A Pre-Registered Neutrino-Sector Prediction from the Finite-QEC Substrate: Normal Ordering at the Oscillation Floor, No Dirac Phase, and an Invisible $m_{\beta\beta}$

David Elliman
Neuro-symbolic Ltd
dave@neusym.ai
https://neusym.ai

Prediction note: 6 July 2026

Abstract

This note pre-registers the finite-QEC framework's neutrino sector in the same style as the programme's five earlier registered predictions. The claims, all fixed by substrate constants plus one declared experimental scale ($\Delta m^2_{31} = 2.515 \times 10^{-3}\,\mathrm{eV}^2$), are: (i) **no intrinsic leptonic Dirac CP violation** — the PMNS lift is the real single-polar frame transport, so $J_\ell = 0$ exactly and $\delta_{\mathrm{CP}} \in \{0, \pi\}$; a sector-selection theorem in the same engine that produces the quark $J \neq 0$ shows *why*: CP transport rides colour intermediates, and the lepton register is colourless; (ii) **normal ordering at the oscillation floor** — the Koide neutrino circulant gives $m = (0.79,\ 8.72,\ 50.16)$ meV, $\sum m_\nu = 59.7$ meV, just 0.7 meV above the absolute normal-ordering minimum; (iii) **CP lives in the Majorana sector** — the orientation pointer $i\Omega_{R1}$ is Hermitian and diagonal in the mass basis, feeding leptogenesis and $0\nu\beta\beta$, never long-baseline δ_{CP}; and (iv) **an invisible** $m_{\beta\beta}$ — the parity branches give $m_{\beta\beta} = 1.76$ or 3.07 meV (envelope < 4.2 meV), below the LEGEND-1000/nEXO discovery band, so the framework predicts a *null* at next-generation $0\nu\beta\beta$; $m_\beta = 9.06$ meV. At registration the CP-conserving point sits 0.16σ from the NuFit 6.0 normal-ordering best fit ($\delta_{\mathrm{CP}} = 177^\circ \pm 20$), while the leading-order mixing angles carry recorded 2.2–3.7σ strains that are declared openly and are not the registered surfaces. Five original kill surfaces were frozen. A later outcome addendum supersedes the joint dark-energy clause: the registered law $w(a) = -1 + a/28$ has since been killed internally, so K5 is now a historical preregistration clause rather than an active rescue of the mass-sum prediction. The lab-facing predictions — normal ordering, Dirac-CP null, m_β, and $0\nu\beta\beta$ null — are unchanged. Grade: conditional, not Locked — three named assumptions are carried explicitly.

Amendment, version 2 (2026-07-07) — pre-data sharpening under the registration's own rules; no bar moved, no retune. Since the original registration (v1, 2026-07-06), the one assumption this note flagged for the $0\nu\beta\beta$ prediction — the $\sigma \to$parity *convention* that split $m_{\beta\beta}$ into two branches — has been *derived*, so the two-branch envelope collapses to a single number. The derivation is entirely upstream and parameter-free: the mass-state labels are fixed by the charged-lepton sector (the Koide singlet is the *heaviest* state, since the same circulant reproduces $m_\mu/m_e = 206.8$ and $m_\tau/m_\mu = 16.8$), so the parity reference is ν_3; the global orientation is $s = +1$, pinned by the observed quark $J > 0$; and the resulting parity pattern $(+, -, +)$ gives

$$m_{\beta\beta} = 0.65 \text{ meV},$$

which is *exactly the floor* of the true continuous-phase envelope [0.650, 4.172] meV (the dominant residue $a_2 = 2.411$ meV exceeds $a_1 + a_3 = 1.761$ meV, so the antialigned pattern *is* the minimum; the v1 figure $[0.7, 4.2]$ was rounding-coarse). **Nothing else changes:** normal ordering, $\sum m_\nu = 59.7$ meV, the Dirac-CP null, and all five kill surfaces $K1$–$K5$ (plus $W1$) stand exactly as registered in v1; the next-generation-$0\nu\beta\beta$ *null* ($K3$) is only *strengthened*, the predicted $m_{\beta\beta}$ now sitting a factor $\sim$15 below the LEGEND-1000/nEXO floor rather than $\sim$3–6. The two-branch text of the body below is the *original v1 registration*, preserved unaltered as the honest pre-derivation form. Receipts (all committed before this version): the label + parity + envelope derivation and its machine-precision Koide seal, in the programme archive under `item87_*` gates (2026-07-07 mass arc).

Amendment, version 3 (2026-07-08) — dark-energy linkage superseded. The v1 kill surface K5 deliberately tied the neutrino mass-sum reading to the programme's then-registered dark-energy law $w(a) = -1 + a/28$. That dark-energy branch has now been killed internally by the selector-H_0/acoustic-clock consistency test and retained only as a frozen reopen record. Therefore K5 below is preserved as the historical preregistration clause, not as a current active kill surface. The neutrino-sector claims that do not rely on that background — normal ordering, Dirac CP conservation, $m_\beta = 9.06$ meV, and the $0\nu\beta\beta$ null with $m_{\beta\beta} = 0.65$ meV after the v2 parity closure — are unchanged. Future cosmological stress on $\sum m_\nu = 59.7$ meV must be evaluated in the live CMB/halo background or a valid reopened dark-energy branch, not in the killed CPL line.

1 Purpose and provenance

This is the sixth registered prediction of the finite-QEC substrate programme, following the now-frozen dark-energy registration $w_0 = -27/28$, $w_a = -1/28$ [1], the sterile-neutrino X-ray target, the primordial tensor null, the K04 fossil non-mobility null, and the G-lock value. The methodology is unchanged: the claim, the numbers, and the kill conditions are frozen and timestamped *before* the deciding data exist, and no parameter may be adjusted afterwards. The instrument-level theory behind the record/response split used below is developed in [2, 3].

The chain being registered was closed inside the programme's canon in four dated steps during 2026-06/07: the portal-phase branch was pinned ($\Phi = 1/3$ rad; the competing $2\pi/3$ branch is a CP zero and the $2\pi/9$ branch fails the observed Δm^2 ratio by a factor 6.6); the "orientation-odd Dirac sign" hypothesis was closed *negatively* (there is no Dirac sign to correlate); the consolidated prediction ledger was built as a single executable gate; and, on the registration date, a sector-selection theorem supplied the mechanism for the Dirac null (Section 2). Every number in this note is recomputed by the registration gate (`neutrino_sector_prereg_gate.py`, exit 0 on 2026-07-06) from substrate inputs; nothing is quoted from memory.

Grade — conditional, not Locked. Three load-bearing assumptions are named and carried openly: (a) the real single-polar QEC frame transport is the *leading* PMNS lift; (b) the discrete-sign (CP-conserving) Majorana reading of the R15 orientation pointer; (c) the $\sigma \to$parity convention mapping the global orientation to the $0\nu\beta\beta$ parity branch. A fourth, historical honesty item: an earlier PMNS mechanism in this programme (the "38°/6° plus sterile-shadow" construction) was explicitly falsified inside the programme on 2026-06-05 and withdrawn; the present lift is a later, different derivation sharing no step with it.

2 The derivation chain

2.1 The Dirac null and its mechanism

The PMNS matrix is the real orthogonal single-polar frame transport

$$U \;=\; R_{23}(\pi/4)\; R_{12}(\pi/4)\; \exp\Big[\tfrac{2}{9}\, A(-\tfrac{1}{2}, \tfrac{1}{2}, 1)\Big], \tag{1}$$

where $A(\mathbf{n})$ is the antisymmetric generator about axis $\mathbf{n}$, the two $\pi/4$ factors are the bipartite baseline, the coefficient 2/9 is the programme's universal topological index, and the axis is the forced mean-cycle tangent. There are no adjustable parameters. A real U has Jarlskog invariant $J_\ell = 0$ *exactly*, hence $\delta_{\mathrm{CP}} \in \{0, \pi\}$: no intrinsic CP asymmetry in vacuum oscillations, for any baseline and any Majorana phases (verified to 10^{-16} over a grid of propagation phases in the gate).

The registration-date theorem says why this is not an accident of Eq. (1). In the same 256-dimensional walk engine that produces the quark sector's *nonzero* Jarlskog invariant ($J_q = \pm 4.3 \times 10^{-5}$, sign set by the global orientation), the rephasing-invariant CP phase is carried entirely by *colour* intermediates: running the identical walk on the colourless register (the lepton analogue) gives $J \equiv 0$ identically — not small, *sector-forbidden.* Leptonic Dirac CP conservation is a selection rule of the engine, not a tuning.

2.2 Where CP actually lives

The orientation-bearing object is the R1 cochain Ω_{R1}; $i\Omega_{R1}$ is Hermitian with eigenvalues $\{0, \pm\sqrt{3}\}$ and is *diagonal in the circulant/mass basis.* Its phases therefore attach to mass eigenstates: they are Majorana phases. A Majorana rephasing cancels from all oscillation probabilities but survives in lepton-number-violating responses — it feeds the leptogenesis sign and the $0\nu\beta\beta$ residue. The framework's statement is structural: **leptonic CP violation exists, and it is Majorana/recovery orientation, never a PMNS Dirac phase**.

2.3 The mass spectrum

The neutrino Koide circulant [4] ($R_\nu = 1$, $\Phi = \delta_\nu = 1/3$ rad — the raw defect fraction 3/9) fixes the mass *ratios* $m_i \propto (1 + \cos(\Phi + 2\pi n_i/3))^2$. One experimental scale, $\Delta m^2_{31} = 2.515 \times 10^{-3}$ eV^2 (declared at registration; the NuFit 6.0 variants $2.534/2.513 \times 10^{-3}$ bracket it), then fixes everything:

$$m = (0.79,\ 8.72,\ 50.16)\ \text{meV}, \qquad \sum m_\nu = 59.7\ \text{meV}, \qquad \frac{\Delta m^2_{21}}{\Delta m^2_{31}} = 0.02997. \tag{2}$$

Normal ordering is forced, the lightest state is nearly massless, and the sum lands 0.7 meV above the absolute normal-ordering floor (58.99 meV at NuFit 6.0 central values): the framework does not merely prefer the floor, it *sits on it.* The mass-squared ratio lands within 0.6–1.4% of the two NuFit variants with zero tuned parameters.

Quantity	Framework	NuFit 6.0 IC19	NuFit 6.0 IC24	tension
θ_{12} [°]	32.163	$33.68^{+0.73}_{-0.70}$	$33.68^{+0.73}_{-0.70}$	2.2σ
$\sin^2\theta_{12}$ (JUNO)	0.2834	0.3092 ± 0.0087		3.0σ
θ_{23} [°]	46.006	$48.5^{+0.7}_{-0.9}$	$43.3^{+1.0}_{-0.8}$	-2.8σ / $+2.7\sigma$
θ_{13} [°]	8.929	8.52 ± 0.11	8.56 ± 0.11	3.7σ / 3.4σ
$\delta_{\rm CP}$ [°]	180 (or 0)	177^{+19}_{-20}	212^{+26}_{-41}	0.16σ / 0.78σ
$\Delta m^2_{21}/\Delta m^2_{31}$	0.02997	0.02956	0.02981	+1.4% / +0.6%
ordering	normal	NO preferred		consistent
$\sum m_\nu$ [meV]	59.7	see text		floor +0.7

Table 1: The registered chain against the 2026-07-06 landscape. The angle rows are zero-parameter leading-order postdictions whose 2.2–3.7σ strains are recorded openly; they are *not* the registered surfaces (see the weak-kill layering, W1). The $\delta_{\rm CP}$ row is the registered null: the CP-conserving point sits inside the current 1σ band, and NuFit 6.0's own conclusion is that normal ordering is consistent with CP conservation within 1σ.

2.4 The $0\nu\beta\beta$ and direct-mass residues

With real U_{ei} and the sign-pointer parities (the $(1,1,1)$ singlet ν_2 is the reference; ν_1, ν_3 take opposite parity, the choice set by the global orientation σ),

$$m_{\beta\beta} = \Big|\sum_i U_{ei}^2\, m_i\, \eta_i\Big| = \begin{cases} 1.76\ \text{meV} & \sigma = +\ (+,+,-) \\ 3.07\ \text{meV} & \sigma = -\ (-,+,+) \end{cases} \qquad \text{envelope } [0.7,\ 4.2]\ \text{meV}, \tag{3}$$

where the constructive extreme $(+,+,+)$ gives 4.17 meV. The kinematic mass is $m_\beta = 9.06$ meV. Both are far below current and next-generation direct reach: KATRIN's current limit is 450 meV [5], and the LEGEND-1000/nEXO discovery band is $m_{\beta\beta} \sim 10$–20 meV [6, 7]. The framework therefore predicts a **null result at next-generation** $0\nu\beta\beta$ — a positive detection there kills this branch outright (Section 4).

3 The numbers against the registration-date landscape

Table 1 places the chain against the global fit NuFit 6.0 [8] (both variants: IC19 without, IC24 with SK atmospheric data) and against JUNO's first measurement [9], which improved the solar-sector precision by a factor 1.6 after 59 days of data.

Three features of Table 1 deserve plain statement.

The null is currently favoured, not merely allowed. For normal ordering the global-fit best value is $\delta_{\rm CP} = 177°$ — 0.16σ from the CP-conserving 180° the framework demands. The often-quoted "3.6σ evidence for leptonic CP violation" belongs to the *inverted*-ordering branch, which the framework independently excludes: the joint claim (NO *and* CP conservation) is exactly the corner of parameter space the current fit prefers.

The angles carry real strain, declared here. The leading-order lift lands all three angles within 0.4–2.5° of the global fit with no free parameters, but at current precision that corresponds to 2.2σ (θ_{12}, rising to 3.0σ against JUNO's first data), 3.4–3.7σ (θ_{13}), and an octant-split $\pm 2.7\sigma$ (θ_{23} — the two NuFit variants disagree with *each other* by more than either disagrees with the framework's 46.0°, and the framework's second-octant call becomes a clean discriminator when the octant resolves). These strains are recorded non-silently at registration. No NLO correction is invoked to soften them here; if one is ever derived it must come from the substrate with zero new parameters and be published *before* the deciding data (rigidity rules, Section 5).

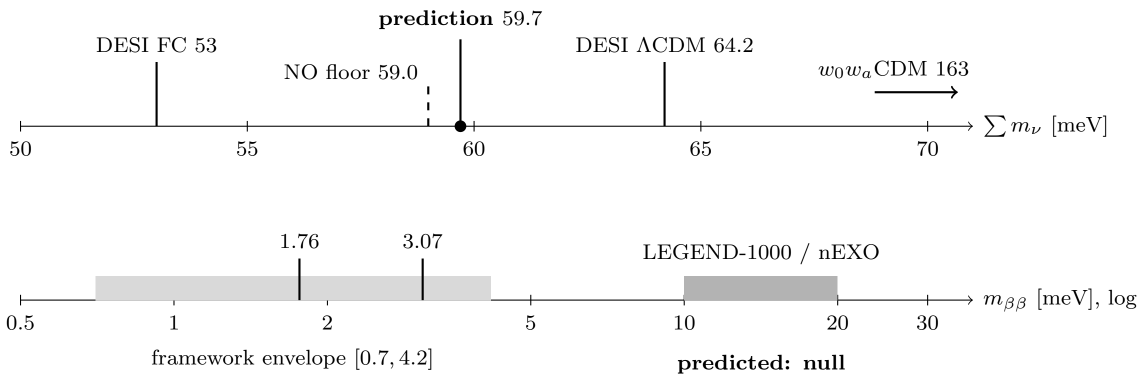

Figure 1: The two mass-sector surfaces at registration (all numbers as cited in the text; positions to scale). *Top:* $\sum m_\nu = 59.7$ meV sits 0.7 meV above the absolute normal-ordering floor — below the DESI DR2 ΛCDM Bayesian bound, above the Feldman–Cousins bound that already breaches the entire NO floor, and comfortably inside the w_0w_aCDM bound that applies in the framework's own registered background. *Bottom:* the $m_{\beta\beta}$ parity branches and envelope sit a factor ~ 3–10 below the next-generation discovery band: the registered prediction there is a null.

The mass sum sits exactly where cosmology is squeezing. Figure 1 (top) shows the position: the DESI DR2 ΛCDM bounds are 64.2 meV (Bayesian) and 53 meV (Feldman–Cousins) at 95% [10] — the latter already *below the entire normal-ordering floor*, a boundary-effect strain that paper itself notes is shared by all NO physics, not by this framework specifically. In the w_0w_aCDM background the same data give 163 meV. At v1 registration the framework had no freedom here: it had registered $w(a) = -1 + a/28$ as its dark-energy law on 2026-06-29 [1], so the mass sum was deliberately judged in that background — where 59.7 meV is comfortable. That was the joint kill surface K5. The later internal kill of the dark-energy branch supersedes this cosmology clause: K5 is retained as a record of the original preregistration, while the lab-facing neutrino predictions remain live.

4 Registered kill surfaces

Frozen on 2026-07-06; never adjusted after data. "Robust" means a result the relevant community itself treats as established (multi-experiment or $\geq 5\sigma$ single-experiment, systematics audited).

K1 (ordering). Inverted mass ordering established at $\geq 5\sigma$ (JUNO or the global fit) $\Rightarrow$ the Koide-circulant branch is falsified.

K2 (Dirac phase). With NO established, the *intrinsic* Dirac phase excluded from both 0 and π at $\geq 5\sigma$ (matter effects removed) $\Rightarrow$ the real frame-transport PMNS and the Dirac-CP null are falsified. New physics would then be a new complex PMNS response operator — a monitored R1/Majorana orientation environment — not another fitted phase.

K3 ($0\nu\beta\beta$). A robust detection with $m_{\beta\beta} \geq 10$ meV $\Rightarrow$ the Majorana parity branch is falsified (the framework band tops at 4.17 meV). No nuclear-matrix-element or alternative-mechanism rescue is permitted.

K4 (direct mass). A robust kinematic $m_\beta \geq 30$ meV $\Rightarrow$ falsified (predicted 9.06 meV).

K5 (joint cosmology; historical after v3). With the background fixed to the registered law $w(a) = -1 + a/28$ [1], a $\sum m_\nu$ posterior excluding 59.7 meV at $\geq 3\sigma$ would have falsified the two registered legs *jointly*. After the 2026-07-02 internal kill of the dark-energy branch, this is no longer an active kill surface. It is retained to show the original preregistration. Future cosmological tests of the neutrino mass sum must be framed in the live background model or in a valid reopened dark-energy branch.

W1 (weak-kill layering). Any single PMNS angle excluded at $\geq 5\sigma$ from $(32.163, 46.006, 8.929)^\circ$ falsifies the single-polar *lift* (the angle layer) — while the Dirac null (sector-selection theorem) and the mass/ordering branch (Koide circulant) remain separately falsifiable via K1–K5. The layers are independent enough that the programme commits to reporting them separately.

Positive surfaces (what survival looks like): NO at $\geq 3\sigma$; $\delta_{\rm CP} \to \pi$ as long-baseline precision improves; a $0\nu\beta\beta$ null through the 10–20 meV band; dynamical-background cosmology converging on $\sum m_\nu \approx 60$ meV; and θ_{23} resolving to the second octant near 46°.

5 Rigidity

No parameter of this sector may be adjusted after 2026-07-06. The inputs 2/9 (topological index), the transport axis, the $\pi/4$ baselines, $\Phi = 1/3$, and the parity convention are substrate-fixed; the single experimental scale Δm^2_{31} is declared above. Not allowed: introducing a Dirac phase post hoc; an inverted-ordering rescue branch; $m_{\beta\beta}$ rescues via matrix-element freedom; retuning angles, parities, or the envelope after data; quietly dropping the K5 joint clause if the dark-energy leg falls. Allowed: an NLO angle correction *only if* derived from the substrate with zero new parameters and published before the deciding data (the leading-order angles remain the registered reference); combining K1–K5 verdicts across experiments in the annual review.

6 Decision timeline

Surface	Experiment	Horizon	Registered outcome
K5 $\sum m_\nu$ (joint)	DESI DR3+/CMB	continuous, nearest	converge on ≈ 60 meV
K1 ordering	JUNO (3σ: $\sim$6 yr data)	$\sim$2032; fits earlier	normal
K2 $\delta_{\rm CP}$	DUNE / Hyper-K	2030s	π (or 0); no intrinsic CPV
K3 $m_{\beta\beta}$	LEGEND-1000 / nEXO	2030s	null at 10–20 meV
K4 m_β	beyond KATRIN	long-term	9.06 meV, unreachable soon
W1 angles	JUNO + LBL precision	continuous	strains resolve or sharpen

JUNO has already delivered its first factor-1.6 precision step after 59 days [9]; its ordering determination needs $\sim$6 years of nominal exposure. The cosmological surface moves fastest and is the reason this registration is timestamped now: the $\sum m_\nu$ squeeze and the dark-energy background question are being decided by the same surveys, and the framework's two registered legs meet exactly there.

7 What failure or survival would mean

If K1–K5 trigger, the loss is structural, not parametric: there is no version of this framework with an inverted hierarchy, a tunable Dirac phase, or a heavy $m_{\beta\beta}$. The sector either is what Eqs. (1)–(3) say or the leptonic closure of the programme is wrong. Conversely, the survival scenario is highly specific: normal ordering at the floor, CP conservation emerging from today's $177^\circ \pm 20$, a next-generation $0\nu\beta\beta$ null that most seesaw-leptogenesis expectations do not prefer, and cosmology settling on $(w_0, w_a, \sum m_\nu) = (-27/28, -1/28, \approx 60 \text{ meV})$ jointly. Nature has little room to satisfy this by accident: each surface is generic in the wider theory space, but their conjunction is not.

Reproducibility. All numbers in this note are recomputed by `neutrino_sector_prereg_gate.py` (exit 0, 2026-07-06), which mirrors the consolidated canon gate `leptonic_cp_majorana_only_prediction.py` and writes the machine-readable registration record; the sector-selection mechanism is certified by `leptonic_dirac_cp_sector_selection_gate.py`. The pre-registration JSON accompanies this note.

References

[1] David Elliman. A Pre-Registered Dark-Energy Prediction from a Finite Substrate Clock: $w_0 = -27/28$, $w_a = -1/28$, 2026. Zenodo, concept DOI 10.5281/zenodo.21049001, published 2026-06-29.

[2] David Elliman. Records and Responses: Dressing-blind theorems for monitored quantum instruments, 2026. Zenodo, DOI 10.5281/zenodo.21189012.

[3] David Elliman. From Counts to Observables, 2026. Zenodo, DOI 10.5281/zenodo.21210527.

[4] Yoshio Koide. A new view of quark and lepton mass hierarchy. *Phys. Rev. D*, 28:252, 1983.

[5] KATRIN Collaboration. Direct neutrino-mass measurement based on 259 days of KATRIN data. 2024. arXiv:2406.13516; $m_\beta < 0.45$ eV (90% CL).

[6] LEGEND Collaboration. The Large Enriched Germanium Experiment for Neutrinoless $\beta\beta$ Decay: LEGEND-1000 Preconceptual Design Report. 2021. arXiv:2107.11462; discovery sensitivity $T_{1/2} \sim 10^{28}$ yr, $m_{\beta\beta} \sim 10$–20 meV.

[7] nEXO Collaboration. nEXO: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity. *J. Phys. G*, 49:015104, 2022. arXiv:2106.16243.

[8] Ivan Esteban, M. C. Gonzalez-Garcia, Michele Maltoni, Ivan Martinez-Soler, João Paulo Pinheiro, and Thomas Schwetz. NuFit-6.0: updated global analysis of three-flavor neutrino oscillations. *JHEP*, 12:216, 2024. arXiv:2410.05380.

[9] JUNO Collaboration. First measurement of reactor neutrino oscillations at JUNO. 2025. arXiv:2511.14593; 59.1 days of data, $\sin^2\theta_{12} = 0.3092 \pm 0.0087$, $\Delta m^2_{21} = (7.50 \pm 0.12) \times 10^{-5}\,\text{eV}^2$.

[10] DESI Collaboration. Constraints on Neutrino Physics from DESI DR2 BAO and DR1 Full Shape. 2025. arXiv:2503.14744; $\sum m_\nu < 0.0642$ eV (ΛCDM, 95%), < 0.053 eV (Feldman–Cousins), < 0.163 eV ($w_0 w_a$CDM).