Threshold Compression in the Charged-Lepton Spectrum

A Muon-Scale Postdiction and the Open Tau Suppression Factor

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The electron, muon, and tau are three versions of the same charged-lepton family — the same electric charge, very different masses. The muon is about 207 times heavier than the electron; the tau is about 17 times heavier than the muon.

That second number is the puzzle. If nature simply repeated the same mass jump, the tau would be about 207 times heavier than the muon, not 17. The hierarchy is steep but uneven.

This paper makes one claim and poses one question.

The claim concerns the first jump, electron to muon. Each deeper refinement generation is held in a smaller closure region, and mass scales like inverse area, so a smaller region means a much larger mass. With the geometric value $\kappa = 8/3$ inherited from the Role-4 model, this gives a muon about 207 times heavier than the electron — close to what is measured, with no quantity adjusted to fit. That is a genuine postdiction.

The question concerns the second jump, muon to tau. The same localisation step, repeated, would make the tau far too heavy — around 21.9 GeV against the measured 1.78 GeV. So the tau must carry an extra suppression. Writing that suppression as a single factor Θ_2, its value is about 0.081.

That number is not a prediction. It is read off the measured tau mass; it records how much suppression is needed, not why. But the programme already has a place to put it. A companion result holds that a particle's mass depends not only on how tightly it is localised but on how *effectively* it couples — on what fraction of its aligned internal states actually complete an irreversible commitment. The live idea is that the saturated third sector is less effective in exactly this sense: only a restricted subset of its states anchor, so its coupling is gated down to about a twelfth of the ground value. Threshold proximity — the tau sitting just inside the binding edge — is the natural cause of that gating. The task this paper hands forward is to derive that reduced effective coupling rather than infer it from the answer.

So: the muon looks structural. The tau marks the next theorem — and that theorem now has an address.

Table of Contents

Abstract

The charged-lepton mass spectrum is decomposed into two factors: an exponential localisation hierarchy and a sector-dependent coherence–anchoring ratio. Under the inherited Role-4 localisation law $L_g = L_0 e^{-\kappa g}$ with $\kappa = 8/3$, the sector-rescaling relation $M_g = \tilde{E}_g / L_g^2$ yields a first mass ratio

$$m_\mu / m_e \approx e^{16/3} = 207.13$$

under the assumption $\tilde{E}_1/\tilde{E}_0 \approx 1$. With the measured electron mass this gives $m_\mu \approx 105.84$ MeV against the observed 105.66 MeV — a postdiction high by 0.17% and free of adjustable parameters once κ is fixed by the Role-4/CP^2 geometry.

The same step iterated to the third sector overpredicts the tau mass by more than an order of magnitude. The discrepancy is absorbed into a single suppression factor

$$\Theta_2 = (m_\tau / m_e) / e^{32/3} \approx 0.081.$$

This factor is inferred from the observed tau mass and carries no independent evidential weight; it is the central open object of the paper. Rather than treat it as an invented compression constant, it is identified with a derived quantity from the *Void Coupling as Phase Coherence* result, in which mass depends on a coherence–anchoring ratio $m \propto p_{eff} / K_c$ with effective coupling $p_{eff} = g_c \cdot p_{align}$. Writing the full mass as $m_g \sim e^{(2\kappa g)} \cdot (p_{eff,g} / K_{c,g})$ makes the sector factor Θ_g the ratio $(p_{eff,g}/K_{c,g})/(p_{eff,0}/K_{c,0})$, which decomposes into a coupling part and an anchoring part. Whether the suppression sits in the coupling, the anchoring depth, or both is left open; the paper investigates the coupling reading, $\Theta_2 \approx p_{eff,2}/p_{eff,0} \approx 0.081$, motivated by its connection to threshold gating of aligned states, under which the saturated third sector is roughly twelve times less anchoring-effective than the ground sector. The outstanding problem is to derive this reduced effective coupling — the alignment probability $p_{align,2}$ of the third sector — from the Role-4 eigenmode structure, with threshold proximity (small $\Delta_g = C_g - C_{crit}(g)$) the candidate cause of the gating. The capacities C_g that fix any threshold route remain underived.

---

1. Introduction

The Standard Model records the charged-lepton masses but does not explain them. The measured values are approximately

$m_e = 0.511$ MeV, $m_\mu = 105.66$ MeV, $m_\tau = 1776.86$ MeV,

with ratios

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

Any structural account must explain three facts:

1. why there are three charged leptons;
2. why their masses are strongly hierarchical;
3. why the hierarchy is non-uniform.

The first two are approached in the Role-4 and refinement programmes. The present paper addresses the third, and separates it sharply into one settled component and one open one.

The central position is:

The first charged-lepton jump is a parameter-free consequence of exponential localisation tightening. The second jump is the same tightening opposed by a reduced anchoring-effective coupling in the saturated third sector — a suppression whose value is known and whose home is a derived object, but whose magnitude is not yet derived.

---

2. Inherited Structure

The following results are inherited.

H1 — Three refinement generations

The charged leptons correspond to three admissible refinement sectors,

$g = 0, 1, 2,$

identified with e, μ, τ. The fourth sector, $g = 3$, is not stably bound.

H2 — Localisation tightening

The Role-4 lepton model gives a sector localisation scale

$L_g = L_0 e^{(-\kappa g)},$

with hierarchy parameter $\kappa \approx 2.67$. The geometrically natural value supplied by the Role-4/CP^2 geodesic exponent is

$\kappa = 8/3.$

H3 — Mass-scale relation

The sector-rescaling theorem gives

$M_g = \tilde{E}_g / L_g^2,$

so that

$M_g = (\tilde{E}_g / L_0^2) e^{(2\kappa g)}.$

If $\tilde{E}_g$ were sector-independent, adjacent mass ratios would be

$M_{g+1} / M_g = e^{(2\kappa)}.$

H4 — Binding thresholds

The Role-4 binding condition is

$C_g > C_{crit}(g),$

with approximate critical values

$C_{crit}(0) \approx 1.192$, $C_{crit}(1) \approx 7.736$, $C_{crit}(2) \approx 18.822$, $C_{crit}(3) \approx 34.435$.

The critical sequence rises steeply and monotonically. Binding holds for $g = 0, 1, 2$ and fails for $g = 3$.

H5 — Coupling and anchoring as separate mass controls

The *Void Coupling as Phase Coherence* result separates two contributions to mass that the localisation law alone leaves merged in the sector energy $\tilde{E}_g$. Mass is set by an effective coupling divided by an anchoring depth,

$$m \propto p_{eff} / K_c, \quad p_{eff} = g_c \cdot p_{align},$$

where K_c is the anchoring depth (the number of successful irreversible micro-events required before commitment completes), g_c is the coupling strength (how often irreversible micro-events become available), and p_{align} is the alignment probability (the fraction of phase-aligned states that actually anchor). The product $p_{eff} = g_c \cdot p_{align}$ is the **anchoring-effective coupling**: a mode may be strongly phase-coupled yet weakly anchoring-effective if only a restricted subset of its aligned states completes commitment.

(Notation: the coupling paper writes the coupling strength as g; it is written g_c here to avoid collision with the refinement-generation index g.)

---

3. The Muon Estimate: Exponential Localisation

Begin from

$$M_g = (\tilde{E}_g / L_0^2)\, e^{(2\kappa g)}.$$

For the first step, assume

$$\tilde{E}_1 / \tilde{E}_0 \approx 1,$$

so that the sector energies cancel and

$$m_\mu / m_e \approx e^{(2\kappa)}.$$

With $\kappa = 8/3$,

$e^{(2\kappa)} = e^{(16/3)} = 207.13$.

Therefore

$m_\mu \approx 207.13\, m_e$.

Using $m_e = 0.51099895$ MeV,

$m_\mu^{est} \approx 105.84$ MeV,

against the observed

$m_\mu^{obs} \approx 105.66$ MeV.

The estimate is high by about 0.17%.

This is the paper's one semi-quantitative result. Once $\kappa = 8/3$ is fixed independently by the Role-4/CP^2 geometry, the electron-to-muon ratio contains **no adjustable quantity**: the agreement to two parts in a thousand is a genuine postdiction, not a fit. The only assumption local to this step is $\tilde{E}_1/\tilde{E}_0 \approx 1$, and the success of the estimate is itself evidence that the first sector energy is essentially uncompressed.

---

4. The Tau Residual

If the same step repeated without correction, then

$m_\tau / m_\mu \approx e^{(2\kappa)} \approx 207$,

predicting

$m_\tau \approx 207\, m_\mu \approx 21.9$ GeV,

against the observed

$m_\tau \approx 1.777$ GeV.

The localisation law alone overpredicts the tau by more than an order of magnitude. The third generation is therefore not a third identical localisation step.

Introduce a sector factor

$\Theta_g = \tilde{E}_g / \tilde{E}_0$,

so that

$m_g = m_e \cdot e^{(2\kappa g)} \cdot \Theta_g$,

with $\Theta_0 = 1$ by normalisation. The muon result requires $\Theta_1 \approx 1$, consistent with §3. The tau fixes

$\Theta_2 = (m_\tau / m_e) / e^{(4\kappa)}$.

Numerically, with

$m_\tau / m_e \approx 3477.2$, $e^{(4\kappa)} = e^{(32/3)} \approx 4.29 \times 10^4$,

we obtain

$\Theta_2 \approx 0.081$.

It must be stated plainly what this number is and is not.

Θ_2 is **defined** as the residual the localisation law leaves behind. It is computed from the measured tau mass by inversion. It is not a prediction, it agrees with observation by construction, and it carries **no independent evidential weight**. The near-exact reproduction of the tau mass that follows from substituting $\Theta_2 = 0.081$ back into $m_\tau = m_e\, e^{(32/3)}\, \Theta_2$ is an identity, not a confirmation. Its only content is the *value* it places on the residual: the third sector must be suppressed by a factor of roughly twelve relative to a pure localisation step.

The observed lepton hierarchy is thus organised as

$1 : e^{(2\kappa)} : e^{(4\kappa)} \cdot \Theta_2 = 1 : 207 : 3477$,

instead of the uncompressed

$1 : 207 : 42900$.

The first colon is derived. The second is not. The paper's question is exactly:

Why does the third sector carry a suppression factor of about 0.081?

---

5. The Sector Factor as a Coherence–Anchoring Ratio

The sector factor $\Theta_g = \tilde{E}_g / \tilde{E}_0$ was introduced in §4 as a black box — whatever the localisation law fails to capture. The coupling result H5 gives it an identity. Identifying the sector energy with the anchoring-effective coupling per anchoring depth,

$\tilde{E}_g \propto p_{eff,g} / K_{c,g}$,

the full mass becomes a product of two controls,

$m_g \sim e^{2\kappa g} \cdot (p_{eff,g} / K_{c,g})$,

the geometric localisation factor $e^{2\kappa g}$ and the coherence–anchoring ratio. Normalising to the ground sector,

$\Theta_g = (p_{eff,g} / K_{c,g}) / (p_{eff,0} / K_{c,0})$,

with $\Theta_0 = 1$ automatically. This is consistent with §4: the operational residual $\Theta_2 = (m_\tau/m_e)/e^{4\kappa} \approx 0.081$ is unchanged, but it is now the *value of a known structural object* rather than an invented compression factor. The localisation factor is not double-counted — $e^{2\kappa g}$ is the L_g^2 geometry, while K_c is a separate anchoring count living inside $\tilde{E}_g$.

The residual therefore decomposes into an anchoring part and a coupling part,

$\Theta_2 = (p_{eff,2} / p_{eff,0}) \cdot (K_{c,0} / K_{c,2})$,

and both factors pull the mass in the same direction (a larger $K_{c,2}$ or a smaller $p_{eff,2}$ each lowers the mass). Which factor carries the suppression is not settled by the coupling relation alone, and is left open here. Three routes are admissible:

- **Route A — anchoring-depth suppression.** $\Theta_2 = K_{c,0} / K_{c,2}$: the coupling is sector-stable and the third sector simply requires more successful events to commit ($K_{c,2} \approx 12 K_{c,0}$).
- **Route B — coupling suppression.** $\Theta_2 = p_{eff,2} / p_{eff,0} = (g_{c,2} \cdot p_{align,2}) / (g_{c,0} \cdot p_{align,0})$: the anchoring depth is sector-stable and the third sector is gated down in anchoring-effective coupling ($p_{eff,2} \approx 0.081\, p_{eff,0}$).
- **Route C — mixed.** Both ratios depart from unity and share the factor of ≈ 12 between them.

The present paper **investigates Route B**, motivated by a stated connection rather than by evidence: it ties directly to threshold gating of aligned states (§5.1), where proximity to the binding edge naturally suppresses the alignment probability p_{align} while leaving the anchoring depth K_c untouched. Route A, by contrast, reproduces the earlier capacity problem in new symbols — it requires a derived account of why K_c grows by a factor of twelve at the third sector — and the anchoring programme has emphasised that anchoring behaviour varies strongly between modes, so sector-stable coupling is no more given than sector-stable anchoring. The choice to investigate Route B is therefore strategic, not evidential: $K_{c,2} \approx K_{c,0}$ is assumed, not shown, and the decomposition between p_{eff} and K_c remains an open object. Under Route B,

$\Theta_2 \approx p\_eff,2 / p\_eff,0 \approx 0.081$,

i.e. the saturated third sector is roughly twelve times less anchoring-effective than the ground sector.

5.1 Threshold compression as the source of reduced alignment

Threshold proximity is not a separate hypothesis here — it is the *mechanism* for a reduced p_align. Near the binding edge, only a restricted subset of phase-aligned states completes irreversible commitment, which is exactly a suppressed alignment probability. Writing the threshold margin

$\Delta\_g = C\_g - C\_crit(g)$,

the link is

small $\Delta\_g \Rightarrow$ reduced p_align,g $\Rightarrow$ reduced p_eff,g $\Rightarrow \Theta\_g \ll 1$.

The link can be stated at two strengths. The weaker form is a qualitative principle that is easy to defend; the stronger is a quantitative limit that sharpens it.

Saturation–Gating Principle (weak form). A refinement sector that saturates the available closure register cannot maintain full participation of all anchoring-effective channels: some channels are committed to maintaining the completed closure and are therefore unavailable for further irreversible commitment. Hence

saturation $\Rightarrow$ p_align < 1.

This much is hard to dispute: a saturated register is by definition one in which the available structure is fully spent, so it cannot also be fully free to commit. It already does the qualitative work the paper needs — it explains why the third (saturated) sector is gated at all and why the first two (unsaturated) sectors are not, preserving p_align,0 $\approx$ p_align,1 ≈ 1 and the muon postdiction. It does not, on its own, say how far below unity p_align,2 falls.

Alignment Saturation Conjecture (sharp form). As a refinement sector approaches its binding threshold, $\Delta\_g = C\_g - C\_crit(g) \to 0^+$, the fraction of phase-aligned states capable of completing irreversible commitment decreases monotonically, p_align(g) ↓, so that the anchoring-effective coupling vanishes at the edge:

$\lim\_\{\Delta\_g \to 0^+\}$ p_eff(g) = 0.

The sharp form is the one that turns "p_align,2 < 1" into a route to a number. Two riders make it usable rather than merely suggestive.

First, **a plateau condition is required for consistency with the muon.** Bare monotonicity is too weak: a $p\_align(\Delta)$ that declined gradually across the whole range would gate the muon as well, destroying the 0.17% postdiction. The conjecture is therefore only admissible in the stronger form in which the decline is confined to a neighbourhood of the threshold — $p\_align(\Delta) \approx 1$ on a plateau away from the edge, falling sharply only as $\Delta \to 0^+$ — so that p_align,0 and p_align,1 sit on the plateau while p_align,2 has entered the fall. This is the §6 "sharp near threshold, flat at Δ_1" specification, now attached to a named object.

Second, **the conjecture fixes shape, not value.** The $\Delta \to 0^+$ limit and monotonicity constrain the qualitative behaviour and guarantee suppression of a near-critical sector, but they do not by themselves yield p_align,2 ≈ 0.081. The tau sits at finite margin (representative $\Delta_2 \approx 3.2$), so the specific value still requires the functional form $p\_align(\Delta)$ and the capacities that locate Δ_2 on it. The conjecture is thus necessary scaffolding for the coupling route, not a derivation of the number.

The inherited and robust fact (H4) is that C_crit(g) rises steeply: 1.19, 7.74, 18.82, 34.44. If sector capacity C_g grows more slowly than the threshold, Δ_g shrinks with g, driving the third sector toward criticality and the fourth past it. This remains conditional on the undecided capacity-growth ordering. Representative capacities $C_0 \approx 18$, $C_1 \approx 29$, $C_2 \approx 22$ give $\Delta_0 \approx 16.8$, $\Delta_1 \approx 21.3$, $\Delta_2 \approx 3.2$ — illustrative only, non-monotonic, and entirely dependent on the chosen C_2. They show the picture is *available*, not that it holds.

The honest statement is therefore:

The reading this paper investigates places the tau residual in a reduced anchoring-effective coupling, p_eff,2 ≈ 0.081 p_eff,0, in the saturated third sector. Threshold proximity — if the capacities place the third sector near-critical — is the candidate cause of that reduction, by gating the alignment probability p_align,2. Both the gating and the near-criticality that would drive it remain to be derived.

This is stronger than the pure threshold story: it gives Θ_2 a home in an existing object, $\Theta_2 \leftrightarrow$ g_c,2 · p_align,2, instead of an invented functional.

---

6. The Open Target: Deriving the Coherence–Anchoring Ratio

The object to be derived is the sector coherence–anchoring ratio, and in the reading investigated here its coupling numerator:

$\Theta\_g = (p\_{eff,g} / K\_{c,g}) / (p\_{eff,0} / K\_{c,0})$, with the live hypothesis $\Theta_2 \approx p\_{eff,2} / p\_{eff,0} = g\_{c,2} \cdot p\_{align,2} / (g\_{c,0} \cdot p\_{align,0})$.

The constraints any derivation must meet are tight, and inherited from the muon's success:

- **Normalisation.** $\Theta_0 = 1$.
- **Muon tolerance.** $\Theta_1 \approx 1$ to within the precision the muon postdiction already enjoys — a few parts in a thousand, not a few tens of percent. The second sector must be essentially un-gated.
- **Tau target.** $\Theta_2 \approx 0.081$ — the third sector roughly twelve times less anchoring-effective than the ground sector.

The decisive advantage of the coupling reading is that p_eff = g_c · p_align is a computable object, not a free functional. The target is therefore concrete: derive the alignment probability p_align,2 (equivalently the anchoring-effective coupling g_c,2 · p_align,2) for the saturated third sector from the Role-4 eigenmode structure, and show that it falls to ≈ 0.081 of the ground value while p_align,1 stays within the muon tolerance of unity. A successful derivation converts Θ_2 from an inverted residual into a predicted ratio.

A pure-threshold reading, $\Theta_g = F(\Delta_g)$ with $F(\Delta_0) = 1$, $F(\Delta_1) \approx 1$, $F(\Delta_2) \approx 0.081$, and $F \to 0$ as $\Delta \to 0$, is the **anchoring-branch realization** of this target: it is what the coupling route reduces to if the suppression is routed through the threshold margin via p_align(Δ). Its difficulties are instructive. Linear and square-root margin dependences vary far too gently to fall from ≈1 to ≈0.081, so the required dependence is sharp near threshold (e.g. an essential-singularity profile $\exp(-a/\Delta)$ or a high inverse power); and sharpness alone is not enough, since the same steepness that suppresses the tau tends to suppress the muon unless Δ_1 sits in the flat region. With the representative margins of §5, no single-parameter monotonic F meets the muon tolerance and the tau target at once. This is not a defect of the hypothesis but a specification of what the alignment law p_align(Δ) — and the capacities that fix Δ — must produce.

No coupling value, alignment probability, or compression constant is reported here as derived. The coherence–anchoring ratio, its coupling and anchoring components, and the capacities feeding any threshold route are all Open. A structural candidate for the *value* of the residual — one that converts it from an arbitrary factor into a counting problem — is recorded in §6.1.

6.1 A Structural Candidate: Boundary–Interior Channel Reduction (Conjectural)

The coupling route leaves the alignment probability p_align,2 — equivalently Θ_2 under Route B — an unexplained residual. We record here a structural candidate for its origin, as a target rather than a result.

Boundary–Interior Reduction Conjecture. The minimal closure algebra possesses

$N_{loop} = 14$

independent reversible closure generators. We conjecture that these split as

$14 = 2_\partial + 12_{int}$,

with two generators belonging to the closure boundary structure and twelve to the interior. The boundary generators enforce closure orientation and endpoint matching: because they define the admissible closure frame itself, they are not available as independent anchoring-effective commitment channels. The number of anchoring-eligible interior channels is therefore

$N_{int} = 12$.

This would be stronger than a bare counting observation if the boundary/commitment distinction (step 2 below) can be established — resting the split on a physical difference between closure-frame generators and commitment generators rather than on the arithmetic $14 - 2$ alone.

Saturated-Sector Gating Conjecture. The third charged-lepton sector corresponds to saturation of the available refinement register. In the saturated regime most anchoring-eligible channels are already consumed by maintenance of the completed closure structure; we conjecture that exactly one anchoring-effective interior channel remains available for irreversible commitment. The first two sectors, being unsaturated, retain their full interior set, so $p_{align,0} \approx p_{align,1} \approx 1$ and the muon postdiction is preserved. The alignment ratio of the saturated sector is then

$p_{align,2} / p_{align,0} = 1 / N_{int} = 1/12$.

If the coupling strength and anchoring depth are approximately sector-stable ($g_{c,2} \approx g_{c,0}$, Route B), this carries straight through to

$\Theta_2 = 1/12 \approx 0.0833$.

Numerical comparison. The residual inferred from observation is $\Theta_2^{obs} \approx 0.0810$; the boundary–interior estimate gives 0.0833, a discrepancy of about 2.8%. This is more than an order of magnitude coarser than the muon postdiction (≈ 0.17%) — large enough that 1/12 cannot be taken as the answer, small enough to motivate the counting problem.

Relation to the Alignment Saturation Conjecture (§5.1). This is a *discrete* model for the *continuous* gating conjectured in §5.1: it realises the suppression of p_{align} at the saturated sector as an integer channel count rather than a smooth function $p_{align}(\Delta)$. The two are not equivalent. A pure channel count predicts exactly 1/12 (up to corrections) independent of the precise margin Δ_2; continuous threshold gating predicts a value tied to where Δ_2 falls on the alignment curve. The 2.8% gap is precisely where they are distinguished: the count overshoots, which either flags a correction to the counting or favours a continuous value near — but not at — 1/12. Reconciling the two (does the integer count emerge as the $\Delta \to 0$ limit of the gating curve, or is it an independent quantisation?) is itself part of the open problem.

Status (Conjectural). The argument is not a derivation. Three independent steps remain unproven:

1. that the fourteen closure generators split naturally into two boundary and twelve interior generators;
2. that boundary generators cannot act as anchoring-effective commitment channels;
3. that saturation leaves exactly one anchoring-effective interior channel.

The proposal is recorded as a structural target. Its value is conceptual localization, not numerical agreement: instead of treating Θ_2 as an arbitrary suppression factor, it reduces the residual to a concrete counting problem inside the closure algebra,

Θ_2 =? 1/12,

connecting the suppression to the closure algebra, the saturation interpretation of the third generation, and the void-coupling framework. A successful derivation would also have to show steps 1–3 return Θ_0, $\Theta_1 \approx 1$ for the unsaturated sectors and account for the residual 2.8%. If the three conjectural steps can be discharged, the tau suppression would emerge from the same closure structure that determines the refinement hierarchy itself.

---

7. Structural Decomposition

The result of §§3–6 is a clean factorisation of the charged-lepton mass:

charged-lepton mass = localisation hierarchy × coherence–anchoring ratio.

The electron is the ground refinement state. The muon is the first fully tightened state, its mass set by localisation alone with the coherence–anchoring ratio near unity. The tau is the saturation generation — it completes the available refinement structure, and the reading investigated here attributes its shortfall to a reduced anchoring-effective coupling, its mass being localisation tightening **opposed** by that suppression.

That asymmetry is the explanation of the non-uniform hierarchy:

e → μ is dominated by localisation tightening, un-gated; μ → τ is localisation tightening opposed by a suppression the investigated reading places in the effective coupling.

The first transition is accounted for. The second is decomposed and given a home, but not yet derived.

---

8. Relation to the Refinement-Efficiency Programme

The hierarchy programme establishes the chain

refinement → efficiency → reuse → completion density → mass.

The present paper adds the spacing layer. Refinement efficiency accounts for why the masses are ordered; Role-4 localisation accounts for why the ordering is large; a reduced anchoring-effective coupling in the saturated sector is proposed to account for why the second jump is much smaller than the first. Combined:

refinement depth → localisation tightening → mass amplification → coupling gating in the saturated sector → observed lepton spacing,

with the final arrow still carrying the open coherence–anchoring derivation of §6.

---

9. Status Table

Claim	Status
Three charged-lepton sectors	Inherited
No fourth charged-lepton sector	Inherited (Role-4 binding analysis)
Localisation law $L\_g = L_0 e\^(-\kappa g)$	Inherited
$\kappa = 8/3$	Conditional (Role-4/CP^2 geometric value)
Mass scaling $M\_g = \tilde{E}\_g / L\_g^2$	Inherited
Coherence–anchoring mass relation m ∝ p_eff/K_c, p_eff = g_c·p_align	Inherited (Void Coupling paper)
Muon estimate m_μ ≈ e^(16/3) m_e	Conditional — parameter-free postdiction to 0.17% given $\kappa = 8/3$ and $\Theta_1 \approx 1$
Tau suppression factor $\Theta_2 \approx 0.081$	Open (inferred from observation; no independent weight)
Decomposition Θ_2 = (p_eff,2/p_eff,0)·(K_c,0/K_c,2)	Conditional (on Ẽ_g ∝ p_eff/K_c)
Suppression sits in coupling (Route B) vs anchoring (A) / mixed (C)	Open (Route B investigated; motivated by gating; K_c,2 ≈ K_c,0 assumed, not shown)
Threshold proximity as cause of reduced p_align,2	Conjectural
Saturation–Gating Principle (saturation ⇒ p_align < 1)	Conjectural (weak form; defensible)
Alignment Saturation Conjecture (p_eff → 0 as $\Delta\_g \to 0^+$, monotone, plateau)	Conjectural (sharp form; sharpening of the Principle)
Sector capacities C_g (fix any threshold route)	Open (underived)
Derivation of p_align,2 / effective coupling g_c,2·p_align,2	Open (principal target)

Claim	Status
Boundary–Interior split $14 = 2_\partial + 12_{int}$	Conjectural (§6.1, target)
Saturated-sector gating → Θ_2 =? 1/12 (2.8% high)	Conjectural (derive-or-discard; not evidence)
Full tau mass derivation	Open

---

10. What Would Decide the Paper

The paper becomes substantially stronger if the following can be derived, in order of increasing strength.

Principal target — the effective coupling. The most direct route to Θ_2 runs through the coherence–anchoring ratio, not through anchoring capacity alone. Derive the alignment probability $p_{align,2}$ — equivalently the anchoring-effective coupling $g_{c,2} \cdot p_{align,2}$ — of the saturated third sector from the Role-4 eigenmode structure, and show that it falls to ≈ 0.081 of the ground value while $p_{align,1}$ stays within the muon tolerance of unity. This converts Θ_2 from an inverted residual into a predicted ratio, and it has the advantage that $p_{eff} = g_c \cdot p_{align}$ is an existing, computable object rather than a free functional.

Supporting — the cause of the gating. Establish *why* $p_{align,2}$ is suppressed. The candidate cause is threshold proximity, stated at two strengths in §5.1. The weak **Saturation–Gating Principle** (saturation ⇒ $p_{align} < 1$) is the defensible first step: proving it establishes that the saturated sector is gated at all and the unsaturated sectors are not, securing the qualitative mechanism. The sharp **Alignment Saturation Conjecture** ($p_{eff} \to 0$ as $\Delta_g \to 0^+$, monotone, plateau form) is the sharpening that makes the value reachable. Both in turn require the capacities: the margins $\Delta_g = C_g - C_{crit}(g)$ are undefined until the C_g are known, so a **Capacity Growth Theorem** — whether C_g rises more slowly than the steeply growing $C_{crit}(g)$ — is what decides the near-criticality of the tau and locates Δ_2 on the alignment curve. The cleanest sub-target is the *inequality* $dC/dg < dC_{crit}/dg$ over the bound sectors, which settles near-criticality without requiring the capacities in closed form.

Weaker but valuable. Prove that any admissible gating must sharply suppress near-threshold sectors — establishing the *mechanism* (reduced p_{align} near the edge) even short of the *number*.

A failure would occur if:

1. the Role-4 eigenmode calculation does not produce a suppressed $p_{eff,2}$ (or, equivalently, a suppressed $\tilde{E}_2$);
2. the suppression is forced into the anchoring depth $K_{c,2}$ in a way that cannot be derived, reproducing the capacity problem rather than resolving it;
3. the muon tolerance and the tau target cannot be met by one alignment law without an independent fit per sector;
4. the no-fourth-generation margin ($\Delta_3 < 0$) fails under improved dynamics.

11. Conclusion

This paper makes one semi-quantitative claim and isolates one open problem.

The claim is the muon. With the inherited localisation law and the geometric value $\kappa = 8/3$, the electron-to-muon jump follows with no adjustable quantity:

$m_\mu / m_e \approx e^{16/3} = 207.13$, $m_\mu^{est} \approx 105.84$ MeV (0.17% high).

This is a genuine postdiction if $\kappa = 8/3$ is fixed independently by the Role-4/CP^2 geometry.

The open problem is the tau. The repeated localisation step overpredicts the tau by more than an order of magnitude, so the third sector carries a suppression factor

$\Theta_2 = (m_\tau / m_e) / e^{32/3} \approx 0.081$.

This number is read off the measured tau mass. It is not explained — but it is no longer homeless. Identifying the sector energy with the coherence–anchoring ratio of the *Void Coupling* result, $m_g \sim e^{2\kappa g} \cdot (p_{eff,g}/K_{c,g})$, makes Θ_2 the value of a derived object rather than an invented constant. The reading this paper investigates places the residual in the effective coupling: the saturated third sector is roughly twelve times less anchoring-effective than the ground sector, $\Theta_2 \approx p_{eff,2}/p_{eff,0} = g_{c,2} \cdot p_{align,2}/(g_{c,0} \cdot p_{align,0}) \approx 0.081$, because only a restricted subset of its aligned states completes commitment. Threshold proximity is the candidate cause of that gating, and the capacities that would establish near-criticality remain underived; but the bottleneck is now a computable quantity, the alignment probability $p_{align,2}$, rather than a free functional.

The honest summary is therefore narrow and strong: the muon looks structural; the tau exposes the next theorem — and that theorem has an address.

The next task is precise:

Derive the anchoring-effective coupling of the saturated third sector — the alignment probability $p_{align,2}$, equivalently $g_{c,2} \cdot p_{align,2}$ — from the Role-4 eigenmode structure, and show it falls to ≈ 0.081 of the ground value while $p_{align,1}$ stays within the muon tolerance of unity. Threshold proximity is the candidate cause; settling it requires the capacity-growth ordering dC/dg vs dC_{crit}/dg.

Until that coupling is derived, the tau remains localised, not derived. If it succeeds, the charged-lepton hierarchy moves from a single muon-scale postdiction to a genuine semi-quantitative derivation of the spectrum.