The Suppression Unit 1/12 from the Closure Cell

Sourcing the Tau and Down-Type First-Step Suppression from the W_7 Spectrum

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

Nature's basic matter particles come in three "generations" — copies that behave identically but get progressively heavier. The electron, the muon, and the tau are one such family of three; the down, strange, and bottom quarks are another. Why the heavier copies weigh what they do is one of physics' long-standing puzzles.

This programme proposes an answer: each step up a generation should multiply the mass by a fixed amount, about 207. The first lepton step (electron to muon) matches that almost perfectly. But two other steps fall far short — the muon-to-tau step and the down-to-strange step each come in at only about one-twelfth of what the simple rule predicts. Something extra is suppressing them. Naming that extra factor "1/12" just because it fits the numbers would not be a real explanation; the question is whether 1/12 has an independent reason to be there.

The point of this paper is to show that 1/12 was already sitting inside the theory's own machinery, for reasons that have nothing to do with particle masses. The programme builds everything on a single small structure — a "closure cell" shaped like a wheel: a central hub with six spokes out to a ring of six points. That shape has a natural set of vibration patterns, like the tones of a drum. The most basic admissible tone turns out to carry exactly one-twelfth of the structure's total vibrational weight. (It happens that the wheel also has twelve connections, so this looks like "one connection in twelve" — but the real content is the tone's share, not the edge count.) So 1/12 is a property of the shape itself, fixed before any mass is looked at.

The cleanest way to test this is the tau, because it is a "clean" particle — it carries none of the extra complications that quarks do. The prediction of 1/12 matches the measured tau to within about 3%. The strange quark, which does carry those extra complications (colour charge, and the fact that quarks are never found alone), comes out about 14% off. Crucially, that 14% is now something to *explain* — the expected size of the quark's extra baggage — rather than a number to fudge.

The honest caveat: 1/12 can be read two ways inside the structure (the most basic tone has a twofold symmetry, which could count as one unit or as two, giving 1/12 or 1/6). The paper argues that committing to a real, irreversible outcome picks just one — like a fork in the road offering two paths while any actual journey takes only one. The theory's own rules force this single-path reading: once a fact is committed, the cell's symmetry leaves only one consistent choice, and either way the chosen path carries the same weight, so the 1/12 outcome does not depend on which path is taken. The one loose thread is that this symmetry rule is itself a built-in assumption of the wider framework rather than something proved from scratch — so the result is as solid as that assumption, which is used throughout the rest of the theory, rather than a number fitted to the data.

In short: the suppression factor 1/12 is not a fudge fitted to the masses. It is built into the geometry of the theory's fundamental cell, the tau confirms it cleanly, and the strange quark's small deviation becomes a prediction rather than a problem.

Contents

---

1. The Factors to Be Sourced

After the universal localization ladder R_0 is removed, two generation steps retain a residual suppression factor:

Tau: $\Theta_\tau = (m_\tau/m_\mu) / R_0 = 16.82 / 207.13 = 0.0812$.

Down: $\Theta_d = (m_s/m_d) / R_0 = 20 / 207.13 = 0.0966$.

Both lie near 0.09. The programme has carried two candidate rational values for them — 1/12 (= 0.0833) and 1/10 (= 0.100) — neither previously sourced from structure. This paper asks whether either value is forced by objects already derived in the closure programme, before any fermion mass is consulted.

---

2. Imports

Import C1 — The admissible closure cell [Proven, conditional on the wheel-family restriction]

The unique admissible cell of the wheel family is W_7 (Theorem J), with Laplacian spectrum

$\mathrm{spec}(L) = \{0, 2, 2, 4, 4, 5, 7\}$,

7 vertices and 12 edges.

Fallback: if minimality is later established only within a sub-family, the spectrum and edge count are unchanged; only the uniqueness scope narrows. The numerical results below depend on the spectrum, not on global minimality.

Import C2 — Primitive spectral normalization [Proven]

The primitive admissible eigenvalue is exactly

$\lambda^* = 2$,

with twofold degeneracy (Theorem E, condition 2). This is the condition that selects W_7; it is a defining feature of admissibility, not a measured quantity.

Fallback: if a future operator audit shifts the normalization to $\lambda^* = 2 + \varepsilon$, the suppression unit below rescales to $(2 + \varepsilon)/24$; the structural form $\lambda^*/\mathrm{trace}(L)$ is unaffected.

Import C3 — Edge / generator count corroboration [Proven]

The edge count $|E| = 12$ of W_7 coincides with the internal-generator count of the D_7 census ($14 = 2\partial + 12\_\mathrm{int}$) and with $\mathrm{trace}(L)/2 = 24/2$. The integer 12 is therefore attested three independent ways within the closure architecture.

Fallback: if the D_7 census denominator verdict revises 14 = 2 + 12, the edge-count attestation via the spectrum stands alone and is sufficient.

Import L1 — Localization ladder [Conditional]

$R_0 = e^{(16/3)} = 207.13$,

the universal per-generation localization factor, shared across charged fermions.

Import O1 — Observed suppression factors [Observation]

$\Theta_\tau = 0.0812$ (lepton masses, clean); $\Theta_d = 0.0966$ (current / MS-bar quark masses; band 0.092–0.099 across scale).

Import C4 — Reversible transport versus irreversible commitment [Conditional]

The substrate distinguishes reversible transport, which explores a sector without generating a committed fact, from irreversible commitment, which closes a fact into an admissible history [⟨cite: irreversible-commitment / proto-time papers⟩]. Used only in Section 3.3.

Fallback: if the distinction is not formally carried in the cited corpus at the grain Section 3.3 requires, the denominator argument there falls back to the Laplacian identity trace(L) = 2|E| (available unconditionally) and the numerator argument reverts to [Open]; the central value 1/12 (Section 3, eigenvalue reading) is unaffected, only its ontological derivation weakens.

Import C5 — Commitment Criterion (one record per commitment) [Conditional]

An irreversible commitment lays down one record / one fold per commitment, independent of the number of candidate branches [⟨cite: Commitment Criterion paper⟩]. The paper is explicit that branch-neutral cost is an operational/constraint argument, not a first-principles derivation. Used only in Section 3.3, Route 2.

Fallback: if the one-record rule does not hold at the grain Section 3.3 requires, Route 2 lapses and the numerator rests on Route 1 (structural, [Conditional] on C6) alone; the central value is unaffected.

Import C6 — Committed-surface stabiliser is D_6 [Inherited / Conditional on H1]

Committed W_7 surfaces carry the full D_6 stabiliser — $\mathrm{Stab}(W_7^t) = D_6$ for every post-projection committed surface — so that a D_6 reflection is an admissibility condition on committed states. This is established by the K=7 architecture papers, where it enters as a structural input (H1), not as a theorem derived inside them; the constitutive transport paper uses and strengthens it, assigning the full D_6 stabiliser to committed states and the reduced D_3 symmetry only to the *transient* transport excitation, not to committed surfaces. The reflection that exchanges the $\lambda^* = 2$ pair ($k \to -k$, Appendix A) lies in D_6, so it is available on committed surfaces. [⟨cite: K=7 architecture papers (H1); constitutive transport paper⟩]. Used only in Section 3.3, Route 1, and distinct from the transport-sector D_7.

Status note. C6 is therefore *supported but not discharged*: the stabiliser is D_6 by inheritance from H1, and the relevant reflection is confirmed to lie in it, but H1 is a postulate of the architecture rather than a derived result. Route 1's forcing of 1/12 is consequently exactly as strong as H1 — clean and structural, conditional on the K=7 architecture establishing H1, and no stronger.

Fallback: if H1 were withdrawn (committed-surface stabiliser not D_6, or the swapping reflection not in it), Route 1's reflection-definiteness would lapse, the λ^* modes would be freely superposable, and *both* routes would weaken (the ensemble average would undercut Route 2 — see "Route coupling," Section 3.3). The eigenvector identification (Appendix A) is unaffected; only the exclusivity reading depends on H1.

---

3. The Primitive-Mode Weight Fraction

The closure operator L on W_7 distributes its spectral weight across the modes of spec(L). The total weight is

$$\mathrm{trace}(L) = \Sigma\, \lambda_i = 0 + 2 + 2 + 4 + 4 + 5 + 7 = 24.$$

The primitive admissible mode is the binary-commitment mode at $\lambda^* = 2$. Its share of the total closure weight is

$$\Theta\_\mathrm{cell} = \lambda^* / \mathrm{trace}(L) = 2 / 24 = 1/12.$$ [Proven, given C1–C2]

Because $\mathrm{trace}(L) = 2|E|$ for any Laplacian, and $\lambda^* = 2$ by C2, this collapses identically to

$$\Theta\_\mathrm{cell} = \lambda^* / (2|E|) = 2 / (2\cdot 12) = 1/|E| = 1/12.$$

So the suppression unit is one over the edge count of the closure cell. The collapse to $1/|E|$ is a consequence of the $\lambda^* = 2$ normalization (C2): the "exactly 2" condition is load-bearing, and it is the same condition that selects W_7 in the first place. The value 1/12 is therefore not a fit but a fixed point of the admissibility spectrum.

3.1 The form, read as equipartition

Since $\lambda^* = 2$ and $\mathrm{trace}(L) = 2|E|$, the primitive eigenvalue is exactly one edge-unit of spectral weight. The relation $\Theta\_cell = \lambda^*/\mathrm{trace}(L)$ is therefore not a claim about a ratio of eigenvalues but a single, nameable premise:

Equipartition premise. Maintenance load is distributed uniformly per edge-unit of the closure operator's spectral weight, and the maintained commitment occupies one such unit.

Stated this way the open step (Section 12) sharpens: either a determined operator L satisfies per-edge-unit equipartition or it does not, and the question becomes computable once L is frozen rather than a matter of interpretation.

3.2 The one fit-choice that remains: multiplicity

The primitive mode is twofold degenerate (C2 — two independent binary commitment modes). "The weight in the primitive mode" is therefore ambiguous between two readings:

eigenvalue alone: $\lambda^* / \mathrm{trace}(L) = 2/24 = 1/12$ (one edge-unit), primitive sector: $\lambda^* \cdot \mathrm{mult} / \mathrm{trace}(L) = 4/24 = 1/6$ (two edge-units).

The value 1/12 is structural; the *choice* between the eigenvalue reading (1/12) and the sector reading (1/6) is a separate question. Taken in isolation, the only thing selecting 1/12 over 1/6 is that the tau lands on 1/12 — a fit-selection of the reading, even though the value itself is sourced. This is the one seam through which the mass-first failure mode could re-enter. Section 3.3 closes it: the reading is fixed by the cell's D_6 symmetry under commitment (Route 1) and the commitment ontology (Route 2), conditional on an inherited architectural postulate rather than on the masses. The qualification, before that argument, is precise: **the suppression value is sourced from the spectrum; the eigenvalue-versus-sector reading is resolved structurally in Section 3.3, conditional on H1.**

Section 3.3 proposes a resolution drawn from the commitment ontology rather than from the masses. It remains conjectural, but it moves the reading from "fit-selected" to "selected by a single, stated ontological premise."

3.3 Irreversible commitment and the numerator [Conditional on H1]

The occupancy ambiguity of Section 3.2 separates cleanly into two factors that are usually conflated:

the **denominator** — against what whole is the primitive weight measured? the **numerator** — how much of the primitive sector does the maintained object occupy?

A resolution must address both, and it must address them from *different* parts of the structure, because they are different questions. The commitment ontology supplies one each.

The denominator: closure completion. The VERSF substrate distinguishes reversible transport from irreversible commitment (Import C4). A committed fact is not merely transported through a sector; it must be closed into an admissible history and returned to a closure-neutral state, which requires carrying the commitment around the *complete* closure structure. The normalising whole is therefore the entire cell, and the natural denominator is the total spectral weight of the completed loop,

$$\mathrm{trace}(L) = 2|E| = 24.$$

This is what closure completion fixes. Note what it does *not* fix: it tells us the whole is the 12-edge loop, but it says nothing about how much of the primitive sector a single commitment occupies. Closure completion is a statement about the denominator only.

Fallback: if the reversible/irreversible distinction is not formally carried in the cited corpus, the denominator argument reduces to the handshake identity $\mathrm{trace}(L) = 2|E|$, which holds for any Laplacian and is therefore available regardless; only the *ontological* motivation for using the whole loop would weaken, not the value.

The numerator: which is the non-default claim. The primitive sector is twofold degenerate, and by Theorem E (condition 3) the two modes are the **two independent binary commitment modes**. The question is how much of this sector a single committed fact occupies — one mode (numerator $\lambda^* = 2$) or the whole sector (numerator $\lambda^* \cdot \mathrm{mult} = 4$).

Honesty requires naming the default first. In ordinary linear algebra a degenerate eigenspace is *additive*: any combination $a\psi_1 + b\psi_2$ is an equally valid eigenvector at λ^*. Superposition is what degeneracy means. So the mathematically default numerator is the full sector — the 1/6 reading — and the claim that a commitment occupies *one* mode is the non-trivial one that must be earned. "A fork has two roads but a journey uses one" is the right intuition, but a superposable sector's natural state is to be on both roads at once until something forces a selection. The burden is therefore to say what forces it.

Two routes do, with different statuses.

Route 1 — structural exclusivity (cleaner; [Conditional on H1]). The primitive $\lambda^* = 2$ sector is the $k = \pm 1$ Fourier pair on the W_7 rim hexagon (Appendix A). It is the two-dimensional E_1-type irreducible representation of $\mathbf{D_6}$ — the hub-plus-hexagon spatial symmetry of the cell, on which C_6 acts by 60° rotation — not the D_7 of the transport/generation irrep structure (see disambiguation below). Being a *nontrivial* irrep, it contains no nonzero D_6-invariant vector.

This forces a distinction that the argument turns on: **surface symmetry is not state symmetry.** H1 makes D_6 the symmetry of the committed geometric *surface*; it does not make committed *states* D_6-invariant. Indeed they cannot be — a D_6-invariant vector in E_1 is the zero vector. What surface-D_6 supplies is that the reflections are *available admissibility operations*. An admissible committed state is required to be definite under one such operation — a Z_2 condition, not a D_6 condition. Commitment selects a reflection axis and the state becomes an eigenstate of that reflection, breaking $D_6 \to Z_2$. So "reflection-definite" means definite under the one commitment-selected reflection, the axis choice is itself part of what commitment does, and the surface retains D_6 while the state does not. There is no tension between full surface stabiliser and single-mode occupancy: the former is a property of the geometry, the latter of the state.

The numerator then follows, and it is robust to the details. On E_1 every state definite under a maximal abelian subgroup is a single eigenvector at $\lambda^* = 2$: the rotation eigenstates ($k = \pm1$) and the reflection eigenstates (cos, sin) are all 1-dimensional and all carry spectral weight $\lambda^* = 2$. The only object carrying weight $\lambda^* \cdot \text{mult} = 4$ is the full *unresolved* 2D sector. So the dichotomy is clean — *any resolution into a 1D pure state gives numerator $\lambda = 2 \rightarrow 1/12$; only the unresolved 2D sector gives $4 \rightarrow 1/6$** — and the 1/12 outcome does not depend on whether the admissibility condition is reflection-definiteness, rotation-definiteness, or which axis is chosen. The axis-dependence of the *state* does not propagate to the *weight*. Exclusivity is therefore a representation-theoretic fact about the E_1 action, not a dynamical postulate.

Route 1 is [Conditional on H1]: the D_6 reflections are admissibility operations on committed states because committed W_7 surfaces carry the full D_6 stabiliser, $\text{Stab}(W_7^t) = D_6$, a structural input of the K=7 architecture (Import C6, inherited from H1). The swapping reflection lies in D_6 (Appendix A). The objection that commitment might break the symmetry and lose the reflection is met by the constitutive transport paper, which confines the reduced D_3 symmetry to the *transient* transport excitation and assigns full D_6 to the committed surface. Granted H1, Route 1 forces 1/12. The eigenvector identification itself (the $k = \pm1$ reading) is [Proven], Appendix A.

Disambiguation — two distinct group actions. The cell carries two symmetry groups doing different jobs, and they must not be conflated. **D_6** is the visible spatial symmetry of the hub-plus-hexagon cell and the stabiliser of committed surfaces; it acts on the λ^* eigenspace and is the group Route 1 invokes. **D_7** is the transport/generation irreducible-representation structure that sources assignment-2 and generation-3 elsewhere in the programme ($|D_7| = 14 = 2\partial + 12\_\text{int}$); it is a different action and is *not* the group acting on the λ^* modes. The earlier reflex to check the eigenspace against D_7 was a category error: the spectrum's visible symmetry is D_6.

A coincidence to watch, not to use. $|D_6| = 12$ equals $|E(W_7)| = 12$ equals the denominator. This is true for the hub-plus-hexagon specifically and is tempting to promote to "the denominator is the order of the symmetry group." It is not earned: "order of the stabiliser" and "edge count of the cell" are distinct objects that happen to coincide here, exactly the kind of substitution the demoted "edge-unit" was. Flagged as a numerical coincidence, used as nothing.

Route 2 — operational record-cost (available now; [Conditional]). The Commitment Criterion (Import C5) states that an irreversible commitment lays down **one record per commitment, independent of how many candidate branches exist** — one fold, one realized history, even when the candidate space is larger than two. Applied here: the committed fact pays one primitive record-cost, not the cost of the whole available sector. The numerator is therefore one record.

The branch-count independence is the strength of this route. It does not say "select one of two" (which would invite a 50/50 average over both); it says the commitment cost is *one record regardless of branch count*. The numerator is one because that is what a commitment costs — it would be one even if the primitive sector were threefold or fourfold degenerate. This severs the numerator from the multiplicity entirely, rather than relying on the multiplicity being exactly 2.

Route 2 carries two explicit dependencies, both of which the Commitment Criterion itself flags or which the spectral translation introduces:

- **Record–mode identification.** Route 2 yields λ* only if *one record = one primitive mode = λ of spectral weight** — record-counting and mode-counting must be the same counting. This is a natural but non-trivial join between the operational and spectral pictures, and it is the analogue of the demoted "edge-unit": the seam where one must confirm that record-cost and spectral weight share a unit. Note this same join is required by Route 1 (which also counts one eigenvector's *weight* λ*, not a dimensionless eigenvector); it is the one assumption common to both routes and is removed by neither. [Assumption, mild]
- **Single branch versus ensemble.** A committed fact lays down one record, but a mass is a standing property, not a single event. Route 2 closes only if mass tracks the record-cost of *one realized branch* rather than an ensemble average over branches — otherwise the 1/6 sector re-enters as the average. C6 secures this: committed surfaces are reflection-definite (full D_6 stabiliser, H1), so the standing committed state is a single E_1 eigenstate, not an ensemble over both branches, and the average is not the admissible standing object. By the weight-invariance of Route 1, any such 1D resolution carries weight λ* = 2 regardless of axis, so the closure does not depend on which eigenstate is realized. The gap is closed conditional on H1, the same import Route 1 rests on. [Conditional on H1]

The Commitment Criterion is explicit that branch-neutral cost is an operational/constraint argument, not a derivation from first principles, so Route 2 is a *support*, not a proof.

Route coupling. The two routes are not independent, and the dependence favours the result. Route 1's reflection-definiteness is what makes "one realized branch" the admissible object rather than the ensemble average — so the single import C6 (committed-surface stabiliser is D_6, inherited from H1) does double duty: it makes Route 1 a structural forcing of 1/12 *and* closes Route 2's single-branch-versus-ensemble gap. Both routes stand on the same inherited foundation. The contrapositive is the live risk: absent H1, the λ* modes are freely superposable, the ensemble average is the natural standing object, and both routes weaken at once. The result's strength is therefore pinned to H1 — a structural commitment used independently across the K=7 architecture and constitutive transport papers, not a postulate introduced here.

The product. Closure completion fixes the denominator (the whole loop, 2|E|); commitment fixes the numerator (one record / one mode, λ*). The suppression unit is their quotient:

Θ_commit = (one committed record) / (completed loop) = λ* / 2|E| = 2 / 24 = 1/12,

whereas a reversible or ensemble-averaged process, keeping the full sector, would read

$\Theta\_reversible$ = (full primitive sector) / (completed loop) = λ*·mult / 2|E| = 4 / 24 = 1/6.

The two factors come from two distinct pieces of the commitment ontology, not from a single substituted quantity.

```
reversible / unresolved          irreversible / committed
   full sector available            one record laid down
   numerator = λ*·mult = 4           numerator = λ* = 2
        ↓                                  ↓
        ╲          denominator = completed loop = 2|E| = 24
         ╲________________________________╱
                        ↓
        1/6 (sector / ensemble)   1/12 (one realized record)
```

The principle in one line: **the multiplicity counts alternatives, not simultaneous occupancy.** Two roads are available; a committed journey uses one.

A caution on a coincidence. Because $\lambda^* = 2$ and trace(L) = 2|E|, the primitive eigenvalue is numerically equal to "one edge-unit," and the phrase "the primitive eigenvalue occupies one edge-unit" reads as natural for that reason alone. That coincidence is not the argument. The numerator is λ^* because a commitment lays down one record (Route 2) or selects one reflection branch (Route 1), not because λ^* happens to equal 24/12. "Edge-unit" is therefore a derived numerical consequence of $\lambda^* = 2$ on this particular cell, not a primitive, and carries no independent weight.

Status. The reading question of Section 3.2 reduces to a single proposition with two candidate proofs:

An irreversible commitment occupies one primitive branch (one record), not the full twofold-degenerate sector.

The eigenspace is identified (Appendix A, [Proven]): the $\lambda^* = 2$ sector is the $k = \pm 1$ hexagon pair, the E_1 irrep of D_6. Route 1 forces the one-branch reading — any 1D resolution of E_1 carries weight λ^*, so the numerator is λ^* — conditional on the committed-surface stabiliser being D_6 (C6). C6 is inherited from H1 rather than derived: the K=7 architecture assigns committed surfaces the full D_6 stabiliser as a structural input, the swapping reflection lies in it, and the reduced D_3 is confined to transient transport. Route 1 is therefore a structural forcing of 1/12, conditional on H1 and no stronger. Route 2 (Commitment Criterion, C5) supports the same conclusion operationally, and its single-branch-versus-ensemble gap is closed by the same C6/H1 reflection-definiteness. What remains open is narrow: the record–weight identification (a mild assumption shared by both routes) and the status of H1 itself as an architectural postulate rather than a theorem. The reading is thus structurally forced, conditional on an inherited architectural postulate, rather than selected by the masses. [Conditional on H1 (Route 1); on C5 + H1 (Route 2); the only un-inherited residual is the mild record–weight identification.]

---

4. The Colourless Test: The Tau

The tau is the clean test, because it carries no colour and no partial-closure structure to modify the cell coupling. If a sector couples to the closure cell and inherits its primitive-mode weight fraction, the prediction is

$m\_\tau/m\_\mu = R_0 \cdot \Theta\_cell = 207.13 / 12 = 17.26$.

Observed: $m\_\tau/m\_\mu = 16.82$. Error: +2.6%.

The tau did not set the spectral value Θ_cell — it was fixed by the spectrum (C2) before the mass was consulted (the *reading* of that value, 1/12 over 1/6, is settled separately in Section 3.3). The 2.6% agreement is therefore a genuine post-selection check on the value, not a fit. [Conditional — see the coupling caveat in Section 7.]

---

5. The Down Quark and the Residual

Applying the same unit to the down-type first step:

$m\_s/m\_d = R_0 \cdot \Theta\_cell = 17.26$.

Observed: $m\_s/m\_d \approx 20$. Error: −13.7% (the unit underpredicts by ~14%).

Unlike the tau, the down quark is a colour triplet and a partial closure. Under the interpretation that 1/12 is the bare cell coupling, the 14% residual is the signature of exactly that extra structure — the colour and partial-closure corrections the tau does not carry. The residual is therefore reclassified: not a fit error to be absorbed, but a predicted correction term to be derived. [Open — the correction is not yet computed.]

The asymmetry between the two tests is consistent with, and a target for, the colour/partial-closure correction: the same unit fits the structureless lepton to 2.6% and the structured quark to 14%. The direction is suggestive but not yet established — $\Theta\_d = 0.0966$ sits *above* the bare 1/12 = 0.0833, so the correction must reduce the suppression by ~14%, and neither that sign nor that magnitude is computed here. Until the correction is derived with the right sign and size this is a consistency check, not evidence (the corresponding failure mode is F3).

---

6. Why Not 1/10

The fitting candidate 1/10 is not adopted, on the following grounds:

Value	source	tau error	down error
1/12	$\lambda^*/\text{trace}(L) = 1/\lvert E(W_7)\rvert$ [C1–C2]	+2.6%	−13.7%

Value	source	tau error	down error
1/10	none (phase window 12/360 read from the answer)	+23.2%	+3.6%

1/10 fits the down quark better but fits the tau badly (+23%) and has no structural origin — the 12°/360° window that produced it cannot be sourced from the corpus and is set by the answer it targets. 1/12 fits the tau cleanly, fits the down quark to within its structural correction, and is forced by the spectrum. Structural sourcing outranks marginal fit, so 1/10 is rejected.

---

7. Spectral-Measure Screen

To guard against selecting 1/12 because it happens to match, the natural spectral measures of W_7 are screened together, not one at a time:

measure	value	in range 0.08–0.10?
λ*/trace = 1/|E|	0.0833	yes — and interpretable
sum(primitive)/trace = 4/24	0.167	no (high)
1/λ_max = 1/7	0.143	no (high)
λ*/λ_max = 2/7	0.286	no (high)
1/trace = 1/24	0.0417	no (low)

λ*/trace is the unique measure landing in the target range, and the only one with a direct reading — the fraction of closure weight carried by the primitive commitment mode. The tau, which did not set this value, lands on it at 2.6%. The match is therefore not one choice among several near-equal candidates.

---

8. The Edge-versus-Colour Decomposition [Open]

Two structural decompositions of Θ_d now coexist:

Colour route (prior work): Θ_d = 3 × (1/31). Edge route (this paper): Θ_d = 1/12 directly.

These cannot both be the primary factorisation. The tau adjudicates: it has no colour and still lands on 1/12, so the edge route does genuine non-colour work and cannot be an artifact of colour multiplicity. This places the colour-3 decomposition of the down *first step* under suspicion of being a forced factorisation (Θ_d written as 3 × X to expose a colour factor that the tau shows is not required for the suppression itself).

The sharp open question:

Does the down quark suppress by 1/12 (cell coupling, as the tau does), with colour entering elsewhere in the mass relation — or by 3 × 1/31, with the tau's colourless 1/12 a separate coincidence?

The tau's clean 1/12 is the strongest available evidence for the former. Resolving this fixes where colour acts in the quark mass relation and is the immediate structural target. [Open]

---

9. What Is Solved and What Is Not: Value versus Activation

This paper sources the *value* of the suppression unit. It does not source its *activation* — which generation steps wear the unit and which do not. The two are independent, and conflating them is what the common-branch reading got wrong.

Where 1/12 turns on and off, across the steps:

e → μ (lepton, gen 1→2): Θ ≈ 1 — pure localization, coupling off μ → τ (lepton, gen 2→3): Θ ≈ 1/12 — coupling on d → s (quark, gen 1→2): Θ ≈ 1/12 — coupling on

The activation pattern is not by sector — leptons appear both off (first step) and on (second step). It is not by generation number — on at the lepton second step but on at the quark first step. It cuts across both classifications, which is precisely the cross-cutting membership that had no discriminant in the common-suppression-branch treatment.

The honest position is therefore a clean split: this paper answers **"what is the unit"** (1/12 = one edge-unit of the W_7 spectrum, Section 3) and leaves **"which steps wear it"** exactly as open as before. Splitting one problem into a solved half and an open half is the progress; claiming the activation came along with the value would not be.

A conjecture for the activation question, held separate from the result: a step may wear the unit only when its primitive mode is still *available* rather than already consumed by a lower step — a saturation reading, which would connect activation to the Saturation Theorem rather than to the spectrum of a single cell. This is a direction, not a claim. [Conjectural]

---

10. Scope — What This Does Not Establish

- It does not revive a universal suppression branch. The up first-step factor (≈2.84), the up second step (≈0.657), and the down second step (≈0.217) do not sit on 1/12. The unit applies to the down *first step* and the tau specifically.
- It does not derive the coupling. "A sector inherits the cell's primitive-mode weight fraction" is a hypothesis with a structural anchor, not a theorem (Section 12).
- It does not derive the down residual. The 14% is a target, not a computed correction.
- It does not derive H1. The eigenvalue-versus-sector reading (1/12 vs 1/6) is resolved structurally (Section 3.3, [Conditional on H1]), but H1 — the D_6 committed-surface stabiliser the resolution rests on — is inherited as an architectural postulate, not derived here.
- It does not resolve the colour/edge decomposition (Section 8).
- It does not source the activation pattern (Section 9).

---

11. Falsification Conditions

F1 — Coupling without source. No derivation links maintenance suppression to the primitive-mode weight fraction; "1/12 = cell coupling" then remains an interpreted coincidence and the tau match is fortuitous.

F2 — Tau drift. A refined lepton-sector treatment moves Θ_τ away from 1/12 by more than the spectral normalization can absorb, removing the clean colourless anchor.

F3 — Residual fails to source. The down 14% residual cannot be obtained from colour and partial-closure structure, so 1/12 is the wrong unit for the down sector and the edge route fails there.

F4 — Normalization shift. A future operator audit gives $\lambda^* \neq 2$, moving the suppression value off 1/12 and breaking the tau and down matches together.

F5 — Sector reading forced. A determined operator selects the two-unit primitive sector (1/6) rather than the single edge-unit (1/12), against the tau; the value would then be sourced but the wrong reading, and the tau match would be a coincidence of magnitude.

---

12. The Next Derivation: One Gate, Three Questions

The missing step is the derivation, from a determined closure operator L, of the relation

$\Theta_{suppression} = \lambda^* / \mathrm{trace}(L)$ = one edge-unit of spectral weight, rather than its current status as an interpreted match. Because $\lambda^* = 2$ and trace(L) = 2|E|, this is the equipartition premise of Section 3.1, and it resolves into three specific questions — all decided by the *same* object, a determined L:

1. **Uniformity.** Is maintenance load distributed uniformly per edge-unit of L's spectral weight? (The equipartition premise itself.)
2. **Occupancy.** Does the maintained commitment occupy one mode (→ numerator λ^*, giving 1/12) or the full two-mode primitive sector (→ numerator $\lambda^*\cdot$mult, giving 1/6)? This is the multiplicity choice of Section 3.2.
3. **One-branch occupancy under commitment.** Does an irreversible commitment occupy one primitive branch (one record) rather than the full twofold-degenerate sector? This is structurally forced, conditional on an inherited architectural postulate. Route 1 (structural): the $\lambda^* = 2$ sector is the E_1 irrep of D_6 (the k = ±1 hexagon pair), with no D_6-invariant vector; committed surfaces carry the full D_6 stabiliser (C6, inherited from H1), so a committed state breaks D_6 to a Z_2 and occupies one 1D resolution — and every 1D resolution carries weight λ^*, forcing 1/12 conditional on H1, independent of axis. Route 2 (operational, Commitment Criterion C5) yields the same conclusion, its ensemble gap closed by the same C6/H1. Both rest on H1 (committed-surface stabiliser is D_6), a postulate of the K=7 architecture, plus the mild record–weight identification. The denominator (the completed loop, 2|E|) is the closure-completion half of the same question. The residual derivation target is therefore not the occupancy itself but the status of H1: whether the D_6 committed-surface stabiliser can be derived rather than postulated.

All three are properties of a fixed operator, not of a chosen spectrum. They land on the same operator-determinacy slot left [Open · Gate: audit-determinacy] by the Stable-Spectrum Gap Functional work. If the frozen audit there determines L uniquely, questions 1–3 become computable, and this paper's central relation moves from [Conjectural] to [Conditional] on that single gate.

The activation question of Section 9 — *which* steps wear the unit — is deliberately not folded into this gate. It is a separate animal, conjecturally a saturation/availability question rather than a single-cell spectral one, and is held apart so the determinacy gate is not overloaded with a problem it cannot decide.

Status Ledger

Claim	Grade
W_7 spectrum and |E| = 12 (C1, C3)	[Proven] (wheel-family scope; inherits Theorem J cap)
$\lambda^* = 2$ normalization (C2)	[Proven]
Suppression *value* = λ^*/trace(L) = 1/|E| = 1/12 (eigenvalue reading)	[Proven] given C1–C2

Claim	Grade
Eigenvalue reading (1/12) vs sector reading (1/6)	structurally forced, [Conditional on H1] (Section 3.3)
λ* = 2 sector is the k = ±1 hexagon Fourier pair (Appendix A)	[Proven]
One-branch occupancy — Route 1 (D_6 reflection-conjugate pair)	forces 1/12, [Conditional on H1] via C6
One-branch occupancy — Route 2 (Commitment Criterion one-record rule)	[Conditional] on C5; ensemble gap closed by C6/H1
C6 (committed-surface stabiliser is D_6)	[Inherited / Conditional on H1] — supported by K=7 architecture + constitutive transport papers; reflection confirmed in D_6; not derived
H1 (D_6 committed-surface stabiliser) itself derived rather than postulated	[Open] — the remaining derivation target
Tau match to 2.6% via cell coupling	[Conditional] on the coupling
Down 14% residual = colour/partial-closure correction	[Open]
Sector inherits the unit (coupling / equipartition)	[Conjectural] → [Conditional] if operator-determinacy gate discharges
Edge-vs-colour primary decomposition	[Open]
Activation pattern (which steps wear the unit)	[Open]; saturation reading [Conjectural] (Section 9)
Universal suppression branch	[Rejected] (Section 10)

Conclusion

The suppression unit shared by the tau and down-type first steps is sourced from structure derived for closure reasons alone:

$\Theta\_cell = \lambda^* / \mathrm{trace}(L) = 1/|E(W_7)| = 1/12$,

forced by the λ* = 2 admissibility normalization. The tau, which did not set this value, is reproduced to 2.6%. The down quark departs by 14%, in a direction consistent with — and a target for — its extra colour and partial-closure structure, converting that residual from a fit error into a derivation target.

The fitting candidate 1/10 is not adopted: it fits the down quark marginally better but fits the tau at +23% and has no structural source. The result's scope is deliberately narrow. It sources the *value* of the unit, not its *activation* — the question of which generation steps wear it cuts across both sector and generation and remains open, conjecturally a saturation question. One choice also sits inside the value: the eigenvalue reading (1/12) over the sector reading (1/6). Section 3.3 settles it: the $\lambda^* = 2$ sector is the E_1 irrep of D_6 (the $k = \pm1$ hexagon pair, Appendix A), a nontrivial irrep with no invariant vector, so a committed state cannot be D_6-invariant; commitment selects a reflection axis, breaking $D_6 \to Z_2$ in the *state* while the *surface* retains D_6. Every 1D resolution of E_1 carries weight $\lambda^* = 2$, so the one-branch reading gives 1/12 independent of the axis, and only the unresolved 2D sector gives 1/6. This reading is structurally forced conditional on H1 (committed surfaces carry the full D_6 stabiliser, a structural input of the K=7 architecture, with the swapping reflection in it and the reduced D_3 confined to transient transport) — Route 1 via that surface-D_6, Route 2 (the Commitment Criterion's branch-count-independent one-record rule) supporting the same conclusion with its ensemble gap closed by the same H1. The result is exactly as strong as the architectural postulate H1 and no stronger; the only un-inherited residual is the mild record–weight identification shared by both routes. The result does not revive the common suppression branch, derive the coupling, or resolve whether the down sector's primary decomposition is the edge unit 1/12 or the colour product 3 × 1/31 — the last adjudicated, on present evidence, in favour of the edge unit by the tau's match.

The remaining derivation target has shifted accordingly. With the one-branch reading inherited from H1, the open question is no longer "which reading" but whether H1 itself — the D_6 committed-surface stabiliser — can be derived rather than postulated, alongside the operator-level statement that maintenance suppression equals one committed record over the completed loop (uniformity and the record–weight identification), which lands on the operator-determinacy gate already isolated by the Stable-Spectrum Gap Functional work. Until H1 is derived and that gate discharges, the strongest defensible claim is the one the tau underwrites at 2.6%: the closure spectrum contains a suppression unit of 1/12 whose one-branch reading is forced by the architecture's own D_6 stabiliser postulate, the tau is its clean colourless test, and the strange quark carries that same unit plus a colour and partial-closure correction.

Appendix A — The $\lambda^* = 2$ Eigenspace Is the k = ±1 Hexagon Pair [Proven]

W_7 is the hub vertex joined to a 6-cycle rim (C_6). Label the hub value h and the rim values $r_0,\ldots,r_5$. The graph Laplacian acts as

$(Lx)_{hub} = 6h - \Sigma_j r_j$, $(Lx)_j = 3 r_j - h - r_{(j-1)} - r_{(j+1)}$,

since the hub has degree 6 and each rim vertex has degree 3 (two rim neighbours plus one spoke).

Rim Fourier modes, k ≠ 0. Take $r_j = \exp(i\cdot 2\pi kj/6)$ for $k = 1,\ldots,5$, with $h = 0$. For $k \neq 0$ the rim values sum to zero, so the hub equation reads $6\cdot 0 - 0 = 0$ and is satisfied. The rim equation gives

$(Lx)_j = 3\, r_j - (\exp(-i2\pi k/6) + \exp(+i2\pi k/6))\, r_j = (3 - 2\cos(2\pi k/6))\, r_j,$

so each $k \neq 0$ mode is an eigenvector with eigenvalue

$\lambda(k) = 3 - 2\cos(2\pi k/6)$:

$k = \pm 1 \rightarrow 3 - 2\cos 60° = 3 - 1 = 2$ (twofold), $k = \pm 2 \rightarrow 3 - 2\cos 120° = 3 + 1 = 4$ (twofold), $k = 3 \rightarrow 3 - 2\cos 180° = 3 + 2 = 5$ (single).

Symmetric subspace (hub + uniform rim). The $k = 0$ rim mode (all $r_j = s$) mixes with the hub. In the (h, s) basis the Laplacian reduces to

$M = [\ 6\ {-6}\ ]\ [\ {-1}\ 1\ ]$,

with trace 7 and determinant $6 \cdot 1 - (-6)(-1) = 0$, hence eigenvalues 0 and 7.

Total. $\{0, 7\} \cup \{2, 2, 4, 4, 5\} = \{0, 2, 2, 4, 4, 5, 7\}$, matching spec(L) of Import C1 exactly.

*Identification of λ.** The primitive eigenvalue $\lambda^* = 2$ is the $k = \pm 1$ pair: the two lowest nontrivial Fourier modes on the rim hexagon, carried into the wheel spectrum with hub-value zero. As a D_6 representation this sector is the two-dimensional E_1 irrep, on which C_6 acts by 60° rotation. Being nontrivial, E_1 contains no nonzero D_6-invariant vector; a D_6-invariant state in this sector is the zero vector.

Reflection action and weight invariance. A D_6 reflection of the hexagon sends $j \rightarrow -j$ (about a chosen axis), hence $k \rightarrow -k$, exchanging the $k = +1$ and $k = -1$ modes. In real form it fixes $\cos(2\pi j/6)$ (reflection-even, +1) and negates $\sin(2\pi j/6)$ (reflection-odd, −1). More generally every state in E_1 that is definite under a maximal abelian subgroup is 1-dimensional and carries spectral weight $\lambda^* = 2$: the rotation eigenstates ($k = \pm 1$) and the reflection eigenstates (cos, sin) all have eigenvalue 2. The only object carrying $\lambda^* \cdot \text{mult} = 4$ is the full *unresolved* 2D sector. Hence any 1D resolution — by any reflection axis or by rotation-definiteness — gives numerator $\lambda^* = 2$, and the weight does not depend on which resolution or axis is chosen; only retaining the unresolved sector gives 4. This is the eigenvector content used by Section 3.3, Route 1. Whether a 1D resolution is an admissibility condition on committed states depends on Import C6 (the committed-surface stabiliser), inherited from the architectural postulate H1; note that, because the committed state breaks D_6 to the Z_2 generated by the selected reflection, full surface-D_6 and single-mode state occupancy are consistent, not in tension. Appendix A establishes only the eigenvector identification, the E_1 structure, and the weight invariance, which are [Proven] independent of C6/H1.