Article

Not peer-reviewed version

Soddy Curvature of the Koide Lepton Triple and the Strange-Quark Mass

Andrew M. Brilliant *

Posted Date: 10 April 2026

doi: 10.20944/preprints202604.0693.v1

Keywords: Koide relation; Descartes circle theorem; Soddy curvature; charged-lepton masses; strangequark mass; renormalization-group running

Article

Soddy Curvature of the Koide Lepton Triple and the Strange-Quark Mass

Andrew M. Brilliant

Applied Dynamics Research, Sapporo, Japan; ab@ad-research.org

Abstract

The Koide relation $Q = (\sum m_\ell)/(\sum\sqrt{m_\ell})^2 = 2/3$ for the charged leptons has held to one part in 10^5 for over forty years without an accepted derivation. Treating lepton mass square roots as Descartes-circle curvatures, the smaller ("outer") root of the Descartes quadratic, $\mathcal{F} \equiv k_4^- = e_1 - 2\sqrt{e_2}$ where $e_1 = \sum_\ell \sqrt{m_\ell}$ and $e_2 = \sum_{\ell<\ell'}\sqrt{m_\ell m_{\ell'}}$, admits the algebraically simple closed form $\mathcal{F} = e_1 - \sqrt{p_2}$ (with $p_2 = \sum_\ell m_\ell$) when Koide holds exactly. Kocik [8] first observed a Descartes-like geometric reading of Koide using a generalized formula for circles in general position; the standard mutually-tangent reading we use here is mathematically valid as an algebraic identity, though distinct from Kocik's original construction. We compute this curvature for the observed lepton triple and report a numerical observation: $\mathcal{F}^2 = 95.113$ MeV is numerically consistent with the strange-quark $\overline{\text{MS}}$ mass at the natural lepton-sum scale $\mu_\star = m_e + m_\mu + m_\tau$ within current lattice precision (residual $+0.04$ MeV against a total uncertainty of ± 0.69 MeV, i.e., about $+0.06\sigma$). Under a Koide-conditioned null model with a log-uniform prior on triples, the observed match arises in 0.47% of valid samples (95% CI [0.39%, 0.57%]); we report this as a model-conditional descriptive frequency under our chosen Monte Carlo construction, not as a p-value or hypothesis test. The hit fraction is robust to four alternative prior choices (range 0.23%–0.47%, fold variation 2.04). Of eight scale prescriptions tested, only $\mu_\star = \sum m_\ell$ and the numerically near-identical $m_\mu + m_\tau$ produce hits; the next-best distinct alternative misses by about $\pm 2.3\sigma$. We present this as a numerical observation, not a derivation or a prediction. We make no claim about an underlying mechanism. We disclose the order of operations leading to the observation and report scale, input, prior, and filter sensitivities along with the error budget decomposition.

Keywords: Koide relation; Descartes circle theorem; Soddy curvature; charged-lepton masses; strange-quark mass; renormalization-group running

---

1. Introduction

The Koide relation [1,2]

$$Q \equiv \frac{m_e + m_\mu + m_\tau}{\left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right)^2} = \frac{2}{3} \tag{1}$$

holds for the observed charged lepton pole masses to one part in 10^5. Using the PDG 2024 values [11] ($m_e = 0.5109989$ MeV, $m_\mu = 105.6584$ MeV, $m_\tau = 1776.93 \pm 0.09$ MeV), one computes $Q = 0.666664$. Although the relation is not part of the standard Standard-Model curriculum, it has been the subject of continuous theoretical investigation since 1981, with derivation attempts, extensions to the neutrino and quark sectors, and geometric reinterpretations appearing across the intervening four decades; we cite the Rivero–Gsponer review [4] as an entry point and [3,5–7,9] as representative further reading. We treat (1) as an empirical input.

In 2012, Kocik [8] observed that the Koide relation admits a geometric reading in terms of intersecting circles. Specifically, Kocik generalized the Descartes circle formula to circles meeting at a common angle, and showed that Koide's $Q = 2/3$ condition fits naturally within that generalized framework once the fourth curvature is set to zero (a degenerate "line" circle); the standard mutually-tangent Descartes case fails to recover Koide because of the 2 versus 2/3 coefficient discrepancy. He does not compute the non-trivial fourth curvature, and he does not connect any geometric quantity to a quark mass. The present work takes a different geometric step: we identify $k_i \equiv \sqrt{m_i}$ for $i \in \{e, \mu, \tau\}$ with three of four curvatures of a *mutually tangent* Descartes-Soddy configuration governed by the *ordinary* Descartes circle theorem. We then show (Proposition 1) that under exact Koide the algebraically simple combination $e_1 - \sqrt{p_2}$ formed from the leptons coincides with the geometrically determined outer Soddy curvature $e_1 - 2\sqrt{e_2}$ of that ordinary-Descartes configuration. The mutually-tangent reading is mathematically valid as an algebraic identity, even though it differs from Kocik's original generalized-angle construction; Proposition 1 is the manuscript's contribution, motivated by but not contained in [8].

Building on Kocik's geometric framing of the Koide relation, we introduce a distinct ordinary-Descartes algebraic variant in which the lepton square roots are treated as mutually tangent curvatures, compute the resulting outer Soddy curvature for the observed lepton triple, denote it $\mathcal{F}$, and report a numerical observation: $\mathcal{F}^2$ is numerically consistent with the strange-quark $\overline{\text{MS}}$ mass at the natural lepton-sum scale $\mu = \sum m_\ell$ within current lattice precision. We make no theoretical claim. The contribution of this paper is the introduction of this ordinary-Descartes variant, the computation of $\mathcal{F}$, and the disclosure of the resulting numerical match, together with a Monte Carlo characterization of the observation's behavior under a family of Koide-conditioned toy null models, including sensitivity to alternative priors, scale prescriptions, and filter cutoffs.

2. The Koide Relation as a Descartes Condition

We summarize the ordinary-Descartes construction used in this paper for completeness and to fix notation. As discussed in Section 1, this construction is distinct from the generalized-angle reading in [8].

The Descartes circle theorem [10] states that four mutually tangent circles in the plane, with curvatures k_1, k_2, k_3, k_4, satisfy

$$(k_1 + k_2 + k_3 + k_4)^2 \;=\; 2\,(k_1^2 + k_2^2 + k_3^2 + k_4^2). \tag{2}$$

Given three curvatures k_1, k_2, k_3, equation (2) is a quadratic in k_4 with two real solutions

$$k_4^{\pm} \;=\; (k_1 + k_2 + k_3) \;\pm\; 2\sqrt{k_1 k_2 + k_2 k_3 + k_1 k_3}, \tag{3}$$

corresponding respectively to the inner and outer Soddy circles tangent to the original three. By convention, the larger-curvature solution k_4^+ corresponds to the smaller circle (the inner Soddy circle, nestled in the gap between the original three), and the smaller-curvature solution k_4^- is conventionally called the "outer" Soddy circle. When k_4^- is negative, it corresponds to a large circle enclosing the original three; when k_4^- is positive, as is the case for the observed lepton triple ($k_4^- \approx 9.75\ \text{MeV}^{1/2}$), it corresponds instead to a fourth circle externally tangent to the three but smaller than k_e^{-1} in radius. The "outer" label refers to the smaller of the two algebraic roots of the Descartes quadratic, not to literal enclosure. The relevant quantity for our purposes is this curvature $k_4^- = e_1 - 2\sqrt{e_2}$.

Proposition 1 (Koide and the outer Soddy curvature). *Identify $k_i \equiv \sqrt{m_i}$ for $i \in \{e, \mu, \tau\}$. Let $e_1 = \sum_i \sqrt{m_i}$, $e_2 = \sum_{i<j} \sqrt{m_i m_j}$, and $p_2 = \sum_i m_i$. The outer Soddy curvature of the Descartes configuration (k_e, k_μ, k_τ) is*

$$k_4^- \;=\; e_1 - 2\sqrt{e_2}. \tag{4}$$

The algebraically simpler expression

$$\mathcal{F} \;\equiv\; e_1 - \sqrt{p_2} \;=\; \left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right) - \sqrt{m_e + m_\mu + m_\tau} \tag{5}$$

 doi:10.20944/preprints202604.0693.v1

coincides with k_4^- *if and only if the Koide relation* $Q \equiv p_2/e_1^2 = 2/3$ *holds exactly.*

Proof. The Soddy form follows from solving the quadratic (2) for k_4 given k_1, k_2, k_3, using the identity $e_1^2 = p_2 + 2e_2$ to simplify the discriminant; the smaller-curvature root is $k_4^- = e_1 - 2\sqrt{e_2}$. Setting $e_1 - 2\sqrt{e_2} = e_1 - \sqrt{p_2}$ yields $4e_2 = p_2$. Combined with $e_1^2 = p_2 + 2e_2$, this gives $e_1^2 = p_2 + p_2/2 = (3/2)p_2$, equivalently $p_2/e_1^2 = 2/3$, which is the Koide condition. Conversely, if $Q = 2/3$ then $e_2 = e_1^2/6 = p_2/4$, so $2\sqrt{e_2} = \sqrt{p_2}$ and the two expressions agree. □

Proposition 1 states the precise algebraic content of the standard-Descartes reading: under exact Koide, the dimensionally simple combination $e_1 - \sqrt{p_2}$, formed without reference to circles, is in fact the smaller Descartes-Soddy root of the configuration whose three input curvatures are the lepton mass square roots. The observed leptons satisfy $Q = 0.666664$, so the equality of (4) and (5) holds to one part in 10^5. In what follows we adopt $\mathcal{F}$ to denote the exact outer Soddy curvature k_4^- for all numerical computations; the algebraic surrogate $e_1 - \sqrt{p_2}$, which would be exactly equal under exact Koide, is reported once in the next section for comparison with k_4^-. Figure 1 shows the geometric configuration schematically.

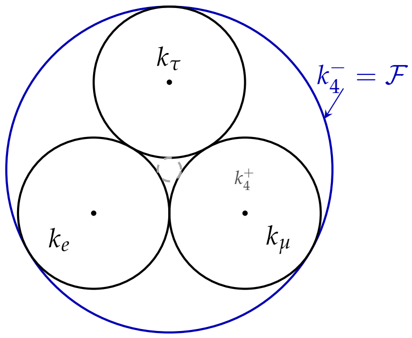

Figure 1. *Sign-convention schematic for the two roots of the Descartes quadratic; the drawing depicts the enclosing-circle orientation that corresponds to negative* k_4^-*, not the observed-lepton case.* The blue k_4^- circle is shown in the enclosing ($k_4^- < 0$) configuration solely to illustrate the conventional "outer" nomenclature; it is not the observed lepton geometry. The three input circles (solid black) carry curvatures $k_i = \sqrt{m_i}$ for $i \in \{e, \mu, \tau\}$. The two Soddy circles tangent to all three are shown: the inner Soddy circle k_4^+ (dashed gray, in the central gap) and the smaller-curvature root k_4^- (blue), conventionally called the "outer" Soddy circle. For the observed lepton triple, $k_4^- \approx 9.75\ \mathrm{MeV}^{1/2}$ is positive, so the actual blue circle would be a small fourth circle externally tangent to the three input circles, not an enclosing one; the "outer" label is then a name for the smaller algebraic root of the Descartes quadratic, not a topological description. When the Koide relation $Q = 2/3$ holds exactly, the (algebraic) outer Soddy curvature k_4^- equals the surrogate $\mathcal{F} = e_1 - \sqrt{p_2}$. *The figure is a topological schematic, not drawn to scale.* The actual lepton curvatures span nearly two orders of magnitude (a factor of about 59, since $\sqrt{m_\tau/m_e} \approx 59$, equivalently $\sqrt{m_e/m_\tau} \approx 0.017$); the electron circle (smallest curvature, $k_e \approx 0.715\ \mathrm{MeV}^{1/2}$) is by far the largest, with the muon ($k_\mu \approx 10.28$) and tau ($k_\tau \approx 42.15\ \mathrm{MeV}^{1/2}$) circles much smaller. The figure illustrates the tangency structure that supports Proposition 1; the algebraic identification of $\mathcal{F}$ with k_4^- does not depend on the relative sizes of the input circles.

Although Kocik [8] gives the geometric identification of Koide as a Descartes condition, to our knowledge the outer Soddy curvature $\mathcal{F}$ has not previously been computed and compared with hadronic observables. We discuss the relation of the present work to existing Koide–quark literature, in particular Rivero's algebraic "waterfall" [7], in Section 6.

3. The Outer Soddy Curvature of the Lepton Triple

We have defined the exact comparison quantity as the outer Soddy curvature $\mathcal{F} \equiv k_4^- = e_1 - 2\sqrt{e_2}$ in Equation (4). Equation (5) gives the Koide-exact algebraic surrogate $\mathcal{F}_{\mathrm{sur}} \equiv e_1 - \sqrt{p_2}$, which we quote once for comparison. Two further remarks before computing $\mathcal{F}$.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 April 2026 doi:10.20944/preprints202604.0693.v1

The function $\mathcal{F}$ in Kocik's framework.

$\mathcal{F}$ is not chosen from a catalog of symmetric polynomials. Once the Koide relation is read as a Descartes condition (Proposition 1), the Descartes-determined quantities of dimension $[\text{mass}]^{1/2}$ are the two Soddy curvatures $k_4^{\pm} = e_1 \pm 2\sqrt{e_2}$. The choice between them is binary (inner versus outer), not a continuous selection. We adopt the outer (smaller-curvature) root k_4^-. The integer power n in any subsequent comparison is fixed at $n = 2$ by dimensional consistency: $\mathcal{F}$ has units of $\text{MeV}^{1/2}$, so only $\mathcal{F}^2$ is dimensionally commensurate with a quark mass, and $\mathcal{F}^n$ for $n \neq 2$ is not unit-invariant. The remaining selection freedom — the choice of comparison quark and the choice of evaluation scale — contributes to the look-elsewhere effect that we quantify explicitly in Section 5. The structural narrowing provided by the Descartes framework is real but partial, and we do not claim that $\mathcal{F}$ is the unique observable one could extract from the construction.

Numerical value.

At the PDG 2024 central charged-lepton pole masses,

$$\mathcal{F} \;=\; 9.7526\ \text{MeV}^{1/2}, \qquad \mathcal{F}^2 \;=\; 95.1134\ \text{MeV}. \tag{6}$$

The propagated uncertainty from PDG lepton inputs is $\sigma_{\mathcal{F}^2} = 0.010$ MeV (Section 5), dominated by the tau mass uncertainty. The value reported in (6) is computed as the exact outer Soddy curvature $k_4^- = e_1 - 2\sqrt{e_2} = 9.752607\ \text{MeV}^{1/2}$. The algebraic surrogate (5), $e_1 - \sqrt{p_2} = 9.752823\ \text{MeV}^{1/2}$, agrees with k_4^- to one part in 10^5, the precision at which Koide holds for the observed leptons; the corresponding difference in the squared quantity is 0.004 MeV. We emphasize that this is an approximation error (a systematic consequence of using $e_1 - \sqrt{p_2}$ as a stand-in for the exact Soddy curvature when the observed leptons satisfy Koide only approximately), not a measurement uncertainty; the two should be tracked separately. Numerically, 0.004 MeV is smaller than the lepton-side measurement uncertainty of 0.010 MeV and very small compared to the lattice-side uncertainty of ≈ 0.69 MeV that dominates the comparison, but its character is different. In what follows, $\mathcal{F}$ denotes the exact outer Soddy curvature k_4^- of (4) throughout, and the algebraic surrogate (5) is reported once for comparison, here. Unless explicitly stated otherwise, all numerical results in subsequent sections use (4), not (5).

4. The Observation

Evaluate the strange-quark $\overline{\text{MS}}$ mass at the natural lepton-sum scale,

$$\mu_\star \;\equiv\; m_e + m_\mu + m_\tau \;=\; 1883.1\ \text{MeV}, \tag{7}$$

by running the FLAG 2024 $N_f = 2{+}1{+}1$ estimate [12] $m_s^{\overline{\text{MS}}}(2\ \text{GeV}) = 93.44 \pm 0.68$ MeV (Table 11 of Ref. [12]) via four-loop QCD [13] with $\alpha_s(M_Z) = 0.1180$. We treat the quoted $m_s(2\ \text{GeV})$ as an $n_f = 4$ $\overline{\text{MS}}$ mass and evolve it to $\mu_\star$ at fixed $n_f = 4$ using four-loop renormalization-group running, since both 2 GeV and $\mu_\star \approx 1.88$ GeV lie above m_c. Because both endpoints of the evolution lie above m_c, the evolution from 2 GeV to $\mu_\star$ uses fixed $n_f = 4$ and is independent of the charm threshold at the quoted precision. Threshold matching enters only in obtaining low-scale α_s from $\alpha_s(M_Z)$, where the b threshold is crossed at m_b (decoupling $n_f = 5 \to 4$). This gives

$$m_s^{\overline{\text{MS}}}(\mu_\star) \;=\; 95.07 \pm 0.69\ \text{MeV}, \tag{8}$$

where the uncertainty includes the lattice input, the propagation of α_s uncertainty, and four-loop truncation, with the lattice contribution dominating; we decompose this in Section 5. The two sides of the comparison are in different renormalization schemes: $\mathcal{F}$ is constructed from charged-lepton pole masses (the scheme in which the Koide relation holds to one part in 10^5), while $m_s^{\overline{\text{MS}}}(\mu_\star)$ is the FLAG lattice $\overline{\text{MS}}$ average run via four-loop QCD. We make no claim of scheme-equivalence between the two

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 April 2026 doi:10.20944/preprints202604.0693.v1

sides; the comparison is reported as a numerical relation between objects defined in different schemes, not as a derived equality. Comparing,

$$\mathcal{F}^2 - m_s^{\overline{\text{MS}}}(\mu_\star) \;=\; +0.042\ \text{MeV}, \tag{9}$$

or about $+0.061$ standard deviations of the lattice uncertainty. For context, the PDG 2024 quark-mass review [11] quotes $m_s^{\overline{\text{MS}}}(2\ \text{GeV}) = 92.93 \pm 0.63$ MeV from the latest $N_f = 2{+}1{+}1$ calculations (combining stat and systematic errors in quadrature); its final lattice-QCD estimate in the $N_L = 4$ theory, obtained by combining both the $N_f = 2{+}1$ and $N_f = 2{+}1{+}1$ averages, is $m_s^{\overline{\text{MS}}}(2\ \text{GeV}) = 92.74 \pm 0.54$ MeV [Equation (60.5) of Ref. [11]]; and the PDG 2024 quark summary table [11] lists $m_s = 93.5 \pm 0.8$ MeV at confidence level 90%. These four numbers—the FLAG Table 11 estimate used above, the PDG review $N_f = 2{+}1{+}1$ average, the PDG review final $N_L = 4$ estimate, and the PDG summary-table entry at 90% CL—are not interchangeable and we do not average across them. Using the PDG review $N_f = 2{+}1{+}1$ central value (92.93 MeV) in place of FLAG shifts the comparison residual by approximately $+0.5$ MeV (residual $\approx +0.9\sigma$); using the PDG review final $N_L = 4$ estimate (92.74 MeV) shifts it by approximately $+0.7$ MeV (residual $\approx +1.1\sigma$). Both remain well within 1σ of the lattice uncertainty; the numerical match is robust to this choice of low-scale input at the level of the conclusions we draw.

We deliberately avoid framing (9) as a precision-significance claim. The residual is small in units of the lattice uncertainty, but the meaningful question is not how many standard deviations separate $\mathcal{F}^2$ from $m_s^{\overline{\text{MS}}}(\mu_\star)$; it is whether the structural setup that produced $\mathcal{F}$ would, under a reasonable null, generically produce numbers near m_s at the natural scale. We address that question in Section 5.

Methodological note on order of operations.

We disclose the order of operations leading to this observation, because it bears on the look-elsewhere assessment in Section 5. The author was investigating Kocik's geometric reading [8] of the Koide relation and computed the fourth (outer Soddy) curvature of the Koide-saturating lepton triple, as an ordinary-Descartes quantity, as a geometric question, with the original motivation being whether the missing fourth curvature might admit a particle interpretation. The numerical agreement $\mathcal{F}^2 \approx m_s^{\overline{\text{MS}}}(\mu_\star)$ was noticed only after the geometric computation was complete, on comparison with a previously consulted reference value for the strange-quark mass. The function $\mathcal{F}$ was therefore not selected from a class of candidate symmetric polynomials of lepton masses; once one takes the ordinary-Descartes reading, it is determined up to the binary inner/outer choice. The choice of evaluation scale, however, deserves separate scrutiny: $\mu_\star = \sum m_\ell$ is the characteristic mass scale of the lepton triple in the same dimensional sense that $\mathcal{F}$ is the Descartes curvature, but it is not the only candidate. We address that selection freedom empirically in Section 5.3.

5. Robustness

We report robustness checks across six axes: (i) propagation of input uncertainties and decomposition of the error budget, (ii) continuous scale sensitivity around $\mu_\star$, (iii) discrete alternative scale prescriptions, (iv) a tight null-model Monte Carlo conditioned on the specific numerical comparison, (v) sensitivity of the null result to the prior, (vi) a loose null-model test admitting expanded selection freedom, together with the sensitivity of both nulls to the lower-cutoff filter on $\mu_\star$. Central values and lattice-input propagation are reproducible from the version-pinned GitHub repository at commit `bb8b1ae` [14]. The $\alpha_s(M_Z)$ uncertainty propagation and the four-loop-versus-three-loop truncation comparison reported in the error budget below are performed by the same running algorithm; their outputs, along with those for the scale-prescription alternatives (`results/scale_alternatives.json`) and the filter-sensitivity scans (`results/filter_sensitivity.json`), are archived as precomputed artifacts at the cited commit [14].

All robustness scans in this section use $\mathcal{F}^2 \equiv (k_4^-)^2$ as the comparison quantity, where k_4^- is the exact outer Soddy curvature defined in (4). The algebraic surrogate (5) was reported once in the previous section to quantify the Koide-approximation offset; it is not used in any robustness scan.

5.1. Input Sensitivity and Error Budget

Propagating PDG 2024 lepton uncertainties through the exact curvature (4) yields $\sigma_{\mathcal{F}^2}^{\text{lepton}} = 0.010$ MeV, dominated by $\sigma_{m_\tau} = 0.09$ MeV (the surrogate (5) gives the same value to the precision quoted). This is to be combined with the running uncertainty on $m_s^{\overline{\text{MS}}}(\mu_\star)$. We decompose the latter by varying each input independently and re-running:

Source	σ contribution (MeV)
FLAG lattice $m_s(2\text{ GeV}) = 93.44 \pm 0.68$	0.692
$\alpha_s(M_Z) = 0.1180 \pm 0.0009$	0.051
Charm threshold $m_c(m_c) = 1.273 \pm 0.005$ GeV (not used for 2 GeV $\to \mu_\star$ evolution)	0.000
Four-loop truncation (4L vs 3L)	0.039
Quadrature sum	0.695

The lattice input dominates by an order of magnitude over all other sources. Decomposing the total variance: the FLAG $m_s(2\text{ GeV})$ input accounts for approximately 99.1%, α_s uncertainty for 0.54%, four-loop truncation for 0.31%, and the lepton-side propagation for 0.02%. The lepton side of (9) is essentially exact at present precision.

One systematic absent from the above decomposition is the QED scheme conversion on the lepton side. $\mathcal{F}$ is constructed from charged-lepton pole masses, in which the Koide relation holds to one part in 10^5, while $m_s^{\overline{\text{MS}}}(\mu_\star)$ is a lattice $\overline{\text{MS}}$ quantity. Converting the lepton pole masses to $\overline{\text{MS}}$ at leading QED order ($m_\ell \to m_\ell(1-\alpha/\pi)$ uniformly) scales $\mathcal{F}$ linearly and shifts $\mathcal{F}^2$ by approximately $-\mathcal{F}^2 \cdot \alpha/\pi \approx -0.22$ MeV, moving the comparison residual from $+0.042$ MeV to approximately -0.18 MeV. Under α prescriptions ranging from $\alpha(m_e)$ to $\alpha(m_\tau)$ the shift ranges from -0.221 to -0.226 MeV; the scheme-insensitive band for the residual is approximately $[-0.18, +0.04]$ MeV, entirely within 0.3σ of the dominant lattice uncertainty. The scheme in which Koide holds most precisely is the pole-mass scheme as conventionally defined; we report the conversion as a directional scheme-dependence estimate rather than a symmetric uncertainty, and note that it does not alter the conclusions drawn from (9).

5.2. Continuous Scale Sensitivity

We compute $m_s^{\overline{\text{MS}}}(\mu)$ for $\mu/\mu_\star \in [0.5, 4.0]$ and compare against the constant $\mathcal{F}^2$. The agreement degrades monotonically away from $\mu_\star$, with residual reaching -36σ at $\mu = 0.5\,\mu_\star$. The intervals over which $|\mathcal{F}^2 - m_s^{\overline{\text{MS}}}(\mu)|$ remains within one and two standard deviations of $\sigma_{m_s}(\mu)$ are

$$1\sigma \text{ window:} \quad \mu \in [1845,\ 1921]\text{ MeV} \quad \approx \quad \mu_\star(1 \pm 2.0\%), \tag{10}$$

$$2\sigma \text{ window:} \quad \mu \in [1808,\ 1996]\text{ MeV} \quad \approx \quad \mu_\star(1 \pm 5.0\%). \tag{11}$$

The 1σ window is narrow (about $\pm 2.0\%$), confirming that the relation is picking out a specific scale rather than holding broadly across hadronic energies. For concreteness: running $m_s^{\overline{\text{MS}}}$ from the FLAG reference scale of 2 GeV to $\mu_\star \approx 1883$ MeV shifts the mass by $+1.63$ MeV ($+1.7\%$), approximately 39 times larger than the comparison residual; the scale prescription is therefore not cosmetic. Once the prescription $\mu_\star = \sum m_\ell$ is fixed, the agreement is not the result of continuously tuning μ. The prescription itself, however, remains part of the post-hoc selection space, and we test that freedom separately in Section 5.3. Figure 2 shows the historical PDG and FLAG determinations of $m_s^{\overline{\text{MS}}}(2\text{ GeV})$ for context; the caption notes that this plot is not an independent validation.

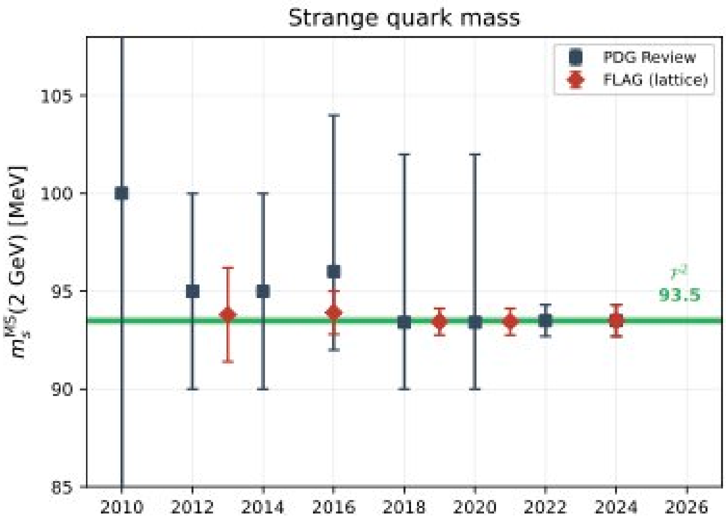

Figure 2. Historical PDG (squares) and FLAG lattice (diamonds) determinations of $m_s^{\overline{\mathrm{MS}}}(2\ \mathrm{GeV})$ from 2010 to 2024, with current uncertainties, *shown for context only*. The horizontal green line marks the value of $m_s^{\overline{\mathrm{MS}}}(2\ \mathrm{GeV})$ obtained by evolving the lepton-derived target $\mathcal{F}^2$ from $\mu_\star$ back to 2 GeV using the running setup described in Section 4; numerically this equivalent value is ≈ 93.50 MeV, which sits approximately 0.06 MeV above the FLAG 2024 Table 11 central value of 93.44 MeV, consistent with the small positive residual $+0.042$ MeV of Equation (9) expressed at the 2 GeV reference scale. *This plot is not an independent validation of the relation.* It is a visual context diagram: the green line is not derived from an independent measurement but from the same FLAG running setup used to produce Equation (9). The plot is a post-hoc assembly ($\mathcal{F}$ was computed in 2025–2026 against then-current data, not predicted in advance). Future updates to PDG and FLAG averages will continue to test the value directly, since the lepton-side prediction is essentially stationary at present precision. The plotted historical data points and their source attributions are tabulated in `results/ms_history.csv` in the accompanying repository [14].

5.3. Alternative Scale Prescriptions

The continuous scan above shows that the relation is sharp at $\mu_\star$, but does not address whether $\mu_\star$ is the unique *natural* scale among candidates derivable from lepton masses alone. We test eight scale prescriptions in total: six alternative lepton-derived prescriptions, the baseline $\mu_\star = \sum m_\ell$, and the conventional non-lepton reference scale of 2 GeV. We compute $m_s^{\overline{\mathrm{MS}}}(\mu)$ at each and compare against the constant $\mathcal{F}^2 = 95.1134$ MeV:

Prescription	μ (MeV)	Hit (within 1σ)?
$\mu_\star = m_e + m_\mu + m_\tau$ (baseline)	1883.10	yes ($+0.06\sigma$)
$m_\mu + m_\tau$	1882.59	yes (trivially equivalent)
m_τ	1776.93	no (-2.31σ)
2 GeV (conventional reference, not lepton-derived)	2000.00	no ($+2.46\sigma$)
$(m_e + m_\mu + m_\tau)/3$	627.70	non-perturbative
$\sqrt{m_\mu m_\tau}$	433.30	non-perturbative
$(m_e m_\mu m_\tau)^{1/3}$	45.78	non-perturbative
$3/(1/m_e + 1/m_\mu + 1/m_\tau)$ (harmonic)	1.53	non-perturbative

Of the four prescriptions in the perturbative regime ($\mu \gtrsim 1$ GeV), only $\sum m_\ell$ and the trivially equivalent $m_\mu + m_\tau$ produce hits. The next-best lepton-derived alternative (m_τ alone) misses by -2.31σ, and the conventional non-lepton reference scale of 2 GeV misses by $+2.46\sigma$. Four further candidates fall below the perturbative regime and we report them as non-perturbative rather than excluding them. Within the perturbative regime, the match is specific to $\mu_\star = \sum m_\ell$; it is not generic across reasonable lepton-derived prescriptions.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 April 2026 doi:10.20944/preprints202604.0693.v1

5.4. Tight Null-Model Test

The relevant question is: under a Koide-conditioned prior on lepton triples, how often does $\mathcal{F}^2$ land within current lattice precision of $m_s^{\overline{\text{MS}}}$ at the natural scale of that triple? We construct the null as follows.

Sample candidate triples (m_1, m_2, m_3) with m_3 drawn log-uniformly from $[10^0, 10^4]$ MeV, m_1 drawn log-uniformly from $[10^{-3}, 10^{-1}] \cdot m_3$, and m_2 solved analytically from the Koide condition $Q(m_1, m_2, m_3) = 2/3$ in the physical range $m_1 < m_2 < m_3$. Writing $a = \sqrt{m_1}$, $x = \sqrt{m_2}$, and $c = \sqrt{m_3}$, the Koide condition $Q = 2/3$ rewrites as the quadratic $x^2 - 4(a+c)\,x + (a^2 - 4ac + c^2) = 0$, whose physical branch (the one yielding $m_1 < m_2 < m_3$) is $x = 2(a+c) - \sqrt{3(a^2+4ac+c^2)}$, with $m_2 = x^2$. Discard triples for which no real m_2 exists in this range. Verify $|Q - 2/3| < 10^{-5}$ for each retained triple. The sampler oversamples and continues until $10^5 = 100{,}000$ Koide-valid triples accumulate (rejections from the no-real-m_2 check are absorbed by oversampling at the sampler level). Of these 100,000 Koide-valid triples, 74,515 are subsequently excluded because the natural scale $\mu_\star^{(i)} \equiv m_1 + m_2 + m_3$ falls below 1 GeV, where four-loop QCD running is not reliable, leaving 25,485 retained valid triples for the analysis. The $\mu_\star \geq 1$ GeV cut is methodological, not post hoc: the test is undefined for triples where running is unreliable. The repository implementation also applies an upper bound of $\mu_\star \leq 50$ GeV as a perturbative-regime guardrail for very heavy triples; this upper bound does not affect the baseline prior-A sample (whose heaviest triples have $\mu_\star \lesssim 10$ GeV) but is disclosed for completeness. We verify in Section 5.6 that the result is stable when this cutoff is varied.

For each valid triple, compute $\mathcal{F}^{(i)}$ and $\mu_\star^{(i)}$, run $m_s^{\overline{\text{MS}}}(\mu_\star^{(i)})$ from the FLAG 2 GeV input, and record a hit if $|\mathcal{F}^{(i)\,2} - m_s^{\overline{\text{MS}}}(\mu_\star^{(i)})| < \sigma_{m_s}(\mu_\star^{(i)})$. Here $\sigma_{m_s}(\mu_\star^{(i)})$ denotes the propagated FLAG Table 11 input uncertainty only; α_s, truncation, and lepton-side contributions are excluded from the null acceptance window (they are reported in the error budget of Section 5 but represent subdominant corrections that do not materially affect the hit fraction). Of 25,485 valid triples, 121 produced hits. The hit fraction is

$$p_{\text{tight}} = 0.4748\%, \qquad \text{95\% Clopper–Pearson interval (sampling noise only): } [0.394\%, 0.567\%]. \tag{12}$$

The interval quantifies binomial sampling uncertainty in the Monte Carlo only; it does not capture model-specification uncertainty from the choice of prior, which we address separately in Section 5.5. The distribution of $(\mathcal{F}^{(i)\,2} - m_s^{\overline{\text{MS}}}(\mu_\star^{(i)}))/\sigma_{m_s}$ across the 25,485 valid samples spans roughly ± 50 standard deviations; the observed lepton triple sits within the central $\pm 1\sigma$ window, which is the outcome quantified by (12).

5.5. Prior Sensitivity

The hit fraction (12) depends on the choice of prior over Koide-saturating triples. The log-uniform prior used above is one reasonable choice but not the only one. We re-run the tight null under three additional priors selected to span the space of reasonable alternatives:

Prior	Hit fraction	95% CI
A: log-uniform $m_3 \in [10^0, 10^4]$ MeV (baseline)	0.47%	[0.39, 0.57]
B: log-uniform $m_3 \in [10^{-1}, 10^5]$ MeV (wider range)	0.32%	[0.26, 0.39]
C: log-uniform $m_1/m_3 \in [10^{-4}, 10^{-2}]$ (smaller ratio)	0.43%	[0.35, 0.52]
D: linear-uniform $m_3 \in [10, 10^4]$ MeV and $m_1/m_3 \in [10^{-3}, 10^{-1}]$ (stress test)	0.23%	[0.20, 0.27]

All four priors give hit fractions in the range 0.23%–0.47%, with maximum-to-minimum fold variation of 2.04. The deliberately adversarial linear prior (D), which heavily weights heavy triples and is not natural for mass-hierarchy problems, lowers the hit fraction rather than raising it. No reasonable prior we tested raises the hit fraction above 1%. We conclude that the magnitude of the hit fraction is stable to the choice of prior within the class we have tested, varying by a factor of two across the four priors but remaining sub-percent throughout.

5.6. Loose Null-Model Test and Filter Sensitivity

To bound the residual look-elsewhere effect across a broader claim, we test $\mathcal{F}^2$ against each of $\{m_u, m_d, m_s, m_c, m_b\}$ at the natural scale $\mu_\star^{(i)}$ of each triple, counting a triple as a hit if any of the 5 comparisons agrees within a common 1σ convention. We use FLAG 2024 1σ uncertainties for s, and PDG 2024 quark summary-table values for u and d (quoted at 90% CL); as with c and b, this CL distinction does not affect the reported result because no triple produces a hit on u or d under either convention. For c and b, the repository implementation uses the PDG 2024 quark summary-table central values ($m_c = 1273$ MeV, $m_b = 4183$ MeV) with their quoted uncertainties, which are stated at 90% CL rather than 1σ; a fully consistent treatment would divide these by 1.645 before forming a Gaussian-equivalent 1σ acceptance window. In practice this distinction does not affect any reported result: no triple in the 25,485-sample produces a hit on c or b under either convention, because $\mathcal{F}^2 \approx 95$ MeV is well outside both the charm range ($m_c \sim 1.3$ GeV) and the bottom range ($m_b \sim 4.2$ GeV) at any reasonable scale. We restrict to $n = 2$ because $\mathcal{F}$ has units of MeV$^{1/2}$, so only $\mathcal{F}^2$ carries mass dimension one and is dimensionally commensurate with a quark mass; comparisons $\mathcal{F}^n$ for $n \neq 2$ are not unit-invariant and would not constitute a dimensionally consistent look-elsewhere accounting. Across the same 25,485 valid triples, the overall hit fraction is

$$p_{\text{loose}} = 0.4748\%, \qquad \text{95\% Clopper–Pearson interval (sampling noise only): } [0.394\%, 0.567\%], \quad (13)$$

identical to p_{tight} (both yield 121 hits on the same 25,485-triple sample). The breakdown by quark species is shown in Table 1: of the five candidate species, only m_s produces hits, because $\mathcal{F}^2 \approx 95$ MeV is well below m_c and m_b at $\mu_\star^{(i)}$ and well above m_u and m_d. The dimensionally consistent broadening from one species to five does not introduce additional matches; the result is robust under this expansion.

Table 1. Loose null hit counts by quark species over 25,485 valid Koide triples. Cells show the number of triples for which $|\mathcal{F}^2 - m_q^{\overline{\text{MS}}}(\mu_\star^{(i)})| < \sigma_{m_q}(\mu_\star^{(i)})$. The observed lepton triple lies in the m_s row.

Quark species	Hit count
m_u	0
m_d	0
m_s	**121**
m_c	0
m_b	0
Total	121

Both null tests use a lower cutoff of $\mu_\star^{(i)} > 1$ GeV and an upper cutoff of $\mu_\star^{(i)} \leq 50$ GeV to ensure four-loop QCD running is reliable. To verify that this cutoff is not coupling the sample to the comparison observable, we re-run the tight null under three alternative cutoffs:

μ_{min}	Tight-null hit fraction
0.7 GeV (extrapolated; running not strictly trusted)	0.41%
1.0 GeV (baseline)	0.47%
1.5 GeV	0.58%
2.0 GeV (excludes the observed lepton triple)	0.00%

Across the three nontrivial cutoffs, the hit fraction varies by a factor of 1.4, well below the factor-of-two threshold at which we would consider the filter to be doing significant work. The 2 GeV cutoff excludes the observed lepton triple ($\mu_\star^{\text{lepton}} \approx 1.88$ GeV < 2 GeV) and yields zero hits among the 17,801 Koide-valid triples retained above this cut, because $\mathcal{F}^2$ drifts outside the 1σ window of $m_s^{\overline{\text{MS}}}(\mu)$ at those larger scales. We conclude that the result is stable under reasonable variations of the lower cutoff.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 April 2026 doi:10.20944/preprints202604.0693.v1

5.7. Interpretation

Before interpreting the numbers reported in this section, we state explicitly what they are and are not. The hit fractions are conditional frequencies under specific Monte Carlo constructions over an explicitly chosen family of toy null models. They are not p-values, do not constitute hypothesis tests against any physical null, and should not be interpreted as evidence against chance in any objective sense. We report them for reproducibility, and to characterize the observation's behavior within a family of toy nulls that bracket several plausible modeling choices; readers should weight any conclusion drawn from these numbers by the credibility they assign to the null family itself, which we make no claim is uniquely natural or exhaustive of all reasonable constructions.

With that framing in place: the tight-null Monte Carlo, conditioned on Koide-saturating triples and using our baseline prior, returns a hit fraction of 0.47%, which remains in the range 0.23%–0.47% across all four priors we tested. The loose null, which broadens the comparison from m_s alone to $\mathcal{F}^2$ against all five quark species (restricted to $n = 2$ for dimensional consistency), returns the same 0.47% on the same sample, because $\mathcal{F}^2$ is too small to match m_c or m_b and too large to match m_u or m_d at the relevant scales. The continuous scale-sensitivity scan shows a narrow $\pm 2.0\%$ window around $\mu_\star$ in which agreement with $m_s^{\overline{\text{MS}}}(\mu)$ holds, and the discrete scale-prescription test shows that no other lepton-derived scale prescription in the perturbative regime produces a hit. Both nulls are stable under variation of the prior and the filter cutoff at the levels reported in subsections 5.5 and 5.6. These numbers describe the observation's behavior within the constructions we built; we draw no further inferential conclusion from them and make no significance claim.

6. Caveats and Non-Claims

We make the following non-claims explicit, because the intuitive reading of a numerical match between lepton-derived and quark-derived quantities tends to import claims we are not making.

Not a derivation.

We have not derived $\mathcal{F}^2 = m_s^{\overline{\text{MS}}}(\mu_\star)$ from any Lagrangian, symmetry, or theoretical framework. The observation is post-hoc in the sense that all of m_e, m_μ, m_τ, and m_s were known when the comparison was made.

Not a prediction.

The relation was not written down in advance of the strange-quark measurement and then confirmed; the masses involved were already established. The Monte Carlo characterization reported in Section 5 is independent of authorial chronology: it asks how often the relation arises within a Koide-conditioned toy null family, not whether it was anticipated.

Not a mechanism.

We do not propose a flavor-symmetry mechanism, a modular construction, a flavon alignment, or a grand-unified embedding that would generate $\mathcal{F}^2 = m_s^{\overline{\text{MS}}}(\mu_\star)$. We do not claim relevance to any specific flavor model. We hand the question of mechanism to model builders without attempting to anticipate their answer.

Not the only algebraic route to a comparable number.

Rivero [7] showed independently in 2011 that the quark triple (s, c, b) with $-\sqrt{m_s}$ satisfies a Koide-like relation $Q \approx 2/3$, and that chained Koide conditions (the "Koide waterfall") connect lepton and quark sectors algebraically; this construction recovers m_s in the same numerical neighborhood as the Soddy curvature reported here. The two routes are conceptually distinct: Rivero's construction is purely algebraic and chains Koide conditions across triples, whereas the present construction is geometric and operates on a single Koide-saturating triple via the Descartes circle theorem. We note the convergence without assigning evidentiary weight to it: independent algebraic and geometric constructions land in a similar numerical neighborhood, but that fact alone does not distinguish structure from coincidence.

 doi:10.20944/preprints202604.0693.v1

We also note the alternative reading: the convergence of two independent routes could indicate that $\mathcal{F}^2 \approx m_s$ is not structurally difficult to achieve from lepton-mass combinations near 1 GeV, a consideration that should inform the evidentiary weight assigned to either route individually.

Not a charm or bottom relation.

Restricting to the dimensionally consistent comparison $\mathcal{F}^2$ versus quark mass, we find $\mathcal{F}^2 \approx$ 95 MeV is well outside both the charm range ($m_c \sim 1.3$ GeV) and the bottom range ($m_b \sim 4.2$ GeV) at any reasonable scale, so no $\mathcal{F}^2 \approx m_c$ or $\mathcal{F}^2 \approx m_b$ relation is suggested. Numerical proximities involving $\mathcal{F}^n$ for $n \neq 2$ noted in earlier drafts of this work (e.g., $\mathcal{F}^3 \approx m_c^{\overline{\text{MS}}}(m_b)$) are dimensionally invalid: $\mathcal{F}$ has units of MeV$^{1/2}$, so $\mathcal{F}^n$ with $n \neq 2$ does not carry mass dimension and cannot be compared with a quark mass in a unit-invariant way. We exclude all such comparisons from this presentation and restrict the present note to the strange sector.

7. Discussion and Conclusions

If future work were to establish that the observation is not coincidental, possible interpretations would fall into three broad classes. First, the Koide relation may have a geometric origin that has not been recognized, and the Soddy connection may be its geometric content; in that case the present observation extends the same geometric structure into the strange sector via the natural mass scale of the lepton triple. Second, the lepton-sum scale $\mu_\star = \sum m_\ell$ may be the natural matching scale of some flavor mechanism that produces m_s from leptonic input, in which case the observation would be a datum that any such mechanism should seek to reproduce. Third, the observation may be a coincidence of the kind that Koide's relation itself might be: an unexplained numerical regularity that has remained stable as measurements have improved, and that any future explanation of Koide will eventually need to account for. We do not argue for any of the three. We report the observation for the record without claiming inferential force beyond what the model-conditional frequencies in Section 5 support.

Author's tentative interpretation.

The author's own reading of this numerical match, offered as speculation rather than argument, is the following. The Koide relation has held to one part in 10^5 for over four decades without an accepted derivation, and $\mathcal{F}^2$ now matches $m_s^{\overline{\text{MS}}}$ at the natural lepton-sum scale within current lattice precision. These are two precise numerical regularities—one internal to the lepton sector, and one near-equality between a lepton-derived quantity and a quark mass within current lattice precision—for which no mechanism is known. In the author's view, regularities of this kind are best understood as shadows of underlying algebraic structure: as representations in a geometric frame of relations that are themselves primarily algebraic, rather than as products of the geometric frame in which they happen to appear. On this reading, the Descartes–Soddy construction presented in Section 2 and Proposition 1 is a visualization of an algebraic identity, not an explanation for it, and any reading that would treat a circle-packing configuration as producing mass would confuse the shadow with what casts it. The author makes no case for this reading and does not ask the reader to accept it; the three possibilities enumerated above are offered as a neutral map of interpretations, and any preference among them is extra-scientific at the present level of evidence. We report the observation and its geometric identification in the possibility that, if such an underlying algebraic structure exists and is eventually identified by others, the present numerical match may serve as an additional constraint to narrow the search for it. We neither propose a mechanism nor claim to understand why this numerical match holds; we report only that the match is present and is stable under the sensitivity scans of Section 5.

Conclusions.

We have introduced an ordinary-Descartes reading of the Koide lepton triple, computed the resulting outer Soddy curvature $\mathcal{F}$, and reported the numerical observation that $\mathcal{F}^2 = 95.113$ MeV is consistent with the strange-quark $\overline{\mathrm{MS}}$ mass run to the natural lepton-sum scale $\mu_\star = \sum m_\ell$ within current lattice precision (residual $+0.042$ MeV against ± 0.69 MeV). The observation is stable across four prior choices, eight scale prescriptions, and the non-excluding lower-cutoff choices tested; at the 2.0 GeV lower cutoff the observed lepton triple is removed by construction and the hit fraction falls to zero. We make no claim of mechanism, derivation, or prediction. The result is presented as a precise numerical regularity whose theoretical status, if any, remains open.

Computational Note

Central values, lattice-input propagation, and the null-model Monte Carlos reported in this paper are reproducible from the version-pinned GitHub repository at commit `bb8b1ae` [14]. Four-loop $\overline{\mathrm{MS}}$ running is performed via CRunDec [13] with $\alpha_s(M_Z) = 0.1180$ and flavor thresholds at $m_c(m_c) = 1.273$ GeV and $m_b(m_b) = 4.183$ GeV. Lepton inputs are PDG 2024 charged-lepton pole masses. Quark inputs are the FLAG 2024 $N_f = 2{+}1{+}1$ Table 11 estimate for m_s (as used in Section 4) and the PDG 2024 quark summary-table values for m_c and m_b. For each run, the null-model sampler oversamples raw candidates until 10^5 Koide-valid triples are obtained; after applying perturbative-regime cuts 1 GeV $< \mu_\star \leq 50$ GeV, 25,485 retained triples enter the baseline analysis. The sampler uses a fixed random seed for reproducibility, and each configuration is replicated across four priors and four filter cutoffs. The $\alpha_s(M_Z)$ uncertainty propagation and the four-loop versus three-loop truncation comparison reported in the error-budget decomposition of Section 5 are performed by the same running algorithm; their outputs are archived as precomputed artifacts in `results/error_budget.json` at the cited commit [14]. Total runtime for the full robustness analysis is under five minutes on a laptop.

Author Contributions: Conceptualization, formal analysis, software, investigation, visualization, writing—original draft preparation, and writing—review and editing: A.M.B. The author has read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: All original computations reported in this paper—the central values, error budget, robustness tables, and null-model statistics in Sections 3–5—are derived from published sources [11,12] and are reproducible with the accompanying code [14]. Figure 2 displays published historical PDG and FLAG determinations of m_s for context; the input data table is archived as `results/ms_history.csv` in the accompanying repository [14].

Acknowledgments: The author thanks J. Kocik for the geometric reinterpretation of the Koide relation that motivated this work, and acknowledges A. Rivero's independent algebraic work [7] on the strange-charm-bottom Koide tuple, which converges with the present geometric route. Computational infrastructure and institutional support were provided by Applied Dynamics Research, Sapporo.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Y. Koide, "Fermion-Boson Two Body Model of Quarks and Leptons and Cabibbo Mixing," Lett. Nuovo Cim. **34**, 201 (1982).
2. Y. Koide, "New view of quark and lepton mass hierarchy," Phys. Rev. D **28**, 252 (1983); Erratum: Phys. Rev. D **29**, 1544 (1984).
3. R. Foot, "A note on Koide's lepton mass relation," arXiv:hep-ph/9402242 (1994).
4. A. Rivero and A. Gsponer, "The strange formula of Dr. Koide," arXiv:hep-ph/0505220 (2005).
5. Y. Sumino, "Family gauge symmetry as an origin of Koide's mass formula and charged lepton spectrum," JHEP **05**, 075 (2009), arXiv:0812.2103.
6. W. Rodejohann and H. Zhang, "Extension of an empirical charged lepton mass relation to the neutrino sector," Phys. Lett. B **698**, 152 (2011), arXiv:1101.5525.
7. A. Rivero, "A new Koide tuple: Strange-charm-bottom," arXiv:1111.7232 [hep-ph] (2011).
8. J. Kocik, "The Koide Lepton Mass Formula and Geometry of Circles," arXiv:1201.2067 [physics.gen-ph] (2012).
9. W. Królikowski, "Two options for the Koide-like formula generalized to neutrinos," arXiv:1404.5705 (2014).

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 10 April 2026 doi:10.20944/preprints202604.0693.v1

10. H.S.M. Coxeter, "The problem of Apollonius," Amer. Math. Monthly **75**, 5–15 (1968).
11. S. Navas et al. (Particle Data Group), "Review of Particle Physics," Phys. Rev. D **110**, 030001 (2024).
12. Y. Aoki et al. (Flavour Lattice Averaging Group, FLAG), "FLAG review 2024," Phys. Rev. D **113**, 014508 (2026), arXiv:2411.04268.
13. B. Schmidt and M. Steinhauser, "CRunDec: A C++ package for running and decoupling of the strong coupling and quark masses," Comput. Phys. Commun. **183**, 1845 (2012); F. Herren and M. Steinhauser, "Version 3 of `RunDec` and `CRunDec`," Comput. Phys. Commun. **224**, 333 (2018).
14. A.M. Brilliant, "koide-soddy: Reproducibility code for 'Soddy curvature of the Koide lepton triple and the strange-quark mass'," GitHub repository, commit `bb8b1ae`, accessed 2026-04-09. https://github.com/AndBrilliant/koide-soddy/tree/bb8b1ae.