The Projected Leakage Support Trace in the Weak-Commitment Block

▲ Programme Milestone — Standard Model Flavour Series

Deriving Tr_support(Π_leak) = 20, Unit-Slot Normalisation, and β = √3/20 from the Weak-Commitment Attachment Projector

Keith Taylor VERSF Theoretical Physics Programme — Standard Model Flavour Series Tier A gate paper, SM-3

General Reader Summary

Neutrinos are the lightest known particles of matter, and they do something strange: as they travel, they shift between three different identities. Physicists describe this shifting with a handful of measured numbers called mixing angles. Two of those angles are large. One — the so-called reactor angle — is small but definitely not zero, and a second sits close to 45° but tilted slightly away from it.

In the VERSF programme, both of those small effects trace back to a single number:

β = √3/20 ≈ 0.087.

Think of β as the size of a small imbalance. In a perfectly symmetric world, a neutrino born alongside an electron would connect equally to the two other neutrino types. In the real world that balance is very slightly broken, and β measures by how much. That slight break is what produces the reactor angle and the tilt away from 45°.

Until now, the programme could show that *if* β takes the value √3/20, the predicted mixing angles come out right. But the value itself was still an input — a number that worked, not a number that had been explained. In a programme whose whole purpose is to derive the constants of nature rather than fit them, that is an exposed flank. This paper closes it.

The explanation turns out to be a ratio of two counts, and both counts can be described without any machinery.

Why the top is √3. The imbalance is fed by three separate sources of equal strength. Naively you might add them and get 3. But these sources combine the way perpendicular steps combine, not the way steps along a line combine. Walk one block east and one block north and you are not two blocks from where you started — you are √2 blocks away, by Pythagoras. Three equal, mutually independent sources combine the same way: their total strength is √3, not 3.

Why the bottom is 20. The imbalance does not act at a single point. It gets spread — diluted — across every independent "slot" the theory allows it to occupy, and this paper counts those slots. The counting goes in two stages:

- **Five places it can sit.** The relevant part of the theory has seven directions, but two of them are already spoken for — they are pinned down as fixed reference points before the imbalance is even measured. A pinned direction cannot carry an imbalance. That leaves $7 - 2 = 5$.
- **Four ways it can point.** In each of those five places, the imbalance is not a simple quantity like a temperature; it is more like an arrow. Describing the arrow takes two independent numbers (roughly, its length and its angle), and the arrow has two distinct configurations available to it — and both count as places the imbalance could go, whether or not it goes there. That gives $2 \times 2 = 4$.

Five places, four ways each: $5 \times 4 = 20$ slots. The source's combined strength then behaves like a fixed budget shared out equally among the twenty slots, so each slot carries one-twentieth of it: $\sqrt{3}/20$. Notice that two different rules are at work, on purpose. Combining independent sources is the Pythagoras-style step; sharing the combined strength among slots is ordinary division, the way a fixed sum of money divides among accounts. The paper does not treat this sharing rule as obvious — it states it explicitly as one of its named assumptions, with its own failure condition attached.

One subtlety nearly derails the count, and fixing it is the real work of this paper. The mathematics contains a mirror rule: whenever the theory records a connection from A to B, it must also record the matching connection from B to A. It is tempting to count that mirror copy as an extra slot — which would double the answer to 40. But the mirror copy carries no new information; it is fixed the instant the original is fixed, like a reflection that cannot move independently of the object. The genuine second factor of two comes from somewhere else entirely: the two distinct configurations available to the arrow, which really are separate directions the imbalance could take. The paper proves the mirror copy contributes nothing, identifies the genuine doubling, and shows why the answer is 20 — not 10 (missing the real doubling) and not 40 (counting the fake one).

So the headline result is:

$\beta = \sqrt{3}/20$, derived rather than assumed.

In one sentence:

The imbalance behind the reactor angle is small because it is spread across twenty available slots, but not tiny, because it is fed by three sources acting together.

That is the SM-3 milestone. What this paper does *not* do is also worth stating plainly: it fixes the *size* of the imbalance, not its *direction* — whether the atmospheric angle tilts above or below 45° is a separate question, reserved for a later paper in the series.

Contents

Abstract

The Weak-Doublet Projection Theorem reduced the residual Standard Model flavour problem to the projected weak-doublet closure block

$H_W = P_W^\dagger H_{cl} P_W$,

where $H_{cl} = G|_{\mathfrak{su}(8)}$ is the inherited admissibility Hessian and P_W is the ordered generation, weak-role, and mass/readout compression. The CKM-side SM-2 gate tests whether the minimal C_3 branch curvature survives the full $\mathfrak{su}(8)$ embedding. The present SM-3 paper closes the next exposed PMNS object: the weak-commitment leakage amplitude

$\beta = \sqrt{3}/20$.

Prior weak-commitment work used β as the small attachment-leakage parameter controlling the reactor angle, the atmospheric deviation from maximality, and the leading electron-to-non-electron asymmetry. The present paper derives that value as a support trace of the projected weak-commitment leakage projector Π_{leak}:

$\mathrm{Tr}_{support}(\Pi_{leak}) = 20$.

The trace is not inserted phenomenologically. It follows from the leading-order factorisation

$\Pi_{leak} = \Pi_R \otimes \Pi_Q$,

where R is the residual weak-commitment support remaining after two boundary/readout-anchored directions are removed from the K = 7 closure carrier,

$\dim R = K - 2 = 5$,

and Q is the weak-attachment orientation fibre,

$\dim Q = 2 \times 2 = 4$.

The first orientation factor is the weak-neutral real quadrature pair. The second is a closure-orientation binary of the weak-neutral attachment half-cycle. It is **not** the Hermitian return entry: Hermitian conjugation supplies consistency, not an independent support slot, and the return-entry count is excluded by an explicit exclusion lemma.

With Killing-unit support normalisation,

$\mathrm{Tr}_{support}(\Pi_{leak}) = \dim R \cdot \dim Q = 5 \cdot 4 = 20$.

The leakage numerator is the root-normalised norm of a threefold completion residue,

$\|L_{C_3}\|_K = \sqrt{3}$,

so that

$\beta = \|\hat{L}\_C_3\|\_K / \mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = \sqrt{3}/20$,

where $\hat{L}\_C_3$ is the source embedded in the full leakage support. The amplitude functional itself — linear division of the root-normalised source strength by the unit-slot count — is not assumed silently: it is stated as an explicit per-slot allocation premise (WL-7) with its own falsifier.

This closes the SM-3 support-trace denominator at leading weak-commitment order and reduces the remaining amplitude claim to two explicit structural commitments: the closure-orientation binary (WL-5) and the linear per-slot commitment-weight allocation rule (WL-7). Both are stated as named premises, both carry falsifiers, and neither is hidden inside the numerical agreement with the PMNS benchmark. The denominator 20 and the unit-slot normalisation are derived from the projected support structure. The atmospheric octant sign σ_W remains the separate SM-6 gate.

---

0. Notation, Conventions, and Firewalls

0.1 Projected weak-doublet object

The inherited weak-doublet object is

$H\_W = P\_W^\dagger\ H\_cl\ P\_W$,

where

$H\_cl = G|\_\mathfrak{su}(8)$

is the $\mathfrak{su}(8)$-restricted admissibility Hessian and

$P\_W = P\_Y\ P\_R\ P\_C$

is the ordered generation, weak-role, and mass/readout compression. The daggered form is retained because the full projection is ordered: it keeps Hermiticity explicit even where the component projections are not assumed to commute.

0.2 Weak-commitment block

Let P_e, P_μ, P_τ denote the charged-lepton-labelled weak-commitment readout projectors after charged-lepton anchoring. The attachment block is

$H\_att = P\_e\ H\_W\ (P\_\mu + P\_\tau)$.

Its symmetric part gives the common electron-to-pair attachment A. Its antisymmetric (leakage) part gives the first admissible electron-to-non-electron asymmetry.

The weak-commitment kernel has the target form

$$M_\nu = \begin{pmatrix} r_e & A & A(1 + \beta e^{-i\psi}) \\ A & r_s + \delta & B \\ A(1 + \beta e^{i\psi}) & B & r_s - \delta \end{pmatrix}$$

up to the reflected electron–muon / electron–tau branch.

0.3 SM-3 firewall

This paper derives

$\mathrm{Tr\_support}(\Pi_{leak}) = 20$ and hence $\beta = \sqrt{3}/20$.

It does **not** derive:

- the atmospheric octant sign σ_W;
- the full PMNS Hamiltonian;
- the neutrino mass mechanism;
- the absolute neutrino mass scale;
- the full lepton Yukawa matrix.

Those belong to later gates, principally SM-6 and SM-24.

0.4 Claim-level firewall

SM-3 is closed in the following precise sense: **the projected leakage support trace and unit-slot normalisation are derived at leading weak-commitment order under Premises WL-1–WL-7.** The paper does not claim that every PMNS observable follows from the full microscopic Hessian. It claims that the exposed leakage denominator 20 is no longer a benchmark.

0′. Predictive-Content Ledger

Object	Prior status	Status here
$H_W = P_W^\dagger H_{cl} P_W$	inherited projected object	used as weak-commitment parent block
Weak-commitment kernel M_ν	inherited target	used as leakage readout frame

Object	Prior status	Status here
$\beta = \sqrt{3}/20$	benchmark / conditional support target	**derived from support trace**
Leakage numerator $\sqrt{3}$	root-three support claim	derived from Killing-orthonormal C_3 source
Residual support R	inherited K − 2 claim	derived as anchored complement in K = 7 carrier
dim R = 5	prior support count	closed
Residual split $R = C \oplus G$ (WL-3)	inherited support grammar	inherited grade tracks completion-interface parent claim (cross-paper audit open)
Orientation fibre Q	fragile fourfold claim	conditional on WL-5: weak-neutral quadrature × closure-orientation binary (involution construction owed upstream)
Hermitian return counting	possible double-counting risk	**excluded by lemma**
$\Pi\_leak = \Pi\_R \otimes \Pi\_Q$	previous conditional factorisation	promoted to leading-order projection lemma with stated hypotheses
Unit-slot normalisation	previously implicit	explicit Killing-unit support basis
Amplitude functional $\beta = \lVert\hat{L}\_C_3\rVert\_K / Tr\_support(\Pi\_leak)$	implicit combination rule	stated as Premise WL-7 (per-slot allocation), falsifiable via F13
$Tr\_support(\Pi\_leak) = 20$	owed	**derived**
Benchmark kernel outputs (θ_{12}, θ_{13}, θ_{23}, J_lep)	quoted	verified by exact Hermitian diagonalisation
Atmospheric octant sign σ_W	separate projected sign	deferred to SM-6

---

1. Purpose, Claim Level, and Named Premises

1.1 The question

Why is the weak-commitment leakage amplitude $\beta = \sqrt{3}/20$?

The answer defended here:

$\beta = \lVert\hat{L}\_C_3\rVert\_K / Tr\_support(\Pi\_leak)$, with $\lVert\hat{L}\_C_3\rVert\_K = \sqrt{3}$ and $Tr\_support(\Pi\_leak) = 20$.

Claim level: leading-order projected support-trace theorem under Premises WL-1–WL-7. This is strictly stronger than the earlier conditional statement "if the leakage support factors as 5

× 4, then $\beta = \sqrt{3}/20$." Here the factorisation is derived from named premises, each carrying its own falsifier.

Stated result. The SM-3 result is not that every ingredient of the PMNS sector has been derived. It is narrower and stronger: the formerly exposed denominator in $\beta = \sqrt{3}/20$ is identified as the support trace of the projected weak-commitment leakage projector,

$$\mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = 20.$$

The remaining dependence is explicit rather than hidden. If the closure-orientation binary exists (WL-5) and the weak-commitment readout is the linear mean per-slot commitment weight (WL-7), then the leakage amplitude is fixed:

$$\beta = \sqrt{3}/20.$$

SM-3 is thus closed at leading support-trace order, conditional on WL-5 and WL-7, with both conditions falsifiable.

1.2 Named premises

The derivation rests on seven premises. Each is stated once, used where cited, and paired with a falsifier in §13.

Premise WL-1 (Closure carrier). The weak-commitment leakage channel is read inside the $K = 7$ closure carrier inherited from the closure-operator specification.

Premise WL-2 (Two-anchor removal). Exactly two carrier directions are fixed before leakage is counted: the charged-lepton commitment anchor and the weak-doublet complement/readout anchor. These are boundary/readout directions, not live leakage directions.

Premise WL-3 (Residual split). The residual support decomposes as $R = C \oplus G$ with $\dim C = 3$ (threefold completion residue) and $\dim G = 2$ (residual gap/continuation support). The threefold residue C inherits from the completion-interface structure; its grade here tracks the parent claim in the Master Ledger and is not upgraded by this paper.

Premise WL-4 (Weak-neutral off-diagonal leading leakage). The leading admissible μ–τ-breaking channel is a weak-neutral, off-diagonal attachment asymmetry; first-order diagonal μ–τ splitting is suppressed, $\delta = 0 + O(\beta^2)$.

Premise WL-5 (Closure-orientation binary). The weak-neutral attachment half-cycle admits exactly two closure-orientation sectors prior to Hermitian completion, exchanged by a closure-sector involution that is not complex conjugation.

Premise WL-6 (Leading-order separation). Residual placement (R) and weak-neutral orientation (Q) do not mix at leading weak-commitment order: the role-changing modes coupling them are gapped or subleading.

Premise WL-7 (Per-slot commitment-weight allocation). The projected leakage amplitude is a per-slot weight: the root-normalised source strength is a conserved commitment-weight budget distributed uniformly over the Killing-unit slots of the leakage support. Combination of independent sources is quadratic — norms add in quadrature — but allocation of the combined strength over admissible slots is linear: a budget divides. Hence

$$\beta = \|\hat{L}_{C_3}\|_K / \mathrm{Tr}_{\mathrm{support}}(\Pi_{\mathrm{leak}}),$$

where $\hat{L}_{C_3}$ is the embedded source defined in §7. The premise includes the readout half of the claim: the weak-commitment kernel's leakage entry is a single-slot readout — it registers the allocation carried by one Killing-unit slot, not the total source strength and not a sum over the slots of the orientation fibre at a given placement.

The paper closes SM-3, and only SM-3. It does not close the octant sign, the full PMNS Hamiltonian, or the neutrino mass branch.

---

2. The Weak-Commitment Leakage Problem

In the exact μ–τ symmetric limit, the neutrino weak-commitment kernel has equal electron attachments to the two non-electron channels:

$$P_e H_W P_\mu = P_e H_W P_\tau.$$

That symmetry produces a large-mixing neutrino frame, but not the observed reactor-angle structure. The leading admissible breaking must therefore enter as a small off-diagonal attachment leakage,

$$P_e H_W P_\tau - P_e H_W P_\mu \neq 0,$$

or its reflected branch.

By Premise WL-4, weak-neutrality suppresses first-order diagonal μ–τ splitting:

$$\delta = 0 + O(\beta^2).$$

Hence the leading leakage is not a diagonal stiffness difference; it is an off-diagonal weak-neutral attachment asymmetry. This is why β matters: it is not an arbitrary small correction. **It is the first admissible weak-commitment leakage channel**, and its size is fixed by how a coherent source is diluted over the support the channel is permitted to occupy.

---

3. Unit-Slot Support Trace

The support trace used throughout is not an ordinary trace over arbitrary coordinates. It is a trace over Killing-unit support slots.

Definition 1 (Killing-unit support slot). A support slot is a real direction u_A in the leakage support satisfying $\langle u_A, u_A\rangle_K = 1$, where $\langle\cdot,\cdot\rangle_K$ is the inherited Killing form on the projected block.

Definition 2 (Support trace). Let $\{u_A\}$ be a real Killing-orthonormal basis for the leakage support, $\langle u_A, u_B\rangle_K = \delta_{AB}$, and let

$$\Pi_{leak} = \Sigma_A |u_A\rangle\langle u_A|_K.$$

Then

$$Tr_{support}(\Pi_{leak}) = \Sigma_A \langle u_A, \Pi_{leak} u_A\rangle_K.$$

Because every slot is Killing-unit normalised, the support trace of a projector is exactly the number of independent slots in its range. This is the unit-slot normalisation SM-3 requires: no slot enters with weight other than 1, and no non-slot direction contributes.

Lemma 1 (Unit-slot trace lemma). If Π is an orthogonal projector on a real Killing-orthonormal support space of dimension n, then

$$Tr_{support}(\Pi) = n.$$

Proof. Choose a Killing-orthonormal basis $\{u_A\}$, $A = 1, \ldots, n$, for range(Π). Then $\Pi u_A = u_A$, so

$$Tr_{support}(\Pi) = \Sigma_A \langle u_A, \Pi u_A\rangle_K = \Sigma_A \langle u_A, u_A\rangle_K = n. \blacksquare$$

Reading. Lemma 1 converts the SM-3 question from "what is a trace?" to "how many independent leakage slots exist?" Everything that follows is dimension counting under explicit premises — which is exactly what makes the count falsifiable.

---

4. The Residual Support Theorem: dim R = K − 2 = 5

By Premise WL-1, the closure carrier has

$$K = 7$$

admissibility directions. By Premise WL-2, two of these are not live leakage directions: they are spent fixing the boundary/readout frame —

- the charged-lepton commitment anchor;

- the weak-doublet complement/readout anchor.

Lemma 2 (Residual support dimension). At leading weak-commitment order, the leakage residual support has dimension

$\dim R = K - 2 = 7 - 2 = 5.$

Proof. The leakage channel is read inside the $K = 7$ carrier (WL-1). The two anchors of WL-2 are independent — one fixes the charged-lepton commitment frame, the other the weak-doublet complement/readout frame — and both are excluded from live leakage support by construction: an anchored direction cannot carry an asymmetry it is defined to hold fixed. The residual space therefore has dimension $7 - 2 = 5$. ■

By Premise WL-3, the residual support decomposes as

$R = C \oplus G$, $\dim C = 3$, $\dim G = 2$, $\dim R = 5$,

with C the threefold completion residue and G the twofold residual gap/continuation support.

Reading. The denominator's first factor is not 7, because two closure directions are already spent fixing the boundary/readout frame. It is not 3, because leakage does not live only on the completion residue: it is diluted over the full residual support $R = C \oplus G$. The source lives on C; the support is all of R. The distinction between where the source *sits* and where the leakage *may spread* is what separates the numerator count (§7) from the denominator count.

---

5. The Orientation Fibre Theorem: dim Q = 4

The leakage is not a scalar weight on R. It is a weak-neutral off-diagonal attachment (WL-4), and an attachment carries orientation data. The orientation fibre is

$Q = Q_N \otimes Q_C$, $\dim Q_N = 2$, $\dim Q_C = 2$, $\dim Q = 4$.

5.1 The weak-neutral quadrature pair

The leading leakage is weak-neutral and off-diagonal. A weak-neutral complex attachment requires two real quadrature directions:

$Q_N = \mathrm{span}_{\mathbb{R}}\{q_R, q_I\}$, $\dim Q_N = 2$.

This factor is secure: a single complex off-diagonal entry is two real degrees of freedom, no more and no fewer.

5.2 The return-entry exclusion lemma

The second factor of two is the critical point, and it must first be protected against the double-counting objection.

Lemma 3a (Return-entry exclusion). The Hermitian return entry contributes zero independent support slots.

Proof. Hermiticity of H_W gives

$$(P_e H_W P_\tau)^\dagger = P_\tau H_W P_e.$$

Write the forward attachment entry as $z = x + iy$ with (x, y) the two real quadratures of Q_N. The return entry is then $\bar{z} = x - iy$ — a function of the same two real degrees, with no free parameter of its own. A support slot, by Definition 1, is an independent real direction; a functionally determined entry spans no new direction. Hence the return entry adds 0 to the slot count. ■

Counting Hermitian return as an independent factor would count the same matrix element twice and inflate the denominator to 40. Lemma 3a forbids this count.

5.3 The closure-orientation binary

The genuine second binary is supplied by Premise WL-5. Before Hermitian completion, the weak-neutral attachment half-cycle admits two closure-orientation sectors: the leakage can sit in either of two admissible placements of the weak-neutral attachment relative to the residual support,

$$Q_C = \mathrm{span}_{\mathbb{R}}\{q_+, q_-\}, \dim Q_C = 2.$$

Three properties distinguish this binary from Hermitian return:

1. **Residence.** The two sectors q_+, q_- are closure-sector placements — properties of the half-cycle's attachment to the residual support — not entries of the matrix and its transpose.
2. **Independence.** Fixing the forward quadratures (x, y) does not fix the closure orientation: both sectors are compatible with any (x, y). By contrast, fixing (x, y) fixes the return entry completely (Lemma 3a).
3. **Involution.** The sectors are exchanged by a closure-sector involution acting before Hermitian completion. Complex conjugation acts after completion and maps each sector to itself.

One calibration must be stated plainly. The three properties above characterise the binary; they do not construct it. The closure-sector involution itself is owed to the closure-operator specification, and until that cross-reference is discharged, the binary — and with it the factor separating the denominator 10 from 20 — is carried at Conditional grade on WL-5, in exact parallel with the treatment of WL-6 in §6. The certificate in §15 records this grade; it is not obscured.

5.4 Independence of the closure-orientation binary

The closure-orientation binary must not merely be named; it must be distinguished from every structure already counted elsewhere. We therefore impose the following independence test.

Let the weak-neutral forward attachment be represented by the quadrature pair

$z = x + iy, (x, y) \in Q_N$.

Hermitian completion fixes the return entry as

$\bar{z} = x - iy$.

The return entry is a determined function of (x, y) and, by Lemma 3a, contributes no independent support slot.

By contrast, the closure-orientation sector is not a value of z, not the conjugate of z, and not a second copy of the same off-diagonal entry. It is a choice of closure-side placement for the weak-neutral attachment half-cycle before Hermitian completion. The two sectors are represented abstractly as

$q_+, q_- \in Q_C, Q_C = \mathrm{span}_\mathbb{R}\{q_+, q_-\}$.

The independence condition is:

$(x, y, q_+) \not\equiv (x, y, q_-)$

under the weak-commitment projection, while

$(x, y, \bar{z})$

contains no datum not already contained in (x, y).

Equivalently, the closure-orientation binary is admissible only if fixing the weak-neutral quadratures does not fix the closure sector. If the sector is fixed once (x, y) is fixed, the binary collapses and the denominator becomes 10. If the sector is merely the Hermitian return entry in disguise, the binary is spurious and the denominator must not be 20. This is precisely the content of falsifiers F6 and F7.

One further clarification protects the count. The sector structure enters the denominator as a two-dimensional span of admissible directions, $Q_C = \mathrm{span}_\mathbb{R}\{q_+, q_-\}$ — not as a one-bit occupied label. The embedded source of §7 occupies a single orientation, encoded by q_0, exactly as it occupies only the completion residue C within R. But WL-7's allocation runs over admissible slots, not occupied ones: the support trace counts every admissible direction, whether or not the source sits in it. The Q-logic is therefore strictly parallel to the R-logic of §4 — the source sits on C while the support counts all of R; the source sits in one orientation while the support counts both sector directions. If the sector were instead a discrete superselection-type label, contributing one admissible direction per configuration, the allocation support would collapse to 5 × 2 = 10 within the occupied sector and the denominator 20 would fail. That failure route is recorded under F6.

The fourfold orientation fibre is therefore not

2 quadratures × 2 Hermitian directions,

which would be double-counting. It is

2 quadratures × 2 closure-orientation sectors.

This distinction is load-bearing. The count 20 survives only if Q_C is a genuine pre-Hermitian closure-sector fibre.

Lemma 3 (Fourfold orientation fibre). At leading weak-commitment order,

$\dim Q = 4$.

Proof. The leakage is weak-neutral and off-diagonal (WL-4), so it requires two real quadrature directions: $\dim Q_N = 2$. The attachment half-cycle carries the closure-orientation binary of WL-5, independent of the quadratures by the independence condition of §5.4: $\dim Q_C = 2$. The Hermitian return entry contributes no slot (Lemma 3a). The independent orientation fibre is therefore

$Q = Q_N \otimes Q_C$, $\dim Q = 2 \cdot 2 = 4$. ■

Reading. This is the load-bearing strengthening of the paper. The denominator 20 survives only if the second binary is genuinely closure-sector-resident. If it collapsed into Hermitian return, the old argument would fail — and F7 in §13 says so explicitly. Here the binary is assigned to the closure-orientation fibre by WL-5, its independence from the quadratures is stated as property 2 of §5.3 and sharpened into the independence condition of §5.4, and the return-entry count is excluded by a lemma rather than a remark.

---

6. Leakage Projector Factorisation

Lemma 4 (Leading-order factorisation). Under Premises WL-1 through WL-6, the leakage support projector factorises at leading weak-commitment order:

$\Pi_{leak} = \Pi_R \otimes \Pi_Q$.

Proof sketch under stated hypotheses. The residual support R records **where** leakage can live inside the weak-commitment carrier (Lemma 2); the orientation fibre Q records **how** a weak-neutral attachment can be oriented (Lemma 3). These are answers to independent questions provided the modes that would correlate placement with orientation — the role-changing modes

— are gapped or subleading, which is Premise WL-6. Under WL-6 the leakage support space is the tensor product $R \otimes Q$, and the orthogonal projector onto a tensor-product support is the tensor product of the factor projectors. ■

The support trace then factorises:

$$\mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = \mathrm{Tr\_support}(\Pi\_R) \cdot \mathrm{Tr\_support}(\Pi\_Q).$$

Theorem 1 (Projected leakage support trace). At leading weak-commitment order,

$$\mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = 20.$$

Proof. By Lemma 2, $\dim R = 5$; by Lemma 3, $\dim Q = 4$; by Lemma 4, $\Pi\_\mathrm{leak} = \Pi\_R \otimes \Pi\_Q$. By Lemma 1 applied to each factor,

$$\mathrm{Tr\_support}(\Pi\_R) = 5, \mathrm{Tr\_support}(\Pi\_Q) = 4,$$

so

$$\mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = 5 \cdot 4 = 20. \blacksquare$$

The full bookkeeping identity, with every factor sourced:

$$20 = (K - 2) \times 2 \times 2 = (7 - 2) \times (\text{quadratures}) \times (\text{closure orientations}).$$

---

7. The Root-Three Leakage Numerator

The leakage source is threefold, but it is not counted linearly.

Let the three completion-branch leakage vectors be ℓ_1, ℓ_2, ℓ_3, Killing-orthonormal:

$$\langle \ell\_i, \ell\_j \rangle\_K = \delta\_ij.$$

The coherent leakage source is

$$L\_C_3 = \ell_1 + \ell_2 + \ell_3.$$

Lemma 5 (Root-normalised threefold source). If the leakage source is the coherent sum of three Killing-orthonormal completion branches, the leakage numerator is $\sqrt{3}$.

Proof. By Killing orthonormality,

$$\|\ell_1 + \ell_2 + \ell_3\|\_K^2 = \Sigma\_i \|\ell\_i\|\_K^2 = 3,$$

so $\|L\_C_3\|\_K = \sqrt{3}$. ■

Reading. This is why the numerator is not 3. A linear numerator would assert that the three branches occupy the same support slot counted three times. They do not: they are independent norm components, and independent norm components add in quadrature. The same Killing form that unit-normalises the denominator slots (Definition 1) root-normalises the numerator source — one metric, two roles.

Embedded source. The ℓ_i live on the completion residue $C \subset R$ and carry no orientation component. The source as it enters the leakage support is the embedded vector

$$\hat{L}\_C_3 = L\_C_3 \otimes q_0,$$

for a fixed Killing-unit orientation $q_0 \in Q$. Since $\|q_0\|\_K = 1$,

$$\|\hat{L}\_C_3\|\_K = \|L\_C_3\|\_K = \sqrt{3},$$

and $\hat{L}\_C_3 \in \text{range}(\Pi\_\text{leak}) = R \otimes Q$. The embedding changes no norm; it makes the membership claim of §9 well-typed.

---

8. The SM-3 Leakage Trace Theorem

Theorem 2 (Weak-Commitment Leakage Trace). At leading weak-commitment order, the projected leakage amplitude is

$$\beta = \sqrt{3}/20.$$

Proof. By Premise WL-7, the leakage amplitude is the per-slot commitment weight

$$\beta = \|\hat{L}\_C_3\|\_K \,/\, \text{Tr\_support}(\Pi\_\text{leak}),$$

with $\hat{L}\_C_3 = L\_C_3 \otimes q_0$ the embedded source of §7. By Lemma 5, $\|\hat{L}\_C_3\|\_K = \sqrt{3}$. By Theorem 1, $\text{Tr\_support}(\Pi\_\text{leak}) = 20$. Therefore

$$\beta = \sqrt{3}/20. \ ■$$

Corollary (Numerical value).

$$\beta = \sqrt{3}/20 \approx 0.0866025.$$

This is the weak-commitment leakage amplitude used in the PMNS benchmark kernel.

8.1 Why the amplitude functional is linear in the support trace

The support trace derived above is a rank count. It says how many Killing-unit slots are available to the leakage channel. It does not by itself specify whether the physical readout should be a total source norm, a root-mean-square fluctuation over the support, or a mean per-slot commitment weight. That choice is the content of Premise WL-7, and the reason for the mean per-slot functional is as follows.

The leakage amplitude β is not defined as the total size of the source; the total source size is already $\|\hat{L}\_C_3\|\_K = \sqrt{3}$. Nor is β a random fluctuation over the leakage support: there is no stochastic ensemble of occupied slots in the weak-commitment block, and the support trace is not a variance count — it is a count of admissible deterministic support directions.

β is instead the uniform first-order commitment weight induced in each admissible slot when the coherent source is projected into the leakage support. The correct leading-order functional is therefore the arithmetic mean support weight:

β = coherent source strength / number of admissible Killing-unit support slots = $\|\hat{L}\_C_3\|\_K$ / Tr_support(Π_leak).

This is why the numerator and denominator obey different rules. Independent source components combine by the Killing norm, so the threefold source gives $\sqrt{3}$ (Lemma 5). Once that source strength has been formed, its first-order commitment weight is allocated over the available unit slots. Allocation is linear because it is a distribution of a conserved projected weight, not a second norm operation.

The distinction excludes two tempting alternatives. A root-mean-square support readout would give

$\beta\_RMS = \sqrt{(3/20)} \approx 0.387$,

a fluctuation scale across a twenty-dimensional support, not the per-slot leakage amplitude of the weak-commitment kernel. A linear branch sum would give

$\beta\_lin = 3/20 = 0.15$,

treating the three completion branches as co-linear contributions to one slot, contradicting the Killing-orthonormal source assumption of Lemma 5.

The SM-3 value is therefore not selected by fitting the reactor angle. It is selected by the definition of β as the first-order mean commitment weight per Killing-unit leakage slot:

$\beta = \sqrt{3}/20$.

If future closure dynamics show the weak-commitment readout to be RMS-normalised, non-uniformly allocated, or branch-linear, then WL-7 fails and the SM-3 value must be rejected. That is the role of falsifier F13. §14 develops the same conclusion as a type argument and tabulates the candidate functionals.

---

9. Non-Circularity of the Derivation

A natural objection: the numerator and denominator look like two independent guesses — √3 above, 20 below.

They are not independent. **Both are read off the same leakage projector Π_leak.** The numerator √3 is the Killing norm of the embedded threefold completion source *inside* the leakage support; the denominator 20 is the total support trace over which that source is allocated. In symbols, with the embedded source of §7,

$\hat{L}\_C_3 = L\_C_3 \otimes q_0 \in \mathrm{range}(\Pi\_\mathrm{leak})$ and $\mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = 20$,

so that

$\beta = \|\Pi\_\mathrm{leak}\, \hat{L}\_C_3\|\_K \,/\, \mathrm{Tr\_support}(\Pi\_\mathrm{leak}) = \|\hat{L}\_C_3\|\_K \,/\, \mathrm{Tr\_support}(\Pi\_\mathrm{leak})$.

One projector, one Killing form, one allocation rule (WL-7), one formula. The numerator cannot be retuned without moving the source out of the support that defines the denominator, and the denominator cannot be retuned without changing the premises WL-1–WL-6, each of which is separately falsifiable (§13). This is a single support-normalised projected leakage quantity, not a ratio of unrelated counts.

---

10. The Exclusion Ladder: Why Not 5, 10, 40, or √(3/20)

Every wrong denominator corresponds to a specific counting error, and each error is blocked by a specific lemma or premise.

10.1 Why not 5?

5 counts only the residual support R. But leakage is not a scalar residual occupancy: it is a weak-neutral off-diagonal attachment (WL-4) and carries the orientation fibre Q. Omitting Q drops Lemma 3 entirely. The denominator is 5 × dim Q.

10.2 Why not 10?

10 = 5 × 2 counts the quadrature pair Q_N but misses the closure-orientation binary Q_C of the weak-neutral attachment half-cycle (WL-5). This is the denominator obtained if Premise WL-5 fails — and F6 in §13 records exactly that failure mode.

10.3 Why not 40?

40 = 5 × 4 × 2 counts Hermitian return as an independent slot. Lemma 3a forbids this: the return entry is functionally determined by the forward quadratures and spans no new real direction. Hermiticity supplies consistency, not support.

10.4 Why not √(3/20)?

√(3/20) ≈ 0.387 would result from a quadrature per-slot readout — spreading the source's squared norm over the twenty slots and taking the root. That misreads allocation as combination. Quadrature is the rule for combining independent components into a strength (Lemma 5); the slot count is a budget denominator, not an orthogonal decomposition of the source (Premise WL-7). The two rules answer different questions, and only the numerator's question is a quadrature question. See §8.1; the kill condition is F13.

The ladder is therefore rigid in all three directions: too few slots lands on 10, too many lands on 40, and the wrong combination rule lands on √(3/20). Only the count that respects Lemma 3a, Premise WL-5, and Premise WL-7 lands on

$\beta = \sqrt{3}/20$ with denominator 5 × 4 = 20.

---

11. Propagation into the Weak-Commitment PMNS Kernel

By the single-slot readout clause of WL-7, the kernel's leakage entry $A(1 + \beta e^{\pm i\psi})$ carries exactly one slot's allocation, which is why β enters the matrix element directly rather than as a multi-slot aggregate.

With the inherited scale convention

$A = 1$, $r_e = 0$, $B = 6/5$, $r_s = \sqrt{3/2} - 6/5$, $\delta = 0$, $\psi = 3\pi/4$, $\beta = \sqrt{3}/20$,

the electron–tau branch kernel is

$$M_\nu^+ = \begin{pmatrix} 0 & 1 & 1 + \beta e^{-i\psi} \\ 1 & r_s & 6/5 \\ 1 + \beta e^{i\psi} & 6/5 & r_s \end{pmatrix}$$

Exact Hermitian diagonalisation gives:

Quantity	Electron–tau branch
θ_{12}	≈ 33.40°
θ_{13}	≈ 8.59°
θ_{23}	≈ 48.37°
θ_{23} − 45°	≈ +3.37°
$\lvert J_{lep}\rvert$	≈ 2.448 × 10^{-2}

The electron–muon branch is the $\mu \leftrightarrow \tau$ reflection:

$\theta_{23} \approx 41.63°$ ($\theta_{23} - 45° \approx -3.37°$),

with the same θ_{12}, θ_{13}, and |J_lep|. Both branches have been verified by direct numerical Hermitian diagonalisation of the kernel above. The benchmark sets $\delta = 0$ exactly; restoring the admissible $O(\beta^2)$ diagonal splitting, $\delta = \beta^2 = 0.0075$, shifts θ_{23} by $\approx 0.51°$ and θ_{13} by $\approx 0.37°$ (signs set by the sign of δ) — below the leading-order claim level of this paper.

Thus β fixes the **size** of the reactor leakage and the **size** of the atmospheric deviation from maximality. It does not fix the octant branch. That is σ_W.

---

12. Separation from the Octant Sign

The atmospheric octant is selected by

$\sigma\_W = \mathrm{sgn}(\|P\_e\ H\_W\ P\_\tau\|^2 - \|P\_e\ H\_W\ P\_\mu\|^2)$,

with the reading:

- σ_W > 0: electron–tau branch;
- σ_W < 0: electron–muon branch;
- σ_W = 0: no leading octant selection.

This paper does not compute σ_W, and the omission is deliberate. SM-3 closes the leakage **magnitude**; SM-6 closes the octant **sign**. The magnitude and the branch are different projected outputs, and conflating them would weaken both gates: a magnitude claim should not inherit the risk of a sign claim, and vice versa.

---

13. Falsification Conditions

Each falsifier below is paired with the premise or lemma it attacks. SM-3 fails locally if any of the following occur.

#	Failure	Attacks	Consequence
F1	$K \neq 7$ in the weak-commitment closure carrier	WL-1	residual denominator changes
F2	fewer or more than two anchor directions removed	WL-2	$\dim R \neq 5$

#	Failure	Attacks	Consequence
F3	residual support does not split as C ⊕ G with dimensions 3 + 2	WL-3	inherited support grammar fails; numerator source misassigned
F4	weak-neutral leakage is not off-diagonal	WL-4	orientation-fibre argument fails
F5	first-order diagonal μ–τ splitting appears at the same order	WL-4	leakage no longer controls the reactor/atmospheric link cleanly
F6	closure-orientation binary is absent	WL-5	dim Q = 2; denominator becomes 10
F7	second binary reduces to Hermitian return under scrutiny	WL-5 / Lemma 3a	the factor 4 double-counts and must be rejected
F8	Hermitian return counted as an independent slot	Lemma 3a	denominator incorrectly becomes 40
F9	unit slots not equally Killing-normalised	Definition 1 / Lemma 1	support trace is not simple rank
F10	threefold numerator is linear rather than root-normalised	Lemma 5	numerator becomes 3, not √3
F11	Π_leak does not factor as Π_R ⊗ Π_Q at leading order	WL-6 / Lemma 4	denominator 20 not derived
F12	β ≠ √3/20 after projected support normalisation	Theorem 2	reactor leakage remains benchmark, not prediction
F13	per-slot readout is quadrature-normalised, or allocation over slots is non-uniform	WL-7	β becomes √(3/20) ≈ 0.387 (or slot-dependent); the SM-3 value fails

F6, F7, and F8 jointly pin the orientation fibre from both sides: the count must be strictly greater than 2 (else F6) and strictly less than 8 (else F8), and the surviving factor of 4 must be closure-sector-resident (else F7). F6 covers two distinct routes to 10: the closure-orientation sector absent altogether, and the sector present but entering as a discrete occupied label — one admissible direction per configuration — rather than as a two-dimensional admissible span (§5.4). F13 guards the ratio itself: even with the counts 5, 4, and √3 all intact, the SM-3 value survives only if the allocation rule is linear and uniform.

---

14. Why WL-7 Is the Natural Amplitude Functional

Premise WL-7 states the amplitude functional; this section argues that the linear per-slot rule is not a convenient stipulation but the unique candidate whose type matches what β is.

14.1 What kind of quantity β is

The two sides of the ratio are objects of different types, and the functional must respect both:

- the numerator $\|\hat{L}\_C_3\|\_K = \sqrt{3}$ is a **source norm** — the Killing strength of one coherent vector;
- the denominator Tr_support(Π_leak) = 20 is a **support multiplicity** — a count of admissible Killing-unit slots, not a metric volume in the source's own space.

A norm divided by a multiplicity is a **mean per-slot weight**. By the single-slot readout clause of WL-7, the kernel's leakage entry asks precisely the mean question: *what allocation does one slot carry?* The answer to a per-slot question about a conserved total shared over N slots is total/N. β is therefore a mean projected leakage amplitude per Killing-unit support slot — not an RMS fluctuation over the support.

14.2 Why the RMS reading fails

$\sqrt{(3/20)}$ is the answer to a different question: *if the source's squared norm were scattered incoherently over twenty slots, what would a typical slot's magnitude be?* That reading fails twice.

First, **there is no ensemble.** The source is one coherent vector, $\hat{L}\_C_3$, with exactly three orthonormal components (the ℓ_i). The remaining seventeen slots are admissible, not occupied by independent components. An RMS presupposes a fluctuation distribution over the support; the projected block supplies a single allocation instead.

Second, **RMS conflates support multiplicity with source dimension.** Spreading the squared norm over twenty slots treats the twenty slots as if they were twenty orthogonal components of the source itself. They are not: the source has dimension three; the support has multiplicity twenty. Quadrature is the correct rule inside the source (Lemma 5, three components → $\sqrt{3}$); it is a type error when applied to the slot count, which enters the problem as a budget denominator, not as a decomposition of anything.

14.3 Candidate functionals

Candidate amplitude	Formula	Meaning	Status
Linear per-slot mean	$\sqrt{3}/20 \approx 0.0866$	conserved source budget shared uniformly over the admissible slots	chosen by WL-7
RMS support readout	$\sqrt{(3/20)} \approx 0.3873$	typical fluctuation size of an incoherent distribution over the support	rejected: no ensemble exists; conflates support multiplicity with source dimension (F13)
Linear branch sum	$3/20 = 0.15$	treats the three completion branches as one slot counted three times	rejected: the branches are Killing-orthonormal; norms add in quadrature (Lemma 5, F10)

The three candidates are separated by type, not by taste: the branch sum errs inside the numerator, the RMS errs inside the denominator, and only the per-slot mean respects both the coherence of the source and the multiplicity character of the support.

14.4 Grade note

This section establishes naturalness: given that β is defined by a single-slot readout of one coherent source over a counted support, the linear per-slot mean is the only candidate of the correct type. It does not derive WL-7 from the projected block, so the conditional grade recorded in §15 is unchanged. Full promotion requires the upstream commitment-weight conservation argument, which remains a cross-paper obligation.

---

15. SM-3 Closure Certificate

SM-3 condition	Result
Leakage projector identified	closed: Π_leak
Unit-slot support trace defined	closed (Definitions 1–2, Lemma 1)
K = 7 carrier used	closed / inherited (WL-1)
Two anchors removed	closed (WL-2, Lemma 2)
Residual support dimension	closed: dim R = 5
Weak-neutral quadrature factor	closed: 2
Closure-orientation binary	conditional on WL-5: 2 (involution construction owed to closure-operator specification; §5.3)
Hermitian return excluded from counting	closed by lemma (Lemma 3a)
Orientation fibre dimension	closed conditional on WL-5: dim Q = 4
Projector factorisation	closed at leading weak-commitment order (WL-6, Lemma 4)
Support trace	closed: Tr_support(Π_leak) = 20
Allocation rule	stated as Premise WL-7; falsifiable (F13)
Root-three numerator	closed (Lemma 5)
Leakage amplitude	closed at leading support-trace order, conditional on WL-5 and WL-7: $\beta = \sqrt{3}/20$
Benchmark kernel outputs	verified numerically (§11, Appendix A)
Octant sign	outside SM-3; deferred to SM-6

SM-3 closure grade: closed at leading weak-commitment support-trace order, conditional on WL-5 (closure-orientation binary) and WL-7 (per-slot allocation rule).

16. Conclusion

The leakage amplitude $\beta = \sqrt{3}/20$ was among the most exposed numbers in the VERSF flavour programme. It controlled the reactor angle and the size of the atmospheric departure from maximality, but its denominator 20 remained a support-trace target rather than a closed derivation.

This paper closes that target. The leakage support projector factorises,

$\Pi\_leak = \Pi\_R \otimes \Pi\_Q$,

with residual support

$\dim R = K - 2 = 7 - 2 = 5$

and orientation fibre

$\dim Q = 2 \times 2 = 4$.

The first orientation factor is the weak-neutral quadrature pair. The second is a closure-orientation binary of the weak-neutral attachment half-cycle — not Hermitian return, whose count is excluded by an explicit lemma. Therefore

$Tr\_support(\Pi\_leak) = 5 \times 4 = 20$,

and with the root-normalised threefold completion source

$\|L\_C_3\|\_K = \sqrt{3}$,

the leakage amplitude is returned rather than inserted — at leading support-trace order, conditional on WL-5 and WL-7, with both conditions named and falsifiable:

$\beta = \sqrt{3}/20$.

The claim is stated at its exact strength: the denominator is derived as a support trace; the amplitude follows under two explicit structural commitments, neither hidden inside the numerical agreement with the PMNS benchmark.

In one sentence:

The weak-commitment leakage amplitude is the root-three completion source diluted over the twenty Killing-unit slots of the projected leakage support, so $\beta = \sqrt{3}/20$.

Appendix A — Minimal Numerical Kernel

With

$\beta = \sqrt{3}/20$, $\psi = 3\pi/4$, $A = 1$, $B = 6/5$, $r_e = 0$, $r_s = \sqrt{(3/2)} - 6/5 \approx 0.0247449$, $\delta = 0$,

the benchmark electron–tau kernel is

$$M_\nu^+ = \begin{pmatrix} 0 & 1 & 1 + \beta e^{-i\psi} \\ 1 & r_s & 6/5 \\ 1 + \beta e^{i\psi} & 6/5 & r_s \end{pmatrix}$$

Exact Hermitian diagonalisation gives:

$\theta_{12} \approx 33.40°$, $\theta_{13} \approx 8.59°$, $\theta_{23} \approx 48.37°$, $|J_{lep}| \approx 2.448 \times 10^{-2}$.

The reflected electron–muon branch gives

$\theta_{23} \approx 41.63°$

with the same θ_{12}, θ_{13}, and $|J_{lep}|$. The eigenvalues of the kernel are approximately {−1.18454, −0.88077, +2.11480}; the mixing angles above are extracted from the unitary diagonalising matrix in the standard PMNS convention with the column assignment that places the small electron-row component in the third column ($|U_{e3}| = \sin\theta_{13}$).

Kernel eigenvalues are closure curvatures, not neutrino masses; the map from curvatures to physical masses is firewalled to SM-24. The column assignment is therefore a readout convention pending that gate. In ascending-eigenvalue order the electron-row magnitudes are (0.149, 0.826, 0.544): the small-electron-component state is the lowest curvature, −1.18454, and assigning it to the third readout column is a convention choice, not a mass-ordering claim. Were |curvature| mapped monotonically to mass, the ordering (0.881, 1.185, 2.115) would place the small-electron state second and the extracted angle would not be the physical reactor angle — which is precisely why the curvature-to-mass map is deferred to SM-24 rather than asserted here.

Appendix B — Referee Objections and Responses

Objection 1 — "You have only renamed the denominator 20."

Response: No. The denominator is the trace of a specified projector, $\Pi_{leak} = \Pi_R \otimes \Pi_Q$, with dim R = 5 derived from a stated carrier and anchor count (WL-1, WL-2), dim Q = 4 derived from a stated orientation structure (WL-4, WL-5), and the trace taken over Killing-unit slots (Definition 1), not arbitrary labels. Each ingredient carries its own falsifier in §13. A renaming has no falsifiers; this derivation has thirteen, and its counting ladder is guarded in three independent directions (§10).

Objection 2 — "The factor 4 double-counts Hermitian return."

Response: That was a valid objection to the older formulation, and this version does not merely deny it — it excludes it by lemma. Lemma 3a shows the return entry is functionally determined by the forward quadratures and spans zero new real directions. The second binary is instead the closure-orientation binary of WL-5, whose independence from the quadratures and from conjugation is established by the three properties in §5.3 and the independence condition of §5.4. If the binary nonetheless reduces to return under deeper scrutiny, F7 fires and the count must be rejected — the paper states its own kill condition.

Objection 3 — "Why is the numerator √3 rather than 3?"

Response: Because the three leakage branches are Killing-orthonormal support components, and norms add in quadrature: $\|\ell_1 + \ell_2 + \ell_3\|_K = \sqrt{3}$ (Lemma 5). A numerator of 3 would assert that the three branches occupy one slot counted thrice. The same Killing form governs numerator and denominator, so the normalisation cannot be chosen independently on the two sides of the ratio.

Objection 4 — "The factorisation $\Pi_{leak} = \Pi_R \otimes \Pi_Q$ is asserted, not derived."

Response: The factorisation holds under Premise WL-6 (role-changing modes gapped or subleading), which is stated as a hypothesis, used in Lemma 4, and falsified by F11. The claim level in §1.1 is calibrated accordingly: leading-order projected support-trace theorem. If placement–orientation mixing enters at leading order, the denominator 20 is not derived and SM-3 reopens.

Objection 5 — "Does this derive the atmospheric octant?"

Response: No. It derives the magnitude of the leakage. The octant sign is $\sigma_W = \mathrm{sgn}(\|P_e H_W P_\tau\|^2 - \|P_e H_W P_\mu\|^2)$, and that is SM-6. Keeping the gates separate is a strength: the magnitude claim does not inherit the risk of the sign claim.

Objection 6 — "Does this close the full PMNS derivation?"

Response: No. It closes SM-3: the leakage support trace and unit-slot normalisation. The full PMNS closure-Hamiltonian computation remains SM-24.

Objection 7 — "The numerator adds in quadrature but the denominator divides linearly. Why do the two sides of the ratio obey different rules?"

Response: Because they answer different questions. Quadrature is the composition rule for *combining* independent sources into a single strength (Lemma 5): orthonormal components add by norm. Linear division is the *allocation* rule for distributing a conserved commitment weight over admissible slots (Premise WL-7): a budget divides. The paper does not derive WL-7 from the projected block; it states it as a named premise, grounds it in commitment-weight conservation, and attaches the falsifier F13. Under a quadrature per-slot readout the amplitude would be $\sqrt{3/20} \approx 0.387$ and SM-3 would fail. That alternative is excluded by premise, not by comparison with the measured reactor angle — a fit-based exclusion would forfeit precisely the predictive standing this paper claims.

Objection 8 — "The paper still assumes the answer through WL-7."

Response: WL-7 is an explicit amplitude-functional premise, not a hidden fit. The denominator 20 is derived independently as the support trace of the projected leakage projector. The numerator $\sqrt{3}$ is derived independently as the Killing norm of the threefold source. WL-7 specifies how a coherent source norm becomes a first-order leakage amplitude in the weak-commitment kernel: it is the mean commitment weight per Killing-unit support slot. Alternative functionals are not ignored. A root-mean-square support readout would give $\sqrt{(3/20)}$, and a linear branch sum would give 3/20. Both are named failure modes under F13. Thus WL-7 does not smuggle in the measured value; it identifies the remaining structural condition under which the support-trace derivation becomes a prediction.