A Charged-Lepton Ratio Law from Localization and Saturated Participation

A Convergence of Two Derivation Chains: Role-4 (CP^2) Localization and Closure-Architecture Participation

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The Standard Model contains three charged leptons — the electron, muon, and tau — with masses that differ enormously:

$m_\mu / m_e \approx 206.77$ $m_\tau / m_\mu \approx 16.82$

The Standard Model records these ratios through three independent Yukawa couplings, but supplies no reason for the particular numbers.

This paper's central claim is one of **convergence**. The charged-lepton masses follow from a single scaling law, $M_g = \tilde{E}_g / L_g^2$ (the Sector Rescaling Theorem), fed by two independently developed structures:

- a **geometric chain** — CP^2 geometric structure fixes the localization exponent $\kappa = 8/3$, so $L_g = L_0 e^{-\kappa g}$; and
- the **rescaled energies** $\tilde{E}_g$ from the Role-4 variational model, whose ratios decide which steps are pure geometry and which are suppressed.

For the electron–muon step the rescaled energies are nearly equal ($\tilde{E}_1 / \tilde{E}_0 \approx 0.999$), so μ/e is essentially pure geometry, $e^{16/3}$. For the muon–tau step the rescaled energy collapses ($\tilde{E}_2 / \tilde{E}_1 \approx 0.085$), and this *is* the tau suppression. The striking point is that the same suppression is produced **independently** by the closure ontology, where the K = 7 census of fourteen generators splits $14 = 2_\partial + 12_{int}$ to give a participation factor $\Pi_2 = 1/12 \approx 0.083$. Two unrelated derivations — a variational energy-collapse and a closure channel-count — land on the same number to within 1.6% (0.0846 vs 0.0833). The single structural statement they both express is this: *saturation leaves roughly one-twelfth of the mass-generating effectiveness available at the tau sector.* They may even be one object described at two levels — an energy ratio and a channel count — though that identity is a conjecture the agreement supports rather than proves.

The resulting law (numbers unchanged) is

$m_e : m\mu : m\tau = 1 : e^{16/3} : (1/12)\, e^{32/3}$,

where the 1⁄12 in the tau term is the suppression $\Theta_2 = \tilde{E}\_2 / \tilde{E}\_1$, read off the closure census. The electron–muon ratio follows to **0.17%**; the full hierarchy to about **3%**. The novelty is not the formula but the meeting points: geometry and the variational energies generate the spectrum, and the one non-geometric feature in it — the tau suppression — is confirmed by a second, independent ontological route.

A word on status, in the programme's house sense. Every input is now **derived, not fitted**: κ from CP^2 geometry, the scaling law from the Sector Rescaling Theorem, the rescaled energies from the Role-4 variational model, and the tau suppression independently from the closure census. But *derived does not mean proven* — each rests on a chain with conditional steps (the Sector–Curvature Correspondence for κ; the variational model's capacity inputs for the $\tilde{E}\_g$; the boundary/interior split and surviving-channel count N_survive = 1 for the closure route). One thing is now retired: the earlier worry that the second-step energy ratio was an unverified assumption set to 1. The variational model *computes* $\tilde{E}\_2 / \tilde{E}\_1 \approx 0.085$ directly, and it coincides with the closure 1⁄12. So this is the programme's first concrete charged-lepton engineering relation — quantitative, falsifiable, and resting on a doubly-derived tau suppression — with two open items: confirming that the CP^2 κ carries no generation index (which lifts the full law to [Conditional]), and reconciling the two derivations of the suppression (variational 0.0846, closure 0.0833) against the observed 0.0812.

The relation is also positioned against the empirical Koide relation (§9): it is a different *kind* of claim — a structural reconstruction, not an algebraic coincidence — and, tellingly, it does **not** reproduce Koide exactly, which converts that comparison into a quantitative cross-check rather than a rhetorical one.

Contents

The Convergence in Plain Terms

The Standard Model contains three charged leptons — the electron, the muon (about 207 times heavier), and the tau (about 3,500 times heavier) — and offers no reason for those particular numbers. This paper assembles a reason from two lines of work that were never designed to meet.

Start with one idea: a particle's mass is tied to how tightly it is localized. Squeeze a particle into a smaller region and it gets heavier — specifically, mass goes as one over the size squared. So if each generation of lepton sits in a tighter region than the last, the masses climb.

The first chain — geometry. How much tighter each step gets is set by a single number, κ. That number is not chosen to match the leptons; it falls out of the geometry of a particular curved space (called CP^2), which fixes $\kappa = 8/3$. Feed that into the size-to-mass rule and the muon comes out 207 times the electron — matching the measured value to about one part in six hundred, with nothing adjusted. The electron-to-muon step is, in effect, pure geometry.

The tau problem. Apply the same geometric squeeze a second time and the tau comes out more than ten times too heavy. So something must hold the tau back. The interesting question is what, and whether it can be derived rather than fitted.

The second chain — and a coincidence that probably isn't one. Two quite different calculations, from different parts of the programme, were each asked how strongly the tau is suppressed. They gave:

an energy calculation: about 1⁄11.8 (the tau sector's effective energy collapses near a binding threshold) a counting calculation: exactly 1⁄12 (of twelve interior "channels," only one survives at the tau)

These are not similar stories. One is about energy draining away as a threshold is crossed; the other is about counting how many channels remain. Yet they land on almost the same number — they differ by about 0.0013, under two percent. The measured suppression is close to both, about 1⁄12.3.

The natural conclusion is that these are not two separate effects to be multiplied together, but **one suppression described in two languages** — an energy ratio and a channel count — that may turn out to be the same underlying thing. The single fact both express can be stated in one line:

Saturation leaves roughly one-twelfth of the mass-generating effectiveness available at the tau.

That two independent routes find the same one-twelfth is the heart of the paper. A number that has only been fitted appears once; a number that is real tends to show up by more than one road.

Putting it together. The three masses then follow a single law,

$m_e : m\mu : m\tau = 1 : e^{16/3} : (1/12)\, e^{32/3}$, reproducing the electron–muon ratio to 0.17% and the full spread to about 3%. Every number in it is derived — κ from geometry, the one-twelfth from the two converging routes — rather than tuned to the data.

What is honest to say, and what is not. "Derived" is not the same as "proven." Each route still rests on a step or two that the programme has yet to close, so the law's standing is conditional, not settled. And the last few percent — the gap between the derived ~1⁄12 and the measured 1⁄12.3 — is a refinement question, not a crack in the structure: the size of the effect is now fixed by two independent arguments, and only its final digit is open. That is a far stronger place to stand than having a single route to the number.

---

1. Inherited Inputs (Named Imports and Structural Assumptions)

In plain terms. This paper does not build everything from scratch. It borrows a small number of results from earlier parts of the programme and combines them in a new way. This section lists exactly what is borrowed and — crucially — labels how secure each borrowed piece is. The two key numbers are *derived* from earlier work, not guessed; but "derived" still ranges from "derived along a chain that is fully nailed down" to "derived along a chain with a step or two left to prove," and the labels say which. That matters, because a conclusion can only be as solid as its shakiest ingredient, so everything is laid on the table before any sums are done. Two of the items are not results at all but quiet *assumptions*, pulled into the open here so a reader can see them.

Each input below carries an epistemic grade and an explicit fallback. By inheritance-cap discipline, no downstream result is graded above the weakest input it depends on (§9).

⟦Import L1 — Localization Hierarchy from CP^2 Geometry⟧

In plain terms. Each heavier lepton sits in a tighter region than the one before, by a fixed factor. The size of that factor is set by a single number, κ. The point that matters here is that κ is not chosen to fit the muon — it falls out of the geometry of a particular curved space (CP^2).

Statement: $L_g = L_0\, e^{(-\kappa g)}$, with $\kappa = 8/3$ and $g \in \{0, 1, 2\}$.

Derivation (geometric chain). CP^2 geometry fixes a holonomy count K_hol = 4; the curvature ratio for CP^n at n = 3 is $(n-1)/n = 2/3$; the effective sector count is therefore $K\_eff = K\_hol \cdot (n-1)/n = 4 \times 2/3 = 8/3$, identified with κ through the Sector–Curvature Correspondence. Equivalently $\kappa = \ln(\varepsilon\_loose / \varepsilon\_spread) = 8/3$. The chain is:

CP^2 geometry → K_hol = 4 → K_eff = 4 × (2⁄3) = 8⁄3 → κ = K_eff → $L_g = L_0\, e^{(-\kappa g)}$.

Source: *Toward a Lepton-Sector Mass Derivation*, §6.4 ("κ Emerges from CP^2 Geometric Structure"); κ = 8⁄3 is the value used throughout the Threshold-Compression and tau-suppression sequence.

Grade: **[Conditional]** — κ = 8⁄3 is *derived from CP^2 geometry, not fitted to any lepton mass*. This retires the earlier back-solving concern: the exact κ reproducing μ⁄e is 2.66580, and the geometric value 8⁄3 = 2.66667 sits 0.0325% above it, so the 0.17% in the muon ratio is the residual between an independently derived geometric value and observation — a genuine postdiction, not a circular fit. Conditionality attaches to the chain, principally the Sector–Curvature Correspondence that identifies $\kappa = K_{eff}$.

Note (generation-independence). The derivation $K_{eff} = 4 \times 2/3$ carries **no generation index** — it is a fixed geometric quantity. If this single value is the localization exponent at *every* step rather than only the first, then ⟦Assumption U⟧ is discharged by the geometric chain itself and the running-κ alternative of §7 is excluded upstream. This is worth confirming explicitly in the Role-4 source, as it would shift the binding constraint off U.

Note (two κ routes). A separate closure-based route yields κ ≈ ln(14) ≈ 2.639 (*Structural Completeness in the Role-4 Lepton Sector Model*). The present paper adopts the CP^2 value κ = 8⁄3 because it is the value inherited by the charged-lepton localization programme and the threshold-compression sequence. Reconciling the two derivations remains an open programme question (§9).

Fallback: if Role-4 fixes only the first localization step, κ is not guaranteed constant across generations — see ⟦Assumption U⟧.

⟦Import M2 — Mass–Localization Scaling (Sector Rescaling Theorem)⟧

In plain terms. A particle's mass is tied to how tightly it is localized — tighter region, heavier particle. This import says mass grows as the *inverse square* of the localization size, and that exponent of 2 is what turns κ = 8⁄3 into the value matching the muon. It is now derived rather than assumed: the Role-4 variational model gives the inverse-square law, and the muon's rescaled energy turns out close to the electron's, so the mass ratio comes purely from the geometry.

Statement: $M_g = \tilde{E}_g / L_g^2$, where $\tilde{E}_g$ is the rescaled variational/eigenvalue energy of generation g. The scaling **structure** is what the Sector Rescaling Theorem establishes; the rescaled energies themselves are *computed* in the Role-4 variational model: $\tilde{E}_0 = 9.864$, $\tilde{E}_1 = 9.854$, $\tilde{E}_2 = 0.834$. Hence $m_g / m_0 = (\tilde{E}_g / \tilde{E}_0)\, e^{(2\kappa g)}$. The light-sector ratio $\tilde{E}_1 / \tilde{E}_0 = 0.999$ leaves μ⁄e essentially pure geometry; the tau ratio $\tilde{E}_2 / \tilde{E}_1 = 0.085$ is the source of the tau suppression (see ⟦Import P3⟧ and §6).

Source: *Toward a Lepton-Sector Mass Derivation* (Sector Rescaling Theorem; Role-4 variational/eigenvalue model).

Grade: **[Conditional]** — *derived, not assumed*. The scaling structure, and so the exponent 2, is derived by the Sector Rescaling Theorem; the rescaled energies are computed by the variational model rather than posited. This retires the previous revision's framing, in which $\tilde{E}_2/\tilde{E}_1$ was treated as an unverified assumption set to 1: the model returns $\tilde{E}_2/\tilde{E}_1 \approx 0.085$ directly, from the tau-sector capacity $C_2 = 22$ exceeding the critical value $C_{crit}(2) = 18.822$ and collapsing the rescaled energy. Conditionality therefore attaches to the variational model's own inputs (the sector capacities and the critical-capacity criterion) and to its reconciliation with the closure route (§6) — not to the existence of the L^{-2} law nor to any assumed energy constancy. For the e–μ step the computed $\tilde{E}_1/\tilde{E}_0 = 0.999$ is consistent with the data to better than 0.1%.

Fallback: the predicted ratios depend directly on the computed $\tilde{E}_g$; revisions to the variational capacity inputs propagate linearly into the suppression, and hence into the tau mass.

⟦Import P3 — Saturated Participation Factor from the Closure Ontology⟧

In plain terms. The tau is held back by a suppression factor of 1⁄12. The 12 is not picked to fit the tau; it is the count of "interior" channels that falls out of the closure architecture's census of fourteen — which itself traces back, step by step, to the geometry's ontology.

Statement: $\Pi_2 = N_{survive}/N_{anchor} = 1/12$ for the saturated third charged-lepton sector. The denominator is the interior (anchoring) channel count $N_{anchor} = N_{int} = 12$; the numerator is the surviving-channel count $N_{survive} = 1$ from the saturated maintenance projection. ("Saturated" means the closure channels are maxed out so that a single channel survives maintenance, and the participating fraction falls — saturation produces suppression, not enhancement.)

This participation factor is the **closure-ontology reading of the same tau suppression** $\Theta_2 = \tilde{E}_2/\tilde{E}_1$ that the variational model computes directly (⟦Import M2⟧): $\Pi_2 = 1/12 = 0.0833$ against the variational 0.0846 — agreement to 1.6%. The two are *not* multiplicative corrections to be combined into $e^{16/3} \times (1/12) \times (\tilde{E}_2/\tilde{E}_1)$; they are two independent derivations of one suppression, and the law carries it once.

Derivation (closure-ontology chain).

geometric ontology → closure architecture → $K = 7$ → fourteen-generator census ($N_{loop} = 14$) → boundary/interior split $14 = 2_\partial + 12_{int}$ → anchoring space $N_{anchor} = 12$ → saturated participation → $\Pi_2 = N_{survive}/N_{anchor} = 1/12$ → tau suppression.

Convergence note. The variational model reaches the same suppression by a wholly different route — the tau-sector capacity exceeding its critical value, collapsing $\tilde{E}_2$ to give $\tilde{E}_2/\tilde{E}_1 \approx 0.085 \approx 1/11.8$. That a channel census (1⁄12) and an energy-collapse calculation (1⁄11.8) independently produce the same suppression, to 1.6%, is the second convergence of the paper (§5) and the strongest evidence that the tau suppression is structural rather than fitted.

Source: the participation and alignment-gating papers of the closure-channel programme.

Grade: **[Conditional]** — *derived, not fitted; and derived does not mean proven*. The participation factor is not introduced phenomenologically: the denominator N_anchor = 12 is *derived* from the boundary/interior decomposition 14 = 2_∂ + 12_int, and the participation object is built on top of it. The existence and value of the 12 are therefore no longer the live question. The remaining explanatory burden has shifted to the **numerator**: why does exactly *one* channel survive — N_survive = 1 — out of the twelve? That is the open item the saturated-maintenance projection must deliver, and the residual conditionality of Π_2 is essentially the conditionality of N_survive = 1. (The deepest closure-algebra realization of the split continues to mature, but it is no longer where the weight of the open question lies.) This is the gauge-paper posture: *derived from the ontology*, with the grade conditional on a single, sharply located step.

Fallback: if the boundary/interior split or the surviving-channel count is revised, the denominator and numerator move with them. Exact agreement on τ/μ would require an effective denominator of 12.317 rather than 12; the integer census therefore predicts a definite residual rather than absorbing it (§6).

⟦Assumption U — Uniform κ across generations⟧

In plain terms. The same "squeezing" number κ is taken to apply to both jumps — electron→muon and muon→tau. This matters because, with only three particles, the data alone cannot tell whether κ stays fixed while a separate factor suppresses the tau, or whether κ simply changes between generations.

Statement: a single κ governs both the e→μ and the μ→τ localization steps.

Grade: **[Conjectural]** as a standalone premise — but likely **discharged** by ⟦Import L1⟧: the CP^2 derivation K_eff = 4 × 2⁄3 carries no generation index, so if K_eff is the localization exponent at every step, uniformity is a consequence of the geometry rather than a free assumption. Pending explicit confirmation of that point in the Role-4 source, U is graded [Conjectural]; on confirmation it becomes [Conditional] alongside L1.

Why it matters. This assumption attributes the *entire* tau deficit to participation rather than to a generation-dependent (running) κ. The two readings are observationally equivalent at the level of two ratios (§7); only an upstream geometric result (the generation-independence note above), or a fourth data point, can break the degeneracy.

⟦Assumption S — Saturation-onset shape⟧

In plain terms. The suppression is taken to switch on at the tau and not before. The participation construction makes this a consequence of how many channels survive at each generation; what still has to be shown is that the switch happens at the third lepton and not the second.

Statement: $\Pi_0 = \Pi_1 = 1$ and $\Pi_2 = 1/12$ — full participation for the electron and muon, with suppression onsetting at the tau.

Grade: **[Conditional]** — and now substantially supported. The onset is no longer a bare step: the variational model gives a mechanism for *why* it falls at the third sector. The tau-sector capacity $C_2 = 22$ exceeds the critical capacity $C_{crit}(2) = 18.822$, which collapses the rescaled energy ($\tilde{E}_2 / \tilde{E}_1 \approx 0.085$); the muon sector stays subcritical, leaving $\tilde{E}_1 / \tilde{E}_0 \approx 0.999$ and no suppression. The closure route concurs, with the surviving count collapsing to $N_{survive} = 1$ at the same sector. What remains to secure is the capacity criterion itself (the values C_g and $C_{crit}(g)$).

Why it matters. The location of the onset (third sector, not second) is the substantive content of this assumption; it is now pinned to the supercriticality condition $C_2 > C_{crit}(2)$ rather than assumed.

---

2. The Electron–Muon Postdiction

In plain terms. The muon is a heavier copy of the electron — about 207 times heavier. This section shows that a single number, $\kappa = 8/3$ (how much more tightly each successive particle is "squeezed" into its region), reproduces that factor of 207 to within about one part in six hundred, with nothing tuned to make it work. Crucially, that κ comes from the geometry of CP^2 — worked out without any reference to lepton masses — so the near-perfect match is a genuine prediction after the fact, not a number bent to fit.

For the first generation step, with no participation suppression ($\Pi_0 = \Pi_1 = 1$):

$m_\mu / m_e = e^{2\kappa} = e^{16/3}$.

Numerically:

$e^{16/3} = 207.127$ observed: 206.768 relative error: 0.17%.

No quantity is adjusted once κ is inherited from ⟦Import L1⟧ and the exponent from ⟦Import M2⟧.

Grade: **[Conditional]** — capped by L1 (κ) and M2 (the exponent) only. This ratio uses κ at the *first* generation step alone, so it does **not** depend on the uniform-κ assumption (U) that gates the second step. The [Conditional] grade reflects only its inheritance from the κ and M2 imports — *not* any weakness in the agreement itself.

Among the Standard-Model reconstruction results currently available in the programme, the μ/e ratio is the strongest quantitative agreement obtained from an independently derived structural parameter: with $\kappa = 8/3$ fixed by CP^2 geometry without reference to lepton data (⟦Import L1⟧), $e^{16/3} = 207.127$ lands against the observed 206.768 — a genuine postdiction at the 0.17% level, with nothing tuned. Both of its inputs carry derivation chains — κ from CP^2 geometry, the scaling law from the Sector Rescaling Theorem (⟦Import M2⟧) — so the result rests on no assumption set by hand. The light-sector rescaled energies are not assumed equal but *computed*: the variational model returns $\tilde{E}_1/\tilde{E}_0 = 0.999$, consistent with the data to better than 0.1% and confirming that μ/e is essentially pure geometry.

---

3. The Tau Participation Correction

In plain terms. The same recipe that works so well for the muon, applied a second time, badly overshoots the tau — it predicts a particle more than ten times too heavy. So something must be holding the third particle back. The paper introduces one suppression factor — a division by 12 — which pulls the prediction back to within about 2.6% of the measured tau. Honesty note: this factor of 12 is not fitted to the tau. It is *derived* — it is the count of "interior" channels in the closure architecture's census of fourteen (14 = 2 + 12). What is still being worked out is not the 12 but why only a single channel survives at the third sector.

Without participation suppression, the same step law gives

$m_\tau/m_\mu = e^{16/3} = 207.13$,

hence $m_\tau \approx 21.9$ GeV — exceeding observation by more than an order of magnitude. This excess is the quantitative motivation for a suppression at the third sector.

Introducing ⟦Import P3⟧ and ⟦Assumption S⟧:

$m_\tau/m_\mu = e^{16/3}/12 = 17.261$ observed: 16.817 relative error: 2.64%.

Grade: **[Conjectural], provisionally** — the tau suppression itself is now doubly-derived (closure 1⁄12 and variational $\tilde{E}_2/\tilde{E}_1 \approx 0.085$, agreeing to 1.6%), so it is no longer the soft point; the residual is the spread among the readings of the suppression, not a missing factor (§6). The binding constraint on the *grade* is Assumption U (uniform κ across steps): this ratio uses κ at the second step, and U is likely discharged by the generation-independence of the CP^2 derivation (⟦Import L1⟧). On confirming that K_eff carries no generation index, U becomes [Conditional] and this ratio rises with it to **[Conditional]**. Until then it is held at [Conjectural] by the running-κ degeneracy (§7).

---

4. The Charged-Lepton Ratio Law

In plain terms. Putting the pieces together — the "squeezing" hierarchy and the tau suppression — gives one compact formula that sets all three masses at once, as proportions. It gets the electron-to-muon step nearly perfectly and the full three-way spread to within about 3%. The suppression in the tau term (the 1⁄12) is the same number two different calculations land on, which is why it appears once, not twice. The formula is written carefully to avoid a notation trap (the 1⁄12 is shown out front, so it cannot be misread).

Combining the localization hierarchy with the tau suppression:

$m_e : m\mu : m\tau = 1 : e^{16/3} : (1/12)\, e^{32/3}$.

(The third entry is the coefficient 1⁄12 multiplying e^(32⁄3); written this way to remove the slash ambiguity of "$e^{32/3}/12$". The 1⁄12 is the closure reading of the suppression $\Theta_2 = \tilde{E}\_2 / \tilde{E}\_1$; the variational reading 0.0846 would put the third entry at $0.0846 \cdot e^{32/3}$.)

Numerically:

1 : 207.13 : 3575 observed: 1 : 206.77 : 3477.

The τ⁄e ratio differs from observation by approximately **2.8%**.

Grade of the full law: **[Conjectural], provisionally → [Conditional] on confirmation** — by inheritance cap, the law cannot exceed its weakest input. Every input is now [Conditional] and *derived*: κ (CP^2), the scaling law and rescaled energies (Sector Rescaling Theorem + variational model), and the tau suppression (doubly-derived, closure and variational). The only sub-[Conditional] input is Assumption U, which the CP^2 derivation likely discharges (⟦Import L1⟧). On confirming κ's generation-independence, every input is [Conditional] and the full law rises to [Conditional] — its residual conditionality then resting on the CP^2/Sector–Curvature chain, the variational model's capacity inputs, and the closure surviving-channel count N_survive = 1.

---

5. Why the Convergence Is the Result

In plain terms. The headline isn't the 3% accuracy on its own — it's *where the two ingredients came from*. They were worked out in completely separate earlier investigations, neither of them aimed at lepton masses, yet they slot together into one formula. That is the kind of coincidence that is hard to arrange by accident. This section is also careful not to oversell: the accuracy is really "one near-perfect result (the muon) plus one merely-good result (the tau)," not a uniform 3% across the board.

The relation is noteworthy precisely because its two ingredients originate in different parts of the programme, along chains that share no common step:

- the **localization factor** e^(2κ) descends from the geometric chain — CP^2 geometry fixing κ = 8⁄3;
- the **participation factor** 1⁄12 descends from the closure-ontology chain — the K = 7 census of fourteen generators and its boundary/interior split.

Set side by side:

Geometric chain: CP^2 geometry → $K\_eff = 4 \times (2/3) = 8/3$ → κ → $e^{16/3}$ → muon ratio

Closure chain: closure ontology → K = 7 → 14 generators → $2\_\partial + 12\_int$ → 1⁄12 → tau suppression

These chains were developed separately and for unrelated reasons, and they meet only at the charged-lepton spectrum. That meeting — not the formula itself — is the result this paper records. A single argument producing two numbers would be unremarkable; two independent ontological branches producing two numbers that *jointly* reproduce the lepton hierarchy is the substantive claim, and the reason the paper is framed as a convergence rather than as a fit.

There is, moreover, a **second convergence**, inside the tau suppression itself. The same suppression Θ_2 is reached two independent ways:

Variational route: $C\_2 = 22 > C\_crit(2) = 18.822$ → $\tilde{E}\_2$ collapses → $\tilde{E}\_2 / \tilde{E}\_1 \approx 0.085 \approx 1/11.8$

Closure route: $14 = 2\_\partial + 12\_int$ → $N\_survive / N\_anchor = 1/12 \approx 0.083$

Numerically, the two derivations and the measurement read:

closure route $\Theta_2 = 0.0833$ (= 1⁄12) variational $\Theta_2 = 0.0846$ (= 1⁄11.8) observed $\Theta_2 = 0.0812$ (= 1⁄12.3)

The two independent derivations differ by just 0.0846 − 0.0833 = **0.0013** — about one and a half percent, and minuscule against the scale of the problem. The suppression itself is an order-of-magnitude effect (a factor of roughly 12), set against a tau mass that stands more than three orders of magnitude above the electron. That an energy-collapse calculation and a channel census — sharing no step — fall within 0.0013 of each other is the central evidence that the tau term is structural, not fitted: a fit gives one number once, whereas here two unrelated derivations independently land on it, and both within a few percent of observation. The earlier temptation to read the rescaled-energy ratio and the participation factor as *separate* multipliers (which would double-count) is exactly what this convergence corrects — they are one suppression, seen twice.

The natural reading is therefore not "closure model + variational model = suppression" but "closure model ↔ variational model" — two descriptions of one object at different levels. The variational model names it at the level of *energy*, $\tilde{E}\_2 / \tilde{E}\_1$; the closure model names it at the level of *channel counting*, $N\_survive / N\_anchor$. The single structural statement both express is:

Saturation leaves roughly one-twelfth of the mass-generating effectiveness available at the tau sector.

Whether the two are *literally* the same object — $N\_survive / N\_anchor$ being the channel-counting shadow of $\tilde{E}\_2 / \tilde{E}\_1$ — is a conjecture the programme can pursue **[Conjectural]**; the numerical agreement to 0.0013 is its first evidence, not its proof. The point that is already secure is the reframing of the problem: the tau suppression is no longer a factor-of-ten puzzle requiring a mechanism, but a few-percent refinement (§6) about an effect whose order of magnitude two independent chains have already fixed.

The paper introduces **no new fitting parameter**. The honest framing of the accuracy is not a uniform "≈3%" but rather *one excellent postdiction (μ/e, 0.17%) plus one good conditional prediction (τ/μ, 2.6%)*; the μ/e channel is an order of magnitude tighter than the τ channel, and that asymmetry is informative (§7).

---

6. The Tau Residual: Two Readings of One Suppression

In plain terms. The tau prediction comes out a little high — by about 2.6%. That gap is now understood not as a missing ingredient but as the small spread between two independent calculations of the same suppression and the measured value. One calculation (a channel count) gives 1⁄12; another (an energy collapse) gives about 1⁄11.8; the measurement wants about 1⁄12.3. All three sit within a few percent of one another. Pinning down which is exact — and closing the last couple of percent — is the open work, but the tau is no longer resting on a guess.

The suppression that would reproduce τ/μ exactly is Θ_2 = 16.817 ⁄ 207.127 = 0.0812 = 1⁄12.32. The two independent derivations bracket the neighbourhood:

reading	Θ_2	as 1⁄x	τ/μ	error vs obs
observed	0.0812	1⁄12.32	16.82	—
closure census (Π_2)	0.0833	1⁄12.00	17.26	+2.64%
variational (Ẽ_2 ⁄ Ẽ_1)	0.0846	1⁄11.82	17.53	+4.24%

Two points follow.

First, the agreement is **structural, not fitted**. A fit would have returned 1⁄12.32 exactly. Instead, two unrelated derivations return 1⁄12.0 and 1⁄11.8 — clustering near, but not on, the observed value. That clustering of *independent* results is the signature of a real suppression mechanism; a single tuned parameter could not produce it.

Second, the residual *is* that spread, and the questions it leaves are now **refinement questions, not structural ones** — the order of magnitude is already fixed by two independent chains. They are:

1. Why does the variational threshold calculation give 0.0846 rather than exactly 1⁄12 = 0.0833 — and is N_survive ⁄ N_anchor the channel-counting shadow of Ẽ_2 ⁄ Ẽ_1, so that the two are one object?
2. Why does nature sit at 0.0812 rather than at either derived value — what sub-leading correction (a further channel, or a refinement of the capacity criterion C_2 vs C_crit(2)) carries the prediction the last ~2–4% onto observation?

Both are few-percent questions. Neither reopens the factor-of-twelve, which is the part the convergence has secured.

A note of correction. An earlier reading of this paper treated the rescaled-energy ratio and the participation factor as *separate* multiplicative corrections, writing $\tau/\mu = e^{16/3} \times (1/12) \times (\tilde{E}\_2 / \tilde{E}\_1)$. That double-counts. The participation factor and the rescaled-energy ratio are two derivations of the **same** suppression, so the law carries it once: $\tau/\mu = e^{16/3} \times \Theta_2$, with $\Theta_2 \approx 1/12$.

Grade: **[Conditional prediction]** — the +2.64% sign and magnitude follow from adopting the closure value 1⁄12; the residual is bounded by the agreement of the two independent derivations, and is revisable by reconciling them or by refining the capacity criterion.

---

7. Underdetermination: Uniform κ versus Running κ

In plain terms. Here is the most important caveat in the paper. The very same masses can be reproduced two different ways: either with the suppression factor this paper uses, or by letting the "squeezing" number change from one generation to the next and dropping the suppression entirely. With only three particles to test against, the data alone cannot tell these two stories apart. So this paper's story is only the privileged one *if* the earlier geometric work independently insists that the squeezing number stays fixed across generations. If it doesn't, the suppression picture is one of two equally good fits — and the paper says so rather than hiding it.

The decomposition "uniform κ + step participation" is **not unique**. The same two ratios are reproduced equally well by a *running* localization exponent with no participation factor at all:

- e→μ step: $\kappa_1 = \frac{1}{2} \ln(m_\mu/m_e) = 2.666$;
- μ→τ step: $\kappa_2 = \frac{1}{2} \ln(m_\tau/m_\mu) = 1.411$, i.e. $\kappa_2/\kappa_1 \approx 0.53$.

These two pictures —

1. $\kappa = 8/3$ uniform, with $\Pi_2 = 1/12$ at the tau, versus
2. κ running (2.666 → 1.411), with no participation factor —

are observationally indistinguishable at the level of two ratios. The participation account is privileged **only if** κ is constant across all generation steps (⟦Assumption U⟧). Two independent considerations now favour this. The geometric chain plausibly supplies it directly: the CP^2 derivation $K\_eff = 4 \times 2/3$ carries no generation index (⟦Import L1⟧), so if K_eff is the exponent at every step, picture 2 (running κ) is excluded at source. And the variational model supplies a *mechanism* for picture 1 that picture 2 lacks: it computes the tau deficit as a genuine collapse of the rescaled energy ($\tilde{E}\_2 / \tilde{E}\_1 \approx 0.085$, from capacity supercriticality) at fixed localization κ — i.e. the suppression is an energy-sector effect, not a change in the geometric exponent. Picture 2 would have to reproduce the same deficit by a running κ with no independent motivation. The degeneracy is therefore resolved upstream (by geometry) and disfavoured downstream (by the absence of a running-κ mechanism), though confirming K_eff's generation-independence in the Role-4 source remains the clean closure.

This is the reason ⟦Assumption U⟧ is flagged separately and not folded silently into the localization import — and why the generation-independence of the CP^2-derived κ is the highest-leverage item in the ledger (§9).

Grade: **[Open]** — degeneracy unresolved within this paper, but with a clear and likely upstream resolution (CP^2 generation-independence) identified.

---

8. Relation to the Koide Relation

In plain terms. There is a famous, almost uncanny numerical pattern among these three masses — discovered by the physicist Yoshio Koide — that holds to about one part in a hundred thousand. Anyone who knows the subject will immediately ask how this paper stacks up against it. The answer has two parts. First, it is a different *kind* of claim: Koide's relation is a striking numerical coincidence spotted in the data, whereas this paper tries to build the masses up from physical structure. Second — and usefully — this paper's formula does *not* reproduce Koide's pattern exactly; it misses by about 0.4%. That miss is a gift, because it hands future work a second, independent check to aim at.

Any reader outside the programme will compare a charged-lepton mass relation to the empirical Koide relation,

$$Q \equiv (m_e + m_\mu + m_\tau) / (\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2 = 2/3,$$

which the observed masses satisfy to $Q = 0.66666$ — about 1 part in 10^5.

Two points must be made plainly.

Kind of claim. The VERSF law does not aim to out-precision Koide. Koide is a tight algebraic coincidence among the measured masses; the VERSF law is a reconstruction of the ratios from independently motivated substrate structure. These are different claims, and "more accurate than Koide" is *not* among the assertions of this paper.

Quantitative cross-check. The VERSF ratio law does **not** reproduce Koide. Evaluating Q on the law's own values ($1 : e^{16/3} : (1/12)e^{32/3}$) gives

$$Q\_VERSF = 0.66929,$$

a **0.39% excess** over 2⁄3. This is consistent with the percent-level residuals in §3–4 and is itself a falsifiable consequence: a corrected participation factor that improves τ/μ would simultaneously pull Q toward 2⁄3. Koide therefore becomes a *second, independent target* for the closure-channel programme — if the corrected Π_2 does not also drive Q → 2⁄3, the two relations are in tension and at least one structural assumption is wrong.

Accordingly, the headline claim is softened to: the first *unified charged-lepton ratio law within the VERSF programme* — not the first such relation in physics.

Grade: **[Conditional]** observation; the cross-check is exact arithmetic on the law's outputs.

---

9. Status and Condition Ledger

In plain terms. This is the scorecard. It lists every claim in the paper next to a plain grade — proven, conditional, conjectural, or open — and follows one strict rule: a conclusion is never graded higher than its weakest ingredient. It then names the three specific unanswered questions ("gates") that would have to be cleared to turn this promising formula into a genuine derivation rather than a suggestive one.

Claim / component	Source or inheritance	Grade
Localization exponent κ = 8⁄3 (Route A)	CP^2 geometry; *Toward a Lepton-Sector Mass Derivation* §6.4	[Conditional] — **derived, not fitted**; conditional on the Sector–Curvature Correspondence
Alternative κ = ln 14 ≈ 2.639 (Route B, **unused**)	closure architecture; *Structural Completeness in the Role-4 Lepton Sector Model*	reconciliation with Route A is [Open]
Mass scaling structure $M_g = \tilde{E}_g / L_g^2$	Sector Rescaling Theorem; *Toward a Lepton-Sector Mass Derivation*	[Conditional] — **derived**
Rescaled energies $\tilde{E}_0$, $\tilde{E}_1$, $\tilde{E}_2$ = 9.864, 9.854, 0.834	Role-4 variational model	[Conditional] — **computed**; conditional on capacity inputs C_g, $C_{crit}(g)$
Interior/anchoring count N_{anchor} = 12 (from 14 = $2_\partial + 12_{int}$)	enumeration / alignment-gating papers	[Conditional] — **derived**; existence and value not the live question
Surviving-channel count $N_{survive}$ = 1	saturated-maintenance projection	[Conditional] — open closure-side item ("why only one channel?")
Tau suppression Θ_2 ≈ 1⁄12	**two derivations**: closure census (0.0833) and variational $\tilde{E}_2 / \tilde{E}_1$ (0.0846)	[Conditional] — **doubly-derived**; agree to 1.6%; obs 0.0812
Identity $N_{survive} / N_{anchor} \leftrightarrow \tilde{E}_2 / \tilde{E}_1$ (one object, two levels)	the two routes' agreement	[Conjectural] — interpretive hypothesis; agreement is evidence, not proof

Claim / component	Source or inheritance	Grade
Uniform κ across steps (Assumption U)	likely discharged by CP^2 generation-independence; favoured by variational mechanism	[Conjectural] → [Conditional] on confirming K_eff carries no generation index
Saturation-onset shape (Assumption S)	capacity supercriticality C_2 = 22 > C_crit(2) = 18.822	[Conditional] — onset now mechanistically located at the tau
$\mu/e = e^{16/3}$ (0.17%)	κ (CP^2), scaling, Ẽ_1 ⁄ Ẽ_0 ≈ 0.999	[Conditional] postdiction — **strongest quantitative agreement in the programme** from an independently derived parameter
$\tau/\mu = e^{16/3} \cdot \Theta_2$ (2.64%)	+ U, suppression (doubly-derived)	[Conjectural] → [Conditional] once U is discharged
Full law $1 : e^{16/3} : (1/12)e^{32/3}$ (≈3%)	both chains	[Conjectural] → [Conditional] once U is discharged
+2.64% τ/μ residual	spread among the readings of Θ_2	[Conditional] prediction
Uniform-κ vs running-κ degeneracy	U / CP^2 / variational mechanism	[Open] pending generation-independence confirmation; disfavoured downstream
Koide Q_VERSF = 0.669 (≠ 2⁄3)	full law	[Conditional] cross-check

All inputs are now *derived* — κ (CP^2), the scaling law (Sector Rescaling Theorem), the rescaled energies (variational model), and the tau suppression (doubly-derived). No import lacks a derivation, and the tau suppression is no longer an assumption. What remains separates into a grade-lifting confirmation, the genuine physical bottleneck, and two narrower questions:

- **Gate (U / generation-independence) — highest leverage.** Confirm in the Role-4 source that K_eff = 4 × 2⁄3 carries no generation index, i.e. the CP^2-derived κ is the localization exponent at every step. This discharges Assumption U, closes the running-κ degeneracy (§7), and lifts τ/μ and the full law to [Conditional]. It is the single most consequential open item for the grade.
- **Gate (Θ_2 refinement) — the remaining physical work.** The factor-of-twelve is fixed by two independent chains; what is left are two few-percent refinement questions. (i) Why does the variational threshold give 0.0846 rather than exactly 1⁄12, and is N_survive ⁄ N_anchor the channel-counting shadow of Ẽ_2 ⁄ Ẽ_1 (one object at two levels of description)? (ii) Why does nature sit at 0.0812 — what sub-leading correction (a further channel, or a refinement of C_2 vs C_crit(2)) closes the last ~2–4%? This is where the tau number can still move; it no longer reopens the order of magnitude.
- **Gate (P3 / surviving channel).** Deliver N_survive = 1 from the saturated-maintenance projection: *why does exactly one channel survive out of the twelve?* This is the closure-side input to the reconciliation above.

- **Gate (κ-route reconciliation).** Reconcile Route A ($\kappa = 8/3$, CP^2) with the closure-based Route B ($\kappa \approx \ln 14$). The ratio law requires Route A; that it is also the better-fitting value should remain explicit until the two derivations are reconciled or one is shown to supersede the other.

Conclusion

In plain terms. To sum up: the programme now has a single formula for the three charged-lepton masses, built from numbers that are all derived, none tuned to fit. The electron-to-muon step is essentially pure geometry. The tau is held back by a suppression that two completely separate calculations — a channel count and an energy collapse — independently put at about one-twelfth, agreeing with each other to better than 2%. That agreement, not the formula, is the real finding. "Derived" is still not "proven": each calculation has a step or two left to secure, and closing the last couple of percent on the tau is the open work.

The programme now contains a single charged-lepton ratio law,

$m_e : m\mu : m\tau = 1 : e^{16/3} : (1/12)\, e^{32/3}$,

assembled from derived structure rather than fitted parameters: $\kappa = 8/3$ from CP^2 geometry sets the localization, the Sector Rescaling Theorem and the Role-4 variational energies set the masses, and the tau suppression is estimated independently two ways. It reproduces the electron–muon ratio to 0.17% and the full hierarchy to about 3%, with the residual concentrated in the tau.

The result's value lies in the **convergences**, not in the formula. The first: a geometric chain ($CP^2 \to \kappa \to e^{16/3}$) and the variational energies generate the spectrum, with μ/e coming out as pure geometry because the light-sector rescaled energies are nearly equal. The second, and the more striking: the one non-geometric feature — the tau suppression — is produced independently by a closure channel-census (1⁄12) and a variational energy-collapse ($\tilde{E}\_2 / \tilde{E}\_1 \approx 0.085$), agreeing to 0.0013. The single statement both express is one structural fact — *saturation leaves roughly one-twelfth of the mass-generating effectiveness available at the tau* — reached at two levels of description, an energy ratio and a channel count, which may yet prove to be one object. *Derived, however, does not mean proven*: the geometric chain rests on the Sector–Curvature Correspondence, the variational energies on their capacity inputs, and the closure route on the surviving-channel count N_survive = 1. With no input now resting on an assumption set by hand, the honest status of the full law is [Conjectural] provisionally, rising to [Conditional] once the CP^2 κ is confirmed to carry no generation index.

The outstanding question is no longer what numerical factors are required, nor whether any was fitted (none was), nor whether the tau suppression rests on an assumption: it is derived twice over. Two open items remain, of different kinds. The grade-lifting one is confirming that the CP^2-derived κ carries no generation index; doing so discharges the last hand-set assumption and lifts the full law to [Conditional]. The other is no longer a structural bottleneck but a pair of refinement questions, since two independent chains have already fixed the order of magnitude of the suppression: why the variational threshold gives 0.0846 rather than exactly 1⁄12 (and whether the channel count and the energy ratio are one object at two levels), and what sub-leading correction carries the prediction the last few percent onto the observed 0.0812. That is the only place the tau number can still move. Settling these, while accounting for the 0.4% pull on the Koide combination (§8), is what would turn the tau channel from a good conditional prediction into a derived one — and graduate the convergence from a unified ratio law into a derivation of the charged-lepton spectrum.