Open problems in the gauged family symmetry programme

Predictive content, obstructions, and paths forward

Working notes for internal discussion

March 2026

Contents

1 The framework in brief

The model is built on three postulates:

(i) An $\mathrm{SU}(3)_F$ family gauge symmetry confines at a scale $\Lambda_F \sim \mathcal{O}(100)\,\mathrm{MeV}$, with $N_f = N_c = 3$ flavours of preons Q^i, $\tilde{Q}_i$.

(ii) The confined spectrum is described by Seiberg's s-confinement: the light degrees of freedom are mesons $M_{ij} = Q^i\tilde{Q}_j$ and baryons $B = \epsilon_{ijk}Q^iQ^jQ^k$, subject to the quantum-modified constraint $\det M - B\tilde{B} = \Lambda^6$.

(iii) The mass ansatz $m_i = \mu\, q_i^2$ (direct branch, from adjugate $\mathrm{adj}(M)_i$) and $\hat{m}_i = \nu/q_i^2$ (inverse branch, from meson M_i), where $q_i = z_0 + z_i$ are the family charges with $z_0 = 1/\sqrt{6}$, $\sum z_i = 0$, $\sum z_i^2 = 1/2$.

The Koide identity $K = 2/3$ for the direct branch is a trace identity:

$$K \;=\; \frac{\sum q_i^2}{(\sum q_i)^2} \;=\; \frac{1}{3/2} \;=\; \frac{2}{3}\,, \tag{1}$$

valid at any direction α in the Cartan plane. The inverse branch does *not* automatically satisfy $K = 2/3$; the empirical relation $K^{-1}(d,s,b) \approx 2/3$ is a non-trivial constraint.

2 Parameter counting and predictive scorecard

2.1 The Standard Model target

The SM has:

- 9 charged-fermion masses: $m_e, m_\mu, m_\tau, m_u, m_d, m_s, m_c, m_b, m_t$
- 4 CKM parameters: $\theta_{12}, \theta_{23}, \theta_{13}, \delta_{\mathrm{CP}}$
- 3 neutrino masses + 4 PMNS parameters (not addressed here)

Total for charged fermions: **13 free parameters**.

2.2 The model's free parameters

Table 1: Free parameters of the framework.

#	Parameter	Fixed from	Value
1	α_{lep}	m_τ/m_μ ratio	0.2222 rad
2	k $(= \mu_{\mathrm{lep}})$	m_τ overall scale	1883 MeV
3	ν	$m_t \cdot q_1^2$	46.8 MeV
4	$\Delta\alpha$	$\arctan(\pi/84)$	0.0374 rad
5	ε (quartic)	*not yet computed*	open

With 4 parameters for 13 observables, the framework claims **9 predictions**. With the quartic ε added as a 5th parameter, the count is 8. The question is how many of these predictions are currently *working* at quantitative precision.

2.3 Scorecard

Exact / proven predictions (0 additional parameters):

1. $K(e,\mu,\tau) = 2/3$ — trace identity, empirically 0.001%.
2. Interleaving order $d_1 > a_3 > d_2 > a_2 > d_3 > a_1$ — proven theorem.
3. $a_i \cdot d_i = \mu\nu$ family-independent — from $\det M = \Lambda^6$.
4. Top = inverse of lightest charge (seesaw structure).
5. Two Higgs doublets from one meson matrix ($H_u = \mathrm{adj}(M)$, $H_d = M$).

Table 2: Quantitative predictions.

Observable	Predicted	PDG 2024	Status
m_e	0.501 MeV	0.511 MeV	2% (from α, k)
m_d	4.59 MeV	4.67 ± 0.48 MeV	1σ (from $\Delta\alpha$)
$K^{-1}(d,s,b)$	$\approx 2/3$	0.667	0.3%
$m_c + m_s$	1084 MeV	1366 MeV	21% tension (unmixed)

Quantitative predictions (from α, k, ν, $\Delta\alpha$):

Table 3: CKM predictions. The self-dual ratio $R = (2+\sqrt{3})^2 \approx 13.928$ satisfies $R^2 - 5R + 1 = 0$.

Observable	Formula	Predicted	PDG 2024	Accuracy
$\lvert V_{cb}\rvert$	$1/(5R)$	0.0406	0.0408	1.1%
$\lvert V_{ub}\rvert$	$2\sqrt{2}/(7R^3)$	0.00355	0.00361	0.06σ
δ_{CP}	$\arctan\sqrt{R}$	75.0°	$73.4 \pm 3^\circ$	0.2σ
J	closed form	3.18×10^{-5}	3.08×10^{-5}	0.09%

CKM predictions (from $R = (2+\sqrt{3})^2$, integers 5, 7 unexplained): These are impressive numerical matches, but the integers 5 and 7 are observed patterns, not derived from the gauge theory. They may encode group-theoretic data (dimensions of representations, Clebsch–Gordan coefficients), but this has not been established.

Not yet predicted (need ε or Kähler metric):

- Individual m_c, m_s (only their sum T_2 is constrained).
- m_u (lightest window, maximally sensitive to mixing).
- PMNS angles (tree-level = identity matrix; need Kähler rotation).
- Neutrino masses (seesaw structure exists but quantitative prediction requires the baryonino mass spectrum).

Structural / qualitative predictions:

- Leptons = composite mesinos: $9 = 3L_i + 3e_i^c$ (anomaly-free).
- SUSY partners near $\mu_{\text{crit}} \approx 4.5\,\text{TeV}$.
- Strong CP solved: involution $M_i \leftrightarrow \Lambda^6/M_i$ acts as CP, forcing $\bar{\theta} = 0$ at tree level.
- Colour sextets absent: antisymmetric obstruction ($\Lambda^3(9) = 84$ contains only 1_C, 3_C, $\bar{3}_C$, never 6_C).

Confirmed negative results:

- α_{lep} is NOT selected by: finite μ, Coleman–Weinberg potential, $N_f = 4$ threshold corrections, or instantons.
- Seiberg–Witten monodromy does NOT produce $\sqrt{21}$. R comes from the family Higgs, not from SW geometry.
- Tree-level baryonino mass matrix is always diagonal $\Rightarrow$ PMNS $\approx \mathbf{1}$ without Kähler corrections.
- Hessian eigenvalues give $K = 0.495$, not $2/3$. SUSY-breaking is essential for the physical mass spectrum.

3 Open Problem 1: The quartic coupling ε

Open Problem 1 (Quartic mixing)**.** *Compute the inter-branch coupling ε that redistributes the unmixed eigenvalues a_i, d_i into the physical quark masses.*

3.1 Why ε is needed

Each family i has two mass parameters:

$$a_i = \mu q_i^2 \quad (\text{direct})\,, \qquad d_i = \nu / q_i^2 \quad (\text{inverse})\,. \tag{2}$$

In the unmixed limit, the 6 masses form an interleaved ladder:

$$d_1 > a_3 > d_2 > a_2 > d_3 > a_1\,. \tag{3}$$

Evaluating the Koide ratio $K = 2/3$ on sliding windows of three consecutive masses gives a quadratic equation in $r = \sqrt{\nu/\mu}$ for each window. There are **four windows but only one free parameter** r, so the system is overconstrained. Explicitly:
The best one can do with a single r is $\max_w |\delta K_w| \sim 7\%$—an order of magnitude worse than the observed $\sim 1\%$ deviations. The constraint surface alone (charge normalisation + $\det M = \Lambda^6$) is therefore **insufficient**.

Table 4: Quadratic solutions $r = \sqrt{\nu/\mu}$ for $K = 2/3$ in each window. No single r satisfies all four. Physical $r = 0.1028$.

Window	r_1	r_2
$W_1 = (d_1, a_3, d_2)$	0.0904	0.00390
$W_2 = (a_3, d_2, a_2)$	1.141	0.00390
$W_3 = (d_2, a_2, d_3)$	3.306	0.01132
$W_4 = (a_2, d_3, a_1)$	0.9436	0.04076

3.2 What the quartic can and cannot do

The inter-branch coupling is parametrised by a 2×2 family Hamiltonian:

$$H_i = \begin{pmatrix} \mu q_i^2 & \varepsilon \\ \varepsilon & \nu/q_i^2 \end{pmatrix}, \tag{4}$$

whose eigenvalues are the physical masses within each family. A key structural feature: $\det(H_i) = \mu\nu - \varepsilon^2$ is **family-independent**. The quartic redistributes the sum $T_i = a_i + d_i$ between the two eigenvalues without changing their product (up to ε^2 corrections).

Scanning ε from 0 to $\sqrt{\mu\nu} = 455\,\text{MeV}$:

- The optimal $\varepsilon \approx 420\,\text{MeV}$ gives $\min\max|\delta K| = 6.7\%$.
- For the charm–strange window (W_2), $\varepsilon = 340\,\text{MeV}$ gives m_s correct but m_c 22% off.
- The quartic **cannot close the gap alone**: the overconstrained system has no solution in a single parameter.

3.3 What else is needed

The companion letter's quark-sector angle shift $\Delta\alpha = \arctan(\pi/84) = 0.0374$ rad modifies *all* the charge ratios. This is the second ingredient. Together, ε and $\Delta\alpha$ provide 2 parameters for 4 constraints—generically solvable (2 equations residual), but the computation has not been done.

3.4 Why the computation is hard

The quartic (**??**) is a proxy for the full non-perturbative dynamics. In the UV description, the coupling arises from integrating out the auxiliary adjoint field Φ of the emergent $\mathcal{N}$=2 structure:

$$W_{\text{quartic}} = -\frac{g^2}{\mu_\Phi}\left[\text{Tr}M^2 - \tfrac{1}{3}(\text{Tr}M)^2\right], \tag{5}$$

where μ_Φ is the adjoint mass and g is the family gauge coupling. At the confinement scale $\Lambda_F \sim 100\,\text{MeV}$, the gauge coupling is strong ($g \sim \mathcal{O}(1)$), so:

1. **Perturbation theory fails**. The ratio g^2/μ_Φ is not a small parameter; the quartic is not a perturbative correction.
2. **The $\mathcal{N}$=2 breaking scale μ_Φ is non-perturbative**. The emergent doublet structure arises from confinement plus isospin, not from a UV $\mathcal{N}$=2 theory that is subsequently broken. The "μ_Φ" is therefore not a Lagrangian parameter but an effective scale determined by the confined dynamics.

3. **Holomorphic vs. physical masses**. The F-term equations give holomorphic masses (superpotential parameters). The physical masses include Kähler corrections $m_{\text{phys}} = m_{\text{hol}}/\sqrt{Z_i}$, where Z_i is the wave-function renormalisation. For leptons, Z_i is generation-universal at leading order (0.04% effect). For quarks, $Z_c/Z_s \approx 1.37$ is needed to match the lattice value $m_c/m_s = 11.78$, compared to the holomorphic prediction of 13.93 (an 18% discrepancy, 86σ).

3.5 Promising paths

(a) **Seiberg–Witten curve for** SU(3)**,** $N_f = 3$. The exact prepotential determines the low-energy effective action including the Kähler metric on the Coulomb branch. However, the physical theory is in the *confining* phase (Higgs branch), not the Coulomb branch. The Seiberg–Witten solution applies only after an analytic continuation whose validity is unclear.

(b) **Lattice family gauge theory**. A lattice computation of SU(3) with 3 fundamental flavours and an adjoint mass term could in principle extract ε from the meson correlator. The obstacle is that the family gauge theory is *chiral* (the preons carry SM quantum numbers), and chiral lattice gauge theories remain technically difficult.

(c) **Large-N extrapolation**. At large N_c, the meson–adjugate mixing scales as $1/N_c$. The physical case $N_c = 3$ is not large, but the scaling may give a useful estimate.

(d) **Deformation theory**. Start from the $\mathcal{N}$=2 point ($\mu_\Phi = 0$, exact SW solution) and deform by soft breaking terms. The Dijkgraaf–Vafa matrix model computes the effective superpotential to all orders in the gluino mass. The obstruction: the DV computation gives the *superpotential*, not the Kähler potential; physical masses still require Z_i.

(e) **Numerical bootstrap**. Use the known constraints (interleaving, det $H_i = \mu\nu - \varepsilon^2$, $K \approx 2/3$ in each window) to fit ε and $\Delta\alpha$ simultaneously. This is a 2-parameter fit to 4 constraints (the waterfall windows) plus 2 constraints (m_c/m_s and the sum rule). It sidesteps the non-perturbative computation but does not explain ε from first principles.

4 Open Problem 2: The Kähler metric

Open Problem 2 (Kähler metric)**.** *Compute the Kähler metric Z_i on the meson moduli space in the confined phase. This is the master obstruction: every quantitative prediction involving quark masses depends on it.*

4.1 Why it matters

The holomorphic masses (F-term, superpotential) are calculable by Seiberg duality. The physical masses are:

$$m_{\text{phys, i}} = \frac{m_{\text{hol, i}}}{\sqrt{Z_i}}, \tag{6}$$

where Z_i is the Kähler wave-function renormalisation. For the lepton sector, the leading-order Kähler metric is $K = M^\dagger M/\Lambda^2$ (canonical normalisation), giving $Z_i \propto 1/\Lambda^2$ which is flavour-universal. This is why the lepton Koide relation is so precise (0.001%): the Z_i cancel in the ratio K.

For quarks, the situation is different. The holomorphic prediction $m_c/m_s = 13.93$ disagrees with the lattice value 11.78 ± 0.025 by 18% (86σ). Since the quark mass ratio is *flavour-independent* under RG (to 5-loop in QCD), this discrepancy must come from flavour-dependent Kähler corrections:

$$\frac{Z_c}{Z_s} \approx 1.37\,. \tag{7}$$

This is a 37% effect, far from the naive expectation of universality.

4.2 Why it's hard

1. The Kähler potential is **not holomorphic** and therefore not protected by non-renormalisation theorems. It receives corrections at every order in perturbation theory and non-perturbatively.

2. In the M-theory description, the Kähler metric on the meson moduli space is related to the DBI action of the M5-brane wrapping the Seiberg–Witten curve. For the $\mathcal{N}$=1 case (genus-0 curve), the DBI analysis of Dasgupta–Hsu–Ooguri–Oh (DHOO) does not apply: their results require $\mathcal{N}$=2 or higher-genus curves.

3. The canonical Kähler metric $K = M^\dagger M/\Lambda^2$ gives $v_{\mathrm{EW}} = 3.09\,\mathrm{GeV}$, two orders of magnitude below the physical value $v = 246\,\mathrm{GeV}$. The needed correction is $Z_{\mathrm{needed}}/Z_{\mathrm{canonical}} \approx 6328$. This large ratio suggests that the canonical normalisation misses important strong-coupling effects.

4.3 What's known

The *overall scale* $k = \mu \approx 1883\,\mathrm{MeV}$ is fitted, not derived. Its ratio to the confinement scale is $\mu/\Lambda_F \approx 24$, which is a non-perturbative number. The fact that $\mu \gg \Lambda_F$ indicates that the Kähler metric enhances the mass scale relative to the naive confinement estimate.

The *relative* Kähler metric (ratios Z_i/Z_j) is what enters mass ratios. For leptons it is approximately universal; for quarks it is not. The quartic mixing ε and the Kähler non-universality Z_c/Z_s are entangled: one cannot determine ε without knowing the Kähler metric, and vice versa.

4.4 Promising paths

(a) **Instanton corrections to Kähler**. On the Coulomb branch, instanton corrections to the Kähler metric are known (e.g., Dorey–Hollowood–Khoze–Mattis). Analytic continuation to the confining phase might give the leading flavour-dependent correction.

(b) **Anomaly-mediated Kähler corrections**. The conformal anomaly determines the leading log-dependent piece of the Kähler metric. This is universal in flavour but might explain the overall scale factor $\mu/\Lambda \approx 24$.

(c) **Superconformal index**. The superconformal index of the UV theory constrains the spectrum of protected operators, which in turn constrains the Kähler metric at special points in moduli space.

(d) **Accept it as a fitted parameter**. The pragmatic approach: treat Z_c/Z_s (or equivalently ε) as a phenomenological parameter and focus on the predictions that are Kähler-independent (ratios within the same sector, the Koide identity, the interleaving theorem, the CKM predictions).

5 Open Problem 3: The origin of α

Open Problem 3 (Cartan angle)**.** *Derive the lepton mixing angle* $\alpha_{\text{lep}} \approx 0.2222$ rad *from the dynamics of the family gauge theory.*

5.1 What α is

The three family charges are:

$$q_i \;=\; z_0 + \frac{1}{\sqrt{3}}\cos\!\Big(\alpha + \frac{2\pi k}{3}\Big)\,, \qquad k = 0, 1, 2\,, \tag{8}$$

with $z_0 = 1/\sqrt{6}$. The angle α parametrises the direction in the Cartan subalgebra of $\mathrm{SU}(3)_F$. The Koide relation $K = 2/3$ holds for *any* α; the specific value $\alpha_{\text{lep}} = 0.2222$ is what determines the mass *ratios* $m_\tau : m_\mu : m_e$.

The notation "2/9" has been used in the literature as a mnemonic, but a rational angle in radians would make $\cos\alpha$ transcendental by Lindemann–Weierstrass. The value is fitted, not exact.

5.2 What has been ruled out

1. **Finite-μ selection**: The superpotential F-flatness conditions do not select α — the moduli space is an entire circle. (Notes 62–64.)

2. **Coleman–Weinberg potential**: The one-loop CW potential on the Cartan direction is *monotonic* in α and cannot produce a non-trivial extremum. (Note 63.)

3. $N_f = 4$ **threshold effects**: Adding a fourth heavy flavour and integrating it out does not select α. (Note 64.)

4. **Instanton effects**: The instanton-generated superpotential in SU(3) with $N_f = 3$ is the $\det M$ term, which is α-independent. (Note 69.)

5. $\mathcal{N}$=2 **Coulomb branch**: The Coulomb-branch analysis is inapplicable to the confining phase. (Note 116.)

5.3 Promising paths

(a) **Kähler modulus of the M5-brane curve**. In the M-theory lift, the family gauge theory corresponds to an M5-brane wrapping a curve Σ in the Calabi–Yau. The angle α is a modulus of Σ that is not fixed by holomorphy. It is fixed by the *Kähler* potential of the compactification—i.e., by the geometry of the internal space. This connects Problem 3 to Problem 2.

(b) **Discrete symmetry**. If α is fixed by a discrete subgroup of the Weyl group (e.g., $\alpha = \pi/6$ corresponds to the $\mathbb{Z}_3$ fixed point of the charge plane, giving $K = 2/3$ at $r_0 = (2+\sqrt{3})^2$), the deviation from $\pi/6$ might be a calculable perturbation. The observed $\alpha_{\text{lep}} = 0.2222$ is 5.5° from the $\pi/12 = 15°$ fixed point.

(c) **Anomaly cancellation at the compactification scale**. Mixed anomalies between $\mathrm{SU}(3)_F$ and gravity (or other gauge factors) might constrain the embedding of $\mathrm{SU}(3)_F$ in a larger group, fixing the Cartan direction.

6 Open Problem 4: The CKM integers

Open Problem 4 (CKM integers)**.** *Derive the integers 5 and 7 appearing in the CKM predictions $|V_{cb}| = 1/(5R)$ and $|V_{ub}| = 2\sqrt{2}/(7R^3)$ from the gauge theory.*

6.1 The pattern

The self-dual ratio $R = (2+\sqrt{3})^2 = 7 + 4\sqrt{3}$ satisfies $R^2 - 5R + 1 = 0$. The CKM predictions use R and small integers:

$$|V_{cb}| = \frac{1}{5R} = 0.01435\ldots \qquad (1.1\% \text{ from PDG})\,, \tag{9}$$

$$|V_{ub}| = \frac{2\sqrt{2}}{7R^3} = 0.001504\ldots \qquad (0.06\sigma)\,, \tag{10}$$

$$\delta_{\text{CP}} = \arctan\sqrt{R} = 75.04° \qquad (0.2\sigma)\,, \tag{11}$$

$$J = \frac{2\sqrt{2}}{5 \times 7^{5/4}} \sqrt{\frac{m_d(m_s - m_d)}{m_s}}\, R^{-15/4} \qquad (0.09\% \text{ from PDG})\,. \tag{12}$$

The integers 5 and 7 might be:

- Dimensions: $\dim(\mathrm{SU}(2)) + 2 = 5$? $\dim(\mathrm{Adj}(\mathrm{SU}(3))) - 1 = 7$?
- Clebsch–Gordan coefficients of the $\mathrm{SU}(3)_F$ representations.
- Artifacts of a deeper algebraic structure (the discriminant $\Delta = 21 = 3 \times 7$, and $5R = R^2 - 1$ from the minimal polynomial).

None of these identifications has been established.

6.2 What would count as a derivation

A satisfactory derivation would:

1. Start from the $\mathrm{SU}(3)_F$ gauge theory (or its M-theory lift).
2. Show that the CKM matrix arises from the misalignment between the direct-branch ($\mathrm{adj}(M)$) and inverse-branch (M) charge directions in the Cartan plane.
3. Derive the angle shift $\Delta\alpha = \arctan(\pi/84)$ and the specific R-dependent entries from group theory.

6.3 Promising paths

(a) $\pi/84$ **from the C_3 polarisation count**. $84 = \binom{9}{3}$ is the dimension of the 3-form representation of SO(9) in 11d SUGRA, and also appears in the E_8 branching $248 = 80 + 84 + \overline{84}$ under SU(9). If the angle shift comes from the M-theory embedding (specifically, from the C_3 coupling of the M5-brane), the 84 would be explained.

(b) **R from the charge-plane fixed point**. At the 15° fixed point of the Cartan plane, the zero-mass ratio is $r_0 = (2 + \sqrt{3})^2 = R$. This is a *derived* property of the charge normalisation. The CKM might probe this fixed point through the angle shift.

(c) **Monodromy of the Seiberg–Witten curve**. The monodromy matrices of the SU(3), $N_f = 3$ SW curve act on the period lattice. Their entries are integers determined by the intersection form. The integers 5 and 7 might appear as matrix elements or traces. So far, this has been checked and *does not work*: no product of monodromies gives $\sqrt{21}$ (note 68).

7 Open Problem 5: The non-perturbative $\mathcal{N}$=2 regime

Open Problem 5 ($\mathcal{N}$=2 breaking dynamics)**.** *Characterise the strong-coupling dynamics of the emergent $\mathcal{N}$=2 structure at the family confinement scale.*

7.1 The nature of the problem

The framework uses an $\mathcal{N}$=2 doublet structure that is *emergent*, not fundamental. The UV theory is $\mathcal{N}$=1 $\mathrm{SU}(3)_F$ SQCD. Upon confinement, the composite spectrum $(M_i, \mathrm{adj}(M)_i)$ has the quantum numbers of an $\mathcal{N}$=2 hypermultiplet, and the inter-branch coupling has the form of an $\mathcal{N}$=2-breaking adjoint mass.

The difficulty is that this "$\mathcal{N}$=2 breaking" is happening at the confinement scale itself, where the gauge coupling is $g \sim \mathcal{O}(1)$. Standard $\mathcal{N}$=2 breaking analyses (Argyres–Plesser–Shapere, or the Seiberg–Witten curve with a mass perturbation) assume that the breaking is a *small perturbation* of an exact $\mathcal{N}$=2 theory. Here the breaking is $\mathcal{O}(1)$.

7.2 What is known

1. **The superpotential is exact**. By holomorphy, the confined-phase superpotential $W = m_i M_i + (\text{quartic}) + \lambda(\det M - \Lambda^6)$ is exact to all orders. The quartic coefficient is the only unknown.

2. **The Kähler potential is not exact**. It receives perturbative and non-perturbative corrections. The physical masses $m_{\text{phys}} = m_{\text{hol}}/\sqrt{Z}$ depend on these corrections.

3. **The $\mathcal{N}$=2 limit is smooth but unphysical**. Setting $\mu_\Phi \to 0$ restores $\mathcal{N}$=2, and the SW solution gives the exact low-energy effective action. But this limit produces massless charged fermions (the doublet is unsplit), so the physical spectrum requires finite breaking.

4. **The deformation is controlled by one parameter**. The adjoint mass μ_Φ (or equivalently the quartic g^2/μ_Φ) is the single parameter controlling the $\mathcal{N}$=2 $\to$ $\mathcal{N}$=1 breaking. The task is to determine its value in the confined phase.

7.3 Promising paths

(a) **Dijkgraaf–Vafa matrix model**. The DV correspondence computes the exact effective superpotential of an $\mathcal{N}$=1 theory obtained by adding $W = \mathrm{Tr}W(\Phi)$ to $\mathcal{N}$=2 SQCD. For a quadratic $W = \mu_\Phi \Phi^2/2$, the matrix model is Gaussian and the computation is tractable. The result gives the gaugino condensate and the effective superpotential, but *not* the Kähler potential.

(b) **Nekrasov partition function**. The Ω-deformed partition function gives the prepotential to all instanton orders. The $\mathcal{N}$=2* theory ($\mathcal{N}$=2 + adjoint mass) has a known Nekrasov partition function. The physical Kähler metric could in principle be extracted from the $\epsilon_1, \epsilon_2 \to 0$ limit, but this is technically challenging.

(c) **Holographic dual**. In the Maldacena limit, the $\mathcal{N}$=2* theory has a known supergravity dual (Pilch–Warner flow). The meson spectrum and its Kähler metric are encoded in the fluctuation equations of the 5d bulk fields. The obstruction is that $N_c = 3$ is far from the large-N limit.

(d) **Phenomenological extraction**. Use the waterfall pattern as input: fit ε to the observed K values and check for consistency with the interleaving theorem and the $\det H_i = \mu\nu - \varepsilon^2$ constraint. This gives a "measurement" of ε that can be compared with any future first-principles computation.

8 Open Problem 6: PMNS and neutrino masses

Open Problem 6 (Neutrino sector)**.** *Derive the PMNS mixing angles and neutrino masses from the baryonino sector of the confined theory.*

8.1 The current status

In the framework, neutrinos are baryoninos: $B_i = \epsilon_{jkl} Q^j Q^k Q^l$, the baryonic composites of the family preons. Their Majorana mass matrix is diagonal at tree level (note 61):

$$M_{\mathrm{Maj}} \;=\; \mathrm{diag}(M_1, M_2, M_3)\,, \tag{13}$$

because the baryonino mass term $B_i \tilde{B}_j$ receives contributions only from the flavour-diagonal contractions. This gives $U_{\mathrm{PMNS}} = \mathbf{1}$ at tree level—no mixing.

The atmospheric and solar mixing angles ($\theta_{23} \approx 45°$, $\theta_{12} \approx 34°$) therefore require:

1. Kähler corrections that rotate the baryonino kinetic terms relative to the mass eigenstates, or
2. SUSY-breaking effects that generate off-diagonal Majorana entries.

Both mechanisms require the Kähler metric (Problem 2).

8.2 A suggestive coincidence

At the permutation point $\nu/\mu = q_1^2 q_2^2$, the "product charge" $p^* = q_1 q_2/q_3 = 0.004019$ gives:

$$\frac{(\mu\, p^{*2})^2}{\Lambda_F} = \frac{(0.072\,\mathrm{MeV})^2}{102\,\mathrm{MeV}} \approx 5 \times 10^{-5}\,\mathrm{MeV} = 0.05\ \mathrm{eV} \approx m_\nu(\mathrm{atm})\,, \tag{14}$$

where $\Lambda_F \approx 102\,\mathrm{MeV}$ is the family confinement scale. This is a single-number coincidence at a *non-physical* value of ν/μ ($\rho_{\mathrm{phys}}/\rho_{\mathrm{perm}} \approx 694$), but it suggests that the seesaw scale is naturally set by Λ_F.

8.3 What's needed

1. The baryonino mass spectrum as a function of α and Λ.
2. The Kähler metric on the baryonino moduli space (to compute the PMNS rotation).
3. The SUSY-breaking contribution to the off-diagonal Majorana entries.

All three require solving Problem 2 (Kähler metric) and Problem 5 (non-perturbative dynamics).

9 Summary: the dependency graph

The open problems are not independent. They form a dependency graph:

$$\begin{array}{rcl}
\textbf{Problem 2}\ \text{(Kähler metric)} & \longleftarrow & \text{master obstruction} \\
\downarrow & & \\
\textbf{Problem 1}\ (\text{quartic } \varepsilon) & & \text{requires Kähler to disentangle} \\
\downarrow & & \\
\textbf{Problem 3}\ (\alpha \text{ origin}) & & \text{Kähler modulus of M5 curve} \\
\downarrow & & \\
\textbf{Problem 4}\ \text{(CKM integers)} & & \text{needs } \alpha \text{ derivation + misalignment} \\
\downarrow & & \\
\textbf{Problem 5}\ (\mathcal{N}{=}2 \text{ dynamics}) & & \text{controls the Kähler metric} \\
\downarrow & & \\
\textbf{Problem 6}\ (\text{PMNS} + \nu) & & \text{requires everything above}
\end{array}$$

The most productive short-term strategy is **path (e) of Problem 1**: a numerical bootstrap that treats ε and $\Delta\alpha$ as free parameters, fits them to the waterfall data, and produces testable predictions for the quark mass ratios. This sidesteps the non-perturbative Kähler computation while extracting the maximal information from the framework.

The most ambitious long-term goal is **path (a) of Problem 3**: deriving α from the M-theory compactification geometry. This would collapse the parameter count from 4 to 3 (or fewer, if $\Delta\alpha$ is simultaneously determined), and would constitute a genuine derivation of the fermion mass hierarchy from string/M-theory.

	Problem	Short-term path	Long-term path
1	Quartic ε	numerical bootstrap	DV matrix model
2	Kähler metric	fitted	Nekrasov/instanton
3	α origin	—	M5 curve geometry
4	CKM integers	—	C_3 polarisation count
5	$\mathcal{N}$=2 dynamics	phenomenological	holographic dual
6	PMNS + ν	—	requires 1–5

The framework is predictive: 4 parameters for 13 observables, with ~ 8 *working at quantitative precision. The obstructions are specific (Kähler metric, non-perturbative* $\mathcal{N}$=2 *dynamics) and the paths forward are identifiable, if technically demanding. The model is not finished, but it is falsifiable: the interleaving theorem, the geometric mean constraint, the waterfall pattern, the SUSY threshold at* $4.5\,\mathrm{TeV}$*, and the absence of colour sextets are all testable predictions.*