

The UV Renormalization-Group Run of Koide-Type Mass Relations:

A Reference Dataset from M_Z to M_{Pl}

A. Rivero*, with computational assistance from Claude (Anthropic)

March 2026

Abstract

We provide a reference dataset and exhaustive analysis of Koide-type mass relations $K = \sum m_i / (\sum \sqrt{m_i})^2 = 2/3$ under Standard Model renormalization-group evolution from M_Z to M_{Pl} . Using PyR@TE 3 two-loop SM beta functions and AHS (Antusch–Hinze–Saad) multi-loop Yukawa data, we track the charge-lattice ratio $\rho \equiv 3K - 1$ for all pure-sector (quark-only or lepton-only) signed Koide triples constructed from the 9 SM Yukawa couplings. Mixed quark–lepton tuples are relegated to an appendix. We confirm that (i) the lepton Koide $K(e, \mu, \tau)$ is RG-stable to 0.003% across 16 decades, (ii) the inverse down-quark Koide $K^{-1}(d, s, b)$ crosses $\rho = 1$ at $\mu \approx 1\text{--}1.6$ TeV, (iii) pure QCD running preserves ρ exactly (the crossing is entirely electroweak), and (iv) the joint probability of finding Koide relations this good in *both* the lepton and quark sectors from random masses is $\lesssim 10^{-3}$ even after full look-elsewhere correction. All code, data, and figures are provided for reproducibility.

1 Introduction and notation

The Koide relation [?] states that the charged-lepton masses satisfy

$$K(e, \mu, \tau) \equiv \frac{m_e + m_\mu + m_\tau}{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2} = \frac{2}{3}, \quad (1)$$

to an accuracy of 0.04% using PDG pole masses. We adopt the **charge-lattice ratio**

$$\rho \equiv \frac{\langle z_i^2 \rangle}{z_0^2} = 3K - 1 \quad (2)$$

where z_0, z_i are the charge-lattice components defined by $m_i = k(z_0 + z_i)^2$ with $\sum z_i = 0$. The Koide condition $K = 2/3$ is equivalent to $\rho = 1$: the variance of the charges equals the squared mean. We use ρ throughout because:

- $\rho > 1$: charges too spread (hierarchy too steep);
- $\rho < 1$: charges too compressed;
- $\rho = 1$: exact Koide.

The percentage deviation is $\delta = (\rho - 1) \times 100\%$.

The **signed Koide** generalizes (??) to

$$K_s(m_1, m_2, m_3) = \frac{m_1 + m_2 + m_3}{(s_1\sqrt{m_1} + s_2\sqrt{m_2} + s_3\sqrt{m_3})^2}, \quad (3)$$

with signs $s_i = \pm 1$ (4 independent patterns; overall sign cancels). The **inverse Koide** is $K^{-1}(m_1, m_2, m_3) \equiv K(1/m_1, 1/m_2, 1/m_3)$.

*Corresponding author. arivero@unizar.es

2 Data and RG evolution

2.1 AHS Yukawa data

All $\overline{\text{MS}}$ Yukawa couplings are from Antusch–Hinze–Saad [?] (Table 2), evaluated at four scales: M_Z , 1 TeV, 10 TeV, and 10^7 GeV using full multi-loop SM running with proper threshold matching at m_b and m_t .

2.2 PyR@TE 2-loop extension to M_{Pl}

We extend the running to $M_{\text{Pl}} = 2.4 \times 10^{18}$ GeV using PyR@TE 3 [?], which generates complete 2-loop SM beta functions. Initial conditions are AHS values at M_Z ; gauge couplings $\alpha_Y = 0.01017$, $\alpha_2 = 0.03378$, $\alpha_s = 0.1181$. Diagonal Yukawa matrices (CKM mixing negligible for Koide ratios). The data is stored as `pyrate_sm_2loop.npz` (3000 points in $t = \ln(\mu/M_Z)$).

2.3 QCD invariance of ρ

Key result: Pure QCD running preserves ρ exactly. The QCD mass anomalous dimension is flavor-universal: $m_q(\mu) = m_q(\mu_0) \cdot c(\alpha_s(\mu))/c(\alpha_s(\mu_0))$ with the *same* function c for all quarks. Therefore every ratio m_i/m_j is a QCD RG invariant, and ρ (which depends only on mass ratios) is QCD-invariant.

Verified numerically with CRunDec [?] (4-loop, $n_f = 5$): $K^{-1}(d, s, b) = 0.66739$ at *all* scales from M_Z to 10^4 GeV under QCD-only running. The AHS full SM running gives $K^{-1}(d, s, b) = 0.66667$ at 1 TeV. **The entire deviation is electroweak + Yukawa**, specifically:

- $\text{SU}(2)_L$ and $\text{U}(1)_Y$ gauge corrections (generation-dependent);
- Top Yukawa: $\frac{3}{2}y_t^2$ in $\beta(y_b)$ but not in $\beta(y_d)$ or $\beta(y_s)$ (isospin partner effect).

3 Non-exotic tuples

3.1 Lepton Koide: $K(e, \mu, \tau)$

Scale	ρ	δ (%)
M_Z	1.00351	+0.176
1 TeV	1.00351	+0.175
10 TeV	1.00345	+0.173
10^7 GeV	1.00346	+0.173
M_{GUT}	1.00347	+0.173
M_{Pl}	1.00347	+0.173

Table 1: Lepton Koide: RG-stable to 0.003% across 16 decades.

The lepton Koide is a near-fixed-point of the SM RG flow (shifting by only 0.004% over 16 decades), a stability first analyzed by Sumino [?]. This is because: (i) gauge corrections are generation-blind (all charged leptons have the same EW quantum numbers); (ii) the Yukawa self-correction $\beta(y_\ell) \supset \frac{3}{2}y_\ell^3/(16\pi^2)$ is generation-dependent but tiny ($y_\tau^2 \sim 10^{-4}$); (iii) unlike quarks, leptons have no heavy isospin partner ($y_\nu \approx 0$ in the SM). The leading correction to ρ is $\sim y_\tau^4/(16\pi^2)^2 \sim 10^{-8}$ per e -fold.

Scale	ρ	δ (%)
M_Z	1.00216	+0.108
1 TeV	1.00001	+0.0003
10 TeV	0.99934	−0.033
10^7 GeV	0.99568	−0.216
M_{GUT}	0.99691	−0.155
M_{Pl}	0.99706	−0.147

Table 2: Inverse down Koide: crosses $\rho = 1$ at $\mu = 1021$ GeV (AHS) / 1588 GeV (PyR@TE 2-loop).

3.2 Inverse down Koide: $K^{-1}(d, s, b)$

This is $K(1/m_d, 1/m_s, 1/m_b) = 2/3$. The crossing at 1 TeV is an electroweak phenomenon (Section ??).

3.3 Inverse upper chain: $K(1/s, 1/c, 1/t)_{(+,+, -)}$

Crosses $\rho = 1$ at $\mu = 2297$ GeV (AHS). The sign flip on $\sqrt{1/y_t}$ corresponds to the top quark being “inverted” in the charge lattice. Crosses near the same TeV scale as $K^{-1}(d, s, b)$.

3.4 Third-generation Koide: $K(t, b, \tau)$

Crosses $\rho = 1$ at $\mu = 2400$ GeV (AHS). The crossing at 2.4 TeV reinforces the TeV-scale clustering.

4 Exotic tuples (quark–lepton mixing)

4.1 $K(u, \mu, \tau)$ — Pati–Salam hint

Crosses $\rho = 1$ at $\mu \approx 3.1 \times 10^6$ GeV. Replacing $e \rightarrow u$ in the lepton Koide. At M_Z , $y_u/y_e \approx 2.5$; but y_u runs down while y_e is nearly stable, and they converge at high scales. The crossing at $\sim 10^6$ GeV is suggestive of Pati–Salam SU(4) unification.

4.2 $K^{-1}(d, s, \tau)$ — SU(5) hint

Crosses $\rho = 1$ at $\mu = 6396$ GeV. Replacing $b \rightarrow \tau$ in the inverse down Koide. In SU(5), b and τ sit in the same $\bar{5}$, so $y_b/y_\tau \rightarrow 1$ at M_{GUT} .

5 Full landscape and rankings

We evaluate ρ for all pure-sector signed triples — quark-only or lepton-only, never mixing the two in a single tuple — along the full RG trajectory from M_Z to M_{Pl} . From 9 Yukawa couplings, we form $2\binom{9}{3} \times 4 = 672$ signed triples (84 all-direct + 84 all-inverse, each with 4 sign patterns). Of the 672, only 168 are pure-sector (the remainder mix quarks and leptons); these are relegated to Appendix ??.

figures/rho_best_planck.pdf

Figure 1: Charge-lattice ratio ρ for the best pure-sector Koide tuples under 2-loop SM RG flow from 10 GeV to M_{Pl} (1-loop below M_Z). Gold band: $|\rho - 1| < 0.2\%$. Markers: AHS multi-loop data. No mixed quark-lepton tuples.





Figure 2: *Left:* Density landscape of pure-sector Koide triples (quark-only in blue, lepton-only in red) from 10 GeV to M_{Pl} (linear ρ scale). *Right:* Histogram of ρ at 1 TeV for the same tuples, confirming the near-uniform background distribution against which the $\rho \approx 1$ tuples stand out.



figures/density_planck_zoom.pdf

Figure 3: Zoom near $\rho = 1$. The convergence of multiple tuples in the TeV region and their divergence at high energy is evident.

figures/panels_abc.pdf

Figure 4: Three-panel summary (pure-sector only). **(a)** Best tuples: $K(e, \mu, \tau)$ (stable to $\sim 0.003\%$), $K^{-1}(d, s, b)$ (crosses $\rho = 1$ near 1 TeV), and $K(1/s, 1/c, 1/t)_{(+,+, -)}$. AHS discrete data (markers) below 10^7 GeV; 1-loop SM extrapolation (dashed) above. **(b)** Next-best pure-sector tuples: $K(u, c, s)_{(+, -, -)}$ and $K(c, s, b)_{(+, -, +)}$, with panel (a) tuples shown faintly for reference. All are $> 4\%$ from exact Koide. **(c)** Histogram of $|\delta|$ at 1 TeV for 270 pure-sector signed triples (blue = direct, red = inverse).

5.1 Best tuples at 10 TeV

Rank	Tuple	ρ	δ (%)
1	$K(1/c, 1/t, 1/s)_{(+,-,+)}$	1.000095	+0.010
2	$K^{-1}(d, s, b)$	0.999076	-0.092
3	$K(e, \mu, \tau)$	1.003517	+0.352
4	$K^{-1}(d, s, \tau)$	1.003534	+0.353
5	$K(1/s, 1/b, 1/\mu)_{(+,-,+)}$	0.992330	-0.767
6	$K(u, \mu, \tau)$	0.992033	-0.797
7	$K(b, \mu, \tau)_{(+,-,+)}$	1.016731	+1.673
8	$K(t, b, \tau)$	0.980309	-1.969

Table 3: Top-8 Koide triples at $\mu = 10$ TeV (PyR@TE 2-loop), from the 9 individual Yukawa couplings only. Chiral combinations (m_u+m_d , $m_u \rightarrow 0$) are not included; see Table ?? for those. Clear gap between ranks 6 ($< 1\%$) and 7 ($> 1.6\%$).

5.2 Top-5 at 1 TeV

Rank	Tuple	ρ	δ (%)
1	$K^{-1}(d, s, b)$	1.000273	+0.027
2	$K^{-1}(t, b, \mu)$	1.000895	+0.090
3	$K(t, b, \tau)$	1.001047	+0.105
4	$K(1/u, 1/s, 1/\mu)$	0.998600	-0.140
5	$K(1/c, 1/t, 1/s)_{(+,-,+)}$	1.001540	+0.154

Table 4: Top-5 Koide triples at $\mu = 1$ TeV (individual Yukawas only; chiral combinations excluded). All have $|\delta| < 0.16\%$.

5.3 Chiral tuples: m_u+m_d and $m_u \rightarrow 0$

The chiral sum m_u+m_d (the GMOR combination governing m_π^2) and the chiral limit $m_u \rightarrow 0$ are physically motivated entries. Since QCD running multiplies *all* quark masses by the same flavor-universal factor, m_u+m_d scales identically and K is again QCD-invariant. The best chiral tuples at the four AHS scales are listed in Table ??.

Rank	Tuple	Scale	ρ	δ (%)
1	$K^{-1}(t, s, u+d)_{(+,-,-)}$	10^7 GeV	0.999329	-0.067
2	$K(s, u+d, 0)$	M_Z	0.993887	-0.611
3	$K(s, u+d, 0)$	1 TeV	0.993055	-0.695
4	$K^{-1}(t, s, u+d)_{(+,-,-)}$	10 TeV	1.006984	+0.698
5	$K^{-1}(t, s, u+d)_{(+,-,-)}$	1 TeV	1.008762	+0.876

Table 5: Best quark-only chiral tuples (involving m_u+m_d or $m_u \rightarrow 0$) at the four AHS scales. $K^{-1}(t, s, u+d)_{(+,-,-)}$ crosses $\rho = 1$ between 10 TeV and 10^7 GeV.

6 Look-elsewhere effect

Any search over many candidate tuples risks finding an accidentally good Koide ratio. We quantify this look-elsewhere effect (LEE) via Monte Carlo.

6.1 Null model and trial counting

The null hypothesis is: Yukawa couplings are unrelated random numbers. We draw 9 masses y_i independently, log-uniformly in $[y_e, y_b]$ (the physical Yukawa range at M_Z), and evaluate K for every possible signed triple. The trial count per random draw is:

- $\binom{9}{3} = 84$ triples, each tested with 4 sign patterns $\Rightarrow 336$ direct evaluations.
- The same 84 triples evaluated on inverse masses $1/y_i \Rightarrow 336$ inverse evaluations.

We record the best (smallest) $|\delta| = |K - 2/3|/(2/3)$ separately for direct and inverse, and for both combined.

Direct–inverse symmetry. For log-uniform random masses, the direct and inverse histograms are statistically identical (Figure ??). This follows from scale invariance of K : since K depends only on mass *ratios*, and the distribution of log-ratios $\log(y_i/y_j)$ is symmetric about zero (triangular, as the difference of two i.i.d. uniforms), inverting all masses $y \rightarrow 1/y$ sends $\log(y_i/y_j) \rightarrow -\log(y_i/y_j)$, which has the same distribution. Thus testing both direct and inverse on random masses is redundant: the effective trial count is 336, not 672. In the real SM, direct and inverse Koide values are of course *different* because the actual mass ratios have no such symmetry.

6.2 Single-sector results

Over 10^6 trials (Figure ??):

- $K(e, \mu, \tau)$ at $|\delta| = 0.17\%$ ($\overline{\text{MS}}$, $\mu = M_Z$): individually **not** significant (63% of random sets have some triple this good, given 336 trials). At pole masses, $|\delta| = 0.001\%$ — the original observation of Koide (1982) — but the LEE-corrected probability remains $\sim 1/539$, comparable to $K^{-1}(d, s, b)$.
- $K^{-1}(d, s, b)$ at $|\delta| = 0.0003\%$ ($\overline{\text{MS}}$, $\mu \approx 1$ TeV): **1-in-539** ($\approx 3.1\sigma$) even after full look-elsewhere correction.

6.3 Joint probability

The real significance is that $K(e, \mu, \tau)$ and $K^{-1}(d, s, b)$ use **completely disjoint** mass sets (3 leptons vs. 3 down-type quarks). To test this, we draw 3 random lepton masses and 6 random quark masses independently, evaluate $|\delta_\ell|$ for the lepton triple and $|\delta_q|$ for the best quark triple (direct or inverse, any signs), and ask how often *both* are as good as the SM:

$$P(|\delta_\ell| < 0.18\% \text{ AND } |\delta_q| < 0.11\%) \lesssim 10^{-3} \quad (4)$$

across tested distributions (log-uniform and SM-like). This is the correct LEE-corrected significance: the probability of finding *two independent* Koide relations this accurate from random masses is $\sim 0.1\%$. See Figure ??.

7 PyR@TE validation

The crossing of $\rho = 1$ is robust across methods; only its exact scale shifts by a factor ~ 1.5 between 2-loop (1588 GeV) and multi-loop (1021 GeV).

figures/look_elsewhere.pdf

Figure 5: Look-elsewhere Monte Carlo (10^6 trials, 9 random masses log-uniform in $[y_e, y_b]$). *Left*: histogram of best $|\delta|$ for direct (blue) and inverse (red) evaluations — the two distributions coincide by the symmetry argument in the text. *Right*: cumulative distribution.

Scale	$\delta_{\text{PyR@TE}}$	δ_{AHS}	Δ
M_Z	+0.108%	+0.108%	0.000%
1 TeV	+0.014%	+0.0003%	0.013%
10 TeV	−0.046%	−0.033%	0.013%
10^7 GeV	−0.157%	−0.216%	0.059%

Table 6: PyR@TE 2-loop vs. AHS multi-loop for $K^{-1}(d, s, b)$. The $\sim 0.01\%$ residual at 1 TeV is from 3–4 loop and threshold matching effects.

figures/joint_probability.pdf

Figure 6: Joint look-elsewhere: best lepton $|\delta|$ vs. best quark $|\delta|$ from independent random masses. Gold star = SM. Three panels show different mass distributions.




Figure 7: PyR@TE 2-loop SM (lines) vs. AHS multi-loop (squares).

8 The 1 TeV coincidence

Multiple structurally different tuples cross $\rho = 1$ near 1 TeV:

Tuple	Crossing scale	Sector
$K^{-1}(d, s, b)$	1021 GeV	quark
$K(1/s, 1/c, 1/t)_{(+,+, -)}$	2297 GeV	quark
$K(t, b, \tau)$	2400 GeV	mixed [†]
$K(d, b, \mu)$	1470 GeV	mixed [†]
$K(u, d, \mu)$	698 GeV	mixed [†]

[†]Mixed quark–lepton tuples; see Appendix ?? for details.

Whether this clustering at the TeV scale has a dynamical explanation (e.g., a compositeness threshold) remains an open question.

9 Summary

1. $K(e, \mu, \tau)$: $\rho = 1.0035$, RG-stable to 0.003% across 16 decades (M_Z to M_{Pl}).
2. $K^{-1}(d, s, b)$: crosses $\rho = 1$ at $\mu \approx 1\text{--}1.6$ TeV.
3. Pure QCD running preserves K exactly (flavor-universal γ_m). The crossing is entirely electroweak.
4. Joint look-elsewhere: finding Koide relations this good in *both* leptons and quarks from random masses is $\sim 10^{-3}$.
5. Seven tuples have $|\delta| < 1\%$ at 10 TeV; a clear gap separates these from the remaining ~ 3250 .

All data and code are available at: <https://doi.org/10.5281/zenodo.XXXXXXX>

References

- [1] Y. Koide, “New viewpoint in lepton mass spectrum,” *Lett. Nuovo Cim.* **34** (1982) 201.
- [2] S. Antusch, V. Hinze, S. Saad, “Running quark and lepton parameters at various scales,” *arXiv:2510.01312*.
- [3] L. Sartore, I. Schienbein, “PyR@TE 3,” *Comput. Phys. Commun.* **261** (2021) 107819, *arXiv:2007.12700*.
- [4] F. Herren, M. Steinhauser, “Version 3 of RunDec and CRunDec,” *Comput. Phys. Commun.* **224** (2018) 333.
- [5] Y. Koide, “Sumino model and my personal view,” *arXiv:1701.01921*.

A Methods: reproducing the analysis

This appendix describes the computational pipeline so that others can reproduce, extend, or adapt the analysis. All code is Python 3 with NumPy, SciPy, and Matplotlib.

A.1 Data generation: PyR@TE 2-loop SM

`run_pyrate_sm.py` solves the 2-loop SM renormalization group equations using PyR@TE 3. Key points:

- PyR@TE 3 generates the full 2-loop beta functions for any non-SUSY gauge theory from a model file. The built-in SM model includes the substitution `g1 : sqrt(5/3) * g1`, so the input coupling is the *un-normalized* hypercharge coupling g' , not the GUT-normalized g_1 .
- Initial conditions: AHS Table 2 values at M_Z . Diagonal Yukawa matrices (off-diagonal CKM elements negligible for Koide ratios).
- The RGE parameter is $t = \ln(\mu/M_Z)$, running from $t = 0$ to $t = \ln(M_{\text{Pl}}/M_Z) \approx 37.8$.
- Solutions are stored via `rge.solutions` (dict with keys `'Yu_{11}'`, `'Yd_{33}'`, etc.) and `rge.tList`.
- Output: `data/pyrate_sm_2loop.npz` containing arrays `mu`, `yu`, `yc`, `yt`, `yd`, `ys`, `yb`, `ye`, `ymu`, `ytau`, `g1`, `g2`, `g3`.

A.2 QCD invariance check: CRunDec

`test_rundec.py` validates that pure QCD running (4-loop, $n_f = 5$) preserves all Koide ratios exactly. Uses the `rundec` Python package. The key insight: the QCD anomalous dimension γ_m is flavor-independent, so all mass *ratios* are RG invariants under QCD alone.

A.3 Koide functions

Four Python functions implement the Koide evaluations:

```
1 def koide(m1, m2, m3):
2     """Standard Koide ratio K = sum(m) / (sum(sqrt(m)))^2"""
3     return (m1+m2+m3) / (np.sqrt(m1)+np.sqrt(m2)+np.sqrt(m3))**2
4
5 def koide_inv(m1, m2, m3):
6     """Inverse Koide: K(1/m1, 1/m2, 1/m3)"""
7     return koide(1/m1, 1/m2, 1/m3)
8
9 def koide_s(m1, m2, m3, s1, s2, s3):
10    """Signed Koide with sign pattern (s1, s2, s3)"""
11    d = (s1*np.sqrt(m1) + s2*np.sqrt(m2) + s3*np.sqrt(m3))**2
12    return (m1+m2+m3) / d if d > 1e-30 else np.inf
13
14 def rho(K):
15     """Charge-lattice ratio: rho = 3K - 1"""
16     return 3*K - 1
```

A.4 Exhaustive scan

`plot_final.py` evaluates ρ for all $2\binom{9}{3} \times 4 = 672$ signed triples (84 all-direct + 84 all-inverse triples, each with 4 sign patterns) at every point along the RG trajectory. The main-text figures show only *pure-sector* tuples (quark-only or lepton-only; 330 of 672). Mixed quark–lepton tuples (the remaining 342) are relegated to Appendix ??.

The earlier `exhaustive_koide_scan.py` (used by `plot_density.py`) employs a larger basis of 21 entries (including $u+d$ chiral sum and the $u \rightarrow 0$ limit) at the 4 AHS scales only. Again, main-text figures show pure-sector tuples at full opacity; exotic entries ($u+d, 0$) appear at reduced opacity; mixed quark–lepton tuples appear only in the supplement.

A.5 Look-elsewhere Monte Carlo

`look_elsewhere_fast.py`: vectorized MC with 10^6 trials. For each trial, draw 9 masses log-uniformly in $[y_e, y_b]$, form all $\binom{9}{3} \times 4 \times 2 = 672$ evaluations (direct + inverse), record the best $|\delta|$.

`joint_probability.py`: the proper test. Draw 3 lepton masses and 6 quark masses *independently*, compute the best lepton Koide (4 sign patterns) and best quark Koide ($\binom{6}{3} \times 4 \times 2 = 160$ evaluations), record whether both beat the SM thresholds simultaneously.

A.6 Hybrid AHS + 1-loop extension

`plot_panels.py` plots honest AHS discrete data points (markers and straight-line connections) below 10^7 GeV, with 1-loop SM beta functions above. This provides a cross-check of PyR@TE and extends the AHS data (which stops at 10^7 GeV) to M_{Pl} .

A.7 File inventory

File	Description
<code>run_pyrate_sm.py</code>	2-loop SM RGEs via PyR@TE 3
<code>rg_ahs_to_planck.py</code>	Hybrid AHS + 1-loop to M_{Pl}
<code>exhaustive_koide_scan.py</code>	All 5320 tuples at AHS scales
<code>test_rundec.py</code>	CRunDec QCD invariance validation
<code>look_elsewhere_fast.py</code>	Vectorized MC (10^6 trials)
<code>joint_probability.py</code>	Joint lepton+quark LEE
<code>plot_final.py</code>	Publication figures (PyR@TE 2-loop)
<code>plot_panels.py</code>	Three-panel figure
<code>plot_density.py</code>	Density landscape at AHS scales
<code>data/pyrate_sm_2loop.npz</code>	NumPy archive: all Yukawas, M_Z to M_{Pl}
<code>note-uv-rg-koide.md</code>	Full analysis note (markdown)

A.8 Dependencies


- Python ≥ 3.8 , NumPy, SciPy, Matplotlib
- `rundec` (pip install rundec) — CRunDec QCD validation
- PyR@TE 3 (<https://github.com/LSartore/pyrate>) — 2-loop SM

B Mixed quark–lepton tuples

The main text restricts attention to pure-sector tuples (all quarks or all leptons). Here we show the complete landscape including mixed quark–lepton tuples such as $K(t, b, \tau)$, $K(u, \mu, \tau)$, and $K^{-1}(d, s, \tau)$. These are structurally less motivated (they mix flavours from different gauge representations) but some cross $\rho = 1$ near the TeV scale.


C Conversation transcript

This dataset was generated in a single interactive session between the author (AR) and Claude (Anthropic’s AI assistant), using the Claude Code CLI tool. The conversation below records the key decisions and directives that shaped the analysis. Tool calls and intermediate outputs are omitted for brevity; only the human directives and key decision points are shown.




figures/density_landscape_supp.pdf

Figure 8: Full Koide landscape at AHS scales, including mixed quark–lepton tuples (gray). Quark-only (blue) and lepton-only (red) at 90% opacity; mixed at 6%.



figures/rho_landscape_zoom_supp.pdf

Figure 9: Zoom near $\rho = 1$ including mixed tuples. Several mixed tuples pass through $\rho \approx 1$ in the TeV region, but the two best remain the pure-sector $K(e, \mu, \tau)$ and $K^{-1}(d, s, b)$.



figures/best_delta_histogram_supp.pdf

Figure 10: Histogram of best $|\delta|$ including all 1050 pure and mixed tuples (cf. 330 pure-sector in the main text). The additional mixed tuples fill in the moderate- $|\delta|$ range but do not produce new outliers near $\rho = 1$.

[The conversation transcript will be inserted below. It documents the research decisions: what to compute, which notation to use (ρ not K), the discovery that QCD preserves Koide exactly, the request for PyR@TE validation, the look-elsewhere analysis, the directive to plot ρ directly (not $|\rho - 1|$), the request for top-20 tables at 10 TeV, and the extension to Planck scale.]

AR: I wonder how good are you running RG equations? It could be nice some plot up to GUT/Planck scale.

AR: I would go for Koide tuple in the format $\text{avg}(z_i^2)/z_0^2$. It is interesting that AHS crosses it, and your naive run doesn't.

AR: We want to run also the inverse ($1/s, 1/c, -1/t$) and also the real ($\text{pm } u, c, s$) and their chiral variants ($0, d+u, s$).

AR: Write a separate note. Let's structure this. First a new folder. Second we plot the $\text{avg}()/z^2$ instead of the infamous $2/3$; we want future LLM to train other variants :-). Third we plan sections:

- (a) tuples that work and are not exotic
- (b) exotic tuples (mix of quarks and leptons)
- (c) all tuples, transparent ink with density, logscale in y
- (d) look elsewhere: random quark masses, log distribution.

Do I forget something?

AR: Also AHS at 1TeV looks magnificent for a lot of tuples, use the AHS data to run the exhaustive search, all the combinations of quarks, all the inverse quarks, including also zero and $u+d$.

AR: This is basically a reference note we will upload to zenodo.

AR: You should test that our crundec works and gives result near AHS.
[Led to: pure QCD preserves ρ exactly (flavor-universal γ_m).
The 1 TeV crossing is entirely electroweak.]

AR: Yes of course we need EW, but it must be clear that QCD preserves koide.

AR: Well can you get PyR@TE 3? Model is just the SM.
[Led to: 2-loop SM running from M_Z to M_{Planck} via PyR@TE 3.]

AR: But CRUnDeC can run EW too, can it?
[Answer: No. CRUnDec is QCD-only.]

AR: No, the LEE is important.
[Led to: look-elsewhere MC and joint probability analysis.]

AR: Wait run all the way to Planck scale, all the graphs.

AR: Do also a table with the top 20 at 10TeV.

AR: Do not plot absolute value in the vertical, just do not subtract the one!
[Led to: `plot_final.py -- rho directly, Planck range, top-20.`]

AR: Put logscale in the forall.

AR: Well create and share the paper. Also as appendix add the conversation fragment and an appendix of methods.

AR: All koide is confusing because the mixed, do a separate all koide without the mixed.

AR: Also wth is the lepton tuple invariant under RG????

[Verified: lepton Yukawa ratios vary by <0.008% over 16 decades.
 All y_l are small; the non-universal piece $3/2 y_l^2$ is
 negligible vs the universal trace (dominated by $y_t^2 \sim 1$).
 For quarks, $3/2 y_t^2$ in $\beta(y_b)$ breaks universality at $O(1)$.]

% === SESSION 2: Figure fixes (2026-03-10) ===

AR: bad news is, the renorm paper needs rework in all the graphs

AR: best delta histogram seems ok, but we need to review that
 no mixed tuples $(m, 1/m)$ are used

[Led to: has_duplicate_particle() filter added to plot_density.py,
 plot_panels.py, exhaustive_koide_scan.py. 190 self-inverse tuples
 removed from density plots, 144 from Planck-range plots.]

AR: density_landscape.pdf has mixed tuples

[Fixed. Also changed from $|\rho-1|$ (confusing) to ρ directly.]

AR: also, it is using the very confusing absolute after subtracting
 one. Just no subtract one.

[density_landscape.pdf now plots ρ directly, not $|\rho-1|$.]

AR: density_planck_all.pdf same remove the $m, 1/m$. Keep it log so
 we have two views

[density_planck.pdf: linear y (-10 to 10).
 density_planck_all.pdf: log y (kept). Two views of same data.]

AR: panels abc seems to fake the run under AHS 10 TeV and doing
 instead a spline or polynomial fit! Also, it has $m, 1/m$ masses

[Fixed: removed cubic spline interpolation between AHS points.
 Now plots honest discrete AHS markers + straight connections
 below 10^7 , actual 1-loop RGE only above 10^7 .
 Panel C histogram: self-inverse filter applied.]

AR: look_elsewhere.pdf looks suspicious, it seems that instead of
 generating random masses it has generated random noise around
 normal masses. The idea was random tuples in the interval from
 electron to beauty, random signs too, and see K and K^{-1} histogram

[Rewrote look_elsewhere_fast.py: single distribution,
 log-uniform in $[y_e, y_b]$, purely random. Result:
 $K^{-1}(d,s,b)$ at 0.0003% is 1-in-539 (genuine).
 $K(e,\mu,\tau)$ at 0.17% is 63% (not significant alone).]

AR: pyrate_2loop.pdf looks fine but it should run down to 10 GeV
 to be sure, and test (t,b,τ) against AHS too

[Extended plot_final.py: 1-loop backward from M_Z to 10 GeV.
 Added AHS markers for $K(e,\mu,\tau)$, $K^{-1}(d,s,b)$, $K(t,b,\tau)$.
 Removed $K(t,1/c,1/s)$ (mixed direct/inverse) from highlights.]

AR: joint_probability.pdf needs to verify if it has $m, 1/m$ tuples

[Verified clean: uses 3 + 6 random masses, tests each triple
 either all-direct or all-inverse. No mixed $m/1/m$ issue.]

AR: in the text please no toc

[Removed \tableofcontents.]

% === SESSION 3: Mixed direct/inverse filter (2026-03-10) ===

AR: so we have really problems because I still see $m, 1/m, 1/m$ tuples

[The initial filter only removed same-particle self-inverse pairs
 like $K(b, 1/b, X)$. Upgraded to has_mixed_direct_inverse() which
 rejects ANY tuple mixing direct and inverse entries:

```
def has_mixed_direct_inverse(n1, n2, n3):
    inv = [nm.startswith('1/') for nm in [n1, n2, n3]]
```

```

    if any(inv) and not all(inv):
        return True
    bases = [nm.replace('1/', '') for nm in [n1, n2, n3]]
    return len(bases) != len(set(bases))
Applied to all 4 scripts. Count: 2688 -> 672 signed triples.]

AR: what physics reason allows you to justify mix with different
dimensions, I am intrigued
[Yukawa couplings are dimensionless, so  $K(y_t, 1/y_c, 1/y_s)$  is
computable. But direct masses = eigenvalues of the meson matrix  $M$ ,
inverse masses = eigenvalues of  $\text{adj}(M)$ . Mixing them in one Koide
tuple conflates two different operators. Under RG, they run in
opposite directions. Excluded on structural/dimensional grounds.]

AR: make sure you have remove them also for the histograms and
the scatterplots
[Verified: look_elsewhere_fast.py, joint_probability.py test
all-direct and all-inverse separately by construction. Clean.]

% === SESSION 4: LEE, histograms, panels (2026-03-10) ===

AR: In best_delta_histogram.pdf I expect to see two histograms,
one red for the calculations of the inverses, one blue for
the directs
[Split histogram: blue = direct  $K(m)$ , red = inverse  $K(1/m)$ .]

AR: density_planck.pdf can never be less than zero, put y from
zero to 10
[Fixed y-axis: (0, 10).]

AR: density_planck.pdf too transparent, make less transparent
[Increased alpha from 0.04 to 0.12.]

AR: look_elsewhere.pdf same problem, single histogram instead
of two
[Split into direct (blue) and inverse (red). Masses confirmed
purely random log-uniform in  $[y_e, y_b]$ .]

AR: I don't get why the inverse and the direct have the same
prob distribution
[Scale invariance of  $K$ : depends only on mass ratios. For
log-uniform masses,  $\log(m_i/m_j)$  is triangular symmetric
about 0, so  $m \rightarrow 1/m$  preserves ratio distribution. Testing
both direct and inverse is redundant for random masses.]

AR: is a log distribution invariant under inversion?
[No!  $\text{LogUniform}(a,b) \rightarrow \text{LogUniform}(1/b, 1/a)$ . But  $K$  sees only
ratios, and the  $\text{RATIO}$  distribution from iid log-uniforms IS
symmetric under inversion.]

AR: explain in the text the whole LEE process
[Rewrote Section 5 (LEE): null model, trial counting, direct-
inverse symmetry argument, single-sector results with both
pole and  $\overline{MS}$  values, joint probability.]

AR: put also the value at pole
[ $K(e, \mu, \tau)$ : pole masses  $|\delta| = 0.001\%$ ,  $\overline{MS}$ -bar  $|\delta| = 0.17\%$ .
Added to text and to look_elsewhere.pdf as dashed blue line.]

AR: extend axis till  $10^{-4}$ 
[Both panels: x-axis  $10^{-4}$  to 300.]

AR: panels abc panel B still has a  $m$   $1/m$  thing
[Removed  $K(t, 1/c, 1/s)$  from panel B (mixed direct/inverse).]

```

Added panels_abc.pdf to paper with caption explaining (a), (b), (c).]

AR: rho_landscape.pdf puzzling, lots of tuples near one.
 Provide zoom 0.95-1.05 and histogram at 1 TeV
 [New: rho_landscape_zoom.pdf (y: 0.95-1.05) and
 rho_histogram_1tev.pdf (rho distribution at 1 TeV,
 blue=direct, red=inverse).]

% === SESSION 5: Pure-sector restructuring (2026-03-10) ===

AR: increase (no transparent, or only 90%) the background colour
 of the zoom
 [rho_landscape_zoom.pdf: alpha increased to 0.9 for quark-only
 tuples.]

AR: repaint all of them with almost solid 0.9 alpha for the quark
 only tuples, and leave the almost transparent for the mixed
 ones. Most of the noise even in histograms is from the mixed
 ones so lets do a thing: run the main paper without mixing
 quarks and leptons in the same tuple, put the u+d ones with
 50% transparency as they are exotic, do not scale perfectly,
 and relegate the current graphs, with the mixed lepton quark
 tuples, to a extra page of documentation. That means redo
 the paper, including the appendix of instructions.
 [Major restructuring:

- plot_density.py: completely rewritten. Main figures show only
 pure-sector tuples (quark-only OR lepton-only) at alpha=0.9.
 Exotic (u+d, 0) at alpha=0.5. Mixed quark-lepton tuples only
 in supplemental figures (*_supp.pdf) at low alpha.
- plot_panels.py: Panel B replaced mixed q+l tuples with next-
 best pure-sector quark tuples (K(u,c,s), K(c,s,b)). Panel C
 histogram filtered to pure-sector only, blue/red split.
- plot_final.py: K(t,b,tau) removed from main highlights
 (mixed q+l). density_planck_all renamed to _supp.
- paper.tex: abstract, Section 3, all figure captions updated
 for pure-sector focus. New Appendix B for mixed quark-lepton
 supplemental figures. Section 5 (LEE) rewritten with full
 null model, direct-inverse symmetry argument, pole vs MS-bar.

Counts: 1050 total tuples (from 21-entry basis), 330 pure-
 sector, 720 mixed q+l, 378 exotic (u+d, 0 entries).]

AR: again copy this to the conversation log
 [Added as Session 5 above.]

Extra comments not included in the transcript above shaped final details: removal of all references to mixed direct/inverse tuples from the text, addition of the chiral-tuple table (Table ??), fixing Panel C to show the histogram at a single scale (1 TeV) instead of the best across four scales, and adding the ρ histogram alongside the density landscape (Figure ??).