Alignment Gating in the Saturated Role-4 Sector

Boundary–Interior Closure Channels, Saturated Participation, and the Tau Suppression Target

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The companion threshold-compression paper showed that the electron-to-muon mass jump follows from geometric localisation: each deeper generation is held in a smaller closure region, and mass rises as that region tightens. With the inherited Role-4 value $\kappa = 8/3$, the muon comes out within about 0.17% of its measured mass, with no quantity adjusted to fit.

The tau is harder. If the same localisation step simply repeated, the tau would weigh about 21.9 GeV. It actually weighs about 1.78 GeV — more than an order of magnitude lighter than the raw law predicts. The companion paper named the missing factor but did not derive it:

$\Theta_2 \approx 0.081$.

That number was a residual — what localisation leaves unexplained, read off the measured tau mass.

This paper asks one question: can that residual be a *counting* effect inside the closure structure, rather than a fitted constant?

The idea rests on a distinction the companion paper introduced. A particle's mass depends not only on how tightly it is localised, but on how effectively its internal states couple to irreversible commitment — on what fraction of its aligned states actually anchor. In the saturated third generation, most of the available closure channels may already be spent maintaining the completed closure structure. The tau would then be structurally tight but coupling-starved.

The closure algebra offers a candidate count. Prior Role-4 work gives a minimal closure structure with fourteen reversible generators. This paper proposes that two of those are *boundary-frame* generators — they define where the closure begins and ends, its orientation and matching — and so are not themselves interior commitment routes. That leaves twelve interior channels eligible to anchor.

The central conjecture is then simple to state:

In the saturated third sector, only one of the twelve interior channels stays anchoring-effective.

That gives $p_{align,2} \approx 1/12 = 0.0833$, against the observed residual $\Theta_2 \approx 0.0810$ — about 2.8% high. That is not exact, and this paper does not pretend it is. The gap is more than fifteen times coarser than the muon postdiction, and the paper is explicit that 1/12 is therefore *not* claimed as the tau's value.

What the closeness does is change the *kind* of problem the tau poses:

The tau residual reduces to a specific counting question — why does the saturated third sector retain exactly one anchoring-effective interior channel out of twelve? — rather than an arbitrary suppression factor.

If that count can be derived from the closure architecture, the tau suppression stops being a fitted number and becomes a consequence of the same structure that fixes the refinement hierarchy. If it cannot, the paper has at least said precisely what would have to fail. Either way, the tau is given an address, not a solution.

One honesty note carried in plain sight: the whole construction hangs on the inherited count of fourteen generators and on its split as 2 + 12. The count itself is inherited from an established closure structure; what is not settled is *how* the fourteen relate to the seven — and on one natural reading that relationship would undercut the split. That tension, and the split, are marked as open.

Contents

Abstract

The companion charged-lepton threshold-compression result isolated the tau suppression factor

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

after the Role-4 localisation law $L_g = L_0 e^{(-\kappa g)}$, $\kappa = 8/3$, postdicted the electron-to-muon ratio at the 0.17% level. The residual Θ_2 was identified with a coherence–anchoring ratio,

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

with $p_{align,2}$ left open. This paper develops a structural candidate for $p_{align,2}$ and grades each step.

Two conjectures are advanced. The **Boundary–Interior Reduction Conjecture** proposes that the inherited minimal Role-4 closure structure, with $N_{loop} = 14$ reversible generators, decomposes into two boundary-frame generators and twelve interior commitment generators, $14 = 2_\partial + 12_{int}$; the boundary generators set closure orientation and endpoint matching and so are excluded from the anchoring-eligible count, leaving $N_{anchor} = 12$. The **Saturated Participation Conjecture** proposes that in the saturated third refinement sector, maintenance of the completed closure register leaves exactly one of the twelve interior channels anchoring-effective. Under sector-stable coupling strength and anchoring depth (the coupling-dominated route, assumed not shown), this yields

$p_{align,2} \approx 1 / 12 = 0.0833$,

a value 2.8% above the inferred residual $\Theta_2 \approx 0.0810$ and correspondingly **not claimed as a prediction**. The result is presented as a derivation target with its dependencies named: the count $N_{loop} = 14$ is inherited from the closure-geometry programme as a structural consequence of

the K = 7 architecture, but its *interpretation* relative to K = 7 is unsettled and on the leading reading would threaten the split; the split 2 + 12, the boundary-exclusion principle, and the one-channel survival rule are each conjectural; and the coupling-dominated route assumes the suppression does not sit in the anchoring depth. If these can be established from the Role-4 eigenmode structure, the tau residual becomes a closure-channel counting effect rather than an arbitrary suppression factor.

---

1. Introduction

The charged-lepton hierarchy contains two problems of different character, and this paper concerns only the second.

The first is the large electron-to-muon jump, $m_\mu / m_e \approx 206.77$. The companion paper argued this jump is captured by the Role-4 localisation law: with $L_g = L_0 e^{-\kappa g}$, $\kappa = 8/3$, and the sector-rescaling relation $M_g = \tilde{E}_g / L_g^2$, the ratio is

$$m_\mu / m_e \approx e^{2\kappa} = e^{16/3} = 207.13,$$

giving a muon estimate high by about 0.17% with no adjustable quantity once κ is fixed by the Role-4/CP^2 geometry. That postdiction is inherited here, not re-argued; its sole conditionality is the geometric origin of $\kappa = 8/3$.

The second problem is the tau. The same step iterated without correction gives

$$m_\tau \sim m_e \cdot e^{32/3} \approx 21.9 \text{ GeV},$$

against an observed $m_\tau \approx 1.777$ GeV. The third sector therefore carries a suppression

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

The companion paper made the status of Θ_2 explicit: it is a residual, computed from the measured mass by inversion, carrying no independent evidential weight. It also gave the residual a home, in the coherence–anchoring mass relation

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

and identified the principal open target as the alignment probability $p_{align,2}$ of the saturated third sector.

This paper takes up that target along one route. It asks whether $p_{align,2}$ can be derived as a *participation count* — a ratio of surviving anchoring channels to available ones — rather than as a value on a smooth threshold curve. The structural question is whether the inherited closure count licenses

$p_{align,2} \approx 1/12$.

The answer offered is conditional and explicitly graded: the route is coherent and lands within 2.8% of the residual, but it rests on an inherited count whose interpretation relative to $K = 7$ is unsettled, and on conjectures whose proofs are the work handed forward. The contribution is to convert "why 0.081?" into "why one channel out of twelve?" — a sharper question with a definite home in the closure algebra — and to state exactly what would refute it.

---

2. Inherited Results

The following are imported and used, not re-derived. Each carries its source's grade; where a source is not yet pinned, that is recorded as owed.

H1 — Role-4 localisation

Charged-lepton generation sectors are indexed $g = 0, 1, 2$, with localisation scales $L_g = L_0 e^{-\kappa g}$. The geometric value carried from the companion paper is $\kappa = 8/3$. **[Inherited; κ Conditional on the Role-4/CP^2 derivation.]**

H2 — Sector mass scaling

The sector-rescaling theorem gives $M_g = \tilde{E}_g / L_g^2$, hence

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

[Inherited.]

H3 — Coherence–anchoring mass relation

The *Void Coupling as Phase Coherence* result separates effective coupling from anchoring depth:

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

where K_c is the anchoring depth (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). **[Inherited — Void Coupling paper.]**

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

H4 — Tau residual as a coherence–anchoring ratio

Normalising to the ground sector,

$\Theta_g = (p\_{eff,g} / K\_{c,g}) / (p\_{eff,0} / K\_{c,0})$, $\Theta_0 = 1$,

with $\Theta_2 \approx 0.081$ for the tau. This decomposes as

$\Theta_2 = (p\_{eff,2} / p\_{eff,0}) \cdot (K\_{c,0} / K\_{c,2})$.

The companion paper left open whether the suppression sits in the coupling, the anchoring depth, or both, and investigated the **coupling-dominated route**

$\Theta_2 \approx p\_{eff,2} / p\_{eff,0}$, and further $\Theta_2 \approx p\_{align,2} / p\_{align,0}$,

on the assumption that g_c and K_c do not supply the dominant sector variation ($g\_{c,2} \approx g\_{c,0}$, $K\_{c,2} \approx K\_{c,0}$). **That assumption is carried, not proved.** This paper works inside the same route and inherits its conditionality. **[Inherited framing; the coupling-dominated reduction is Conditional, not shown.]**

H4a — The three routes are one factor in three places

It is worth making explicit what the routes A, B, C are and are not, because the choice of B is easy to mistake for an evidential preference when it is a strategic one. The decomposition $\Theta_2 = (p\_{eff,2} / p\_{eff,0}) \cdot (K\_{c,0} / K\_{c,2})$ is a single measured factor — the tau's 12.35-fold suppression relative to a bare localisation step — written as a product of two ratios. The three "routes" are not three hypotheses with different physical content; they are three *placements* of that one factor:

Route	Where the factor sits	Required value	Reading
A — anchoring-depth	all in K_c	$K\_{c,2} / K\_{c,0} \approx 12.35$	third sector needs ~12× more successful micro-events to commit
B — coupling	all in p_eff	$p\_{eff,2} / p\_{eff,0} \approx 0.0810$	third sector is ~12× less anchoring-effective (the route investigated)
C — mixed	shared	product = 0.0810 (e.g. ×3.5 depth, ÷3.5 coupling)	factor split between the two

Three things follow, and none of them favours B on evidence.

First, the routes are arithmetically equivalent: each reproduces $\Theta_2 \approx 0.081$ exactly, by construction, because each is the same number reparametrised. No route fits *better*; the 12.35 is not in dispute, only its address is. So the table cannot, and does not, argue that the suppression *is* in the coupling.

Second, the factor the data demands is 12.35, not 12 — visible in Route A's required depth ratio just as in Route B's inverse. This is the same non-integer that the discrete channel count (§5–§6) provably cannot hit, and it appears here independently of the channel picture: whichever route carries the factor, the factor itself is non-integral. The integer 12 is a property the *count* proposes, not a property the *data* exhibits.

Third, the paper's reason for working in Route B is therefore not "B fits best" but "B connects to a structural counting problem (the boundary–interior channels of §3–§5) that A and C do not." Route A would require a derived account of why the anchoring depth grows ~12× at saturation — reproducing the companion paper's capacity problem in new symbols — and the anchoring programme has emphasised that K_c varies strongly between modes, so sector-stable coupling is no more given than sector-stable anchoring. **The choice of B is strategic, not evidential: $K_{c,2} \approx K_{c,0}$ is assumed so that the factor lands in p_{eff} where a channel count can address it; it is not shown that K_c is in fact sector-stable.** Routes A and C are not refuted here; they are set aside, and §9 (F4) records that distinguishing them from B requires the per-sector eigenmode computation of g_c, K_c, p_{align} — a debt this paper does not discharge.

H5 — Minimal closure-algebra count

The Role-4 closure architecture is taken to contain

$N_{loop} = 14$

independent reversible closure generators.

[Inherited from the closure-geometry programme.] This count is load-bearing for everything below. It is not a free numerical input: as recorded in H5a, the closure-geometry programme establishes a K = 7 closure architecture (six boundary constraints and one hub constraint) whose response structure carries fourteen oriented traversals and a nullity-one mode, so $N_{loop} = 14$ is a structural consequence of that architecture. What remains owed is not the existence of the count but the *interpretation* of its relation to K = 7 — and that interpretation bears directly on the construction below.

A consistency obligation between two factors of two. The fourteen and the seven are related as 14 = 2K, and there are two distinct 2's the construction must keep apart. The first is the doubling in 14 = 2K: H5a notes the leading reading is that the fourteen are *oriented* traversals of the seven closure constraints — i.e. each constraint carries a directed pair. The second is dim $\Gamma_\partial = 2$, the boundary frame of §3. The danger is that these may be the *same* two: if the doubling is orientation-pairing, and orientation is precisely what §3 assigns to the boundary frame, then subtracting "2 boundary generators out of 2K" removes a frame already distributed through the count, and the denominator double-counts — the twelve would not survive. Sharpening this, the inherited architecture's *six boundary constraints* sit suggestively close to the frame structure §3 invokes, and whether the boundary frame is drawn from those six (or is orthogonal to them) is part of the same unresolved question. If instead the doubling is closure-preservation structure (a formation/preservation split, on a formation/preservation axis rather than a frame axis), the two 2's are independent and the construction stands. The matter is *neither confirmed nor excluded*, and it is the cheapest flag for the highest-leverage failure: if the 2's coincide, the entire $N_{anchor} = 12$ denominator collapses. The obligation is therefore explicit: **any use of $N_{loop} = 2K$ must show the provenance of its factor of two distinct from the boundary dimension $\dim \Gamma_\partial = 2$ of §3** — and because the leading provenance (oriented traversals) is the branch that most threatens the split, this is urgent rather than precautionary, and is the first thing the next paper must settle.

H5a — Provenance of the Fourteen-Generator Census

The count $N_{loop} = 14$ is inherited from the closure-geometry programme rather than introduced independently here. Earlier work establishes a $K = 7$ closure architecture consisting of six boundary constraints and one hub constraint, together with a fourteen-traversal response structure and a nullity-one mode. The present paper therefore treats $N_{loop} = 14$ as a structural consequence of the established closure architecture rather than as a free numerical input.

[Provenance grade — to be pinned.] A distinction must be held here that the prose above does not, on its own, enforce. That the count is fourteen is the inherited result the construction needs. The *internal* decomposition stated above — six boundary constraints plus one hub, a fourteen-traversal response structure, a nullity-one mode — is the load-bearing detail, and it must be confirmed to read *exactly* this way in the closure-geometry source, with a locator, before it is treated as settled. If the source states the six-plus-one decomposition and the nullity-one mode in these terms, this subsection is a genuine inherited result and Step 0 becomes a derivation task against a fixed structure. If any of that internal detail is instead a reconstruction assembled to make the 2K reading cohere, then it is the softest link in the chain *precisely because it no longer advertises its softness* — the earlier "source citation owed" confessed its own gap, whereas a confidently-stated architecture does not. Until the decomposition is checked against the source paper and cited, the count fourteen is inherited but its stated internal structure is carried as **reconstruction, not confirmed result**, and is not promoted past that grade.

The detailed relationship between the seven closure constraints and the fourteen traversals is not required by any argument below. Future work may clarify whether the fourteen-generator census should be understood as oriented traversals of the seven closure constraints, as closure-preservation structure (a formation/preservation split, in which seven generators realise the closure and seven preserve it against immediate reversal), or as another equivalent realization of the same $K = 7$ architecture.

These readings are not interchangeable for the construction, and the difference is exactly the consistency obligation flagged in H5. If the fourteen are *oriented traversals* of the seven constraints, the doubling is a frame-geometric one — and frame geometry is what §3 assigns to the boundary generators, so the split $14 = 2_\partial + 12_{int}$ risks removing a frame already counted, collapsing the denominator. If instead the doubling is *closure-preservation* — a formation/preservation distinction between forming and holding the closure — then the boundary frame is a separate structure living within one layer, the two factors of two are independent, and the denominator survives. The first reading is the leading one suggested by the inherited architecture and is therefore the one that puts the count most at risk; the second would discharge the obligation favourably, and would additionally tie the doubling to the programme's commitment-and-irreversibility structure concretely, since preservation-against-reversal *is* the irreversibility the anchoring programme tracks. Which realization is correct is open, and is what the next paper's derivation must decide.

A forbidden identification: the surviving channel is not the nullity. The architecture above carries a nullity-*one* mode, and §5/§9 require $N_{survive}$ = *one*. The two 1's invite an identification — that the single surviving anchoring channel of the saturated sector *is* the architectural nullity — and that identification is not merely unproven but **forbidden by F5**. The nullity-one mode is a fixed feature of the $K = 7$ architecture, hence sector-independent: it is the same nullity for $g = 0, 1, 2$. If $N_{survive}$ were sourced from it, the survival count would read 1, 1, 1 across all three generations, gating the muon and electron and destroying the 0.17% postdiction. But F5 demands the count read 12, 12, 1 — a discontinuity that switches on only at saturation. A sector-independent feature cannot produce a sector-dependent step. Therefore the surviving channel of §5 must come from something that *distinguishes* the saturated sector from the unsaturated ones, which the nullity by construction does not. This is recorded as an a-priori exclusion in the same spirit as the monotone-in-occupancy exclusion of F5: a future derivation of $N_{survive} = 1$ must not source it from the closure nullity, on pain of F5. The numerical coincidence of the two 1's is a trap, not a lead.

---

3. The Boundary–Interior Reduction

The minimal closure structure supplies fourteen reversible generators (H5). The construction turns on a single distinction: not every generator that participates in *reversible* closure need be available as an *irreversible* anchoring channel.

Closure requires a frame. Any closed structure must establish (i) an orientation of traversal and (ii) a matching between initiation and completion. These are not interior commitment operations; they are the conditions under which interior commitment is well-defined. The proposal is that the frame is carried by a fixed, small number of generators, distinct from the commitment-carrying interior.

Boundary–Interior Reduction Conjecture

The fourteen reversible generators decompose into two boundary-frame generators and twelve interior commitment generators:

$$\Gamma_{min} = \Gamma_{\partial} \oplus \Gamma_{int}, \quad \dim \Gamma_{\partial} = 2, \quad \dim \Gamma_{int} = 12,$$

equivalently

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

The two boundary generators ∂_-, ∂_+ admit several equivalent readings — entry/exit, inner/outer matching, initiation/termination, orientation/return — and the exact representation depends on the closure formalism. What the conjecture asserts is the functional split: Γ_∂ defines the closure frame; Γ_{int} carries the anchoring-eligible interior routes. Hence the anchoring-eligible count is

$N_{anchor} = 12$.

Grade: Conjectural. Two things must be shown to establish it: that exactly two generators are frame-setting (not one, not three), and that they are identifiable in the Role-4 closure algebra as such. The "two" is currently motivated only by the two frame conditions (orientation, matching); whether each condition consumes exactly one generator, or whether they share or multiply, is not fixed. This is the first place the count 12 could move — and §9 (F2) tracks where Θ_2 lands if it does.

---

4. Boundary Exclusion

Granting the split, the boundary generators must also be *excluded* from the anchoring count for the number 12 to be the right denominator. This is a separate claim from the split itself, and is stated separately.

A commitment channel is a route by which an aligned state completes an irreversible micro-event. A boundary-frame generator instead defines the admissible domain in which such routes occur. The proposal is that these roles cannot coincide.

Boundary Exclusion Principle

A generator whose role is to define closure orientation or endpoint matching cannot simultaneously serve as an independent anchoring-effective commitment channel:

$\Gamma_\partial \cap \Gamma_{anchor} = \emptyset$, hence $\Gamma_{anchor} \subseteq \Gamma_{int}$, $N_{anchor} \leq 12$.

Grade: Conjectural. The principle is structurally natural — frame-setting and commitment are different functional roles — but "different role" does not by itself force "disjoint generator." A generator could in principle participate in both the frame and the interior dynamics; ruling that out is what a proof must do. Note the principle gives an *inequality*, $N_{anchor} \leq 12$: it caps the denominator at twelve but does not by itself forbid the frame generators from leaking partial anchoring weight. The clean count $N_{anchor} = 12$ requires the exclusion to be exact, which is stronger than the role-distinction argument delivers on its own.

The principle does not deny that boundary generators are necessary — without them closure is undefined. It says only that necessity-for-closure is not availability-for-anchoring. Boundary generators define the frame; interior generators carry commitment.

5. Saturated Participation

The split and exclusion fix the *denominator* at twelve. The suppression itself lives in the *numerator*: how many of those twelve interior channels stay anchoring-effective in each sector. This is where the saturation of the third generation enters.

In the refinement-efficiency picture, the generation loads are 1, 2, 4 with cumulative occupancy 1, 3, 7. The third sector closes the $K = 7$ register — it is the saturation generation. The proposal is that channel participation depends on whether the register is saturated.

A reader may reasonably ask why saturation is counted in sevens but participation in twelves — if the register is seven and it is full, why is the denominator not seven? Because two different spaces are in play, one being *filled* and the other *taxed*. Occupancy fills the seven-element closure register (six boundary constraints plus one hub, H5a); that is what reaches $K = 7$ at the third sector. The participation fraction, by contrast, lives on the interior *traversal* space — the twelve anchoring-eligible channels of §3, the 12 of $14 = 2_\partial + 12_{int}$. Saturation is the statement that the seven-constraint register is complete; the consequence the mechanism proposes is that maintaining that completed register then taxes the twelve-dimensional interior, removing all but one channel from anchoring. The seven and the twelve are distinct counts on distinct spaces — one the register that fills, one the channel space that is taxed once it is full — not competing values for the same denominator. (This is a different pair from the two-factors-of-two obligation of H5, which concerns the 2 in 2K versus $\dim \Gamma_\partial$.)

For the first two sectors the register is unsaturated; interior channels remain freely available as alignment routes, so

$p_{align,0} \approx 1$, $p_{align,1} \approx 1$.

This preserves the muon postdiction — the second sector is essentially un-gated, exactly the tolerance the muon's 0.17% agreement demands. For the third sector the register is saturated: most interior channels are spent maintaining the completed closure, and only a residual route remains.

Saturated Participation Conjecture

In the saturated third refinement sector, exactly one of the twelve anchoring-eligible interior channels remains effective for irreversible commitment:

$p_{align,2} = N_{survive}/N_{anchor} = 1/12$, $N_{survive} = 1$, $N_{anchor} = 12$.

Grade: Conjectural — and this is the crux of the paper. The justification currently available is a *consistency* argument, and it must be labelled as such rather than dressed as a derivation: if no channel survived ($N_{survive} = 0$) the sector would be inert rather than massive; if many survived the tau would not be strongly suppressed. So N_survive = 1 is the value *consistent with* a tau that is both massive and strongly suppressed. That is a bracketing argument — it shows one channel is what the answer *needs*, not what saturation *produces*. Promoting it past Conjectural requires deriving N_survive = 1 from the saturation dynamics (the Role-4 eigenmode structure of the completed register), independently of the tau mass it is meant to explain. Until that derivation exists, the "exactly one" is the principal theorem target, and the danger to guard against is precisely that the residual fixes the count rather than the count predicting the residual.

The honest statement: the denominator twelve is a closure-algebra count (Conjectural via §3–§4); the numerator one is, at present, the value consistency requires (Conjectural, consistency-driven). The product 1/12 is therefore a structurally-motivated target, not a derived participation fraction.

---

6. Tau Estimate from Boundary–Interior Gating

Working inside the coupling-dominated route (H4) with $p_{align,0} \approx 1$ and the conjectured $p_{align,2} = 1/12$,

$\Theta_2^{BIG} \approx p_{align,2} / p_{align,0} = 1 / 12 = 0.0833$.

Through the mass formula $m_g \sim m_e \cdot e^{(2\kappa g)} \cdot \Theta_g$, the tau estimate is

$m_\tau^{BIG} = m_e \cdot e^{(32/3)} \cdot (1 / 12)$.

With $m_e = 0.51099895$ MeV and $e^{(32/3)} \approx 4.290 \times 10^4$,

$m_\tau^{BIG} \approx 1827$ MeV ≈ 1.83 GeV, against $m_\tau^{obs} \approx 1776.86$ MeV,

high by about 2.8%. Equivalently, $\Theta_2^{BIG} = 0.0833$ against $\Theta_2^{obs} \approx 0.0810$, a relative gap

$(0.0833 - 0.0810) / 0.0810 \approx 2.8\%$.

The comparison must be read at its true weight. This is more than fifteen times coarser than the muon postdiction's 0.17%, and the muon postdiction is itself only as strong as the inherited $\kappa = 8/3$. A 2.8% agreement on a single ratio, reached through two unproven conjectures and one assumed route, is **not evidence that 1/12 is the tau's value**. Its content is narrower: the leading structural count in the closure algebra lands within a few percent of the residual, which is enough to make the counting problem worth posing and not enough to claim it solved. The number 1/12 is offered to be derived or discarded (§9, §12), not relied upon.

---

7. Interpretation

The construction, taken as a conditional picture, organises the three charged leptons as follows.

The electron is the ground sector — unsaturated, its aligned closure channels participating normally. The muon is the first tightened sector — also unsaturated, its channels effectively available, which is why the bare localisation law works cleanly for the electron-to-muon jump. The tau is the saturated sector — structurally tighter, and so heavier by localisation, but with most interior channels spent maintaining completion rather than generating fresh irreversible commitment.

The asymmetry of the hierarchy then reads:

$e \to \mu$: localisation tightening, un-gated; $\mu \to \tau$: localisation tightening opposed by saturation gating of the effective coupling.

So the tau is not light because it is *less* structured. In the reading investigated here it is light *relative to the raw localisation prediction* because it is over-structured: saturation consumes the channels that would otherwise anchor. Compactly:

The tau is heavy by localisation but suppressed by saturation.

This sentence is an interpretation of the conjectured mechanism, not an established result; it inherits the grade of §5.

---

8. Relation to Threshold Compression

The companion paper offered two readings of the tau residual: a *continuous* threshold-gating reading, in which p_align falls along a smooth curve as the binding margin $\Delta_g = C_g - C_{crit}(g)$ shrinks toward zero, and a *discrete* channel-count reading. This paper develops the discrete reading. The two are genuinely different, and the difference is testable rather than cosmetic.

The continuous reading (Alignment Saturation Conjecture of the companion paper) sets

$\lim_{\Delta_g \to 0^+} p_{eff}(g) = 0$, with $p_{align,2} = F(\Delta_2)$,

so the value depends on *where* the tau's margin Δ_2 falls on the alignment curve. The discrete reading sets

$p_{align,2} = 1/12$,

so the value depends on the *channel count* and is, to leading order, independent of the precise margin.

This independence is exactly where the 2.8% gap stops being a near-miss to be apologised for and becomes a structural statement. H4a already established the sharp version arithmetically: the data demands a denominator of 12.35, and an integer count cannot reach it *by construction*. The discrete reading therefore predicts the integer 12 exactly and rigidly — and predicts *nothing whatever* about the residual 0.35 of denominator slack. The slack is not noise around a good fit; it is structure the discrete count does not contain and cannot, even in principle, supply. *If the discrete reading is the leading term*, then at least one of the following is not an optional refinement but a **forced** ingredient supplying the slack, because a pure integer count provably misses 0.0810:

1. a correction to the count itself (the leading 1/12 receives sub-leading structure);
2. finite-margin leakage — a small continuous correction riding on top of the discrete count, i.e. the two readings are not exclusive but layered;
3. a residual share of the suppression carried by K_c after all (the coupling-dominated route is approximate, not exact).

These three are *sources of the 3%* — they presuppose the discrete count is correct and supply what it cannot. A fourth possibility is of a different kind: that the boundary–interior counting is simply wrong, in which case there is no integer 12 to correct and the forced-slack framing does not apply at all. (1)–(3) answer "where does the 3% come from?"; the fourth denies the premise. The two should not be read as a single menu.

This reframes the relationship between the discrete count and the residual. If the count holds, it is not "close to 0.081"; it is *exactly* the integer 12 and *necessarily* accompanied by additional structure of order 3%. Step 4 (§12) is therefore not polishing a near-miss — it is supplying mandatory structure, and which of (1)–(3) supplies it is a physical claim that distinguishes the discrete and continuous readings, not a cosmetic choice. The honest headline is that the paper claims the integer and owes the decimal — *conditional on the integer being the right leading term at all*.

The cleanest experimentum crucis between the readings is therefore not the tau alone but the *margin-dependence*: if a fourth or heavier hypothetical sector existed, the discrete reading would predict a count-determined jump while the continuous reading would predict a margin-determined one. In the actual three-sector spectrum the distinguishing question reduces to whether the 2.8% is structured (favouring 1, 2) or smooth (favouring 2, 3). This paper does not decide among them; its purpose is to put the residual into a form where the question is well-posed.

A note on whether the discrete count could *be* the $\Delta \to 0$ limit of the continuous curve: this would require the smooth $F(\Delta)$ to terminate at a rational $1/N_{int}$ as $\Delta \to 0^+$, i.e. for the continuous gating to quantise at the edge. That is an attractive unification of the two readings but is itself unproven, and is recorded as part of the open problem rather than as a bridge already built.

9. Failure Modes

The mechanism is refuted, in whole or in the stated part, if any of the following holds.

F1 — No boundary–interior split. If the fourteen generators do not decompose as $14 = 2_\partial + 12_{int}$, the 1/12 count has no foundation. This is the most basic failure and the first thing a proof must secure.

F2 — Wrong boundary number. If the frame is set by a number of generators other than two, the denominator moves, and the fit degrades asymmetrically. If three generators are frame-setting, the denominator is eleven, giving 1/11 = 0.0909, about 12% *above* the residual. If the boundary generators are not excluded at all (they anchor like any interior channel), the denominator is fourteen, giving 1/14 = 0.0714, about 12% *below*. The nearest competitor, though, is on the low side and closer than either: if orientation and matching *share* a single frame generator — the branch §3's own grade note explicitly leaves open ("whether each condition consumes exactly one generator, or whether they share") — the frame is one-dimensional and the denominator is thirteen, giving 1/13 = 0.0769, only about 5% below the residual. So the fit ordering by closeness is twelve (2.9%), then thirteen (5.0%), then eleven and fourteen (~12%). Twelve fits best, but thirteen is a genuine competitor that the *shared-frame* reading of §3 would select, so "the fit is specific to twelve" must be stated as "twelve fits best, with the shared-frame thirteen the nearest alternative" — and the choice between them is exactly the unresolved question in §3 of whether the two frame conditions consume one generator each or share one. A small error in the boundary count is not a small error in Θ_2, and the closest such error (the shared-frame thirteen) is one §3 has not excluded.

F3 — Saturation does not leave exactly one channel. If saturation leaves two interior channels, 2/12 = 0.1667 — roughly double the residual, far too weak a suppression. If it leaves none, the sector cannot produce a finite massive state at all. The theory needs $N_{survive} = 1$ exactly, and (per §5) currently *infers* that from the answer rather than deriving it; F3 is the failure mode the §5 caveat is guarding.

F4 — Suppression sits in anchoring depth. If the residual is dominantly $K_{c,0} / K_{c,2}$ rather than $p_{eff,2} / p_{eff,0}$, the coupling-dominated route (H4) is the wrong route and the boundary–interior alignment count is not the operative mechanism, whatever its numerical coincidence. Distinguishing the routes requires the eigenmode computation of g_c, K_c, p_{align} per sector — the same debt the companion paper records.

F5 — Muon gating. If the mechanism gates the muon (or electron) sector appreciably, $p_{align,1}$ falls below unity and the 0.17% muon postdiction is destroyed. The mechanism must therefore distinguish unsaturated from saturated sectors sharply — which is precisely what the saturation-conditioned numerator (§5) is for, and what any derivation of $N_{survive}$ must respect: it has to return $N_{survive} = N_{anchor}$ (full participation) for $g = 0, 1$ and $N_{survive} = 1$ for $g = 2$.

This has a consequence worth drawing out, because it eliminates a whole class of otherwise-natural derivations in advance. The demand is not for a *value* but for a *discontinuity*: N_survive must read 12, 12, 1 across g = 0, 1, 2 — a step function at register closure, not a curve. Any saturation dynamics that is *smooth* in occupancy — anything in which p_align degrades continuously as the register fills from 3/7 toward 7/7 — would gate the muon at least slightly, and even a slight gating of p_align,1 destroys a postdiction good to 0.17%. So the admissible mechanism must be insensitive to *partial* filling and switch on only at *exact* closure. The target of Step 3 (§12) is therefore not "derive one surviving channel" but "derive a threshold collapse that activates only at K = 7 saturation and not before" — a closure-*completion* effect, not a filling-*fraction* effect. This rules out, a priori, any monotone-in-occupancy participation law, and tells the next derivation where not to look.

---

10. Predictive Status

The construction so far has one target and one answer: it asks why the tau is suppressed by $\Theta_2 \approx 0.081$ and replies that perhaps only one of twelve interior channels survives saturation. That is a structural hypothesis, not yet a test. A test requires the same survival rule to constrain an observable it was not built to fit. This section states plainly what the mechanism would have to do to earn that status, and where it currently stands against each demand. The honest summary is that the paper does not yet pass the strongest internal test it sets itself.

Test 1 — Generation-wide suppression (not the tau alone)

If the gating is tied to third-sector saturation rather than to the tau specifically, an analogue of the suppression should appear across the third generation — bottom and top, not only tau. This is the natural first test, and it is not yet runnable, for two reasons.

First, Θ_2 is defined only after dividing out the bare localisation prediction e^(2κg). Asking whether the bottom or top carries the same residual requires a quark localisation law with its own κ_q, derived as the companion paper derived κ = 8⁄3 for the leptons. No such law exists yet, and the quark mass pattern does not obviously fit a single exponential ladder. **[Open — prerequisite; the quark localisation law is owed before Test 1 is well-posed.]**

Second, run naively the top quark looks hostile to the mechanism — but the threat is route-conditional, and naming the condition is what keeps this section consistent with its own shared-denominator point below. The top is the least-suppressed fermion in the spectrum, with a coupling of order unity, the one fermion that looks fully un-gated. That reading, however, *assumes route B*: it infers from the top's large mass a large p_eff, and under sector-stable g_c a large p_align — hence un-gated, hence in contradiction with a shared 1/12. The contradiction therefore exists only if route B holds for the top. Under the route-A/C reading — the one cross-species comparison forces anyway, and the same reading the shared-denominator escape below relies on — the top's large mass sits in g_c and K_c, and p_align,top can be 1/12 like its siblings with *no* contradiction. So the precise statement is: the top refutes the **route-B generation-wide**

form (heavy = un-gated = not 1/12), but is **consistent with the route-A/C shared-denominator form** (heavy by coupling/depth, p_align still gated). The top is the sharpest threat to one form of the mechanism and no threat at all to the other, and the two cannot both be asserted without saying which. The participation companion takes the further step that dissolves the ambiguity entirely: the quark localisation baseline is not independently validated (it fails two generations early), so the top's mass cannot be read as evidence about p_align at all until a validated κ_q exists — which is why that paper restricts to leptons rather than adjudicating the top here. **[Open — route-conditional: adversarial to the route-B generation-wide form, consistent with the route-A/C form; not adjudicable until the quark baseline is independently validated.]**

Test 2 — Sharp transition, suppression nowhere else

The mechanism requires generations one and two unsaturated and only the third saturated, so the suppression must be absent below the third sector — charm and strange ungated, exactly as the muon is (F5). This is F5 generalised across species. It is sound as a demand and shares Test 1's prerequisite: checking whether charm carries a residual again needs the quark law. Note this is the same content as the F5 step-function requirement (§9): a participation law smooth in occupancy would gate charm partially and die, just as it would gate the muon. **[Conditional on the quark law; equivalent to the F5 closure-completion requirement.]**

Test 3 — Do the numbers 12 and 1 recur independently?

This is the strongest internal test, and the paper does not currently pass it. The honesty obligation is to say so.

The denominator 12 has no established *independent* home — distinct, that is, from the closure count it was drawn from. The count fourteen is inherited (H5a); but 12 = 14 − 2 does not *also* fall out of the refinement census (loads 1, 2, 4; cumulative 1, 3, 7) or out of the flavour-mixing structure in any clean way shown here, so it does not recur anywhere it was not placed. The numerator 1 is in a sharper bind. It is tempting to source N_survive = 1 from the architectural nullity-one mode of H5a — but that identification is **forbidden by F5** (and excluded a priori in H5a and §12): the nullity-one mode is a fixed feature of the K = 7 architecture and is therefore sector-independent; were N_survive drawn from it, the count would read 1, 1, 1 across g = 0, 1, 2, gating the muon and destroying the 0.17% postdiction. The architectural "1" and the survival "1" must be distinct objects, and the survival "1" is consequently left with no independent home either.

The net status is uncomfortable and is recorded as such: at present neither 12 nor 1 recurs anywhere independent of the tau residual they were introduced to reproduce. This is precisely the failure the §5 caveat anticipates — the danger that the residual fixes the count rather than the count predicting the residual — now stated at the level of the numbers themselves. **[Open — failing; independent recurrence of N_anchor = 12 and N_survive = 1 is not yet exhibited, and the architectural nullity is excluded by F5 from supplying the latter.]**

Test 4 — A second observable wanting the same fraction

The decisive test would be a second, independently measured quantity that requires the same 1/12 participation fraction: one mechanism, two observables, far harder to dismiss as coincidence. The paper deliberately declines to manufacture one. Hunting through known ratios for a value near 0.083 would be the §5 disease one level up — a second fit dressed as a second prediction. The legitimate form of Test 4 routes through Test 1: derive the participation fractions for the saturated third quark sectors from the same boundary–interior count, then compare against the quark residuals once the quark law exists. That converts one observable into several from one mechanism without inventing a coincidence. **[Open — the only admissible version is the quark-sector extension, not a search for a numerical match.]**

The cost of generalising: route B does not survive the extension

There is a structural obstruction to Tests 1–4 that must be stated, because it is easy to miss. The lepton construction's central simplification (H4a) was the strategic choice of route B: assume g_c and K_c sector-stable so the suppression lands in p_{align} where a channel count can reach it. That assumption is across generations of *one species* (e, μ, τ). Tests 1–2 compare *different species* within the third generation (tau, bottom, top), where the couplings are manifestly unequal — the top's is of order unity, the tau's is tiny. The moment the construction extends to the quark third generation, g_c and K_c variation is forced, and route B's simplification is no longer available: the quark extension reactivates exactly the route-A / F4 debt the lepton paper deferred. **[Open — the natural generalisation breaks the assumption that made the lepton case tractable.]**

A partial escape exists and cuts both ways. If the shared structure is the p_{align} count (1/12, common to everything in the saturated register) while g_c, K_c remain species-specific, then within the third generation the 1/12 is a common factor — and common factors cancel in mass ratios. So m_t / m_τ does not test the count at all; the 1/12 is testable only in each fermion's absolute suppression against its own bare localisation, which returns to the quark-law prerequisite. There is no shortcut. **[Conditional — the count is testable per-fermion against the quark law, not via cross-species ratios.]**

Two further obligations the mechanism cannot dodge

The neutrino line. The neutrinos form a fourth three-generation tower, presumably saturating at the third sector like the others. If the gating is generation-structural rather than species-specific, ν_τ is obligated to carry the same channel logic. Neutrino masses are the most speculative instance — small, possibly Majorana, ordering unknown — but the obligation is not optional: it is a fourth independent place the same count must apply. **[Open — speculative but non-evadable; the mechanism owes ν_τ.]**

The shared-denominator constraint. Even granting the ratio-cancellation problem, the mechanism makes one sharper structural claim available now: the denominator 12 is shared across tau, bottom, and top because they share the saturated register, while only the numerators (and g_c, K_c) differ. This converts three masses into a system constrained by a common architectural factor rather than three independent fits — more predictive than the single tau number, and a place the boundary–interior count carries weight across the whole generation or fails visibly. **[Conditional on the quark law; the strongest predictive content the count offers.]**

Summary of predictive status

The construction is, at present, a single-target structural hypothesis. Of the tests it sets itself: Tests 1, 2, and 4 are not yet runnable, each awaiting a quark localisation law that does not exist and may not take clean exponential form; the top quark is the sharpest threat to the *route-B generation-wide* form of the mechanism but is consistent with the *route-A/C shared-denominator* form, and is not adjudicable either way until a quark baseline is independently validated; Test 3 is failing in the specific sense that neither structural number recurs independently and the one architectural "1" available is excluded by F5; and the natural generalisation to quarks breaks the route-B assumption on which the lepton result rests. The mechanism therefore earns evidential weight only through a quark-sector companion that derives the localisation law and lets the same count predict the third-generation quark residuals — turning one target into several, or refuting the route-B form on the top quark. Until that paper exists, the present result remains a located hypothesis, not a tested one. **[Open — predictive validation deferred to the quark-sector extension; this paper claims structural location only.]**

---

11. Status Table

Claim	Status
N_loop = 14 (minimal closure count)	**Inherited** from closure-geometry programme; the count is a structural consequence, but its stated internal decomposition (6+1, 14-traversal, nullity-1) is **reconstruction until checked against source**
Internal architecture (6 boundary + 1 hub, 14 traversals, nullity-1)	**Provenance grade to be pinned** (H5a; load-bearing detail, locator owed)
Surviving channel ≠ architectural nullity	**Excluded a priori** (H5a/Step 3; nullity is sector-independent, would force 1,1,1, gating muon — forbidden by F5)
Provenance of the "2" in 2K vs the "2" in dim Γ_∂	**Open — consistency obligation** (H5/H5a; leading reading (oriented traversals) threatens the 12; closure-preservation reading discharges it)
N_loop = 14 as closure-preservation (formation + preservation)	**Conjectural — the favourable branch, not the leading one** (H5a)
Boundary–interior split 14 = 2_∂ + 12_int	**Conjectural** (§3; the "two" not yet forced)

Claim	Status
Boundary generators frame-setting, not commitment channels	**Conjectural** (§4; gives N_anchor ≤ 12, exact count needs exact exclusion)
Interior anchoring count N_anchor = 12	**Conditional** (on the split ∧ exact exclusion; leakage is one-directional against the fit)
Saturated third sector leaves N_survive = 1	**Conjectural — consistency-driven, not derived** (§5; the crux)
Required form of the survival rule	**Step function, not curve** (§9 F5; 12,12,1 across g; rules out smooth-in-occupancy models)
p_align,2 = 1/12	**Conditional/Conjectural** (numerator consistency-driven, denominator count-conjectural)
Routes A/B/C are one factor (12.35) reparametrised; none fits better	**Established** (H4a; arithmetic identity, no route preferred on evidence)
Coupling-dominated route $\Theta_2 \approx$ p_align,2/p_align,0	**Conditional** (H4; K_c,2 ≈ K_c,0, g_c,2 ≈ g_c,0 assumed, not shown; B chosen strategically per H4a)
Θ_2^BIG = 1/12 = 0.0833 (integer 12 exact; misses data's 12.35)	**Conjectural — claims the integer, owes the decimal; not evidence**
Residual 3% structure (12 → 12.35)	**Forced *if* the discrete count is the leading term** (H4a, §8; integer count provably cannot reach 0.0810 — but F1/§8 item 4 may deny the premise)
Tau estimate m_τ^BIG ≈ 1.83 GeV	**Conditional/approximate**
Muon postdiction (0.17%) preserved (p_align,0 ≈ p_align,1 ≈ 1)	**Inherited**, requires F5 to hold
Discrete vs continuous reading reconciled	**Open** (§8)
Predictive status (Tests 1–4)	**Open — single-target**; Test 3 failing (no independent recurrence), top quark route-conditional (threatens route-B form, consistent with route-A/C), route B breaks under cross-species extension; quark companion required (§10)
Full tau mass derivation	**Open**

---

12. What Would Prove the Result

The route becomes a theorem if the following are established, in order of logical dependence.

Step 0 — Settle the interpretation of the fourteen, and separate the two factors of two (prerequisite). The count N_loop = 14 is inherited (H5a) as a structural consequence of the K = 7 closure architecture; what must be settled is *how* the fourteen relate to the seven. Determine whether the fourteen are oriented traversals of the seven constraints (the leading reading) or a closure-preservation doubling, and — per H5's consistency obligation — show that whichever doubling produces the fourteen, its factor of two is *distinct* from the boundary dimension dim $\Gamma_\partial = 2$ of Step 1, on pain of the denominator double-counting the frame. The leading reading is the dangerous one, so this is not a formality. Everything downstream is conditional on it.

Step 1 — Identify the boundary generators. Exhibit explicit generators $\partial_-, \partial_+ \in \Gamma_{min}$ with $\Gamma_{min} = \Gamma_\partial \oplus \Gamma_{int}$, dim $\Gamma_\partial = 2$, and show the "two" is forced by the frame conditions rather than assumed.

Step 2 — Prove boundary exclusion (exactly). Show $\Gamma_\partial \cap \Gamma_{anchor} = \emptyset$ with no partial leakage, so that N_anchor = 12 exactly and not merely ≤ 12. Note the leakage is one-directional against the fit: any inexactness lowers N_anchor below 12, which raises 1/N_anchor above 0.0833 — *further* from the residual 0.0810, never toward it. Exact exclusion is therefore load-bearing for the 2.8%, not a tidiness preference.

Step 3 — Prove the saturation discontinuity, not merely the value. The target is a *step function*, not a number. Derive from the Role-4 eigenmode structure of the completed K = 7 register that N_survive collapses to 1 *at exact closure* while remaining at N_anchor for every unsaturated sector — i.e. N_survive = 12, 12, 1 across g = 0, 1, 2 — *independently of the measured tau mass*. The hard content is the sharpness: per F5, any participation law smooth in occupancy gates the muon and dies, so what must be derived is a closure-*completion* effect that switches on only at saturation, not a filling-*fraction* effect that degrades gradually. Two derivations are excluded a priori: (i) any participation law monotone in occupancy (F5); and (ii) sourcing the surviving channel from the architectural nullity-one mode — sector-independent, hence forcing 1, 1, 1 and gating the muon (the full argument is in §10, Test 3, and H5a). This rules out monotone-in-occupancy models and the nullity identification, and is the step that converts §5's consistency argument into a derivation. It is the principal target.

Step 4 — Supply the forced residual structure. This is not optional polishing of a near-miss. Per H4a and §8, the discrete count predicts the integer 12 exactly and *provably cannot* reach the data's required 12.35, so additional structure of order 3% is mandatory, not cosmetic. Explain why $\Theta_2^{obs} \approx 0.0810$ sits below 1/12 = 0.0833 by identifying which source supplies the slack: a sub-leading correction to the count, finite-margin leakage layered on the discrete count, a residual K_c share, or boundary leakage. The choice is itself a physical claim that distinguishes the discrete and continuous readings.

Until Steps 0–3 are completed the result is a structural candidate, not a derivation. Step 4 is then required — not to improve the fit but because the integer count cannot reach the residual unaided — to upgrade the candidate from one that claims the right integer to one that also accounts for the mandatory 3% it does not itself contain.

---

13. Conclusion

The companion charged-lepton paper closed with a precise open target: derive the alignment probability $p_{align,2} \approx 0.081$ of the saturated third sector. This paper proposes a structural route to that number and grades it honestly.

The inherited minimal Role-4 closure structure supplies fourteen reversible generators. If two are boundary-frame generators and twelve are interior anchoring-eligible generators, and if saturation of the third sector leaves exactly one interior channel anchoring-effective, then

$$p_{align,2} \approx 1/12 = 0.0833,$$

and the tau estimate is about 2.8% high. That does not solve the tau, and the paper does not claim it does — the agreement is fifteen times coarser than the muon postdiction, the count's interpretation relative to K = 7 is unsettled (and on the leading reading would threaten the split), and both the split and the one-channel rule are conjectural, the latter currently inferred from the answer it is meant to predict.

What the construction does is change the nature of the problem. The tau residual is no longer a fitted number floating free of structure; it is a closure-channel participation question with a definite chain of dependencies:

$$N_{loop} = 14 \rightarrow 2_{\partial} + 12_{int} \rightarrow N_{survive} = 1 \rightarrow p_{align,2} = 1/12 \rightarrow \Theta_2,$$

each link of which is named, graded, and assigned a proof obligation (§12). The chain is not established — its first link is a count whose interpretation could undercut the very split that uses it, and its central link is a consistency argument awaiting a derivation — but it is *locatable*, which is the difference between an open problem and a mystery.

The next mathematical task is to prove or refute the boundary–interior split and the one-channel survival rule from the Role-4 eigenmode structure, independently of the tau mass. If both hold, the tau suppression becomes a consequence of closure-channel gating in the saturated Role-4 sector — derived from the same architecture that fixes the refinement hierarchy itself. But derivation is not yet validation: even a clean internal derivation would leave the count fitted to a single number. Evidential weight requires the predictive step of §10 — a quark-sector companion that derives the quark localisation law and lets the same boundary–interior count predict the third-generation quark residuals, turning one target into several, or failing visibly on the top quark, which on present reading is the sharpest threat to the whole picture. If either the internal derivation or the external test fails, §9 and §10 say precisely how. Until then the tau, like the muon before it, has an address rather than a solution — but the address is now a specific room, and §10 says what would confirm someone lives there.