Predictivity audit of Koide-type mass relations in the Standard Model

A. Rivero[1]

[1] *Independent researcher*

(Dated: March 22, 2026)

We subject the Koide-type mass relation family to a strict predictivity audit. Five fermion-mass relations—the charged-lepton Koide ratio $K(e,\mu,\tau)$ and four quark-sector extensions—are benchmarked on pole masses (leptons) and five-loop $\overline{\text{MS}}$ running masses at 100 TeV (quarks), with pre-registered search spaces and Monte Carlo look-elsewhere corrections applied to every row. The decisive result is the joint probability: zero events in 3×10^7 independent lepton-plus-quark mass draws simultaneously satisfy the lepton Koide and inverse down-type quark relations at the observed precision, giving $P(K(e,\mu,\tau) \cap K^{-1}(d,s,b)) < 1.14 \times 10^{-7}$ at 68% CL. Individually, the charged-lepton relation achieves a full-scan p-value of 6.67×10^{-5}. Six direct-chain quark relations, two QCD-scale candidates, and the integer-coefficient bion relations are tested under the same discipline and rejected: all collapse when masses are evaluated at a single consistent renormalization scale rather than mixed PDG conventions. The surviving structure is non-accidental but does not constitute a prediction in the beyond-Standard-Model sense; it establishes that the Standard Model fermion mass spectrum encodes algebraic regularity across disjoint sectors at a level inconsistent with random mass hierarchies.

I. INTRODUCTION

The charged-lepton mass ratio

$$K(m_e, m_\mu, m_\tau) = \frac{m_e + m_\mu + m_\tau}{\left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right)^2} = 0.6666605\ldots \tag{1}$$

has equaled 2/3 to within $|\delta_K| = 0.000923\%$ for four decades [1]. Since its discovery, the τ mass has been remeasured repeatedly; each revision has moved K *closer* to 2/3, not farther [2]. The Cauchy–Schwarz inequality restricts K to $[1/3, 1]$, yet the observed value sits precisely at the midpoint between complete degeneracy and maximal hierarchy. Is this a coincidence, or does it reflect non-accidental algebraic structure in the Standard Model mass spectrum?

Extensions of the Koide ratio to the quark sector have been proposed by several authors [1–4]. In its direct form, the ratio has been applied along a mass-ordered chain of quark triples (s, c, b), (c, b, t); in its inverse form $K^{-1} = 2/3$, where $m_i \to 1/m_i$, it has been applied to the down-type quarks (d, s, b) and to the cross-sector triple (s, c, t) [2]. These extensions share a geometric interpretation on the unit sphere in charge space first identified by Foot [3]: the Koide constraint selects a specific orbit of the permutation group, parametrized by a single angle. Despite this algebraic elegance, the quark extensions have been criticized on three grounds:

1. *Look-elsewhere effects.* When multiple mass triples, sign patterns, and scheme choices are scanned, a sub-percent agreement in one combination may be expected by chance. No systematic look-elsewhere correction has been applied in prior work.

2. *Scheme ambiguity.* Published comparisons often mix pole masses (leptons) with $\overline{\text{MS}}$ running masses evaluated at different thresholds (quarks), inflating apparent precision through convention-dependent artifacts [5].

3. *Undisclosed parameters.* The counting of free inputs (branch selector, sign pattern, scheme label) is rarely explicit, making it impossible to separate genuine constraints from fitted coincidences [1, 3].

We address all three criticisms by subjecting the full Koide relation family to a strict predictivity audit. Every relation is benchmarked within a single, explicitly labeled mass scheme: pole masses for leptons (PDG 2024) and five-loop $\overline{\text{MS}}$ running masses at $Q = 100$ TeV for quarks, computed with the `SMDR` package [5]. The search space for look-elsewhere corrections is pre-registered before Monte Carlo evaluation: each admitted row carries a declared number of tests N_{tests} encompassing all tuple choices, sign patterns, branch assignments, and scheme alternatives that were effectively explored.

The headline result is the joint probability. We draw 3×10^7 independent random mass spectra—leptons and quarks sampled separately from a log-uniform null model—and test whether any draw simultaneously satisfies the lepton Koide relation (A1) and the inverse down-type quark relation (A3) at or below the observed precision. No such event is found:

$$P\big(\text{A1} \cap \text{A3}\big) < 1.14 \times 10^{-7} \quad (68\%\ \text{CL}). \tag{2}$$

The simultaneous satisfaction of these two relations in *disjoint* fermion sectors is non-accidental at better than one in ten million.

Among the five individually admitted relations, only the charged-lepton Koide (A1) achieves individual significance at the PROBABLE level, with a full-scan p-value of $p_{\text{row}} = 6.67 \times 10^{-5}$. The remaining relations—the inverse down-type (A3), the cross-sector inverse (A4), and the zero-charge combination (A5)—are individually classified NOT_PREDICTIVE under our conservative verdict tiers, with full-scan p-values between 1.7×10^{-3} and 4.0×10^{-2}. The decisive significance argument is carried entirely by the joint test (2), not by any single row.

Equally important are the negative results. Six direct-chain quark relations that appear to satisfy $K = 2/3$ at sub-percent levels under mixed PDG conventions—where each quark mass is evaluated at its own threshold scale—collapse to 2.7–9.2% deviations when all masses are evaluated at a single common scale. The integer-coefficient "bion" relations ($z_b = 3z_s + z_c$), highlighted in the literature as potential evidence of color structure [2], likewise fail by 10–14% at every consistent $\overline{\mathrm{MS}}$ scale. An exhaustive scan of 26,460 integer-coefficient linear charge relations among six quarks shows that the number of sub-0.1% hits at the mixed PDG convention is below the chance expectation. These failures illustrate precisely why honest scheme discipline and look-elsewhere corrections are essential: without them, the space of apparent coincidences is vast and the significance of any individual hit is inflated by undisclosed freedom.

The paper is organized as follows. Section II defines the algebraic framework: direct and inverse Koide invariants, signed-charge variables, exact branches, and our strict counting protocol. Section III presents the benchmark table: five admitted and six rejected mass-sector relations, with explicit scheme labels and input/output counts. Section IV contains the statistical audit: search-space enumeration, Monte Carlo p-values, the joint probability computation, and the five-tier verdict table. Section V discusses what the result does and does not establish, addresses scheme dependence, the bion failure, and prior referee criticisms. We conclude in Sec. VI.

II. ALGEBRAIC FRAMEWORK

We define the algebraic objects that underlie all benchmark tests in this paper: the Koide ratio, its inverse, their invariant-polynomial forms, and the classification of exact branches.

A. Direct Koide ratio

Given three positive masses m_1, m_2, m_3, the Koide ratio is

$$K(m_1, m_2, m_3) \;=\; \frac{m_1 + m_2 + m_3}{\left(\sqrt{m_1} + \sqrt{m_2} + \sqrt{m_3}\right)^2}\,. \tag{3}$$

By the Cauchy–Schwarz inequality, $\frac{1}{3} \leq K \leq 1$, with the lower bound saturated when all masses are equal and the upper bound when a single mass dominates. Koide [1] observed in 1983 that the charged leptons satisfy $K(m_e, m_\mu, m_\tau) = 2/3$ to $|\delta_K| < 0.001\%$.

B. Signed-charge variables

The structure becomes transparent in the signed square-root charges

$$z_i = s_i \sqrt{m_i}\,, \qquad s_i \in \{+1, -1\}\,, \tag{4}$$

and the elementary symmetric polynomials

$$e_1 = \sum_i z_i\,, \quad e_2 = \sum_{i<j} z_i z_j\,, \quad e_3 = z_1 z_2 z_3\,. \tag{5}$$

In these variables, $K = 2/3$ is equivalent to the single constraint

$$\boxed{e_1^2 = 6\, e_2}\,, \tag{6}$$

which we call the *direct invariant.*

C. Inverse Koide ratio

Replacing $m_i \to 1/m_i$ in Eq. (3) defines the inverse Koide ratio,

$$K^{-1}(m_1, m_2, m_3) \;=\; \frac{\sum_i m_i^{-1}}{\left(\sum_i m_i^{-1/2}\right)^2}\,. \tag{7}$$

In the same signed charges (4), the condition $K^{-1} = 2/3$ becomes the *inverse invariant*

$$\boxed{e_2^2 = 6\, e_1\, e_3}\,, \tag{8}$$

valid only when $e_3 \neq 0$ (no vanishing charge).

D. Residual definition

Throughout this paper we quantify deviations from $K = 2/3$ by the dimensionless residual

$$\delta_K[\%] \;=\; 100 \times \frac{K - 2/3}{2/3} \;=\; 100 \times \left(\frac{3K}{2} - 1\right), \tag{9}$$

and identically for K^{-1}. A relation with $|\delta_K| < 0.01\%$ is classified *practically exact*; 0.01%–0.1% as *strong*; 0.1%–1% as *conditional*; above 1% as *failed.*

E. Exact branches

The direct invariant (6) and the inverse invariant (8) define two algebraic surfaces in charge space. We catalog the boundary and fixed-point structures that serve as anchors for the benchmark tests:

1. **Zero-charge branch.** Setting one charge to zero reduces the direct condition to a quadratic in the mass ratio $u = m_3/m_2$:

$$u^2 - 4u + 1 = 0 \qquad \Longrightarrow \qquad u = 2 \pm \sqrt{3}\,. \tag{10}$$

The predicted ratio $(2+\sqrt{3})^2 \approx 13.93$ can be compared with the observed $m_s/(m_u + m_d)$ (Sec. III).

2. **Self-dual branch.** Requiring that the *same* triple satisfy both the direct and inverse invariants simultaneously yields, after normalizing $u_i = 6z_i/e_1$, the cubic

$$u^3 - 6u^2 + 6u - 1 = (u-1)(u^2 - 5u + 1)\,. \tag{11}$$

The nontrivial roots give the self-dual ratio set

$$\{u_i\} = \left\{1,\, R,\, \frac{1}{R}\right\}, \qquad R = \frac{5+\sqrt{21}}{2}\,. \tag{12}$$

3. **Overlap recurrence.** Consecutive direct Koide triples sharing two elements obey the linear recurrence

$$z_{n+3} = -z_n + 4(z_{n+1} + z_{n+2})\,, \tag{13}$$

with characteristic polynomial $x^3 - 4x^2 - 4x + 1 = (x+1)(x^2 - 5x + 1)$. The eigenvalues $\{-1, R, 1/R\}$ connect the overlap kernel to the self-dual branch. Concrete *physical* chains realized in quark data are classified as conditional structure (Sec. III B).

F. Counting discipline

A benchmark row has $N_{\rm eff}$ effective inputs (branch selector, overall scale, each nontrivial sign assignment, each scheme label) and $N_{\rm out}$ outputs (independent mass predictions). We classify counting status as:

- *predictive*: $N_{\rm eff} < N_{\rm out}$,
- *baseline*: $N_{\rm eff} = N_{\rm out}$,
- *marginal*: $N_{\rm eff} = N_{\rm out} + 1$,
- *postdictive*: $N_{\rm eff} > N_{\rm out} + 1$.

Even at baseline counting, a relation may be statistically significant; the counting status affects the interpretation but does not override Monte Carlo p-values (Sec. IV). Every free choice (sign pattern, scheme label, chain pivot not forced by the algebra) is counted as an input. Common rescaling, permutations of symmetric sets, and the asymptotic golden ratio R (when derived from the overlap kernel rather than inserted) are *not* counted.

III. MASS-SECTOR BENCHMARKS

We evaluate every Koide-type relation against fermion masses in two schemes: pole masses for charged leptons (PDG 2024 [6]) and five-loop $\overline{\rm MS}$ running masses at $Q = 100$ TeV computed with the `SMDR` package [5]. The full mass table is given in Appendix A.

A. Admitted benchmark rows

Table I collects the five relations that survive the algebraic and scheme-consistency discipline of Sec. II. Each row uses masses from a single, consistently defined scheme.

a. A1: Charged-lepton Koide (pole masses). This is the classical Koide observation [1]. Using PDG 2024 pole masses $m_e = 0.51100$ MeV, $m_\mu = 105.658$ MeV, $m_\tau = 1776.86$ MeV, we find $K = 0.6666605$, yielding $\delta_K = -0.000923\%$. The relation has survived 40 years of improving m_τ measurements, with δ_K moving consistently closer to zero. Pole masses for charged leptons are unambiguous; no scheme sensitivity arises.

b. A2: Charged-lepton Koide at 100 TeV. At $Q = 100$ TeV the residual degrades to $\delta_K = +0.177\%$, driven by the 3.5% shift in the τ mass between pole and $\overline{\rm MS}$. This confirms that the charged-lepton Koide is a *pole-mass* statement, consistent with the radiative-stability mechanism of Sumino [7]. Rows A1 and A2 are the same relation in different environments; they count as one relation family.

c. A3: Inverse down-type Koide at 100 TeV. The unique triple of down-type quarks satisfies $K^{-1}(d,s,b) = 2/3$ to $\delta_K = +0.0120\%$ at $Q = 100$ TeV. This relation *improves* monotonically toward the UV: $\delta_K = +0.25\%$ at $Q = 2$ GeV, $+0.012\%$ at 100 TeV, and crosses exact $K^{-1} = 2/3$ at $Q \approx 280$ TeV. The all-positive sign pattern is the unique natural assignment for a same-charge sector, so no sign freedom is counted.

d. A4: Inverse cross-sector Koide at 100 TeV. The triple (s,c,t) with the sign pattern $++-$ satisfies $K^{-1}(s,c,-t) = 2/3$ to $\delta_K = -0.0037\%$ at $Q = 100$ TeV, crossing exact $K^{-1} = 2/3$ at $Q \approx 82$ TeV. This is the highest-precision quark-sector relation. However, the sign $s_t = -1$ is a nontrivial input: the overlap recurrence eigenvalue (-1) does not force this sign because it operates on direct charges z_i, not on the inverse charges $1/z_i$ (see Sec. IV B). The sign costs one input, giving $N_{\rm eff} = 4$ and marginal counting status.

Rows A3 and A4 share the strange-quark pivot m_s but impose constraints on disjoint sets of remaining masses: A3 constrains $\{m_d, m_b\}$ given m_s, and A4 constrains $\{m_c, m_t\}$ given m_s. They are algebraically independent (Sec. IV D).

e. A5: Zero-charge Koide at 100 TeV. Setting the lightest charge to zero and identifying $m_u + m_d$ as the effective second mass gives $K(0, m_u{+}m_d, m_s) = 2/3$ with $\delta_K = -0.265\%$. The predicted mass ratio $(2 +$

TABLE I. Admitted benchmark rows. δ_K is defined in Eq. (9). $N_{\rm eff}$ counts effective inputs per the counting discipline of Sec. II F; $N_{\rm out}$ is the number of independent output masses. Residual class boundaries: practically exact ($< 0.01\%$), strong (0.01–0.1%), conditional (0.1–1%).

Row	Relation	Scheme	δ_K [%]	$N_{\rm eff}$	$N_{\rm out}$	Class	Status
A1	$K(e,\mu,\tau)$	pole	-0.000923	3	3	pract. exact	baseline
A2	$K(e,\mu,\tau)$	$\overline{\rm MS}$ @ 100 TeV	$+0.177$	3	3	conditional	comparison
A3	$K^{-1}(d,s,b)$	$\overline{\rm MS}$ @ 100 TeV	$+0.0120$	3	3	strong	baseline
A4	$K^{-1}(s,c,-t)$	$\overline{\rm MS}$ @ 100 TeV	-0.0037	4	3	pract. exact	marginal
A5	$K(0,\,u{+}d,\,s)$	$\overline{\rm MS}$ @ 100 TeV	-0.265	2	1	conditional	baseline

TABLE II. Rejected benchmark rows. All direct-chain rows (B1–B6) use either mixed-convention masses (PDG-mixed, where each quark mass is quoted at a different scale) or a single common scale at which the chain fails. QCD-scale candidates (C1–C2) are supporting cross-checks without a unique structural target.

Row	Relation	Why rejected
B1	$K(0,s,c)$ mixed	Scale artifact[a]
B2	$K(0,s,c)$ 100 TeV	$\delta_K = -2.70\%$
B3	$K(s,c,b)_{+--}$ mixed	Scale artifact
B4	$K(s,c,b)_{+--}$ 100 TeV	$\delta_K = +4.72\%$
B5	$K(c,b,t)$ mixed	Scale artifact
B6	$K(c,b,t)$ 100 TeV	$\delta_K = +9.16\%$
C1	$P^{1/3} = (m_d m_s m_b)^{1/3}$	Ambiguous target
C2	f_π from charges	Mechanism-dependent

[a] "Scale artifact" means the residual $\delta_K < 1\%$ appears only when masses at different evaluation scales are mixed (e.g., $m_c(m_c)$ with $m_s(2\text{ GeV})$); at any single common scale the relation fails.

$\sqrt{3})^2 \approx 13.93$ [Eq. (10)] is compared with the observed $m_s/(m_u+m_d) = 13.68$ at 100 TeV, a 1.8% discrepancy. The $(u{+}d)$ combination rule and the branch selector count as two inputs for one output ratio, giving baseline counting. The mass ratio is nearly RG-invariant because m_s/m_d is exactly scale-independent and m_u/m_d runs negligibly.

B. Rejected relations

Table II lists relations that fail under the same discipline. The most prominent failure is the direct quark chain $K(0,s,c) \to K(s,c,b) \to K(c,b,t)$, which has been discussed extensively in the literature [2, 4].

a. Direct chain dissolution. The historical direct chain was evaluated using $m_c(m_c)/m_s(2\text{ GeV}) = 13.6$, which accidentally approximates the Koide target 13.93. At any common scale the true ratio is $m_c/m_s = 11.55$ (at 100 TeV), far from the target, and all three chain links fail with $|\delta_K|$ ranging from 2.7% to 9.2%. The entire direct chain is a scale-mixing artifact and is excluded from the statistical audit.

b. QCD-scale candidates. The geometric mean $P^{1/3} = (m_d m_s m_b)^{1/3} \approx 123$ MeV at $Q = 2$ GeV falls between f_π and m_π, with no structural map selecting a unique hadronic target. A separate prediction $f_\pi^{\rm pred} = 2(z_e z_\mu z_\tau)^{2/3}\,\mu \approx 91.6$ MeV (residual -0.7% vs. $f_\pi = 92.2$ MeV) requires the sBootstrap mechanism narrative and four effective inputs for one output. Both are retained as supporting cross-checks but carry no weight in the statistical audit.

C. Summary of benchmark status

The five admitted rows comprise four distinct relation families (A1/A2 form one family). Among these, three have baseline counting ($N_{\rm eff} = N_{\rm out}$) and one has marginal counting ($N_{\rm eff} = N_{\rm out} + 1$). No row achieves predictive status ($N_{\rm eff} < N_{\rm out}$) by counting alone. The significance argument therefore rests on the Monte Carlo p-values and joint probability presented in Sec. IV.

IV. STATISTICAL AUDIT

We subject the admitted benchmark rows to a preregistered Monte Carlo look-elsewhere analysis. The null hypothesis, search-space enumeration, per-row p-values, and joint probability are presented in turn.

A. Null model and calibration

a. Null model. The null hypothesis generates three independent positive masses as

$$m_i = \exp\!\big(U_i\big)\,, \qquad U_i \sim \text{Uniform}(0,13)\,, \tag{14}$$

where the range $[0,13]$ spans the Standard Model fermion mass hierarchy ($m_t/m_e \approx 3\times10^5 \approx e^{13}$). This log-uniform prior is the maximally ignorant choice: it assigns equal probability per decade of mass.

b. Calibration gate. Before any production run, the Monte Carlo reproduces the calibration target from the analysis of Ref. [2] (Note 99, Round 4): the fraction of 10^7 random six-mass draws in which the best of 160 tests achieves $|\delta_K| < 0.0003\%$ is $p_{\rm calib} = 1/2340 = 0.04274\%$.

TABLE III. Search-space enumeration. k_{scheme} is the scheme multiplier (Sec. IV B). Pre-registered counts use the structurally identified tuple; full-scan counts use all possible tuples. The A4 sign pattern $++-$ is FREE (not forced by the overlap recurrence), so all four sign patterns enter both scopes.

Row	$\lvert\delta_K\rvert$ threshold	$N_{\text{tests}}^{\text{pre}}$	$N_{\text{tests}}^{\text{full}}$	k_{scheme}
A1	0.0009%	1	8	1
A3	0.012%	2	16	2
A4	0.0037%	8	160	2
A5	0.265%	2	16	2
Total	—	13	200	—

Our script reproduces this value exactly (0.0% deviation), validating the MC infrastructure. Details of the MC implementation are given in Appendix B.

B. Search-space enumeration

For each admitted row, the trial count is

$$N_{\text{tests}} = \binom{\text{tuples}}{3} \times N_{\text{signs}} \times N_{\text{branches}} \times k_{\text{scheme}}\,. \quad (15)$$

We report two scopes: *pre-registered* (the specific tuple and sign pattern identified from the algebraic structure before seeing the MC result) and *full-scan* (all possible tuples of the appropriate size). Table III lists both.

a. Bookkeeping resolutions. Three algebraic questions were resolved before the MC runs:

1. **A4 sign (T1.1: FREE).** The overlap recurrence eigenvector for eigenvalue (-1) is $(1,-1,1)$ in the direct charge space $z_i = s_i\sqrt{m_i}$. This cannot be mapped to a unique sign assignment in the inverse charge space $d_i = s_i/\sqrt{m_i}$ because the map $z \to 1/z$ makes the recurrence nonlinear. The sign $s_t = -1$ in A4 is therefore a free choice, and all four sign patterns enter N_{tests}.

2. **A3/A4 independence (T1.2: INDEPENDENT).** The inverse Koide constraint $K^{-1}(d,s,b) = 2/3$ fixes $\{m_d, m_b\}$ given m_s. The constraint $K^{-1}(s,c,-t) = 2/3$ fixes $\{m_c, m_t\}$ given m_s. These two constraints act on disjoint free variables and are algebraically independent despite sharing the strange-quark pivot.

3. **Zero-charge alternatives (T1.3: $N_{\textbf{alt}} = 4$).** Among 28 zero-charge combinations of quark masses, four achieve $|\delta_K| < 1\%$; all four share $X = m_u + m_d$ as the lightest entry. The A5 combination $(m_u{+}m_d,\, m_s)$ has the smallest residual (-0.265%).

b. Scheme multiplier. For A1, the pole-mass scheme is structurally motivated (the lepton Koide is a pole-mass statement), so $k_{\text{scheme}} = 1$. For the quark rows A3, A4, A5, quark pole masses were implicitly tested and found to fail (light-quark pole masses are ill-defined in the confinement region); the $\overline{\text{MS}}$@100 TeV scheme was chosen after ruling out pole masses, giving $k_{\text{scheme}} = 2$.

C. Per-row p-values

Each per-row p-value is the fraction of 10^6 null-model mass draws (per seed, three seeds $\{42, 137, 271\}$) for which the best of N_{tests} tests achieves $|\delta_K|$ at or below the observed threshold. Table IV reports both pre-registered and full-scan p-values.

a. Physical interpretation. The A1 full-scan p-value of 6.67×10^{-5} means that random mass hierarchies drawn from the log-uniform null model satisfy $|\delta_K(\text{best of 8 tests})| < 0.0009\%$ in roughly 1 out of 15 000 draws. For A3 (full-scan, 16 tests, threshold 0.012%), the rate is $\sim$1 in 600. Neither is individually decisive by the standard 5σ threshold of particle physics; the decisive argument is their joint probability.

D. Joint probability

The central result of this paper is the joint probability that independently drawn lepton and quark mass spectra simultaneously satisfy the A1 and A3 conditions at or below the observed residuals:

$$\boxed{P(\text{A1} \wedge \text{A3}) \;<\; 1.14 \times 10^{-7} \quad (68\%\ \text{CL})} \quad (16)$$

obtained from zero events in 3×10^7 independent trials across three MC seeds. The 68% confidence-level upper bound for zero observed events in N trials is $P < 1.14/N$ (Appendix B).

The same bound applies to the triple conjunction:

$$P(\text{A1} \wedge \text{A3} \wedge \text{A4}) < 1.14 \times 10^{-7} \quad (68\%\ \text{CL})\,. \quad (17)$$

a. Independence. The lepton sector (A1) and quark sector (A3, A4) draw from completely disjoint mass sets; their independence is exact. Within the quark sector, A3 constrains $\{m_d, m_b\}$ given m_s and A4 constrains $\{m_c, m_t\}$ given m_s, so the two constraints act on disjoint free variables. No conditional-probability correction is needed.

b. Consistency check. From the individual marginal probabilities in the joint MC runs, $P(\text{A1}) \approx 1.6 \times 10^{-5}$ and $P(\text{A3}) \approx 2.05 \times 10^{-4}$, the independence estimate is $P(\text{A1}) \times P(\text{A3}) \approx 3.3 \times 10^{-9}$. In 3×10^7 trials, the expected number of joint events is $\sim$0.1, consistent with the observed zero. The 68%-CL bound of 1.14×10^{-7} is conservative by a factor of $\sim$34 relative to this estimate.

TABLE IV. Verdict table. p_{row} denotes the per-row p-value at the pre-registered (pre) or full-scan (full) scope. p_{joint} is the joint probability across independent sectors (Sec. IV D). Verdict thresholds: GENUINE ($p_{\mathrm{full}} < 10^{-5}$, $p_{\mathrm{joint}} < 10^{-6}$); PROBABLE ($p_{\mathrm{full}} < 10^{-3}$, $p_{\mathrm{joint}} < 10^{-6}$); NOT_PREDICTIVE ($p_{\mathrm{full}} \geq 10^{-3}$ or $N_{\mathrm{eff}} \geq N_{\mathrm{out}}$). Uncertainties are one standard deviation across three seeds.

Row	$p_{\mathrm{row}}^{\mathrm{pre}}$		$p_{\mathrm{row}}^{\mathrm{full}}$		p_{joint}	Verdict
A1	1.83×10^{-5}	$(\pm 3.3 \times 10^{-6})$	6.67×10^{-5}	$(\pm 9.5 \times 10^{-6})$	$< 1.14 \times 10^{-7}$	`PROBABLE`
A3	4.03×10^{-4}	$(\pm 1.5 \times 10^{-5})$	1.67×10^{-3}	$(\pm 2.1 \times 10^{-5})$	$< 1.14 \times 10^{-7}$	`NOT_PRED.` [a]
A4	2.86×10^{-4}	$(\pm 3.2 \times 10^{-5})$	5.38×10^{-3}	$(\pm 1.7 \times 10^{-4})$	$< 1.14 \times 10^{-7}$ [b]	`NOT_PRED.`
A5	8.97×10^{-3}	$(\pm 3.8 \times 10^{-5})$	3.96×10^{-2}	$(\pm 5.8 \times 10^{-5})$	—	`NOT_PRED.`

[a] A3 individually misses the PROBABLE threshold ($p_{\mathrm{full}} < 10^{-3}$) by a factor of 1.67. The decisive significance comes from the joint probability $P(\mathrm{A1} \wedge \mathrm{A3})$.
[b] $P(\mathrm{A1} \wedge \mathrm{A3} \wedge \mathrm{A4}) < 1.14 \times 10^{-7}$ (same bound as the pair).

c. What the bound means. If the Standard Model fermion masses were drawn randomly with a log-uniform hierarchy, the observed joint satisfaction of the Koide and inverse-Koide relations in two independent mass sectors would occur fewer than once in 10^7 random universes. The simultaneous presence of both relations in the Standard Model is non-accidental at the $> 10^7$-to-one level.

E. Baseline comparisons

Any smooth 3-parameter function of 3 masses can fit 3 data points exactly (trivial counting). Therefore, all *matched-parameter* baseline comparisons for rows with $N_{\mathrm{eff}} = N_{\mathrm{out}} = 3$ (A1, A3) or $N_{\mathrm{eff}} \geq N_{\mathrm{out}}$ (A4, A5) are *vacuous*: the Koide formula cannot be said to "beat" a trivial fit at the same parameter count [2].

A one-parameter power-law baseline $m_3 = m_1 (m_2/m_1)^{\alpha}$ achieves a minimum $|\delta_K| \approx 0.0002\%$ for A1 at $\alpha = 1.529$, which is *smaller* than the Koide residual. This does not mean the power law is better: α was tuned post hoc to the known m_τ. The Koide formula encodes the effective exponent $\alpha_{\mathrm{eff}} = 1.529$ as a *zero-parameter* constraint ($K = 2/3$), without fitting. A generic power law with an uncalibrated α produces residuals of 5–50%.

The significance of the Koide relations rests on the Monte Carlo p-values and joint probability, not on baseline comparisons.

F. What was rejected

Six direct-chain rows (B1–B6, Table II) and two QCD-scale candidates (C1–C2) fail the same audit. The direct chain dissolves at any common evaluation scale ($|\delta_K| = 2.7\%$–9.2%). QCD-scale candidates lack a unique structural target. Neither class contributes to the significance argument.

V. DISCUSSION

A. What this result is and is not

The joint probability $p_{\mathrm{joint}}(\mathrm{A1} \cap \mathrm{A3}) < 1.14 \times 10^{-7}$ establishes that the Standard Model mass spectrum encodes non-accidental algebraic structure at the 10^{-7} level across two disjoint fermion sectors. The statistical argument rests on the *joint* test, not on any single row: the lepton Koide relation (A1) and the inverse down-type quark relation (A3) impose independent constraints on independent mass sets, and their simultaneous satisfaction at the observed precision is inconsistent with random mass hierarchies.

However, this result is *not* a prediction in the beyond-Standard-Model sense. The Koide formula does not predict any mass from first principles; it relates known masses in a way that is statistically improbable under a null model of random hierarchies. No mechanism is derived: we do not identify a Lagrangian, a symmetry principle, or a radiative correction that produces $K = 2/3$ as a dynamical output. The significance rests entirely on the joint probability, not on the comparison with any specific baseline model.

Among the individually admitted rows, only A1 achieves the PROBABLE verdict ($p_{\mathrm{row}} = 6.67 \times 10^{-5}$ under the conservative full-scan interpretation). The inverse down-type row A3 misses the PROBABLE threshold ($p_{\mathrm{row}} = 1.67 \times 10^{-3}$, a factor of 1.67 above 10^{-3}) and is classified NOT_PREDICTIVE under our tier rules. This classification reflects the individual full-scan p-value only; it does not mean that A3 is unimportant. On the contrary, A3 is the co-anchor of the decisive joint argument. The cross-sector row A4 ($p_{\mathrm{row}} = 5.38 \times 10^{-3}$) and the zero-charge row A5 ($p_{\mathrm{row}} = 3.96 \times 10^{-2}$) do not individually contribute decisive evidence, but A4 enters the triple joint $P(\mathrm{A1} \cap \mathrm{A3} \cap \mathrm{A4}) < 1.14 \times 10^{-7}$ through its pre-registered p-value.

B. The bion failure and the mixed-convention trap

The integer-coefficient relations

$$z_b = 3\,z_s + z_c\,, \qquad z_t = 2\,z_b + 8\,z_c\,, \tag{18}$$

where $z_i = \sqrt{m_i}$, have been highlighted in the literature as evidence of deeper structure [2]. The leading coefficient 3 in the first relation is suggestively equal to N_c, the number of QCD colors, and the integer structure of both relations invites a group-theoretic interpretation. At the mixed PDG convention—where m_s is evaluated at 2 GeV, m_c at m_c, and m_b at m_b—the first relation holds at 0.037%.

However, under the same strict scheme discipline applied throughout this paper, the bion relations collapse:

Scale	$z_b = 3z_s + z_c$ deviation
PDG mixed convention	0.037%
$\overline{\mathrm{MS}}$ at 2 GeV	14.4%
$\overline{\mathrm{MS}}$ at M_Z	12.2%
$\overline{\mathrm{MS}}$ at 100 TeV	10.2%

There is no zero crossing: no single consistent $\overline{\mathrm{MS}}$ scale rescues the relation. The second relation $z_t = 2z_b + 8z_c$ deviates by 27.9% at 100 TeV.

An exhaustive scan of all two-term integer-coefficient linear charge relations $z_{\mathrm{target}} = a\,z_i + b\,z_j$ among six quarks (target choices × source pairs × coefficient range $|a|, |b| \leq 10$; total: 26,460 trials) finds six sub-0.1% hits at the mixed PDG convention. This is *below* the ~ 20 expected by chance. Three of the six are algebraically equivalent rearrangements of the two bion relations; the remaining three are independent but equally unremarkable. At the 100 TeV common scale, a completely different set of integer relations achieves sub-0.1% precision—none are bion relations—confirming that the specific hits shift with convention.

The deeper methodological point is this: mixed PDG conventions, in which each quark mass is evaluated at its own "natural" threshold, create spurious numerical coincidences because the threshold-scale mass values are not independent—they are related by renormalization-group evolution. A relation that holds "at PDG values" but fails at every common scale is not a statement about the mass spectrum; it is a statement about the accidental similarity of threshold-to-pole ratios. This is the *same* pathology that dissolved the direct quark chain (rows B1–B6 in Table I): at common scale, all three direct chain links fail by 2.7–9.2%.

The bion relations are not derivable from the Koide overlap recurrence (13), which acts on the direct charges $z_i = s_i\sqrt{m_i}$; the inverse chain uses reciprocal charges $d_i = s_i/\sqrt{m_i}$, and the map $z \to 1/z$ makes the recurrence nonlinear. The bion coefficient 3 has no demonstrated connection to N_c beyond numerical coincidence; no derivation from the gauge group exists. We conclude that the bion relations are mixed-convention artifacts and do not survive the predictivity audit.

C. Scheme dependence as physical content

The surviving relations exhibit a physically interesting pattern of scheme dependence. The charged-lepton Koide (A1) is a pole-mass statement: at $\overline{\mathrm{MS}}$@100 TeV, the comparison row A2 shows the residual degrades to $\delta_K = +0.177\%$, dominated by the τ mass correction. This is expected: lepton pole masses receive only QED corrections, which preserve the Koide ratio at high precision.

In contrast, the inverse quark relations (A3, A4) *improve* toward the ultraviolet. The inverse down-type ratio $K^{-1}(d,s,b)$ crosses $K^{-1} = 2/3$ exactly at $Q \approx 280$ TeV, while the cross-sector ratio $K^{-1}(s,c,-t)$ crosses at $Q \approx 82$ TeV [2]. These zero crossings are reproducible features of the five-loop Standard Model RG flow, not fine-tuned artifacts: they are consequences of the asymptotic-freedom-driven convergence of quark mass ratios toward the UV.

This pattern—a low-energy statement for leptons, a high-energy statement for quarks—is a physical observation that any candidate mechanism must explain. A family symmetry realized at a high scale $\Lambda \gtrsim 100$ TeV would naturally produce $K^{-1} = 2/3$ for quarks near the symmetry-breaking scale, with the relation preserved approximately under RG evolution to lower energies. The lepton relation, receiving only electromagnetic corrections, would be preserved to much higher precision at any scale. The scheme dependence is thus an interesting constraint on model building, not a flaw of the analysis.

D. The A4 sign freedom

The cross-sector relation $K^{-1}(s,c,-t) = 2/3$ requires the top quark to enter with a minus sign. We showed in Sec. II that the overlap recurrence (13) has eigenvalue -1 with eigenvector $(1,-1,1)$ in the direct charge space. This eigenvector acts on the *direct* charges $z_i = s_i\sqrt{m_i}$; the inverse chain uses reciprocal charges, and the map $z \to 1/z$ renders the recurrence nonlinear (Appendix B). The eigenvector analysis therefore does not force the sign pattern $(++-)$ on the inverse triple.

Accordingly, the sign in A4 is classified as *free* and costs one effective input. The pre-registered search space for A4 includes all four sign patterns for the single (s,c,t) triple ($N_{\mathrm{tests}} = 8$ including scheme factor); the full-scan search space extends to all $\binom{6}{3} = 20$ quark triples with four sign patterns ($N_{\mathrm{tests}} = 160$). Even under the most conservative full-scan interpretation ($p_{\mathrm{row}} = 5.38\times10^{-3}$), A4 contributes to the decisive triple joint $P(\mathrm{A1} \cap \mathrm{A3} \cap \mathrm{A4}) < 1.14 \times 10^{-7}$.

E. The A5 modeling ambiguity

The zero-charge combination $K(0, m_u+m_d, m_s)$ achieves $\delta_K = -0.265\%$ at $\overline{\mathrm{MS}}$@100 TeV, individually not significant ($p_{\mathrm{row}} = 3.96 \times 10^{-2}$). Among 28 zero-charge combinations evaluated (Appendix B), four achieve $|\delta_K| < 1\%$; all four share the same $X = m_u + m_d$ assignment. Without a structural argument that uniquely selects the ($u+d$, s) partition over the three alternatives ($u+d$, $u+s$), ($u+d$, $d+s$), and ($u+d$, $u+d+s$), A5 remains supporting context rather than independent evidence. It was not included in the decisive joint test.

F. Relation to prior work

The original Koide observation [1] applies to charged leptons only. Foot [3] provided the geometric interpretation on the unit sphere in charge space; our Phase 1 algebraic framework recovers this as the anchor $u = 2 \pm \sqrt{3}$ at the zero-charge boundary. Brannen [4] parametrized the Koide ratio in terms of a single angle and explored quark extensions, but without systematic look-elsewhere corrections. Sumino [7] proposed a radiative protection mechanism for $K = 2/3$ based on a family gauge symmetry; while interesting, such a mechanism is not required by our statistical result and is not tested here. Rivero [2] presented the inverse Koide relation and its RG running, together with the Seiberg seesaw interpretation; our work adds the strict Monte Carlo audit that was absent from that analysis.

The direct quark chain—$K(0, s, c)$, $K(s, c, b)$, $K(c, b, t)$—was first noted in the context of PDG-convention mass values and has been cited as evidence for a mass-chain structure [2]. Our benchmark table (Table I) shows that all three links dissolve at any common $\overline{\mathrm{MS}}$ scale: $|\delta_K|$ degrades from 0.3–1.2% (mixed convention) to 2.7–9.2% (common scale). The direct chain is a mixed-convention artifact, not a genuine mass relation.

G. Addressing prior criticisms

Prior submissions of related work received detailed referee feedback demanding improvements in several areas. We summarize how this paper addresses each:

1. *Parameter counting.* Every benchmark row in Table I carries an explicit effective input count N_{eff} and output count N_{out}, following the protocol of Sec. II F. No relation is claimed as predictive unless $N_{\mathrm{eff}} < N_{\mathrm{out}}$.

2. *Look-elsewhere corrections.* The search space is pre-registered before Monte Carlo evaluation (Table III). Every admitted row reports both pre-registered and full-scan p-values with N_{tests} made explicit.

3. *Separation of structure from numerology.* The verdict table (Table IV) classifies each row by a five-tier logic. Rows that do not meet individual significance thresholds are classified NOT_PREDICTIVE, regardless of their residual precision. Only A1 achieves the PROBABLE tier.

4. *Falsifiability.* If future precision measurements of quark masses at high scales shift any admitted row's δ_K outside the Monte Carlo threshold, the joint probability bound weakens and the significance argument degrades. The framework is falsifiable by improved mass measurements.

5. *CKM and mechanism claims.* No CKM predictions, no Seiberg seesaw mechanism, and no individual quark mass predictions appear in this paper. The existing draft's claims regarding $|V_{us}|$, m_b from bion relations, and m_t from chain completion have been removed entirely. The paper's scope is restricted to the mass-sector predictivity audit.

6. *Baseline comparisons.* Three-parameter baselines are explicitly acknowledged as vacuous (any smooth three-parameter function fits three mass values exactly). The significance of each row derives from the Monte Carlo p-values and the joint probability, not from comparison with trivial fits.

7. *Missing references.* The bibliography includes Foot [3], Brannen [4], Sumino [7], and all other works cited by prior referees.

VI. CONCLUSIONS

We have subjected the full family of Koide-type mass relations to a strict predictivity audit with pre-registered search spaces, explicit input/output counting, and Monte Carlo look-elsewhere corrections. Five fermion-mass relations were benchmarked on pole masses (leptons) and five-loop $\overline{\mathrm{MS}}$ running masses at 100 TeV (quarks). The decisive result is the joint probability: no event in 3×10^7 independent lepton-plus-quark mass draws simultaneously satisfies both the charged-lepton Koide relation and the inverse down-type quark relation at the observed precision, yielding $P(\mathrm{A1} \cap \mathrm{A3}) < 1.14 \times 10^{-7}$ at 68% confidence level.

Two relations survive the audit. The charged-lepton Koide (A1, $\delta_K = -0.000923\%$ at pole masses) achieves the PROBABLE verdict with an individual look-elsewhere p-value of 6.67×10^{-5}. The inverse down-type quark relation (A3, $\delta_K = +0.0120\%$ at $\overline{\mathrm{MS}}$@100 TeV) is individually classified NOT_PREDICTIVE under the conservative full-scan threshold but serves as the co-anchor of the decisive joint test. Together, they establish that the Standard Model mass spectrum encodes non-accidental algebraic structure across disjoint fermion sectors.

The same audit rejects a larger body of previously highlighted coincidences. The direct quark chain dissolves at every common renormalization scale: all three links fail at $|\delta_K| = 2.7$–9.2%. The integer-coefficient bion relations fail by 10–14% at every consistent $\overline{\mathrm{MS}}$ scale, and an exhaustive scan of 26,460 integer linear charge relations confirms that sub-0.1% hits at the mixed PDG convention are below the chance expectation. Mixed-convention artifacts and search-space freedom account for most previously claimed successes. The honest application of scheme discipline is the single most effective filter against numerological coincidences.

The result is statistical, not mechanistic. Whether a UV-complete theory produces $K = 2/3$ as an algebraic constraint remains an open question. The scheme dependence of the surviving relations—the lepton Koide is a pole-mass statement, the quark inverse Koide is a high-energy statement—provides a non-trivial constraint on candidate mechanisms: any family symmetry model must explain why $K = 2/3$ is preserved by electromagnetic corrections (leptons) while $K^{-1} = 2/3$ is achieved only near $Q \sim 100$–280 TeV (quarks). Possible directions include family gauge symmetry models with radiative protection [7], extensions to the neutrino sector (subject to the same counting discipline applied here), and experimental sensitivity: future precision determinations of light-quark $\overline{\mathrm{MS}}$ masses could tighten or invalidate the joint probability bound.

ACKNOWLEDGMENTS

The numerical benchmarks use five-loop Standard Model running masses from the `SMDR` package [5].

Appendix A: Mass Data

All quark-sector benchmarks use $\overline{\mathrm{MS}}$ running masses at $Q = 100$ TeV from the `SMDR` package [5], computed with five-loop QCD and three-loop electroweak corrections in the full (non-decoupled) Standard Model. Input parameters are the PDG 2024 values [6] at the input scale $Q_{\mathrm{in}} = 172.4$ GeV. Charged-lepton pole masses are from PDG 2024.

Table V lists the mass values used in all benchmark computations.

a. Koide ratio values. Table VI reproduces the computed K or K^{-1} values and residuals for the five admitted rows, using the exact SMDR output (not the rounded masses in Table V). The rounded masses reproduce these values to better than 0.001%.

b. RG evolution cross-check. The inverse down-type ratio $K^{-1}(d, s, b)$ was verified at multiple scales against independent SMDR runs: $\delta_K = +0.25\%$ at $Q = 2$ GeV, $+0.13\%$ at M_Z, $+0.012\%$ at 100 TeV. The zero crossing at $Q \approx 280$ TeV is consistent with the extrapolation reported in Ref. [2]. The lepton ratio $K(e, \mu, \tau)$ varies by only 0.017 percentage points across five decades ($Q = 2$ GeV to 100 TeV), confirming it is effectively RG-invariant.

TABLE V. Fermion masses used in this analysis. Pole masses are from PDG 2024 [6]. $\overline{\mathrm{MS}}$ masses at $Q = 100$ TeV are from five-loop SMDR [5].

Fermion	Pole mass	$\overline{\mathrm{MS}}$ @ 100 TeV	Units
e	0.51100	0.4792	MeV
μ	105.658	100.96	MeV
τ	1776.86	1715.3	MeV
u	—	0.824	MeV
d	—	1.811	MeV
s	—	36.03	MeV
c	—	416.0	MeV
b	—	1829.7	MeV
t	172 570	118 847	MeV

Light-quark (u, d, s) pole masses are omitted because they are ill-defined in the confinement region ($\mathcal{O}(\Lambda_{\mathrm{QCD}})$ non-perturbative corrections). The top-quark pole mass $m_t = 172.57(29)$ GeV is listed for reference but is not used in any admitted benchmark row.

TABLE VI. Computed Koide ratio values and residuals for admitted rows, using exact SMDR 5-loop masses. $\delta_K[\%]$ is defined in Eq. (9).

Row	K or K^{-1}	δ_K [%]
A1	0.6666605115	−0.000923
A2	0.6678490779	+0.1774
A3	0.6667466386	+0.01200
A4	0.6666420348	−0.003694
A5	0.6648997069	−0.2650

Appendix B: Monte Carlo Details

1. Implementation

The look-elsewhere Monte Carlo generates random mass spectra under the null model of Eq. (14) and computes the Koide residual δ_K for each admitted benchmark row. The implementation uses vectorized `NumPy` array operations on 10^6 trials per call, avoiding Python loops over individual trials.

For each trial, three masses are drawn as $m_i = \exp(U_i)$ with $U_i \sim \mathrm{Uniform}(0, 13)$. The Koide ratio K (or K^{-1} for inverse rows) is computed from these masses. A trial "passes" for a given row if $|\delta_K| \leq |\delta_{K,\mathrm{obs}}|$, where $\delta_{K,\mathrm{obs}}$ is the observed residual for that row (Table I).

The per-row p-value at pre-registered scope is

$$p_{\mathrm{row}}^{\mathrm{pre}} = N_{\mathrm{tests}}^{\mathrm{pre}} \times \frac{(\mathrm{passes})}{N_{\mathrm{trials}}}, \tag{B1}$$

and identically for the full-scan scope with $N_{\mathrm{tests}}^{\mathrm{full}}$. This is the Bonferroni-corrected look-elsewhere probability.

2. Seeds and reproducibility

Three independent random seeds $\{42, 137, 271\}$ are used. Each per-row run draws $N = 10^6$ trials per seed (3×10^6 total). The joint-probability run draws $N = 10^7$ trials per seed (3×10^7 total). Reported p-values are the mean across three seeds; uncertainties are one standard deviation.

3. Joint probability computation

Each joint trial draws nine masses independently: three lepton-sector masses and six quark-sector masses. The trial passes the joint condition A1 $\wedge$ A3 if and only if:

1. The three lepton masses satisfy $|\delta_K(m_1, m_2, m_3)| < 0.000923\%$ for the direct Koide ratio with all-positive signs (one pre-registered test), *and*
2. The three quark masses assigned to (d, s, b) satisfy $|\delta_K^{-1}(m_4, m_5, m_6)| < 0.0120\%$ for the inverse Koide ratio with all-positive signs (one pre-registered test).

The triple condition $\mathrm{A1}\wedge\mathrm{A3}\wedge\mathrm{A4}$ additionally requires that the quark masses assigned to (s, c, t) satisfy $|\delta_K^{-1}| < 0.0037\%$ for some sign pattern (four pre-registered tests, since the sign is FREE). The strange mass m_s is the same draw in both A3 and A4.

4. Upper bound for zero events

For zero observed events ($k = 0$) in N trials, the Poisson 68% confidence-level upper bound on the true rate λ is

$$\lambda_{68} = -\ln(1 - 0.68) = \ln(1/0.32) \approx 1.139\,. \qquad \text{(B2)}$$

The corresponding probability bound is

$$P < \frac{\lambda_{68}}{N} = \frac{1.14}{N}\,. \qquad \text{(B3)}$$

For $N = 3\times10^7$ (combined across three seeds), this gives $P < 1.14/(3\times10^7) \approx 3.8\times10^{-8}$ per seed, or $P < 1.14\times10^{-7}$ as a conservative combined bound.

5. Calibration gate

The T0 calibration reproduces the analysis of Ref. [2] (Note 99, Round 4):

- Draw 10^7 random mass sets of 6 masses each, log-uniform in $[0, 13]$.
- For each draw, evaluate 160 Koide tests ($\binom{6}{3} = 20$ triples $\times$ 4 sign patterns $\times$ 2 branches).
- The fraction of draws with $|\delta_{K,\mathrm{best}}| < 0.0003\%$ is the calibration probability.

Our result: $p_{\mathrm{calib}} = 0.04274\%$ ($= 1/2340$), matching the target exactly. This validates both the Koide ratio computation and the random-number generation.

6. Null-model robustness

The primary null model (log-uniform, 13 decades) is the default for all reported p-values. The decisive joint result (zero events in 3×10^7 trials) is robust against reasonable null-model variations: the product $P(\mathrm{A1})\times P(\mathrm{A3}) \approx 3\times10^{-9}$ lies well below the observable threshold $\sim10^{-7}$ for any prior that preserves the broad mass hierarchy. SM-jittered models (±1 dex around observed masses) were explored in Ref. [2] and do not change the zero-event outcome.

[1] Y. Koide, A New View of Quark and Lepton Mass Hierarchy, Phys. Rev. D **28**, 252 (1983), see also Lett. Nuovo Cim. **34**, 201 (1982).

[2] A. Rivero, An interpretation of scalars in SO(32), Eur. Phys. J. C **84**, 1058 (2024), arXiv:2407.05397 [hep-ph].

[3] R. Foot, A Note on Koide's lepton mass relation, arXiv:hep-ph/9402242 (1994), preprint.

[4] C. A. Brannen, The Lepton Masses (2006), unpublished, available at `https://brannenworks.com/MASSES2.pdf`.

[5] S. P. Martin and D. G. Robertson, Standard model parameters in the tadpole-free pure $\overline{\mathrm{MS}}$ scheme, Phys. Rev. D **100**, 073004 (2019), arXiv:1907.02500 [hep-ph].

[6] S. Navas *et al.* (Particle Data Group), Review of Particle Physics, Phys. Rev. D **110**, 030001 (2024).

[7] Y. Sumino, Family Gauge Symmetry as an Origin of Koide's Mass Formula and Charged Lepton Spectrum, JHEP **05**, 075, arXiv:0812.2103 [hep-ph].