

Charged-Lepton Participation Gating in Saturated Closure

A Sector-Restricted Participation Object, Completion Maintenance, and the Tau Suppression Target

Keith Taylor *VERSF Theoretical Physics Programme*

*Scope note. This paper states a sector-restricted result for the charged-lepton Role-4 sector only. The unrestricted "Participation Theorem" — a generation-wide rule covering quarks and neutrinos — is **not** claimed here. The computation in §4 shows a naive generation-wide version is already in tension with quark masses, and that tension is the reason the present statement is restricted rather than general. No claim is made that the participation law extends to the neutrino, up-type, or down-type sectors; per §4.4 those sectors lie outside the validated-baseline domain, so the question there is open in both form and existence.*

General Reader Summary

Two earlier VERSF ideas appear to be pointing at the same mechanism.

The first is refinement: as particles occupy deeper generations, their closure structure becomes more refined, and deeper structure tends to mean greater mass. The second is saturation: the third generation is special not because it is merely "more refined" than the second, but because it *fills* the available closure register completely. It reaches the edge of what the register can hold.

This paper asks what happens at that edge. The proposal is that below saturation, the channels of the closure structure are free to participate in anchoring — the process by which mass is generated. At saturation, most of those channels are no longer free: they are spent merely holding the completed structure closed. The structure has committed its internal degrees of freedom to keeping itself shut. What remains free to anchor is a small residual, and that residual fraction is the *participation factor*.

This would give a structural reason for the tau's suppression. The tau is the saturated charged lepton: structurally tightest, and so heavy by localisation, but with most channels consumed by maintenance, and so suppressed in effective coupling. The two effects pull in opposite directions, and the observed tau mass is what is left.

But there is an honest limit, and this paper draws it in plain sight rather than hiding it. The obvious way to test such an idea is to ask whether the *whole* third generation — bottom and top quarks as well as the tau — shows the same suppression. When the lepton law is applied to

quarks directly, it does not. In fact the law that works so well for the leptons fails for quarks two generations earlier, at the muon-analogue step, before saturation is even reached. So the mechanism cannot be a generation-wide mass rule. It is restricted here to the charged leptons, where the underlying law is known to hold, and the quark question is set aside as a separate problem requiring its own foundation.

The result of the paper is therefore modest and exact: it defines the structural object the tau suppression would have to be — a residual participation fraction after saturated maintenance — proves that *if* saturation activates such a maintenance cost *then* participation drops at saturation (a genuine step being the natural reading, not one the lepton data forces), and states clearly that the numerical value (one channel in twelve) is still imported from the tau, not yet derived, and not yet corroborated by any second observable inside its restricted domain.

Table of Contents

- Abstract
- 1. Introduction
- 2. Inherited Structure
 - H1 — $K = 7$ closure register
 - H2 — Refinement occupancy
 - H3 — Saturation as a threshold
 - H4 — Coherence-anchoring mass relation
 - H5 — Tau residual
 - H6 — Domain validity of the localisation law
- 3. The Participation Object
- 4. Domain and the Quark Tension
 - 4.1 The generation-wide temptation
 - 4.2 The direct check
 - 4.3 What the check actually shows
 - 4.4 The domain restriction
 - 4.5 The cost, stated directly
- 5. Why a Threshold Is Natural, Not Why It Is Forced
 - 5.1 What the muon forces (Established within domain)
 - 5.2 What the muon does not force
 - 5.3 The defensible statement: a threshold is natural, a smooth law is tuned
 - 5.4 Why they are identical — and what that costs
- 6. Completion Maintenance and the Statement
 - Sector-Restricted Participation Collapse (conditional)
 - Proof (of the conditional)
- 7. What Is Proven and What Is Not
 - The nullity trap (a forbidden shortcut)
- 8. Relation to Completion Density
- 9. Predictive Status

- [10. Failure Modes](#)
- [11. What Would Complete the Paper](#)
- [12. Conclusion](#)

Abstract

The charged-lepton threshold-compression analysis isolated a third-generation suppression factor $\Theta_2 \approx 0.081$ after the Role-4 localisation law $L_g = L_0 e^{-\kappa g}$, $\kappa = 83$, postdicted the electron-to-muon ratio to $\approx 0.17\%$. A companion alignment-gating paper proposed that this suppression arises because the saturated third charged-lepton sector retains only a reduced fraction of its anchoring-effective channels. This paper formulates the structural object underlying that proposal and **restricts its domain**.

We define the anchoring-effective participation fraction Π_g of a refinement sector as the fraction of closure channels available for irreversible commitment after closure-maintenance obligations are imposed. The central statement is a **sector-restricted Participation Collapse**: within the charged-lepton Role-4 sector, unsaturated sectors retain full participation ($\Pi = 1$), while saturation of the $K = 7$ register activates a maintenance projection that reduces Π . The muon datum requires the mechanism to be negligible at the second generation; a genuine threshold at exact closure is the *natural* realization of that requirement, while a smooth degradation would gate the muon unless tuned to mimic a threshold — so the threshold reading is natural rather than forced, and the two are observationally identical on charged-lepton data.

The paper is explicit about its limits. **(i)** The theorem is conditional: it proves *if a maintenance projection is absent below saturation and present at it, then Π is full below and reduced at saturation* — the existence of the projection and its location at *exact* closure (rather than within the unsaturated range) is graded Conjectural and carries the load. What the muon datum actually establishes is narrower: that the mechanism is negligible at the second generation ($\Pi_1 \approx 1$). A true threshold is the *natural* realization of that, but is not *forced* by it — a tuned crossover is observationally identical on charged-lepton data (and §5.4/§9 show why this is the same root cause as the single-observable limitation below). **(ii)** Both factors of $\Pi_2 \approx 1/12$ are owed, not just the numerator: the numerator $\dim \text{Ker}(P_{\text{maint}}) = 1$ is **imported from the tau residual, not derived**, and the denominator $\dim \Gamma_{\text{anchor}} = 12$ is itself **doubly conditional** — a conjectural split ($14 = 2_{\partial} + 12_{\text{int}}$) atop an architecture (H1) whose provenance is pending confirmation. **(iii)** A generation-wide version is **already in tension with quark masses**: the lepton localisation baseline fails to postdict second-generation quark ratios by an order of magnitude, and fails in *opposite directions* ($c/u \approx 588$ overshooting a predicted ≈ 207 ; $s/d \approx 20$ undershooting it), so no single quark law *of the same exponential form* can repair both towers and quark " Θ_2 " values are residuals against an invalid baseline. The domain is restricted to sectors where the Role-4 localisation law is independently validated — at present the charged leptons alone. **(iv)** The cost is conceded directly: the restricted mechanism has **no in-domain second observable** and remains a single-target hypothesis, and whether $\dim \text{Ker}(P_{\text{maint}})$ is even computable from the Role-4 algebra is presently unassessed.

1. Introduction

Matter comes in three generations. The electron, muon and tau share charge and identity yet differ in mass by large factors. Earlier VERSF work proposed that this structure arises from refinement of a finite closure register: the first generation occupies the least-refined state, the second a deeper state, and the third saturates the register. Saturation makes the third generation structurally special — not another step in a smooth sequence, but the point at which refinement completes the register.

The charged-lepton threshold-compression paper showed the first jump, electron to muon, follows from Role-4 localisation: with $\kappa = 8/3$ the ratio $m_\mu/m_e \approx e^{(2\kappa)} = e^{(16/3)} = 207.13$ is postdicted to $\approx 0.17\%$. The same law iterated gives a tau far too heavy ($m_\tau \sim m_e e^{(32/3)} \approx 21.9$ GeV against ≈ 1.78 GeV), so the tau carries a residual suppression $\Theta_2 \approx 0.081$. The alignment-gating paper proposed an address for that residual: the saturated third sector has reduced anchoring-effective participation, because most of its channels are spent maintaining the completed closure rather than performing new anchoring.

This paper develops that idea structurally — and then immediately fences it.

The natural temptation is to lift the idea to a generation-wide rule: *saturation gates the third generation*. That would be far more powerful, because it would make a sharp, checkable claim about bottom and top quarks as well as the tau. §4 runs that check directly against measured quark masses and finds it fails — and fails for a deeper reason than third-generation confounding. The result is that the mechanism cannot be stated generation-wide without contradiction. It is therefore stated here for the charged-lepton Role-4 sector only, and the domain boundary is drawn empirically: the mechanism is defined where the Role-4 localisation law is independently validated, which at present means the charged leptons.

The contribution is consequently narrow and exact. It is not a derivation of the tau mass, nor a generation-wide theorem. It is: (a) a clean definition of the participation object Π_g ; (b) a conditional theorem that, given a maintenance projection located at exact closure, saturated maintenance reduces participation there while leaving unsaturated sectors untouched — with the projection's existence and location flagged as the conjectural, muon-underdetermined part (§5); and (c) an explicit statement of what remains owed — the value of $\dim \text{Ker}(P_{\text{maint}})$, still imported from the tau — and of what the restriction costs, namely the loss of any in-domain second observable.

2. Inherited Structure

Imported results and definitions. Each carries its source's grade; where a source is not yet pinned, that is recorded as owed.

H1 — $K = 7$ closure register

[Inherited — pending source confirmation]

The minimal committed closure architecture is taken to be $K = 7$: six boundary constraints plus one hub constraint. This is inherited from the closure-geometry programme. As recorded in the companion alignment-gating paper, the architecture's status *as stated* (six-plus-one, with a fourteen-traversal response structure and a nullity-one mode) is load-bearing and its citable provenance is still being confirmed; it is used here as an inherited input, not re-derived, and is **not** treated as established within this paper.

H2 — Refinement occupancy

[Inherited]

The refinement sequence occupies the register in cumulative loads 1, 3, 7. With sector index $g = 0, 1, 2$:

$$O_0 = 1, O_1 = 3, O_2 = 7 = K.$$

The first sector is partially occupied, the second remains unsaturated, the third saturates the register.

H3 — Saturation as a threshold

[Inherited]

The third charged-lepton sector is the saturation sector, $O_2 = K$. Saturation is a threshold property (the register is *complete*), not a degree of filling. This distinction is the hinge of §4.

H4 — Coherence—anchoring mass relation

[Inherited — Void Coupling paper]

Mass depends on localisation and on effective anchoring:

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

with g_c the coupling availability, p_{align} the aligned-state anchoring fraction, K_c the anchoring depth. (The coupling paper writes the coupling strength g ; written g_c here to avoid collision with the generation index g .)

H5 — Tau residual

[Inherited; a residual, no independent evidential weight]

The tau suppression factor $\Theta_2 \approx 0.081$ is computed from the measured mass by inversion. In the coupling-dominated route it is read mainly as reduced anchoring-effective participation, $\Theta_2 \approx p_{\text{align},2}/p_{\text{align},0}$, on the assumption (carried, not shown) that g_c and K_c do not supply the dominant sector variation. This paper inherits that route's conditionality and asks whether saturated completion naturally produces such a reduction.

H6 — Domain validity of the localisation law

[Established within domain; see §4]

The Role-4 localisation law with $\kappa = 8/3$ is independently validated *only where it postdicts an observed second-generation ratio*. For the charged leptons it does (muon, 0.17%). §4 shows it does **not** for quarks at the corresponding step. The mechanism of this paper is defined only on the validated domain. This is the formal domain restriction, and it is empirical, not stipulated.

3. The Participation Object

Define a sector participation fraction Π_g : the fraction of closure structure remaining available for irreversible anchoring after closure-maintenance obligations are satisfied.

Let Γ_{anchor} be the anchoring-eligible channel space of a sector and $\Gamma_{\text{maint},g}$ the subspace required to maintain closure at refinement level g . The anchoring-effective participation space is the complement

$$\Gamma_{\text{part},g} = \Gamma_{\text{anchor}} \ominus \Gamma_{\text{maint},g}, \quad N_{\text{part},g} = N_{\text{anchor}} - N_{\text{maint},g},$$

and the participation fraction is

$$\Pi_g = N_{\text{part},g}/N_{\text{anchor}}.$$

This separates three things otherwise easily conflated: **available** channels (could in principle anchor), **maintenance** channels (consumed preserving closure), and **participating** channels (still free for new irreversible commitment).

A modeling assumption is built into this definition and must be named. The denominator N_{anchor} carries no sector index: all sector variation is loaded into $\Gamma_{\text{maint},g}$ (zero below saturation, non-trivial at saturation) while the anchoring-eligible space is held fixed across sectors. Call this **anchoring-space stability**:

$N_{\text{anchor},g} = N_{\text{anchor}}$ for all g . [**Assumption — anchoring-eligible dimension sector-independent.**]

This does not threaten $\Pi_1 = 1$: with $\Gamma_{\text{maint},1} = 0$ the ratio is 1 regardless of the denominator. But it means the "12" is specifically the *saturated-sector* anchoring dimension, and that if refinement altered N_{anchor} across generations, Π_2 would shift even with $\dim \text{Ker}(P_{\text{maint}}) = 1$. The assumption is reasonable (the anchoring-eligible space is a property of the closure architecture, which is the same register at every g) but it is an assumption, not a derived fact.

The structural claim is then that $\Gamma_{\text{maint},g}$ is trivial below saturation and non-trivial at saturation:

$$\Gamma_{\text{maint},0} = \Gamma_{\text{maint},1} = 0 \implies \Pi_0 = \Pi_1 = 1, \Gamma_{\text{maint},2} \neq 0 \implies \Pi_2 < 1.$$

This is the participation collapse. Its *value* is deferred (§6–§7); whether its onset is a true threshold or a tuned crossover is taken up next (§5).

4. Domain and the Quark Tension

This section is placed before the theorem deliberately. The theorem's honest scope cannot be stated until the domain is fixed, and the domain is fixed by data.

4.1 The generation-wide temptation

The strong, attractive form of the mechanism would be generation-wide: *saturation gates the third generation*. This would make the participation collapse a sharp claim about every third-generation fermion — tau, bottom, top — and so would supply exactly the second, third and fourth observables a single-target hypothesis lacks. It is worth testing precisely because it would be so powerful.

4.2 The direct check

The lepton paper defines the third-step suppression as $\Theta_2 = (m_{\{g=2\}}/m_{\{g=0\}})/e^{(32/3)}$, with the second-generation ratio predicted as $e^{(16/3)} = 207.13$. Applying the *same* definition to quarks, with no quark-specific law, gives the table below. The value 207.13 in the second column is the **lepton-derived** prediction; it is shown against the quark rows only to exhibit that

the lepton law does not describe them, not as a prediction the quark masses were ever expected to meet.

| Sector | $g=1/g=0$ ratio (predicted 207.13) | $\Theta_2 =$ ($g=2/g=0$)/ $e^{(32/3)}$ | $1/\Theta_2$ | Verdict |
|------------------|---------------------------------------|---|--------------|---|
| Charged leptons | $m_\mu/m_e \approx 206.8 \checkmark$ | 0.0811 | 12.34 | baseline valid; tau suppressed |
| Up-type quarks | $m_c/m_u \approx 588 \times$ | 1.86 | 0.54 | baseline fails; third step <i>enhanced</i> (wrong sign) |
| Down-type quarks | $m_s/m_d \approx 20 \times$ | 0.021 | 48 | baseline fails; over-suppressed, not ≈ 12 |

(Quark masses are scheme- and scale-dependent running quantities; the ratios above use standard reference values and the qualitative verdicts are robust to that dependence — none of the quark second-generation ratios comes within a factor of several of 207.)

4.3 What the check actually shows

The decisive observation is not that the quark *third* generation behaves differently. It is that the localisation baseline **fails two generations earlier**, at the muon-analogue step. For leptons the law postdicts the second-generation ratio to 0.17%; for quarks it misses by an order of magnitude in *both* sectors ($s/d \approx 20$ not 207; $c/u \approx 588$ not 207). Therefore the quark " Θ_2 " is a residual measured against a baseline that was never valid for quarks. The wrong-sign up-type value ($\Theta_2 = 1.86$) and the over-suppressed down-type value ($1/\Theta_2 = 48$) are symptoms of that broken baseline, not independent third-generation effects, and neither tests the participation mechanism.

A naive generation-wide version of the participation collapse is therefore already in tension with quark masses — and the tension is in the baseline, not merely the third step. This is stated plainly rather than absorbed: the mechanism, taken generation-wide, is in conflict with measured quark hierarchies before saturation is even reached.

4.4 The domain restriction

The honest consequence is not denial but restriction. The mechanism is defined only on sectors where the Role-4 localisation law is independently validated (H6). That domain currently contains the charged leptons and excludes the quarks — not by stipulation, but because the law's own postdiction succeeds for the former and fails for the latter.

Why the law fails for quarks is a separate question the paper does not answer and does not need to. Colour confinement, QCD running, Higgs-sector Yukawa structure, and CKM-sector mixing are candidate reasons the quark mass channel is not controlled by Role-4 charged-lepton closure in the same way; these are offered as conjecture, not claim. **[Conjectural — candidate confounders, not established.]** The operative point is independent of which confounder

dominates: wherever the localisation baseline is not independently validated, there is no well-defined Θ_2 for the mechanism to explain, and the sector is out of domain.

Quarks could re-enter the domain only by supplying their own validated localisation law together with a quark anchoring-channel decomposition — and the table forecloses the *simplest* form of that route. The two quark towers fail in **opposite directions**: the up-type second-generation ratio ($c/u \approx 588$) *overshoots* the lepton value 207, while the down-type ratio ($s/d \approx 20$) *undershoots* it. Since a larger κ predicts a larger $e^{(2\kappa)}$ ratio, repairing the up-type tower requires $\kappa_{\text{up}} > 8/3$ and repairing the down-type tower requires $\kappa_{\text{down}} < 8/3$. No single κ_q can do both *within a law of the same exponential form* $L_g = L_0 e^{(-\kappa g)}$, and the gaps are far too large to be reconciled by scheme or running. (A quark law of genuinely different functional form is not caught by this bracketing argument; what is ruled out is a single exponential-form law with one κ , which is the form route B inherits.) Quark re-admission therefore requires **separate validated localisation constants κ_{up} and κ_{down} (and a separate κ_v for neutrinos)**, each postdicting its own sector's second-generation ratio — not one universal quark law of the lepton form. This is the route-B / anchoring-depth debt the companion paper deferred, now made concrete and made harder: extending to quarks reactivates the assumption (sector-stable g_c, K_c) the lepton construction set aside, *and* multiplies the number of localisation constants that must be independently derived. **[Open — quark re-entry requires separately validated κ_{up} , κ_{down} , κ_v ; a single same-form κ_q is ruled out by the opposite-direction failure; not attempted here.]**

The scope is therefore stated explicitly. The present construction concerns the charged-lepton closure architecture only. No claim is made that the corresponding participation law extends to the neutrino, up-type quark, or down-type quark sectors, which may possess distinct closure geometries and distinct maintenance projections — and which, per the foregoing, lie outside the validated-baseline domain entirely, since the Role-4 localisation law they would rest on is not yet independently validated for them. Whether a distinct participation law holds in those sectors is therefore open in both its form and its existence, not merely in its parameters. The neutrinos are included in this exclusion deliberately: they form a fourth three-generation tower the mechanism would be obligated to if it were generation-wide, and the restriction makes that obligation explicitly undischarged rather than silently assumed.

4.5 The cost, stated directly

The restriction is honest, but it is not free. By scoping to the charged leptons, the mechanism returns to a **single in-domain observable** — the tau — with no sibling in scope to corroborate it. There is no heavier charged lepton, so the domain offers no further test until a validated quark (or neutrino) baseline is independently derived. The paper buys correctness about quarks at the price of conceding it has no in-domain second observable. §9 carries this into the predictive-status accounting and does not let the restriction disguise it.

5. Why a Threshold Is Natural, Not Why It Is Forced

It is tempting to say the muon *forces* a discontinuous collapse. It does not, and the paper should not claim it. The muon constrains only one thing, and the rest is a naturalness argument that must be labelled as such.

5.1 What the muon forces (Established within domain)

The electron-to-muon ratio is explained by localisation to 0.17%, with no participation suppression. Therefore whatever the mechanism, **its action at $O_1 = 3$ is negligible**:

$\Pi_1 \approx 1$ to $\approx 0.1\%$. [**Established within domain.**]

That is the entire empirical content the muon supplies. It pins the value at one occupancy point and says nothing about the *shape* of Π between O_1 and O_2 .

5.2 What the muon does *not* force

It does not force a literal step. There is no in-domain data point between $O_1 = 3$ and $O_2 = 7$, so a smooth monotone law that sits at ≈ 1 near $O = 3/7$ and falls to ≈ 0.081 by $O = 1$ (full saturation) fits every constraint just as well as a true discontinuity. The earlier claim that "any monotone-in-occupancy law is excluded a priori" is therefore **withdrawn as overstated**: what the muon excludes is monotone laws with *appreciable gradient at O_1* , not monotone laws as such.

5.3 The defensible statement: a threshold is natural, a smooth law is tuned

The argument that survives is one of naturalness, and it is cleaner than the a-priori claim it replaces. A *generic* smooth degradation gates the muon: if Π falls appreciably as the register fills, it has already fallen by $O_1 = 3$. To reproduce $\Pi_1 \approx 1$ and $\Pi_2 \approx 0.081$ with a smooth law, the crossover must be tuned to stay flat through O_1 and then plunge between O_1 and O_2 — i.e. tuned into a step in all but name. The completion story makes that profile natural rather than tuned: if the maintenance obligation activates on *completion* of the register (a threshold property of $O = K$) rather than on degree of filling, then Π is flat below saturation and drops at it, with no crossover to position by hand. So:

$\Pi_g = 1$ for $O_g < K$, $\Pi_g < 1$ for $O_g = K$

is the **natural** realization, not the **forced** one. A genuine threshold and a finely-tuned steep crossover are observationally identical on the present data.

5.4 Why they are identical — and what that costs

This indistinguishability is not a separate weakness; it is the same fact as the single-in-domain-observable limitation conceded in §9, viewed from the other side. A true step and a tuned crossover differ only in their value at occupancies between O_1 and O_2 , and the charged-lepton domain contains no such occupancy to measure. Therefore the step-function assumption is not merely *unproven* — it is **unprovable within the domain** until a second in-domain anchor exists. The threshold is the natural reading; its *necessity* is not establishable from charged-lepton data alone. Hypothesis 3 (§6) carries exactly this gap as a conjecture: that the projection exists and that its activation sits at *exact* closure rather than somewhere in $O \in (3, 7]$.

6. Completion Maintenance and the Statement

A completed register is not passive; it must be held closed. Below saturation, spare capacity lets the structure absorb local commitments without forcing the whole architecture into maintenance mode. At saturation there is no spare capacity: any anchoring event now engages the completed closure as a whole, which must preserve boundary consistency, interior compatibility, hub coherence, orientation matching, and return stability. These maintenance obligations generate no mass directly; they preserve the conditions under which mass-generating commitment remains possible. The claim is that maintenance consumes participation, splitting the anchoring-eligible space

$$\Gamma_{\text{anchor}} = \Gamma_{\text{maint},2} \oplus \Gamma_{\text{res},2}, \Pi_2 = \dim \Gamma_{\text{res},2} / \dim \Gamma_{\text{anchor}}.$$

Sector-Restricted Participation Collapse (conditional)

Domain. Let the sector lie in the validated-baseline domain (H6) — for this paper, the charged-lepton Role-4 sector.

Hypotheses. 0. (Domain) The Role-4 localisation law is independently validated on the sector. **[Established for charged leptons; §4.]**

1. A refinement sequence occupies a finite K -register with occupancies O_g ; $O_g < K$ unsaturated, $O_s = K$ saturated. **[Inherited H2–H3.]**
2. Anchoring-effective participation requires channels not committed to closure maintenance. **[Definition, §3.]**
3. (a) The mechanism's action is negligible at $O_1 = 3$, i.e. $\Pi_1 \approx 1$. **[Established within domain — the muon, §5.1.]** (b) A non-trivial maintenance projection P_{maint} exists on Γ_{anchor} , and its activation is located at *exact* closure $O = K$ rather than at some

occupancy in $O \in (3, 7]$. [**Conjectural — this is the part the muon does *not* force (§5.2–5.4), and it carries the load.**]

The split matters: 3(a) is data, 3(b) is conjecture, and the proof below uses only that P_{maint} is absent below saturation and present at it. Whether the onset is a true discontinuity or a steep crossover tuned to mimic one is not decided by the conclusion (§5.3) — the theorem establishes the collapse *given* the projection's location, not the location itself.

Conclusion. Then

$$\Pi_g = 1 \ (O_g < K), \ \Pi_s = \dim \text{Ker}(P_{\text{maint}}) / \dim \Gamma_{\text{anchor}} < 1 \ (O_s = K).$$

If, in the $K = 7$ Role-4 architecture, $\dim \Gamma_{\text{anchor}} = 12$ and $\dim \text{Ker}(P_{\text{maint}}) = 1$, then $\Pi_s = 1/12$. [**The dimensions are imported from the tau paper, not derived here — see §7.**]

Proof (of the conditional)

For $O_g < K$ the register is incomplete; by hypothesis 3b no global maintenance projection is active, so $\Gamma_{\text{maint},g} = 0$, $\Gamma_{\text{part},g} = \Gamma_{\text{anchor}}$, and $\Pi_g = 1$.

For $O_s = K$, hypothesis 3b activates P_{maint} on Γ_{anchor} . Channels in $\text{im}(P_{\text{maint}})$ are committed to maintenance and cannot independently anchor; the anchoring-effective subspace is the residual complement $\Gamma_{\text{part},s} = \text{Ker}(P_{\text{maint}})$, giving $\Pi_s = \dim \text{Ker}(P_{\text{maint}}) / \dim \Gamma_{\text{anchor}}$. Since P_{maint} is non-trivial, $\dim \text{Ker}(P_{\text{maint}}) < \dim \Gamma_{\text{anchor}}$ and $\Pi_s < 1$. The numerical specialisation to $1/12$ follows only on the imported dimensions. ■

The proof is valid but its force is entirely in hypothesis 3b. The theorem says: *granting that a maintenance cost is absent below saturation and present at it, participation is full below and reduced at saturation.* It does **not** establish that the projection exists, that its onset is located at exact closure rather than within the unsaturated range, or the size of the cost — and §5 shows the location in particular cannot be fixed from charged-lepton data.

7. What Is Proven and What Is Not

Proven (conditionally, within domain): *if* a maintenance projection is absent below saturation and present at it, *then* participation is full below saturation and reduced at it, located as $\dim \text{Ker}(P_{\text{maint}}) / \dim \Gamma_{\text{anchor}}$. The collapse is **negligible at O_1 and active at O_2 ; a genuine step is the natural realization but is not forced** (§5.3–5.4). The participation object Π_g is given a precise definition.

Not proven, and owed:

1. **Activation and its location.** That a non-trivial P_{maint} exists *and* activates at exact closure rather than within $O \in (3, 7]$ (hypothesis 3b). This is the crux, and per §5 its location is unprovable within the domain.
2. **Denominator — doubly owed.** That $\dim \Gamma_{\text{anchor}} = 12$. This rests on *two* unconfirmed pillars, not one: (a) it is imported from the companion paper, where the split $14 = 2_{\partial} + 12_{\text{int}}$ is itself **conjectural** and subject to the unresolved two-factors-of-two obligation; and (b) the underlying count rests on H1, whose architecture (six-plus-one, fourteen traversals, nullity-one) is **inherited pending source confirmation** and not established within this paper. So even the denominator is doubly conditional — a conjectural split atop an unconfirmed architecture.
3. **Numerator.** That $\dim \text{Ker}(P_{\text{maint}}) = 1$. **Imported from the tau residual, not derived.** This is the value consistency requires, not a computed kernel dimension.
4. **The residual percent (range, partly κ -absorbable).** The integer count gives $\Pi_2 = 1/12 = 0.0833$; the inverted residual sits near 0.0810. Stated as a clean "forced 2.9% of new physics" this overclaims, for two reasons. First, the gap is **κ -treatment-dependent**: $\kappa = 83$ postdicts the muon only to 0.17%, so it is not pinned tighter than that, and because $e^{(32/3)}$ carries 4κ of leverage, κ precision alone shifts the tau target by a few tenths of a percent — moving the gap into roughly the **2.4–2.9%** range and absorbing part of it into κ -anchoring rather than a new mechanism. Second, though, κ -refitting cannot absorb the *whole* gap: sharpening κ to the tau would pull it off the muon — but κ is not a free parameter to retune, because it is *fixed by the muon postdiction*, which is what validates the localisation law in the first place. Retuning κ to the tau would dissolve the very fit that licenses the law, so κ cannot be moved to the tau without removing the law's own warrant. A residual therefore survives any honest κ treatment. The defensible statement is therefore: a residual of order a couple of percent remains after κ precision is accounted, and *that* residual — smaller and less certainly new-physics than a bare 0.0833-vs-0.0810 comparison suggests — is what must be sourced from a sub-leading count correction, finite-margin leakage, a residual K_c share, or κ -anchoring precision. The candidate absorbers now explicitly include κ precision alongside the structural ones.

So the paper does not say " $\Pi_2 = 1/12$ is proven." It says: Π_2 *is* the residual anchoring dimension after the saturated maintenance projection, and computing $\dim \text{Ker}(P_{\text{maint}})$ — independently of the tau mass — is the open target. Whether that computation is *tractable* from the Role-4 algebra is itself presently unassessed (§11), so the claim is that the problem is *located*, with the honest caveat that located-ness has not yet been demonstrated by a feasibility check.

The nullity trap (a forbidden shortcut)

H1's architecture carries a *nullity-one mode*, and §6 needs $\dim \text{Ker}(P_{\text{maint}}) = 1$. The two "ones" will be tempting to identify — sourcing the single surviving channel from the architectural nullity. **That identification is forbidden.** The nullity-one mode is a fixed feature of the $K = 7$ architecture and is therefore *sector-independent*; if $\dim \text{Ker}(P_{\text{maint}})$ were the architectural nullity, it would read 1 for *every* sector, gating the muon ($\Pi_1 = 1/12$) and destroying the postdiction. This is the same fatal structure as smooth degradation: a sector-independent suppression cannot leave $\Pi_1 \approx 1$ while delivering $\Pi_2 \approx 0.081$ (§5.1). The required kernel must instead be a property that distinguishes the *saturated* sector from the unsaturated ones; the

architectural nullity, being sector-independent, cannot supply it. Any derivation of $\dim \text{Ker}(P_{\text{maint}}) = 1$ must exclude this shortcut explicitly. **[A-priori exclusion; the survival-one is not the architectural nullity-one.]**

8. Relation to Completion Density

The completion-density programme already derives generation ordering from refinement structure, with the third generation special because refinement fills the register. This paper adds one distinction within the charged-lepton sector: refinement has two separable effects — it increases structural depth (tending to *raise* mass) and, at saturation, reduces free participation (tending to *suppress* effective coupling). The first would make the tau too heavy alone; the second pulls back. The tau is then not an exception to the hierarchy logic but the signature that the hierarchy has reached closure completion:

electron — unsaturated, full participation; muon — deeper, still unsaturated, full participation; tau — deepest and saturated, reduced participation.

This unification is *organisational*: it ties the tau-suppression and completion-density strands to one object. It is not, by itself, evidential — it adds no observable (see §9). Its value is conceptual economy, and it is claimed as nothing more.

9. Predictive Status

The theorem is structurally clean, but the programme's own discipline requires stating plainly what it does and does not buy.

It adds no new observable. The chain ending at $\Pi_2 \approx 1/12$ still has exactly one empirical anchor: the measured tau mass. The 12 and the 1 are imported (§7), not independently exhibited. Lifting the lepton conjecture to a theorem about a participation *object* generalises the *explanation* without generalising the *test*. The " $\Theta_2 \approx \Pi_2$ " identification of §6 equates two quantities each fixed by the tau; it is an identification, not a second measurement.

Test 3 (independent recurrence of the numbers) still fails. Neither 12 nor 1 recurs anywhere independent of the tau residual they reproduce, and the one architectural "1" available (the nullity) is excluded from being the survival "1" by §7. This is recorded as a present failure, not a deferred check.

The domain restriction does not repair this — and must not be read as doing so. Scoping out the quarks removes the *adversarial* generation-wide test (which the mechanism was failing, §4) but it does not supply a corroborating one. On the contrary, it *removes* the only near-term

external test the idea had. Within the restricted domain there is no second charged lepton, so the mechanism is back to a single target with no sibling. The restriction is the honest move, but its predictive cost is real and is owned here: **the paper concedes it has no in-domain second observable and will not until a validated quark or neutrino baseline is independently derived.**

The step-function assumption is not just unproven but unprovable here — and this is the same limitation. §5 showed a true threshold and a tuned crossover are observationally identical on charged-lepton data, because no in-domain occupancy lies between O_1 and O_2 . That indistinguishability and the no-second-observable concession above are *one fact viewed twice*: both say the domain contains too few anchor points to constrain the shape of Π . So hypothesis 3b's location-at-exact-closure cannot be settled empirically within the domain; it can only be settled by deriving the projection from the algebra (Step 1) or by acquiring a second in-domain anchor (which does not exist). Connecting the two admissions is not bookkeeping — it shows the central conjecture and the predictive limitation share a single root cause.

What would change the status. Only an out-of-sample success: a *separately* validated κ_{up} , κ_{down} , or κ_v re-admitting that sector to the domain (§4.4 rules out one universal κ_q doing this) and then the *same* participation count predicting that sector's third-generation residual; or a derivation of $\dim \text{Ker}(P_{\text{maint}}) = 1$ from the Role-4 eigenmode structure independent of the tau mass (converting the imported numerator into a computed one) — itself contingent on the tractability question §11 leaves open. Until one of these lands, the result is a *located* hypothesis, not a *tested* one, and "located" is itself an assertion pending the feasibility check.

10. Failure Modes

F1 — No maintenance projection. If saturation activates no closure-maintenance projection (hypothesis 3b false), there is no reason for participation to collapse. The theorem's antecedent fails and the construction is empty.

F2 — Appreciable gating below saturation. If participation degrades appreciably by $O_1 = 3$, the muon is gated and the 0.17% postdiction dies (§5.1). Note this does *not* require the mechanism to be a literal step (§5.2): it requires only that the mechanism be negligible at O_1 . A genuine threshold satisfies this naturally; a smooth law satisfies it only if tuned to stay flat through O_1 (§5.3). What F2 excludes is appreciable gradient at the second generation, not smooth laws as a class.

F3 — Wrong residual dimension. If the saturated projection leaves two or more residual channels, the suppression is too weak ($2/12 \approx 0.167$); if none, the saturated sector cannot anchor a massive state. The interpretation needs $\dim \text{Ker}(P_{\text{maint}}) = 1$ exactly, currently inferred from the answer (§7).

F4 — Suppression belongs to anchoring depth. If Θ_2 is dominantly $K_{c,0}K_{c,2}$ rather than $p_{align,2}p_{align,0}$, participation is not the operative mechanism whatever its numerical coincidence. Distinguishing routes needs the per-sector eigenmode computation of g_c, K_c, p_{align} — an undischarged debt.

F5 — No connection to completion density. If the completion-density hierarchy and the tau suppression turn out to use unrelated mechanisms, the §8 unification fails; the tau interpretation might survive in isolation but would not advance the common spine.

F6 — Domain leakage. If the localisation baseline can be shown valid for some sector outside the charged leptons, the generation-wide tension of §4 returns as a live test rather than an out-of-domain exclusion, and the mechanism must then *predict* that sector's residual or be refuted by it. Re-admission is sector-specific: because the up- and down-type towers fail in opposite directions (§4.4), it requires a *separately* validated $\kappa_{up}, \kappa_{down},$ or κ_v — a single universal κ_q of the lepton's exponential form is already ruled out, so F6 cannot be triggered by one same-form quark law repairing both quark sectors at once. The domain restriction is provisional on the baseline's continued failure sector by sector, not a permanent exemption.

11. What Would Complete the Paper

The next derivation must compute P_{maint} and its residual kernel from the Role-4 closure algebra, in this order of dependence:

Step 0 — Confirm the architecture (prerequisite). Pin H1 *as stated* (six-plus-one, fourteen traversals, nullity-one) to a citable closure-geometry result, and resolve the companion paper's two-factors-of-two obligation, since $\dim \Gamma_{anchor} = 12$ rests on it.

Step 1 — Derive the activation and its location. Show that exact saturation ($O = K$) activates a non-trivial maintenance projection P_{maint} on Γ_{anchor} while $O < K$ does not — and, the harder half, that the activation sits at *exact* closure rather than within $O \in (3, 7]$. This is hypothesis 3b, the antecedent; without it the theorem is vacuous, and per §5 the *location* cannot be fixed by charged-lepton data, so it must come from the algebra.

Step 2 — Compute the kernel dimension. Derive $\dim \text{Ker}(P_{maint})$ from the eigenmode structure of the completed register, *independently of the tau mass*, and show it equals 1 — while the same construction returns full participation (no projection) for the unsaturated sectors, securing F2. The derivation must **not** identify the kernel with the architectural nullity (§7).

Step 3 — Account for the residual. Account for the gap between $1/12 = 0.0833$ and the inverted residual near 0.0810. This is *not* cleanly " $\approx 3\%$ of forced new physics": per §7 item 4 the gap is κ -treatment-dependent (range $\approx 2.4\text{--}2.9\%$), and κ -anchoring precision is one candidate absorber alongside a sub-leading count correction, finite-margin leakage, or a residual K_c share. The task

is to identify which combination supplies the residual that survives an honest κ treatment; the choice is a physical claim.

Tractability — presently unassessed. The forward value of the paper rests on Steps 1–2 being *computable* from the Role-4 algebra independently of the tau. Whether the saturated-register eigenmode computation is in fact tractable has not been assessed here. The paper claims the problem is *located* — given a precise object to compute and a precise independence condition — but does not claim it is *known to be solvable*; demonstrating tractability (or its absence) is itself part of Step 1–2 and would convert "located" from an assertion into a supported claim.

Until Steps 1–2 are done the result is a structural candidate within its restricted domain, not a derivation; Step 3 then accounts for the residual that remains once κ precision is taken into account.

12. Conclusion

This paper defines the participation object Π_g and proves, conditionally and within a restricted domain, that a maintenance projection located at saturation reduces anchoring-effective participation there while leaving unsaturated sectors at full participation — converting formerly free channels into closure-maintenance structure. The muon fixes only that the mechanism is negligible at the second generation; a genuine threshold at exact closure is the natural reading of that, not a forced one. Within the charged-lepton Role-4 sector this gives a structural reading of the tau: heavy by localisation, suppressed by saturation, the candidate sign of register completion.

Four honesties are carried in plain sight. The theorem is conditional on an unproven activation whose *location* at exact closure the muon does not fix (hypothesis 3b, §5); both factors of 1/12 are owed, the numerator imported from the tau and the denominator doubly conditional (a conjectural split atop an unconfirmed architecture); a generation-wide version is already in tension with quark masses, the lepton baseline failing for quarks two generations before saturation and failing in opposite directions so no single same-form quark law repairs it — which is why the domain is restricted to the charged leptons and why the mechanism, so restricted, has no in-domain second observable; and whether the deciding computation is even tractable from the Role-4 algebra is unassessed, so "located" is a claim still pending its own feasibility check.

The construction therefore does not solve the tau and does not pretend to. What it does is convert a lepton-specific suppression factor into a named structural object with a definite computation attached:

$K = 7$ register \rightarrow exact saturation \rightarrow maintenance projection P_{maint} \rightarrow residual kernel dim
 $\text{Ker}(P_{\text{maint}}) \rightarrow \Pi_2 \rightarrow \Theta_2,$

each link graded, the load-bearing link (the activation and its kernel dimension) marked as the open target, and the domain boundary drawn where the data draws it. The next mathematical task is to derive P_{maint} and its kernel from the Role-4 algebra independently of the tau mass. If the kernel is one-dimensional, the tau suppression becomes a consequence of saturated completion participation. If it is not, the participation route fails cleanly — and §10 says exactly how.

The right spine is *not* to start with the tau. It is to start with the participation object, fence its domain honestly, and let the tau be the one application the domain currently licenses — application, not premise.