A Closed-Form Neutrino Sector from DFD v3.5: TBM Geometry, a Discrete S_2 Lock, and a Seesaw Scale Closure

Gary Alcock

January 1, 2026

Abstract

DFD v3.5 provides three ingredients that, when combined with a strict no-hidden-knobs rule, appear to close the neutrino sector to a surprising extent: (i) a tribimaximal (TBM) neutrino mixing base from the neutrinos-at-center overlap rule (Appendix K), (ii) a derived heavy Majorana scale $M_R = M_P\alpha^3$ (Appendix P), and (iii) a derived electroweak scale $v = M_P\alpha^8\sqrt{2\pi}$ (Section 13 / Appendix K).

This note pushes a fourth ingredient as hard as possible. TBM singles out a canonical residual transposition S_2 (the $\mu \leftrightarrow \tau$ swap), and the unique smallest positive S_2-equivariant deformation of the identity is $I + P_-$, where P_- projects onto the odd-parity axis. On the doublet this produces eigenvalues $(2, 1)$ exactly, hence a discrete lock $m_2/m_1 = 2$.

The closure step is that the same S_2 doublet structure also forces a normalization factor $1/\sqrt{2}$ in the center-coupling Dirac overlap, turning the Appendix-P ansatz $y_D \sim \sqrt{\alpha}$ into a no-knobs value $y_D = \sqrt{\alpha/2}$. With the seesaw, this removes the remaining continuous scale and yields explicit absolute masses, mass-squared splittings, and $0\nu\beta\beta$ and beta-decay effective masses in terms of α and M_P only.

1 What TBM gives you for free in DFD (and what it does not)

Appendix K of the unified manuscript states the TBM base (when neutrinos are "at center"):

$$U_{\mathrm{TBM}} = \begin{pmatrix} \sqrt{2/3} & \sqrt{1/3} & 0 \\ -\sqrt{1/6} & \sqrt{1/3} & \sqrt{1/2} \\ \sqrt{1/6} & -\sqrt{1/3} & \sqrt{1/2} \end{pmatrix}. \tag{1}$$

TBM fixes the eigenvectors (columns) and therefore fixes a discrete set of residual permutation symmetries of the neutrino mass matrix. However, TBM by itself does not fix the eigenvalues (m_1, m_2, m_3).

The push here is: can TBM's residual symmetry content, plus a strict no-hidden-knobs principle, force a specific doublet split such as $m_2/m_1 = 2$?

2 Why full S_3 invariance cannot split a doublet

Let generation space carry the permutation representation of S_3. The S_3-invariant endomorphisms are the centralizer, spanned by I_3 and $J = \mathbf{1}\mathbf{1}^T$. On the standard doublet subspace $\{x_1+x_2+x_3 = 0\}$ one has $J = 0$, hence every S_3-equivariant operator is proportional to the identity on the doublet. Therefore:

Any non-degenerate doublet spectrum requires breaking S_3 to a proper subgroup.

3 TBM naturally singles out a transposition S_2 (the $\mu \leftrightarrow \tau$ swap)

Consider the transposition that swaps the μ and τ components:

$$S_{\mu\tau} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

Its eigenvectors in the μ–τ plane are the even and odd parity axes

$$v_+ = \frac{1}{\sqrt{2}}(0,1,1), \qquad v_- = \frac{1}{\sqrt{2}}(0,1,-1),$$

with $S_{\mu\tau} v_\pm = \pm v_\pm$. Up to an unphysical rephasing of the τ row, the TBM basis contains exactly this even/odd structure: the third TBM column is v_+ as written above, and a row sign flip converts it to v_- without changing physical mixing probabilities. Thus TBM motivates a canonical residual transposition subgroup $S_2 = \langle S_{\mu\tau} \rangle$.

4 The no-hidden-knobs split: the minimal positive S_2-equivariant deformation is $I + P_-$

Let P_- be the rank-1 projector onto the odd axis v_-:

$$P_- := v_- v_-^T = \frac{1}{2} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -1 & 1 \end{pmatrix}.$$

Impose three no-knobs constraints:

1. Residual symmetry: the splitting operator must commute with $S_{\mu\tau}$.
2. Positivity: it must be positive (mass-like, not tachyonic).
3. Minimality: among nontrivial choices, pick the smallest deformation of I with no continuous coefficient.

The unique candidate satisfying these is

$$\boxed{O := I_3 + P_-}. \tag{2}$$

On v_- one has $Ov_- = 2v_-$, while on the orthogonal complement of v_- one has eigenvalue 1 (because P_- annihilates that subspace). In particular,

$$\lambda_- : \lambda_+ = 2 : 1$$

on the two parity axes.

If the light-neutrino doublet (m_1, m_2) corresponds to the (v_+, v_-) parity sectors under the TBM-motivated S_2, then the minimal no-hidden-knobs split is

$$\boxed{\frac{m_2}{m_1} = 2.}$$

5 A fully explicit neutrino mass matrix (TBM + $m_2/m_1 = 2$ + $m_3/m_2 = \alpha^{-1/3}$)

Assume the DFD hierarchy

$$\frac{m_3}{m_2} = r := \alpha^{-1/3},$$

and the discrete lock above $m_2/m_1 = 2$. Then, up to the overall scale m_1, the spectrum is fixed:

$$m_1 : m_2 : m_3 = 1 : 2 : 2r.$$

Using TBM eigenvectors, the mass matrix is

$$M_\nu = m_1\, P_1 + (2m_1)\, P_2 + (2rm_1)\, P_3,$$

where the TBM projectors $P_i = c_i c_i^T$ are rational matrices. Writing them explicitly:

$$P_1 = \frac{1}{6}\begin{pmatrix} 4 & -2 & 2 \\ -2 & 1 & -1 \\ 2 & -1 & 1 \end{pmatrix}, \qquad P_2 = \frac{1}{3}\begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix}, \tag{3}$$

$$P_3 = \frac{1}{2}\begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}. \tag{4}$$

Therefore, the neutrino mass matrix is fixed in closed form:

$$\boxed{M_\nu = m_1 \left[\frac{1}{6}\begin{pmatrix} 4 & -2 & 2 \\ -2 & 1 & -1 \\ 2 & -1 & 1 \end{pmatrix} + \frac{2}{3}\begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix} + r \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix} \right].} \tag{5}$$

All entries are rational linear combinations of $(1, r)$, with $r = \alpha^{-1/3}$ fixed by the single topological constant α.

6 Parameter-free oscillation invariant (the compression)

From $m_1 : m_2 : m_3 = 1 : 2 : 2r$ one gets

$$\Delta m_{21}^2 = 3m_1^2, \qquad \Delta m_{32}^2 = 4(r^2 - 1)m_1^2,$$

hence

$$\boxed{\frac{\Delta m_{32}^2}{\Delta m_{21}^2} = \frac{4}{3}(r^2 - 1) = \frac{4}{3}\left(\alpha^{-2/3} - 1\right) \approx 34.106787}.$$

7 Seesaw closure from S_2 normalization

Appendix P motivates a center-overlap Dirac Yukawa scale $y_D \sim \sqrt{\alpha}$. In the presence of the TBM-selected S_2 doublet, there is a canonical no-hidden-knobs refinement: if the relevant center-coupled right-handed state is the normalized symmetric combination of a two-state subspace, then any overlap amplitude acquires a factor $1/\sqrt{2}$. Thus one is led to

$$\boxed{y_D = \frac{\sqrt{\alpha}}{\sqrt{2}} = \sqrt{\frac{\alpha}{2}}.} \tag{6}$$

With the DFD theorem $M_R = M_P\alpha^3$ and the seesaw estimate $m_\nu \sim (y_D v)^2/M_R$, the heaviest light-neutrino mass closes as

$$m_3 = \frac{(\alpha/2)\,v^2}{M_P\alpha^3} = \frac{v^2}{2M_P\alpha^2}. \tag{7}$$

Using $v = M_P\alpha^8\sqrt{2\pi}$, this becomes a pure α-power:

$$\boxed{m_3 = \pi\, M_P\, \alpha^{14}.} \tag{8}$$

Given the fixed ratios $m_2/m_1 = 2$ and $m_3/m_2 = \alpha^{-1/3}$, all three light masses follow:

$$\boxed{m_1 = \frac{m_3}{2\alpha^{-1/3}}, \qquad m_2 = \frac{m_3}{\alpha^{-1/3}}, \qquad m_3 = \pi M_P\alpha^{14}.} \tag{9}$$

8 Numerical predictions (manuscript conventions)

Using the manuscript values $\alpha^{-1} = 137.036$, $M_P = 1.22 \times 10^{19}$ GeV, and $v = 246.09$ GeV, Eq. (7) gives:

Quantity	Prediction	Notes
m_1	4.52 meV	from $m_2/m_1 = 2$ and $m_3/m_2 = \alpha^{-1/3}$
m_2	9.04 meV	same
m_3	46.61 meV	from the S_2-normalized seesaw closure
Σm_ν	60.17 meV	fully determined
Δm^2_{21}	6.13×10^{-5} eV^2	equals $3m_1^2$
Δm^2_{32}	2.09×10^{-3} eV^2	equals $4(r^2 - 1)m_1^2$
$\Delta m^2_{32}/\Delta m^2_{21}$	34.1068	equals $(4/3)(\alpha^{-2/3} - 1)$

Beta decay and neutrinoless double beta decay (TBM limit)

In the TBM limit $U_{e3} = 0$,

$$m_{\beta\beta} = \left|\frac{2}{3}m_1 + \frac{1}{3}m_2\right| = \frac{4}{3}m_1, \qquad m_\beta = \sqrt{\frac{2}{3}m_1^2 + \frac{1}{3}m_2^2} = \sqrt{2}\,m_1. \tag{10}$$

Thus

$$\boxed{m_{\beta\beta} = 6.03 \text{ meV}, \qquad m_\beta = 6.39 \text{ meV}.} \tag{11}$$

9 Falsifiers specific to this closure

This closure is deliberately sharp, so it can fail sharply:

- If the measured ratio $\Delta m_{32}^2/\Delta m_{21}^2$ is incompatible with $(4/3)(\alpha^{-2/3} - 1)$ at high precision, the S_2 lock $m_2/m_1 = 2$ is wrong.
- If future cosmology strongly prefers Σm_ν far from ~ 60 meV while α remains fixed, then the S_2-normalized seesaw closure (or the identification of M_R) fails.
- If $0\nu\beta\beta$ bounds push below ~ 6 meV (in the same TBM-limit mapping), the TBM+S_2 closure for the (m_1, m_2) subspace fails.

Pointer

Unified DFD manuscript (Zenodo DOI): `https://doi.org/10.5281/zenodo.18066593`