepton-paired indices

$$c_\alpha, \quad \alpha = e, \mu, \tau,$$

with c_α denoting the read-out labeling of the same three primitive directions c_1, c_2, c_3 . In ordinary particle-physics language, the PMNS matrix compares charged-lepton-paired neutrino read-out directions with neutral mass or closure directions. Structurally, this is an overlap matrix:

$$U_{\alpha i}^{\text{str}} = \langle c_\alpha | n_i \rangle, \quad \alpha = e, \mu, \tau, \quad i = 1, 2, 3.$$

After carrier read-out, this overlap structure is the object represented phenomenologically by the PMNS matrix [5,8–10].

The present paper does not derive the full matrix $U_{\alpha i}^{\text{str}}$. It derives only the normalized solar-sector overlap in the 1-2 neutral subspace. In that subspace, the electron-paired neutral direction may be represented schematically as

$$c_e^{(12)} = \cos \theta_{12}^{\text{str}} n_1 + \sin \theta_{12}^{\text{str}} n_2.$$

This equation defines the normalized 1-2 neutral subspace after projecting out the remaining neutral directions. In the usual PMNS parameterization, the corresponding solar angle is likewise the normalized 1-2 mixing angle, rather than the full matrix element $|U_{e2}|^2$ without qualification.

Thus

$$\sin^2 \theta_{12}^{\text{str}}$$

is the neutral solar-leakage weight in the normalized 1-2 subspace, while

$$\cos^2 \theta_{12}^{\text{str}}$$

is the complementary closure-aligned weight. This is why the notation θ_{12}^{str} is used. It is the pre-read-out structural counterpart of the solar PMNS angle, not an empirical PMNS input inserted into the charged-lepton hierarchy.

Appendix D.3. Primitive Democratic Neutral Exposure

At the primitive neutral-compensation level, the three directions c_1, c_2, c_3 are indistinguishable. No carrier-facing distinction has yet selected one as larger, lighter, heavier, more strongly coupled, or dynamically preferred. By the equal-weight principle, they therefore carry equal primitive weight:

$$c_1 : c_2 : c_3 = 1 : 1 : 1.$$

The primitive solar leakage is one direction selected out of this threefold neutral exposure. Hence the uncorrected democratic value is

$$\sin^2 \theta_{12}^{(0)} = \frac{1}{3}.$$

This should not be read as a measured PMNS input. It is the democratic pre-read-out value before the first protected neutral displacement is imposed. Equivalently, before the first neutral refinement, one solar leakage direction occupies one of three equal compensational directions.

The complementary neutral component at this primitive level is

$$\cos^2 \theta_{12}^{(0)} = 1 - \frac{1}{3} = \frac{2}{3}.$$

However, the charged endpoint residual is not patched at this primitive level. The residual appears only after the charged endpoint tower has already selected and stabilized the Z_5 - Z_6 interface. Therefore the neutral overlap must be evaluated at the corresponding first higher neutral interface, not at the uncorrected democratic level.

Appendix D.4. Why the First Neutral Correction Is a Two-Slot Overlap

A neutral mixing angle is an overlap quantity, not a one-slot occupation quantity. A one-slot count would ask which single neutral continuation slot is occupied. A mixing or compensation angle asks how one neutral read-out direction overlaps another closure direction. It is therefore intrinsically pairwise.

For this reason, the first neutral correction is counted not by single Z_5 slots but by two-slot overlap chambers among the five exposed Z_5 continuation positions. Define the Z_5 two-slot overlap space by

$$\mathcal{C}_5^{(2)} = \{\{i, j\} : i, j \in Z_5, i < j\}.$$

Its size is

$$|\mathcal{C}_5^{(2)}| = \binom{5}{2} = 10.$$

The pairs are counted as unordered because the compensational quantity is a sector-power overlap, not an ordered transition amplitude. The overlap between slots i and j is the same primitive chamber as the overlap between slots j and i . Interchanging the two labels does not create a new sector-power chamber.

The five slots are counted as exposed continuation positions of the embedding environment. They are not quotiented by a global cyclic rotation at this stage, because the overlap chambers represent carrier-facing exposed positions inside the selected interface. The only symmetry reduction used here is the unordered-pair identification $\{i, j\} = \{j, i\}$.

This also explains why the neutral correction begins at Z_5 , rather than at Z_4 . The Z_4 exposure opens the leading charged endpoint,

$$(Z_2^+ \rightarrow Z_3^+) \hookrightarrow Z_4.$$

The neutral sector does not open a separate leading charged endpoint. It patches the residual after the first higher refinement of the endpoint structure has become available. The first such higher neutral overlap environment is the Z_5 two-slot exposure.

Appendix D.5. The Protected-Chamber Displacement

Among the ten pairwise Z_5 overlap chambers, one chamber is protected by the neutral-parent closure condition. It is not available as an ordinary solar-leakage chamber. Instead, it is reserved as the closure unit whose removal distinguishes the leakage-facing part from the closure-aligned part.

Since the primitive democratic solar leakage is $1/3$, removing one protected chamber out of the ten Z_5 two-slot overlaps subtracts

$$\frac{1}{3} \cdot \frac{1}{10} = \frac{1}{30}$$

from the democratic solar leakage. The structurally corrected solar leakage is therefore

$$\sin^2 \theta_{12}^{\text{str}} = \frac{1}{3} - \frac{1}{30} = \frac{9}{30}.$$

Thus the proposed solar-sector prediction is

$$\boxed{\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}}.$$

The interpretation is simple. The neutral continuation begins with a democratic leakage weight of $1/3$. The first Z_5 two-slot refinement contains ten possible overlap chambers. One chamber is protected by closure and is removed from the ordinary leakage count. The result is a slightly reduced solar leakage weight, $3/10$, rather than the uncorrected democratic value $1/3$.

Appendix D.6. Closure-Aligned Complement and Endpoint Projection

The charged endpoint residual is not patched by the solar leakage part itself. It is patched by the complementary neutral component aligned with parent closure. Since

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10},$$

the closure-aligned complement is

$$\cos^2 \theta_{12}^{\text{str}} = 1 - \frac{3}{10} = \frac{7}{10}.$$

This $7/10$ is the neutral component available to close the charged endpoint seam.

However, the charged endpoint residual sees only one compensational family projection of this closure-aligned neutral component. The same threefold compensational-family structure that entered the endpoint projection therefore contributes a further factor $1/3$. Hence the neutral overlap fraction that actually patches the charged endpoint is

$$\Omega_\nu = \frac{1}{3} \cos^2 \theta_{12}^{\text{str}} = \frac{1}{3} \cdot \frac{7}{10} = \frac{7}{30}.$$

Thus

$$\boxed{\Omega_\nu = \frac{7}{30}}.$$

Combining this overlap with the Z_5 - Z_6 interface scale gives the neutral-compensational correction

$$\Delta_\nu = q_5 q_6 \Omega_\nu = q_5 q_6 \frac{7}{30}.$$

This is the correction used in Section 9.

The sequence of counts is summarized as

$$\frac{1}{3} \longrightarrow \frac{1}{3} - \frac{1}{30} = \frac{3}{10} \longrightarrow 1 - \frac{3}{10} = \frac{7}{10} \longrightarrow \frac{7}{10} \cdot \frac{1}{3} = \frac{7}{30}.$$

In words: democratic neutral exposure gives $1/3$; removing one protected chamber from the ten Z_5 two-slot overlaps gives the solar leakage $3/10$; the complementary closure-aligned part is $7/10$; and one compensational-family projection gives the endpoint patch factor $7/30$. This sequence is illustrated schematically in Figure A1.

The neutral-overlap count underlying the compensational factor is shown schematically in Figure A1.




Figure A1. Neutral-overlap closure mechanism underlying the compensational factor. The primitive neutral continuation begins with a democratic $1/3$ solar leakage among three compensational directions. Removing one protected chamber from the ten Z_5 two-slot overlaps gives $\sin^2 \theta_{12}^{\text{str}} = 3/10$. The complementary closure-aligned component is $7/10$. One compensational-family projection then gives the endpoint patch factor $\Omega_\nu = 7/30$, so that the neutral correction is $\Delta_\nu = q_5 q_6 (7/30)$.

Appendix D.7. Why the Factor Patches the Charged Endpoint Residual

The charged endpoint tower has the form

$$\left(\frac{m_e}{m_\mu} \right)_{\text{ch}} = \frac{1}{3^2(4! - 1)} [1 + q_5 + q_5 q_6 + q_5 q_6 q_7 + \dots],$$

where

$$q_n = \frac{2/n}{3(n! - 1)}.$$

The product structure is essential. The factor q_5 is not an independent additive correction unrelated to the leading endpoint. It is a conditional refinement inside the already selected endpoint. Likewise, q_6 refines the already Z_5 -refined endpoint. Thus

$$q_5 q_6$$

marks the first nested interface at which the charged endpoint has become highly stabilized.

The residual left after the charged tower has stabilized has the scale expected from a secondary projection of this Z_5 - Z_6 interface. This motivates writing the neutral correction in the form

$$\Delta_\nu = q_5 q_6 \times \Omega_\nu.$$

The charged calculation determines the interface scale $q_5 q_6$. The neutral compensation calculation determines the overlap factor $\Omega_\nu = 7/30$. Hence

$$\Delta_\nu = q_5 q_6 \frac{7}{30}.$$

This is why the neutrino sector is described as compensational. It does not replace the charged endpoint tower. It does not generate a new independent charged-lepton mass parameter. Rather, it supplies the small neutral component required to make the already selected charged endpoint compatible with the full neutral-parent closure.

A useful analogy is a seam in a three-part closure. The charged endpoint tower constructs the charged side of the seam and nearly closes it. The positive-side nested embedding fixes the chamber architecture in which the seam is measured. The neutral continuation then provides the final closure-aligned patch. The size of the seam is set by $q_5 q_6$, and the fraction of the neutral sector aligned with the seam is $7/30$.

Appendix D.8. Relation to the Solar PMNS Angle

The notation

$$\sin^2 \theta_{12}^{\text{str}}$$

is used because the neutral-overlap factor has the same structural role as the solar component of the PMNS mixing matrix. In the present proposal,

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}, \quad \theta_{12}^{\text{str}} = 33.21^\circ.$$

This value lies close to the empirical solar-neutrino mixing scale summarized in global neutrino-oscillation analyses [5,8–10].

The logical direction is important. The observed solar mixing angle is not inserted into the charged-lepton hierarchy calculation. Instead, the charged endpoint residual identifies the need for a neutral compensational patch, and the neutral-overlap count predicts a solar-sector value close to the empirical PMNS solar angle. Thus the solar angle enters the paper as an external consistency comparison, not as a fitted parameter.

Appendix D.9. What Is and Is Not Claimed

The neutral-overlap argument is a structural counting proposal. It claims that the solar-sector overlap needed to close the charged endpoint residual is not an arbitrary fitted PMNS input, but follows from a simple protected-chamber displacement in the first neutral Z_5 two-slot exposure.

The argument does not claim to derive the full PMNS matrix. A complete neutrino-sector theory would need to derive the atmospheric angle, the reactor angle, the CP phase, mass ordering, and the full neutrino mass structure. The present appendix derives only the particular solar-sector overlap needed for the charged-lepton hierarchy:

$$\sin^2 \theta_{12}^{\text{str}} = \frac{3}{10}, \quad \frac{\cos^2 \theta_{12}^{\text{str}}}{3} = \frac{7}{30}.$$

This is sufficient for the endpoint residual calculation, but it is not yet a complete theory of neutrino mixing.

References

- Li, B. Particle Structure from Codimension-Two Carrier Closure. *Symmetry* **2026**, *18*, 1154. <https://doi.org/10.3390/sym18071154>.
- Koide, Y. A fermion-boson composite model of quarks and leptons. *Phys. Lett. B* **1983**, *120*, 161–165.
- Koide, Y. New view of quark and lepton mass hierarchy. *Phys. Rev. D* **1983**, *28*, 252. <https://doi.org/10.1103/PhysRevD.28.252>. Erratum in *Phys. Rev. D* **1984**, *29*, 1544. <https://doi.org/10.1103/PhysRevD.29.1544>.
- Foot, R. A note on Koide's lepton mass relation. *Mod. Phys. Lett. A* **1994**, *9*, 169–174.
- 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. (Particle Data Group). Review of Particle Physics. *Phys. Rev. D* **2024**, *110*, 030001.
- National Institute of Standards and Technology. 2022 CODATA Value: Electron mass energy equivalent in MeV. In *The NIST Reference on Constants, Units, and Uncertainty*; NIST: Gaithersburg, MD, USA, 2024. Available online: <https://physics.nist.gov/cgi-bin/cuu/Value?mec2mev=> (accessed on 15 June 2026).
- National Institute of Standards and Technology. 2022 CODATA Value: Muon mass energy equivalent in MeV. In *The NIST Reference on Constants, Units, and Uncertainty*; NIST: Gaithersburg, MD, USA, 2024. Available online: <https://physics.nist.gov/cgi-bin/cuu/Value?mmuc2mev=> (accessed on 15 June 2026).
- Maki, Z.; Nakagawa, M.; Sakata, S. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.* **1962**, *28*, 870–880.
- Pontecorvo, B. Neutrino experiments and the problem of conservation of leptonic charge. *Sov. Phys. JETP* **1968**, *26*, 984–988.
- Esteban, I.; Gonzalez-Garcia, M.C.; Maltoni, M.; Martínez-Soler, I.; Pinheiro, J.P.; Schwetz, T. NuFit-6.0: Updated global analysis of three-flavor neutrino oscillations. *J. High Energy Phys.* **2024**, *12*, 216.
- Weinberg, S. *The Quantum Theory of Fields, Volume I: Foundations*; Cambridge University Press: Cambridge, UK, 1995.
- Weinberg, S. *The Quantum Theory of Fields, Volume II: Modern Applications*; Cambridge University Press: Cambridge, UK, 1996.
- Peskin, M.E.; Schroeder, D.V. *An Introduction to Quantum Field Theory*; Westview Press: Boulder, CO, USA, 1995.
- Lawson, H.B.; Michelsohn, M.-L. *Spin Geometry*; Princeton University Press: Princeton, NJ, USA, 1989.
- Nakahara, M. *Geometry, Topology and Physics*, 2nd ed.; Institute of Physics Publishing: Bristol, UK, 2003.
- Georgi, H. *Lie Algebras in Particle Physics: From Isospin to Unified Theories*, 2nd ed.; Westview Press: Boulder, CO, USA, 1999.
- Zee, A. *Group Theory in a Nutshell for Physicists*; Princeton University Press: Princeton, NJ, USA, 2016.
- Hatcher, A. *Algebraic Topology*; Cambridge University Press: Cambridge, UK, 2002.
- Bott, R.; Tu, L.W. *Differential Forms in Algebraic Topology*; Springer: New York, NY, USA, 1982.
- Milnor, J.W.; Stasheff, J.D. *Characteristic Classes*; Princeton University Press: Princeton, NJ, USA, 1974.
- Mermin, N.D. The topological theory of defects in ordered media. *Rev. Mod. Phys.* **1979**, *51*, 591–648.
- Li, B. Topological classification of admissible reconstruction operations. *Int. J. Topol.* **2026**, *3*, 8. <https://doi.org/10.3390/ijt3020008>.
- National Institute of Standards and Technology. 2022 CODATA Value: Muon–electron mass ratio. In *The NIST Reference on Constants, Units, and Uncertainty*; NIST: Gaithersburg, MD, USA, 2024. Available online: <https://physics.nist.gov/cgi-bin/cuu/Value?mmusme=> (accessed on 15 June 2026).
- Froggatt, C.D.; Nielsen, H.B. Hierarchy of quark masses, Cabibbo angles and CP violation. *Nucl. Phys. B* **1979**, *147*, 277–298.
- Altarelli, G.; Feruglio, F. Discrete flavor symmetries and models of neutrino mixing. *Rev. Mod. Phys.* **2010**, *82*, 2701–2729.
- Ishimori, H.; Kobayashi, T.; Ohki, H.; Okada, H.; Shimizu, Y.; Tanimoto, M. Non-Abelian discrete symmetries in particle physics. *Prog. Theor. Phys. Suppl.* **2010**, *183*, 1–163.
- King, S.F.; Luhn, C. Neutrino mass and mixing with discrete symmetry. *Rep. Prog. Phys.* **2013**, *76*, 056201.
- Sumino, Y. Family gauge symmetry and Koide's mass formula. *Phys. Lett. B* **2009**, *671*, 477–480.
- Rosen, G. Koide's lepton mass formula and the cosmic mass scale. *Mod. Phys. Lett. A* **2008**, *23*, 2385–2390.
- Rivero, A.; Gsponer, A. The strange formula of Dr. Koide. *arXiv* **2005**, arXiv:hep-ph/0505220.
- Brannen, C.A. The lepton masses. *arXiv* **2006**, arXiv:hep-ph/0605331.
- Fritzsch, H. Calculating the Cabibbo angle. *Phys. Lett. B* **1977**, *70*, 436–440.
- Carroll, S.M. *Spacetime and Geometry: An Introduction to General Relativity*; Addison-Wesley: San Francisco, CA, USA, 2004.

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.