

# The Structural Origin of Yukawa Operators in VERSF

## Generation-Depth Sectors, Closure-Norm Mediation, and the Constrained Architecture of Fermion Mass

Keith Taylor

VERSF Theoretical Physics Programme

---

### Summary for the General Reader

The Standard Model gets fermion masses right, but it does not explain them. It writes down a table of numbers — the Yukawa couplings — by hand: one for the electron, a bigger one for the muon, a much bigger one for the tau, and so on across quarks and leptons. The numbers span twelve orders of magnitude, they come with a pattern of "mixing" between generations, and the theory offers no reason for any of it. They are inputs, not outputs.

This paper does not produce those numbers either — and it says so plainly throughout. What it does is locate *where the numbers come from* inside VERSF, and then squeeze that location for everything it forces without further fitting. The proposal is simple to state. In VERSF, matter is made of stable knots in the substrate (Persistent Fold Defects), the three generations are three "depths" at which such a knot can stably sit, and the Higgs is the substrate's closure-norm mode — its scalar breathing degree of freedom. A mass term is the one thing that bridges a left-handed knot to a right-handed knot, and the only substrate object that *can* sit in that bridge is the scalar closure-norm mode. So a Yukawa coupling is read as an **overlap amplitude**: how much a left-handed generation-depth state, the right-handed state, and the closure-norm mode share support. Write that overlap  $Y_{ij}$ , and it has exactly the shape a Yukawa matrix element must have.

The central structural content is that reading Yukawas as overlaps does not merely relocate them — it *constrains* them. Some of those constraints are theorems that hold with no further assumption; others hold under one clearly-flagged hypothesis, and the paper keeps the two grades apart throughout. The strongest of the assumption-free ones: because the closure-norm cost is a positive quantity, the overlaps obey a Cauchy–Schwarz inequality, so the cross-generation couplings cannot exceed the geometric mean of the same-generation supports. That is a real prohibition — it forbids strong mixing between two sectors that are each weakly coupled — and it is the structural seed of the long-noticed empirical pattern that mixing angles track the square roots of mass ratios. Its extreme case is a second theorem: when one closure mode dominates, the Yukawa matrix is an outer product, i.e. rank one, so at leading order exactly one generation is heavy and the lighter generations are switched on by additional modes — precisely the staircase pattern the data show, with the spread of masses set by the spectrum of closure

modes. The framework also reproduces the Standard Model's Yukawa parameter count exactly, ties the total mass-squared budget to the total closure-coupling strength, and explains an asymmetry the Standard Model only stipulates: gauge forces are flavour-blind while masses are wildly generation-dependent — forced here, because gauge labels are carried by the defect *class* while flavour is carried by *depth*, so the two are orthogonal substrate structures. These are the theorems; what they do *not* do is fix any number.

One of those results deserves singling out, because it answers a question the Standard Model only writes down and never explains: why are the forces the same for every generation while the masses are so different? An electron and a muon feel electromagnetism identically, yet differ in mass by two hundred. In this picture that is not two facts but one. What a particle *is* — its electric charge, its place in the forces — is carried by the *kind* of knot it is; how *heavy* it is is carried by the knot's *depth*. Those are two independent labels on the substrate, so the forces (which see only the kind) cannot tell generations apart, while the masses (which see the depth) must come out unequal. The universality of the forces and the hierarchy of the masses fall out of the same separation — arguably the cleanest single idea in the paper.

One further step is taken under an explicitly flagged assumption — that the closure-norm mediation is *local in depth*. If that holds, the off-diagonal overlaps fall off with depth separation, so the most-suppressed mixing is between the first and third generations, exactly as the measured quark-mixing matrix shows. The mass *hierarchy* is a separate matter — it lives on the diagonal and needs a preferred origin in depth (the closure-norm mode's own location), each generation graded by its absolute distance from that support, which is the standard split-fermion mechanism. This is the same kind of mechanism that makes wavefunction-localization flavour models work, so the framework is doing something physics already knows can succeed; it is marked conditional, because the locality of the mediator is the assumption doing the work, and the paper does not hide that.

For the charged leptons specifically, companion papers turn this from a shape into numbers — and the rate of the hierarchy is no longer a fitted dial. The factor by which each generation is heavier than the last (about two hundred for electron→muon) now comes out of counting the independent ways the substrate's closure architecture can be traversed: the count is fourteen, and the same fourteen that fixes why there are exactly three generations and no fourth. The bare count slightly undershoots the observed muon-to-electron ratio, and the small gap is carried honestly as a known, bounded correction rather than hidden. A second companion route, from the curvature of the internal space, lands on almost the same number — the two agree because the curvature value, exponentiated, rounds to fourteen — so they are two readings of one structural scale; the small gap between the integer fourteen and that exponential is exactly the "correction" mentioned above, not a separate effect. The muon-to-tau jump is much smaller than the electron-to-muon one (seventeen, not two hundred), and the companion explains this too: the tau sits right at the edge of being able to exist at all, which compresses its mass spacing. The structural number behind all this is fourteen, not thirteen — a thirteen-based rule of thumb gets the first ratio close but mispredicts the second, precisely because the real count is fourteen.

This points to a way of seeing the whole flavour puzzle differently, offered as a conjecture rather than a result. If a fermion's mass is an overlap, then the question "why are the masses so

different?" becomes "why are the substrate's closure-maintenance costs so different?" — and the paper conjectures that heavier generations are heavy because the substrate must work harder to keep more committed, distinguishable structure in existence. In that reading mass is a measure of bookkeeping complexity, not just of how tightly something is squeezed. The paper is careful to mark this as a conjecture and not to lean on it; but if it can be proved, the Standard Model's table of flavour constants would become something computable from the geometry of the substrate, and this paper, with its companions, is the bridge from the force-and-matter sectors into that mass programme. It also names the next problem precisely: compute the "costs" of the substrate's closure modes — the eigenvalue problem for the closure operator — and the individual masses would follow.

What stays open is still the absolute scale — the actual mass of the electron in everyday units — and the values of the individual overlaps, which need the substrate's closure operator computed from first principles. But less is open than before: the *rate* of the hierarchy is now inherited rather than fitted, the count of generations is a theorem, and the charged-lepton ratios are reproduced to a few percent by the companions. The honest result is structural and constrained: Yukawas are overlap operators, that reading is backed by theorems that forbid certain textures, fix the rank pattern and parameter count, separate flavour from gauge structure, and predict the qualitative hierarchy — with the hierarchy's rate now derived — and the remaining numbers left to be computed rather than fitted. In short, the flavour problem is transformed from "Why are the Yukawa couplings what they are?" into "What determines the spectrum of closure-maintenance costs?"

---

## Note on Epistemic Discipline

This paper makes a structural claim and is careful never to let it round up to a numerical one. The central identification — Yukawa matrix  $\leftrightarrow$  closure-depth overlap operator — is labelled [structural identification], conditional on the inherited identification of the Higgs with the closure-norm mode. The consequences are sorted into two grades: those following from the operator structure and the positivity of the mediator alone, marked [structural], and those requiring the additional depth-locality hypothesis, marked [conditional] with the hypothesis named at the point of use. No numerical value of any mass, angle, or matrix element is claimed; all are [open]. Two prior-paper inheritances remain to be confirmed — the closure-norm equilibrium value and the PFD-dictionary class/gauge assignments — and are flagged [open: inheritance to confirm] rather than asserted; a third, the depth-decay exponent  $\kappa_{\text{clos}} \approx \ln(14)$ , has now been confirmed against the companion closure-loop count (convergent with the CP<sup>2</sup>-focusing value 8/3) and is inherited in §7.3.

# Table of Contents

1. **Introduction** — the Yukawa sector as the Standard Model's largest unexplained input.
2. **Epistemic Status** — inherited, identified, structural-consequence, conditional, open.
3. **Inherited Structure** — PFDs, three generation-depth sectors, the closure-norm Higgs mode.
4. **The Mediator Is Forced** — the closure-norm mode is the unique identified substrate scalar that can bridge L and R.
5. **The Overlap Identification** —  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$ , and the scale/structure split of the mass.
6. **Structural Theorems (assumption-free)** — the closure-overlap bound, the single-mode rank theorem and the closure-spectrum hierarchy, hierarchy without tuning, the parameter count, flavour localization, chiral protection, and the total mass budget.
7. **Conditional Consequences (depth geometry)** — the generation-separation metric, banded overlaps, exponential decay with the inherited closure-loop exponent  $\kappa \approx \ln(14)$  (convergent with  $8/3$ ), the kinship to localization flavour models, the charged-lepton cascade realized in the companion Role-4/BCB papers, the closure-mode spectrum as the conjectured source of hierarchy, and the synthesis (mass as closure-maintenance cost).
8. **Relation to Mass Hierarchies and Mixing** — where the numbers must originate; what is and is not inherited.
9. **Why This Matters** — relocation of the flavour sector from parameters to substrate geometry.
10. **Open Problems and Inheritances to Confirm.**
11. **Conclusion.**

---

## Abstract

The Standard Model generates fermion masses through Yukawa couplings  $\mathcal{L}_Y = -Y_{ij} \bar{\psi}_{L,i} H \psi_{R,j} + \text{h.c.}$ , with the matrices  $Y_{ij}$  treated as fundamental inputs whose existence, hierarchy, count, and mixing structure are unexplained. This paper gives a structural account of the Yukawa sector within VERSF and extracts its forced consequences as theorems. Building on inherited results — matter as Persistent Fold Defects, three refinement-stable closure-depth generation sectors, and the Higgs as the closure-norm mode — it argues that the closure-norm mode is the *unique identified* substrate scalar capable of bridging left- and right-handed sectors, and identifies the Yukawa matrix element with the closure-depth overlap amplitude  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$  [structural identification, conditional on the Higgs = closure-norm identification]. The mass then splits cleanly: its overall scale is the closure-norm equilibrium value (the substrate origin of  $v$ ), its flavour structure is the overlap. The identification is shown to carry theorem-level content rather than mere relocation. From positivity of the closure-norm mediator alone: the Closure-Overlap Bound  $|Y_{ij}|^2 \leq \langle \gamma_{L,i} | \hat{H}_c | \gamma_{L,i} \rangle \langle \gamma_{R,j} | \hat{H}_c | \gamma_{R,j} \rangle$  forbids large mixing between individually weakly-coupled sectors and seeds the  $\sqrt{\text{mass-ratio}}$  texture pattern; its equality case is the Single-Mode Closure Limit, in which a dominant closure mode makes  $Y$  an outer product of rank one — one heavy generation at leading order, lighter generations from subleading modes

— generalizing to the spectral expansion  $Y = \sum_n \lambda_n a_n b_n^\dagger$ , in which the fermion mass hierarchy is controlled by the closure-mode spectrum  $\{\lambda_n\}$ , reducing the hierarchy problem to the eigenvalue problem for  $\hat{H}_c$ ; the  $3 \times 3$  overlap reproduces the Standard Model Yukawa parameter count exactly; the total mass-squared budget satisfies  $\sum m_i^2 \propto \text{Tr}(Y^\dagger Y)$ ; and a Flavour Localization result (conditional on the dictionary) shows gauge representation (carried by PFD class) and flavour hierarchy (carried by generation depth) are orthogonal substrate structures, explaining why flavour resides in the mass sector and the gauge sector is flavour-universal. Under an explicitly flagged depth-locality hypothesis on  $\hat{H}_c$ , the off-diagonal overlaps are graded by generation separation, giving a banded matrix and the observed ordering of mixing suppression (generation-1–3 most suppressed, as in CKM); the mass hierarchy is a separate, *inhomogeneous* statement — three stacked rank-one localization blocks, one closure mode per depth,  $Y = \sum_g \lambda_g a_g b_g^\dagger$  with  $\sigma_g \sim \exp(-2\kappa a_g)$ , the rank-three Proposition 6.2' form, reproducing  $M_g \propto e^{(2\kappa g)}$ , the genuine split-fermion mechanism (a separation-only Toeplitz form would give mixing but no hierarchy, and a single block is rank one — one mass, not three — so the three blocks must be summed). The depth-decay exponent is inherited from the companion closure-loop count,  $\kappa_{\text{clos}} \approx \ln(14)$  ( $\dim(\Gamma_{\text{min}}) = 14$ ); the bare loop count gives the per-generation factor  $e^{(2\kappa_0)} = 14^2 = 196$ , 5% below the observed 206.8, the residual  $w \approx 1.055$  being the gap to the focusing value  $e^{(8/3)} = 14.39$  — whether the two are one structural scale (so  $w$  is a rounding artifact) or two convergent-but-independent derivations is open (§7.3, problem #8), and the structural choice of the integer 14 is the empirically worse fit, preferred on structural-unity grounds. The charged-lepton cascade is realized in the companions with threshold-proximity suppression at the tau and the three-generation cutoff a companion theorem ( $C_0 \leq 14\kappa'$ ). No numerical mass, angle, or matrix element is derived; all remain [open], as do two prior-paper inheritances (the closure-norm equilibrium value and the PFD-dictionary class/gauge assignments). The contribution is a constrained structural identification, now theorem-backed: Yukawas are closure-depth overlap operators, that reading forbids textures, fixes the rank pattern and parameter count, separates flavour from gauge structure, predicts the qualitative hierarchy, and locates both halves of fermion mass — scale and structure — in substrate geometry. In one line: mass is a closure-maintenance cost, Yukawas are its overlap expression, and the hierarchy is the spectrum of that cost — controlled, for the charged leptons, by the same closure architecture ( $\dim(\Gamma_{\text{min}}) = 14$ ) that fixes the generation count.

---

## 1. Introduction

The Yukawa sector is the largest single reservoir of unexplained parameters in the Standard Model. Fermion masses arise from

$$\mathcal{L}_Y = -Y_{ij} \bar{\psi}_{L,i} H \psi_{R,j} + \text{h.c.},$$

and after electroweak symmetry breaking

$$m_{ij} = Y_{ij} v / \sqrt{2},$$

where  $v$  is the Higgs vacuum expectation value. The matrices  $Y_{ij}$  reproduce the observed masses and flavour mixing with great accuracy, yet they are introduced by hand. The Standard Model does not explain:

- why Yukawa operators exist at all,
- why there are three generations,
- why the masses are hierarchical, spanning some twelve orders of magnitude,
- why flavour mixing occurs, and why its pattern is what it is.

The purpose of this paper is *not* to compute the numerical Yukawa matrices. It is to identify the substrate-level object that plays the role of the Yukawa sector within VERSF, and — this is the new emphasis — to prove the constraints that identification forces on Yukawa structure. An identification that merely relabelled the couplings would be thin. The identification offered here is shown to forbid certain textures, fix the rank pattern and the parameter count, separate flavour from gauge structure, and predict the qualitative hierarchy and mixing pattern, while leaving the numbers to a future first-principles computation.

## 2. Epistemic Status

This paper contains no new numerical predictions. Its claims are graded as follows.

**[Inherited]** (from earlier VERSF work, not rederived here)

- Matter corresponds to Persistent Fold Defects.
- Three refinement-stable generation-depth sectors exist.
- The Higgs field corresponds to the closure-norm mode.

**[Structural identification]** (conditional on the Higgs = closure-norm identification)

- Yukawa operators correspond to closure-depth overlap operators,  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$ .

**[Structural]** (forced by the operator form and mediator positivity, no further assumption)

- The closure-norm mode is the unique identified substrate scalar mediator (§4).
- Theorem 6.1, the Closure-Overlap Bound, and its texture restriction.
- Proposition 6.2, the Single-Mode Closure Limit (rank-one Yukawa), as the equality case of 6.1.
- Proposition 6.2', the Hierarchical Closure Spectrum:  $Y = \sum_n \lambda_n a_n b_n^\dagger$ , mass hierarchy controlled by the closure-mode spectrum  $\{\lambda_n\}$ .
- Proposition 6.3, hierarchy without tuning (assumption-free part).
- Proposition 6.4, the exact Standard Model Yukawa parameter count.
- Result 6.5, Flavour Localization (conditional on the dictionary) (gauge class  $\perp$  flavour depth).

- Proposition 6.6, chiral protection of zero-overlap sectors.
- Proposition 6.7, the total mass-squared budget  $\Sigma m_i^2 \propto \text{Tr}(Y^\dagger Y)$ .

**[Conditional]** (require the depth-locality hypothesis on  $\hat{H}_c$ , flagged at use)

- Generation separation grades the off-diagonals  $\rightarrow$  banding and the CKM-like mixing ordering (§7.1–7.2).
- The hierarchy as three stacked rank-one localization blocks (one closure mode per depth),  $Y = \Sigma_g \lambda_g a_g b_g^\dagger$  with  $\sigma_g \sim \exp(-2\kappa a_g)$  — the rank-three 6.2' form — giving  $M_g \propto e^{(2\kappa g)}$  (§7.3).
- Banded, exponentially-suppressed overlaps, the mass hierarchy, and the CKM-like mixing ordering (§7), with depth-decay exponent  $\kappa_{\text{clos}} \approx \ln(14)$  now inherited from the companion closure-loop count (convergent with 8/3; corrected to the observed ratio by the bounded asymmetry  $w \approx 1.055$ ), not a free constant.

**[Inherited realization — conditional]**

- The charged-lepton singular-value cascade of Proposition 6.2' is realized in the companion Role-4/BCB papers as closure-loop exponential localization (per-generation factor  $e^{(2\kappa_0)} = 14^2 = 196$  from  $\kappa_0 = \ln(14)$ , 5% below the observed 206.8; the gap  $w \approx 1.055$  is the residual to the focusing value  $e^{(8/3)} = 14.39$ , a rounding artifact *if* the two routes are the same scale — open, #8 — else a residual between two convergent derivations) modulated by a threshold-proximity profile factor at the tau (§7.4); the no-fourth-generation cutoff is a companion theorem ( $\dim(\Gamma_{\text{min}}) = 14$ ,  $C_0 \leq 14\kappa'$ ). The structural integer governing the hierarchy is  $14 = \dim(\Gamma_{\text{min}})$ , not 13; a 13-based rule mispredicts the second jump.

**[Conjectural — reframing of the hierarchy source (§7.5)]**

- Closure-Mode Spectral Conjecture: the eigenvalues  $\lambda_n$  of  $\hat{H}_c$  are closure-maintenance costs, and the hierarchy comes primarily from their spread ("why are closure modes different?" replacing "why are Yukawas different?").
- Conjecture 7.4 (Non-Uniform Hierarchy Growth): infrastructure reuse makes the second generation increment smaller ( $\mathcal{S}_\tau - \mathcal{S}_\mu < \mathcal{S}_\mu - \mathcal{S}_e$ ), offered as a candidate *cause* of the companion's threshold proximity (unifying, not competing), and so as a route to deriving the  $C_2 \approx 22$  input of open problem #8.
- Hierarchy Dominance Conjecture: persistent distinguishability  $\mathcal{S}_D$  dominates the log-mass decomposition (mass as informational rather than geometric complexity). Explicitly *not* a theorem — the five-term decomposition is underived, the dominance currently traces to the 13-based reconstruction (not a sound basis), and it must reduce to the companion's two-term form.

**[Open]**

- Exact fermion masses, CKM angles, PMNS angles, numerical Yukawa elements.
- The substrate origin of the left/right chirality-sector distinction.

## [Open: inheritance to confirm]

- The closure-norm equilibrium value as the substrate origin of  $v$  (§5).
- The PFD-dictionary assignment of gauge representation to defect class, used in Result 6.5.

*(Now confirmed and no longer open: the depth-decay exponent  $\kappa_{\text{clos}} \approx \ln(14)$ , derived in the companion closure-loop count (convergent with  $8/3$ ) — see §7.3.)*

---

## 3. Inherited Structure

### 3.1 Persistent Fold Defects

Matter is represented by Persistent Fold Defects (PFDs), each carrying an invariant tuple

$$\mathcal{J}(D) = (C\_D, \beta_1(D), h\_D, \pi\_D, \chi\_D, \gamma\_D, \ell\_D, \rho\_D),$$

where  $\gamma\_D$  denotes generation depth. [inherited]

### 3.2 Generation Depth

Earlier VERSF work identifies exactly three refinement-stable generation-depth sectors,

$$\gamma_1, \gamma_2, \gamma_3,$$

corresponding to the observed three-generation structure: a fourth sector is not refinement-stable, which is what makes the count three rather than an empirical input. The present paper does not rederive this and inherits both the count and the stability criterion from the companion work, where it is now a theorem: the Structural Completeness companion proves  $(n\_int, n\_role) = (3, 4)$  the unique admissible pair and the no-fourth-generation cutoff structural  $(C_0 \leq 14\kappa' \implies C_3 < C\_crit(3))$  for all admissible  $\eta\_A$ , so the "three" is structurally fixed, not calibrated. [inherited; companion theorem]

### 3.3 Closure-Norm Mode

Earlier work identifies the Higgs sector with fluctuations of the closure norm,

$$\rho_{\text{clos}}(x),$$

with equilibrium value and fluctuation  $\delta\rho_{\text{clos}}$  controlling the energetic cost of closure stabilization. (This closure-norm  $\rho_{\text{clos}}$  is distinct from the transport-offset cochain  $\rho$  of the closure-charge sector; the symbols are kept apart deliberately.) The present paper inherits this identification. [inherited]

---

## 4. The Mediator Is Forced

A fermion mass term connects a left-handed sector to a right-handed sector through a Lorentz scalar. In substrate language the question is:

What object within VERSF can sit in the bridge  $L \rightarrow (\text{scalar}) \rightarrow R$ ?

The answer is not merely *a* candidate but, among the substrate modes the framework has identified, the *only* one. A mass bilinear  $\bar{\psi}_L(\cdot)\psi_R$  requires the inserted object to be a Lorentz scalar (here, as throughout, "Lorentz" is effective-description language: the symmetry is itself emergent from the substrate geometry, and "Lorentz scalar" means a spin-0 object in that emergent description) and a genuine substrate mode, not an external field. The closure-norm mode is the substrate's scalar sector — its radial, gauge-singlet breathing degree of freedom. The vector and tensor substrate modes carry spin and so cannot occupy the spin-0 bridge — that is the operative exclusion; the gauge-connection modes are vectors for the same reason.

**Observation 4.1 (Unique identified scalar mediator) [structural].** Among the substrate modes identified in the framework, the closure-norm mode is the unique Lorentz-scalar singlet that can mediate a left–right mass bilinear. Asked "if not the closure-norm mode, what substrate scalar remains?", the answer is: none currently identified. The closure-norm mode is therefore not merely a candidate for the Yukawa bridge but the only currently identified substrate object capable of occupying it — and the Higgs role is forced upon it by the requirement that the mediator be the substrate's scalar singlet, rather than assigned by choice.

This upgrades §3.3's inherited identification from "the Higgs is the closure-norm mode" to "the closure-norm mode is the only identified thing that could play the Higgs's role in the mass term" — the stronger statement for what follows. The qualifier "identified" is deliberate and honest: it is a no-currently-known-alternative statement, not a proof that no other scalar can exist.

---

## 5. The Overlap Identification

Let  $|\gamma_{L,i}\rangle$  and  $|\gamma_{R,j}\rangle$  denote the left- and right-handed generation-depth states ( $i, j \in \{1, 2, 3\}$ ), and let  $\hat{H}_c$  denote the closure-norm operator associated with the Higgs mode. The natural — and, given §4, essentially forced — bilinear is

$$M = \bar{\psi}_L \hat{H}_c \psi_R + \text{h.c.},$$

which in the generation-depth basis reads

$$M = \sum_{ij} \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle \bar{\psi}_{L,i} \psi_{R,j}.$$

**Identification 5.1 (Yukawa correspondence) [structural identification, conditional on the Higgs = closure-norm identification of §3.3].** The Standard Model Yukawa matrix element is the closure-depth overlap amplitude

$$Y_{ij} \equiv \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle.$$

This is not a theorem in the mathematical sense; it is a correspondence between a substrate object and an effective-field-theory object. Its significance is that it fixes *where* Yukawa structure resides in the VERSF ontology — and, as §6 shows, it lets one prove what that residence forces.

**The scale/structure split.** With  $m_{ij} = Y_{ij} v / \sqrt{2}$  and  $v$  the closure-norm equilibrium value, the mass sector separates into two substrate quantities: the overall **scale** is set by the closure-norm equilibrium (the substrate origin of  $v$ ), and the **flavour structure** is set by the overlaps  $Y_{ij}$ . Both halves are thereby located in substrate geometry. The identification of  $v$  with the closure-norm equilibrium is marked [open: inheritance to confirm], pending the closure-norm-sector paper; the structural split holds regardless of the numerical value.

## 6. Structural Theorems (Assumption-Free)

The overlap reading is not consequence-free relabelling. The following results follow from the operator form and the positivity of the mediator alone — no locality, no dynamics, no fitting. Throughout,  $\hat{H}_c$  is the energetic cost of closure stabilization and is therefore a positive operator,  $\hat{H}_c \geq 0$ ; positivity is the natural and minimal physical input, and it is enough.

### 6.1 The Closure-Overlap Bound

Cauchy–Schwarz for the positive sesquilinear form  $\langle \cdot | \hat{H}_c | \cdot \rangle$  gives, for any states,

$$|\langle x | \hat{H}_c | y \rangle|^2 \leq \langle x | \hat{H}_c | x \rangle \cdot \langle y | \hat{H}_c | y \rangle.$$

**Theorem 6.1 (Closure-Overlap Bound) [structural].** For a positive closure-norm operator  $\hat{H}_c$ , every Yukawa overlap is bounded by the left-handed and right-handed diagonal overlaps of its endpoints:

$$|Y_{ij}|^2 \leq \langle \gamma_{L,i} | \hat{H}_c | \gamma_{L,i} \rangle \cdot \langle \gamma_{R,j} | \hat{H}_c | \gamma_{R,j} \rangle.$$

The right-hand factors are the **left-left and right-right** self-overlaps  $D_i^{\wedge}L \equiv \langle \gamma_{L,i} | \hat{H}_c | \gamma_{L,i} \rangle$  and  $D_j^{\wedge}R \equiv \langle \gamma_{R,j} | \hat{H}_c | \gamma_{R,j} \rangle$  — diagonal *closure supports*, not Yukawa matrix elements (which are themselves L–R overlaps). The theorem is therefore stated rigorously in terms of the self-overlaps.

**Consequence (the texture restriction).** A large off-diagonal Yukawa requires at least one large diagonal closure support:  $|Y_{ij}|$  large forces  $D_i^{\wedge}L$  and  $D_j^{\wedge}R$  both large. Equivalently, the framework *forbids* large mixing between two sectors that each have small closure support — a

Yukawa texture the Standard Model permits but nature does not exhibit. This is a real restriction on admissible Yukawa textures, with no free parameter.

**Aligned corollary (the  $\sqrt{\cdot}$ -mass-ratio pattern) [structural, with alignment].** Under the natural alignment assumption that left and right states of a given generation approximately coincide in the  $\hat{H}_c$ -metric (so the self-overlaps track the physical mass-basis diagonals,  $D_i \approx Y_{ii}$ ), the bound reads

$$|Y_{ij}| \lesssim \sqrt{Y_{ii} Y_{jj}} \sim \sqrt{(m_i m_j) / (v/\sqrt{2})},$$

which is exactly the form of the long-noticed empirical texture relations (Gatto–Sartori–Tonin / Fritsch type), of which  $|V_{us}| \approx \sqrt{(m_d/m_s)}$  is the standard instance — and a strong one: it gives  $\approx 0.225$  against the measured 0.225 in the down sector, the "partial" applying to the up-sector relation  $\sqrt{(m_u/m_c)}$ . The framework structurally produces relations of this type and forbids their violation. One caveat the alignment carries: reaching  $D_i \approx Y_{ii}$  needs  $|\gamma_{L,i}\rangle \approx |\gamma_{R,i}\rangle$  in the  $\hat{H}_c$ -metric, yet *exact* L–R coincidence would collapse the very chirality distinction that makes the mass term nontrivial (§4). So the alignment can only be approximate, and there is an implicit balance — enough L–R overlap to be massive, enough L–R distinction to be chiral; the corollary holds in that intermediate regime. The inequality of Theorem 6.1 itself is assumption-free; only this  $\sqrt{\cdot}$ -mass-ratio reading invokes the (approximate) alignment.

## 6.2 The Closure Spectrum: Single-Mode Limit and the Mass Hierarchy

The equality case of Theorem 6.1 is itself a theorem, and a physically pointed one. Cauchy–Schwarz is saturated exactly when  $\hat{H}_c^{1/2}|\gamma_{L,i}\rangle$  and  $\hat{H}_c^{1/2}|\gamma_{R,j}\rangle$  are parallel for all  $i, j$  — which holds precisely when  $\hat{H}_c$  has a single dominant mode coupling to the depth states,  $\hat{H}_c = \lambda |h\rangle\langle h|$  as restricted to the depth-state subspace (effectively rank one —  $\hat{H}_c$  may carry further modes orthogonal to every  $|\gamma_{L,i}\rangle, |\gamma_{R,j}\rangle$  without touching the Yukawas). In that limit

$$Y_{ij} = \lambda \langle \gamma_{L,i} | h \rangle \langle h | \gamma_{R,j} \rangle = \lambda a_i \bar{b}_j, \quad a_i \equiv \langle \gamma_{L,i} | h \rangle, \quad b_j \equiv \langle \gamma_{R,j} | h \rangle,$$

an outer product.

**Proposition 6.2 (Single-Mode Closure Limit) [structural, conditional on single-mode dominance].** If closure-norm mediation is dominated by a single coherent mode, the Yukawa matrix is rank one at leading order:

$$\text{rank}(Y) = 1, \quad Y = \lambda a b^\dagger \quad (\text{rank one on the depth-state subspace}).$$

Consequences, all immediate:

- A rank-one  $Y$  has a single nonzero singular value, so **exactly one generation is massive at leading order** while the other two are massless — the extreme hierarchy the Standard Model exhibits, in which each charge sector has one dominant member ( $t \gg c \gg u, b \gg s \gg d, \tau \gg \mu \gg e$ ).

- **Lighter generations are switched on by additional closure modes.** Full rank (three massive generations) requires at least three independent contributing modes, and the mass hierarchy becomes the hierarchy of the modes' singular values — a singular-value cascade driven by the spread of the closure-mode spectrum.
- The single-mode limit **saturates the Closure-Overlap Bound** for every entry, tying Proposition 6.2 to Theorem 6.1 as its equality case: maximal coherence of the mediator is exactly the rank-one, one-heavy-generation regime.

One does not have to claim the mediator is single-mode; the proposition states what follows if it is, and the staircase it produces is the observed one.

The single-mode limit is the leading term of a general spectral expansion, and writing out the full spectrum is what turns the rank-one extreme into a hierarchy mechanism. Any positive  $\hat{H}_c$  has a spectral decomposition  $\hat{H}_c = \sum_n \lambda_n |h_n\rangle\langle h_n|$  with  $\lambda_n \geq 0$ , and the overlap inherits it term by term:

$$Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle = \sum_n \lambda_n \langle \gamma_{L,i} | h_n \rangle \langle h_n | \gamma_{R,j} \rangle = \sum_n \lambda_n (a_n)_i (\bar{b}_n)_j,$$

so that

$$Y = \sum_n \lambda_n a_n b_n^\dagger, \quad (a_n)_i \equiv \langle \gamma_{L,i} | h_n \rangle, \quad (b_n)_j \equiv \langle \gamma_{R,j} | h_n \rangle.$$

**Proposition 6.2' (Hierarchical Closure Spectrum) [structural].** The Yukawa matrix is the spectral sum  $Y = \sum_n \lambda_n a_n b_n^\dagger$  of rank-one closure-mode contributions, weighted by the closure-mode spectrum  $\{\lambda_n\}$ . If the spectrum is itself hierarchical,

$$\lambda_1 \gg \lambda_2 \gg \lambda_3,$$

then the leading term  $\lambda_1 a_1 b_1^\dagger$  is the rank-one, one-heavy-generation matrix of Proposition 6.2; the second term switches on the second generation; the third the lightest — a generation-by-generation activation in which the fermion mass hierarchy is controlled by the spectrum  $\{\lambda_n\}$  of the closure-norm operator.

Two points of precision. First, the masses are the singular values of  $Y$ , and these track  $\{\lambda_n\}$  *exactly* only when the mode profiles  $\{a_n\}$  and  $\{b_n\}$  are each orthonormal sets — in which case  $Y = \sum_n \lambda_n a_n b_n^\dagger$  is literally the singular-value decomposition and  $\sigma_n(Y) = \lambda_n |a_n| |b_n|$ . In general the singular values are *controlled by* (bounded through) the  $\lambda_n$  rather than equal to them, and the mode profiles also contribute; the clean statement is that a hierarchical closure spectrum *induces* a hierarchical mass spectrum, with the profiles modulating it. Second, this locates the next paper's question exactly: the masses are not free inputs but images of  $\{\lambda_n\}$ , so the matter-sector hierarchy reduces to *what dynamics fixes the closure-mode spectrum* — the eigenvalue problem for  $\hat{H}_c$  on the generation-depth sectors. Proposition 6.2 (one mode) and Proposition 6.4 (three slots) are the rank-one and full-rank ends of this single spectral statement.

The Standard Model has a naturalness puzzle: gauge couplings are  $O(1)$  while Yukawas span twelve orders of magnitude, and the smallness of the light Yukawas is unexplained. The overlap reading dissolves the puzzle in two stages.

**Proposition 6.3 (Hierarchy without tuning).** (a) [*structural*]. For normalized depth states and bounded  $\hat{H}_c$ , every overlap satisfies  $|Y_{ij}| \leq \|\hat{H}_c\|$ , and overlaps of distinct states are generically strictly below the diagonal supports. Small Yukawas are therefore the generic case, not a tuned one: a small number arising as the overlap of imperfectly-matching states needs no cancellation or fine adjustment, in contrast to a small *fundamental* parameter. The smallness is structural. (b) [*conditional, see §7*]. The *magnitude* of the smallness — twelve orders — is supplied by the exponential suppression of localized-state overlaps in depth separation, under the depth-locality hypothesis of §7. Modest depth differences then produce enormous mass ratios without tuning.

The contrast with gauge couplings is the heart of it: gauge couplings are local geometric data (depth-blind, hence  $O(1)$  and universal), whereas Yukawas are overlaps (depth-sensitive, hence naturally small and hierarchical). The same architecture that makes gauge couplings uniform makes Yukawas hierarchical.

## 6.4 Exact Parameter Count

**Proposition 6.4 (Parameter count) [structural].** A 3-sector  $\times$  3-sector overlap operator has exactly nine complex matrix elements per fermion type, so the closure-depth overlap supplies exactly the Standard Model count of Yukawa parameters per charge sector — a  $3 \times 3$  complex matrix, before the field redefinitions that remove unphysical phases. The framework neither over-determines the count (which would predict relations the data refute) nor under-determines it (which would leave structure unexplained).

## 6.5 Flavour Localization

The closure-norm overlap and the gauge structure live on different substrate coordinates, and saying so sharply explains a Standard Model stipulation. The PFD dictionary assigns gauge representation labels to the defect *class* (the discrete invariants of  $\mathcal{J}(D)$  — hypercharge as a closure-ledger invariant, colour and weak representations as class data), while generation is the *depth* coordinate  $\gamma_D$ . These are independent slots of the invariant tuple.

**Result 6.5 (Flavour Localization) [structural, conditional on the PFD-dictionary class assignment].** Gauge representation labels are carried by PFD class; mass (flavour) structure is carried by generation depth. The two are orthogonal substrate structures. (Labelled a Result rather than a Theorem because it is conditional on the inherited dictionary assignment, not unconditional like 6.1.) Consequently:

- **Gauge couplings are flavour-universal.** They are vector-mode (connection) couplings fixed by class, depth-blind, hence identical across generations of a given class.
- **Flavour structure resides in the mass sector and nowhere else.** Only the mass term is a depth overlap; it is the only place generation-dependence can enter.

These two consequences are one fact, and together they answer a question the Standard Model only records: *why are the gauge interactions generation-universal while the masses are wildly hierarchical?* The single orthogonality — gauge labels on class, flavour on depth — forces both at once: universality of the forces and hierarchy of the masses are the two readings of the same structural separation, not independent facts. This is the cleanest conceptual statement the overlap picture yields, and it is assumption-free given the dictionary's class assignment.

This converts the Standard Model's unexplained gauge/Yukawa asymmetry — flavour-blind forces, flavour-rich masses — into a consequence of the substrate carrying gauge labels and flavour depth on orthogonal coordinates. The dependence on the dictionary's class/gauge assignment is flagged [open: inheritance to confirm]; given that assignment, the orthogonality and its two corollaries are structural.

## 6.6 Chiral Protection

**Proposition 6.6 (Chiral protection) [structural].** A generation-depth sector with vanishing closure-norm overlap is massless: if  $\hat{H}_c$  annihilates it, or its L and R states are  $\hat{H}_c$ -orthogonal, its mass entry is exactly zero. Masslessness is thus not fine-tuning but the structural statement that the L and R states share no closure-norm support. This gives a handle on the bottom of the hierarchy and on protected states; any application to the neutrino sector would require additional seesaw or Majorana machinery and is not claimed here.

## 6.7 The Total Mass Budget

A basis-invariant trace identity links the total mass to the total closure-coupling strength. Note first that  $\text{Tr}(Y) = \sum_i Y_{ii}$  is *not* the right object: it is basis-dependent, and the physical masses are the singular values of  $Y$ , which equal the diagonal entries only when  $Y$  is diagonal and positive. The correct invariant is built from  $Y^\dagger Y$ , whose eigenvalues are the squared singular values.

**Proposition 6.7 (Total mass budget) [structural].** The total fermion mass-squared of a charge sector is proportional to the total closure-depth coupling strength:

$$\sum_i m_i^2 = (v^2/2) \sum_i \sigma_i(Y)^2 = (v^2/2) \text{Tr}(Y^\dagger Y) = (v^2/2) \sum_{ij} |Y_{ij}|^2,$$

where  $\sigma_i(Y)$  are the singular values of  $Y$ . The Frobenius norm  $\text{Tr}(Y^\dagger Y) = \sum_{ij} |Y_{ij}|^2$  is the total squared closure-depth overlap — a single substrate quantity — so the mass-squared budget of each charge sector is fixed by the total support the depth sectors share through the closure-norm mode, independent of basis. (The naive  $\text{Tr}(Y)$  would give the mass sum only in the no-mixing limit and is not used.)

---

## 7. Conditional Consequences (Depth Geometry)

One further hypothesis, stated explicitly as the assumption doing the work, makes flavour into geometry and yields the hierarchy and mixing pattern.

**Hypothesis 7.0 (Depth-locality of the mediator) [the flagged assumption].** The closure-norm operator  $\hat{H}_c$  is local in generation depth: it couples nearby depth sectors more strongly than distant ones, so the overlap  $\langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$  decreases as the depth separation grows.

Everything in this section is [conditional] on Hypothesis 7.0 (with the stronger homogeneity form noted where used). The hypothesis is physically natural — a stabilization energy is generically short-ranged in the stabilization coordinate — but it is an assumption, and the conditional results stand or fall with it.

## 7.1 Generation Separation Governs Mixing, Not the Hierarchy

Define the generation-separation

$$d_{ij} = |\gamma_i - \gamma_j|.$$

**Consequence 7.1 [conditional on 7.0].** Under depth-locality the *off-diagonal* overlaps decrease with separation, so the comparison-data of distinct sectors are graded by  $d_{ij}$ : the larger the depth gap, the smaller the mixing. This is the content that survives — it controls the *mixing* sector (banding, §7.2), where flavour structure is indeed a kind of distance in closure space.

It is essential *not* to over-read this into a separation-only Yukawa matrix  $Y_{ij} = F(d_{ij})$ . A separation-only (Toeplitz) form would have a constant diagonal,  $Y_{11} = Y_{22} = Y_{33} = F(0)$ , and a  $3 \times 3$  symmetric matrix with equal diagonal and decaying off-diagonals has all three singular values clustered near  $F(0)$  — banding without any mass hierarchy. So the pure-separation reading gives the mixing ordering but actively *forbids* the mass hierarchy, and it would also fail to factorize into the rank-one-per-mode form of Proposition 6.2'. The hierarchy therefore cannot come from separation alone; it requires absolute depth-position, which §7.3 supplies. Separation is the right variable for mixing, the wrong one for mass.

## 7.2 Banded, Exponentially-Suppressed Overlaps

**Consequence 7.2 [conditional on 7.0].** Depth-locality makes the overlap matrix *banded*: off-diagonal entries are suppressed relative to diagonal ones, and the suppression grows with separation,

$$|Y_{13}|, |Y_{31}| < |Y_{12}|, |Y_{23}|, |Y_{21}|, |Y_{32}| < |Y_{11}|, |Y_{22}|, |Y_{33}|.$$

The most-suppressed mixing is therefore generation-1-to-3 — exactly the ordering the measured quark-mixing matrix exhibits:  $|V_{ub}| \ll |V_{us}|, |V_{cb}|$ . The framework reproduces the qualitative ordering of mixing suppression from the geometry of depth separation, without fitting.

## 7.3 The Hierarchy: Inhomogeneous Absolute Localization

**Consequence 7.3 (the hierarchy is inhomogeneous — absolute depth-position) [conditional on 7.0].** The mass hierarchy lives on the diagonal, and the diagonal is graded only if there is a

preferred origin in depth — which there is: the localization of the closure-norm mode itself. The mechanism is the standard split-fermion one, built from *one closure mode per generation*. Take closure mode  $g$  localized at absolute depth-position  $a_g$  (distance from the closure-norm support), so its contribution to the overlap is the factorized rank-one block

$$Y^{(g)}_{ij} \sim \exp(-\kappa_{\text{clos}} \cdot a_g) \cdot (\text{profile})_i (\text{profile})_j,$$

a single outer product — rank one, the one-heavy-generation building block of Proposition 6.2. The full Yukawa is the *sum* over the three modes,

$$Y = \sum_g \lambda_g a_g b_{g\dagger}, \lambda_g \sim \exp(-2\kappa_{\text{clos}} a_g),$$

which is exactly the three-term spectral expansion of Proposition 6.2' — rank three, three graded singular values  $\sigma_g \sim \exp(-2\kappa_{\text{clos}} a_g)$ , hence three masses. This is the point a single factorized form *cannot* do alone: one outer product has one nonzero singular value (one mass plus two zeros), and its off-diagonals saturate Cauchy–Schwarz rather than being small — so the graded diagonal of a single mode is not a three-generation spectrum. The hierarchy is three stacked one-generation blocks, one per depth, not one block read three ways.

With the standard convention that heavier generations sit *nearer* the closure-norm support,  $a_g$  *decreases* with  $g$ , so the singular values

$$\sigma_g \sim \exp(-2\kappa_{\text{clos}} a_g) \text{ run } e^{(-6\kappa)}, e^{(-4\kappa)}, e^{(-2\kappa)} \text{ for } g = 1, 2, 3 \text{ — increasing,}$$

giving  $m_1 < m_2 < m_3$  and the consecutive ratio  $\sigma_{\{g+1\}}/\sigma_g = e^{(2\kappa)}$  at unit depth-step, i.e.  $M_g \propto e^{(2\kappa g)}$  (§7.4). (The factor 2 is from left and right states both overlapping the support; it is the same 2 that makes the per-generation factor  $e^{(2\kappa)} = 196$ , not  $e^{(\kappa)}$ .) This is why the inhomogeneous reading is the right one structurally, not merely the consistent one: each mode's block is rank-one and the three sum to the  $Y = \sum_g \lambda_g a_g b_{g\dagger}$  structure of Proposition 6.2', welding §7 to the spectral decomposition. The separation-only form of §7.1 does not factorize this way and is disconnected from §6.2. Generation depth plays, structurally, the role of a discrete extra dimension in flavour space, with the closure-norm support as the brane — the split-fermion / wavefunction-localization lineage, an understood and partially-successful class of models. (The translation-invariant reading would forbid exactly the preferred origin this mechanism needs, which is the §7.1 point.)

A note on  $\kappa_{\text{clos}}$ , now fixed by inheritance. The depth-decay rate  $\kappa_{\text{clos}}$  is no longer a free structural constant; the companion papers derive it, and tie it to the same structural object — the minimal closure algebra  $\dim(\Gamma_{\text{min}}) = 14$  — that controls the generation count elsewhere in the corpus. The featured route is the closure-loop count: the Structural Completeness companion proves  $\dim(\Gamma_{\text{min}}) = 2(n_{\text{int}} + n_{\text{role}}) = 2(3 + 4) = 14$  with  $(n_{\text{int}}, n_{\text{role}}) = (3, 4)$  the unique admissible pair, and the hierarchy-exponent companion derives

$$\kappa_0 = \ln(N_{\text{loop}}) = \ln(14) \approx 2.639,$$

with a first-order additivity theorem (cross-mode Hessian terms vanishing at the ground configuration) justifying the underlying  $\varepsilon_{\text{loose}} = N_{\text{loop}} \cdot \varepsilon_{\text{spread}}$  step rather than asserting it. The bare loop-count prediction for the per-generation factor is therefore  $e^{(2\kappa_0)} = 14^2 = 196$  — explicitly *not* the final observed value. The observed first ratio is  $M_{\mu}/m_e \approx 206.8$ , so the required correction is  $206.8/196 \approx 1.055$ . This factor is the bounded inter-sector weight asymmetry  $w$  of the companion, not a new fitted hierarchy mechanism: applied to the ratio it is  $w \approx 1.055$ , equivalently absorbed into the exponent as  $\kappa = \ln(14 \cdot w_{\kappa})$  with  $w_{\kappa} \approx 1.027$  ( $14 \cdot 1.027 \approx 14.39$ ,  $14.39^2 \approx 207$ ). The companion treats  $w$  as bounded but not yet derived from microdynamics; it is recorded here at its own epistemic status.

A convergent companion route reaches nearly the same value independently:  $CP^2$  geodesic focusing gives  $\kappa = K_{\text{eff}} = K_{\text{hol}} \cdot (n-1)/n = 4 \cdot 2/3 = 8/3 \approx 2.667$ , and  $e^{(8/3)} \approx 14.39 \approx 14 = \dim(\Gamma_{\text{min}})$ . So the focusing exponent and the loop count are two near-coincident readings of one structural scale —  $e^{(8/3)}$  rounds to the loop integer 14 — and are not yet formally unified. The  $8/3$  value happens to match the empirical  $\frac{1}{2} \ln(M_{\mu}/m_e) \approx 2.666$  more closely (to 0.03%), but it is featured here only as a convergence check rather than the primary route, because the  $\ln(14)$  route carries the deeper structural unity (it is  $\dim(\Gamma_{\text{min}})$ , the same integer that yields the no-fourth-generation theorem).

Consequence 7.3 therefore inherits a *derived* exponent,  $\kappa_{\text{clos}} \approx \ln(14)$ , with the convergent reading  $\kappa_{\text{clos}} = 8/3$ . These are *conjectured* to be one structural scale read two ways —  $e^{(8/3)} = 14.39$  close to the loop integer  $14 = \dim(\Gamma_{\text{min}})$  — but they are two independent derivations (one combinatorial,  $2(3+4) = 14$  exact; one geometric,  $K_{\text{eff}} = 8/3$  exact) that are not yet formally unified (open problem #8). On that identification, the residual between them is precisely the  $w \approx 1.055$  of §7.4, in which case  $w$  is the integer-rounding artifact of a single underlying scale and not a separate dynamical effect; until the identification is proven,  $w$  is a residual between two independent derivations that agree to 3%. Either way  $w$  is not a free fitted dial. This promotes the exponent from conditional-with-free-constant to conditional-with-fixed-constant. The inheritance carries the companions' fencing: the loop-count route depends on the uniqueness of  $(n_{\text{int}}, n_{\text{role}}) = (3, 4)$  and the first-order additivity theorem; the focusing route is a theorem given the scale-homogeneity postulate and the Sector-Curvature Correspondence axiom ( $\kappa = K_{\text{eff}}$ ), with the supporting confinement tension  $T_{\text{conf}} = 2e^{(K_{\text{eff}})}$  a closure hypothesis. So Consequence 7.3 is now [conditional on 7.0, with  $\kappa_{\text{clos}} \approx \ln(14)$  derived in the companions] — not unconditionally proven, but no longer carrying a free constant.

A point of honesty about which route the paper features. Empirically, the focusing value  $8/3$  is favored:  $e^{(2 \cdot 8/3)} = 207.1$  against the observed 206.8 (0.03%), whereas the loop count gives  $e^{(2 \ln 14)} = 196$ , a 5% miss. The paper nonetheless features  $\ln(14)$  — on *structural*, not empirical, grounds: 14 is  $\dim(\Gamma_{\text{min}})$ , the same integer that yields the no-fourth-generation theorem and the generation count, so featuring it ties the hierarchy rate to the architecture rather than to a separately-fitted geometric constant. This is a defensible preference but it should be named for what it is: structurally motivated, empirically the worse fit, with the open question (#8) being whether the fundamental object is the integer 14 or the focusing value  $e^{(8/3)}$  that fits better.

## 7.4 The Charged-Lepton Hierarchy: Realization in the Companion Lepton-Mass Model

**Status [conditional realization, inherited from the Role-4/BCB companion papers].** The charged-lepton hierarchy is worked out explicitly in the companion Role-4/BCB papers, and this subsection records how they realize the singular-value cascade of Proposition 6.2'. The structural integer governing the hierarchy is  $\dim(\Gamma_{\min}) = 14$  (the closure-loop count), not 13: a 13-based rule of the form  $1 : (2K-1) \cdot 2^4 : (2K-1) \cdot 2^8 = 1 : 208 : 3328$  (with  $2K-1 = 13$ ) matches  $m_{\mu}/m_e \approx 206.8$  to  $\sim 0.6\%$  on a single fitted integer but mispredicts  $m_{\tau}/m_e \approx 3477$  by  $\sim 4.5\%$  and needs an undetermined correction — close but off by one, exactly because the structural object is 14, not 13. The Level-2 route is not a  $13 \times 16$  product but  $e^{(2\kappa)}$  with  $\kappa$  derived from the closure-loop count, plus threshold compression for the tau.

**The cascade is realized as exponential localization plus threshold suppression.** Identifying the charged-lepton masses with the dominant singular values of the overlap operator  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$  (Identification 5.1, Proposition 6.2'), the companion supplies two facts.

First, the *leading* hierarchy is the single-exponential law  $M_g \propto e^{(2\kappa g)}$ , with  $\kappa$  the inherited closure-loop exponent of §7.3. The bare loop-count value  $\kappa_0 = \ln(14)$  gives a per-generation factor  $e^{(2\kappa_0)} = 14^2 = 196$  — and this is 5% below the observed electron  $\rightarrow$  muon ratio 206.8, a gap worth owning rather than smoothing: featuring the structural integer 14 costs the precise fit (the focusing value  $8/3$  recovers it,  $e^{(2 \cdot 8/3)} = 207.1$ , but that is the demoted route — see §7.3). The residual is not a free fitted dial. The "bounded inter-sector asymmetry"  $w \approx 1.055$  is, to the quoted precision, the gap between the loop integer 14 and the focusing value  $e^{(8/3)} = 14.39$  (equivalently  $w_{\kappa} \approx 1.027$  inside the exponent,  $14 \cdot 1.027 \approx 14.39$ ). *If* the loop count and the focusing value are the same structural scale — the open identification of problem #8 — then  $w$  is the integer-rounding artifact of that identity rather than independent dynamics; until that is proven,  $w$  is the residual between two independent derivations that agree to 3%. Either reading keeps  $w$  from being a fitted parameter. This is the  $\sigma_2/\sigma_1$  step of the cascade: geometric scaling along absolute depth, the inhomogeneous exponential of §7.3.

Second, the muon  $\rightarrow$  tau step is *not* a second copy of the same factor, and this is exactly why a 13-based product overshoots if iterated. If the  $e^{(2\kappa)}$  factor applied uniformly, the second step would again be  $\approx 207$ ; the observed step is  $\approx 16.8$ . The companion attributes the compression to **threshold proximity**: the genus-2 (tau) sector sits just above its binding threshold ( $C_2 = 22$  vs  $C_{\text{crit}}(2) = 18.822$ ), so its rescaled sector eigenvalue  $\tilde{E}_2$  is strongly suppressed relative to  $\tilde{E}_0 \approx \tilde{E}_1$ , and  $M_{\tau}/M_{\mu} = (\tilde{E}_2/\tilde{E}_1) \cdot e^{(2\kappa)} \approx 0.085 \times 208 \approx 17.7$  (against 16.8, 4.9%; the 208 here is the companion's calibrated per-generation factor, §7.3, not the bare 196). In the overlap language of Proposition 6.2' this is a statement about the mode profiles: the third singular value is pulled down by near-threshold binding, so the cascade is *not* a clean geometric progression but a geometric progression modulated by a sector-eigenvalue (profile) factor — exactly the "profiles modulate the singular values" caveat of 6.2', here made concrete.

**Why the loop-count route is the right reading rather than a 13-based product.** A  $13 \times 16$  product and the companion give incompatible accounts of the second jump: the product iterated would make it another large factor needing a  $\sim 4.5\%$  correction, whereas the companion makes it

$e^{(2\kappa)}$  suppressed by a threshold factor of  $\sim 12$ . They cannot both be the mechanism, and the companion is far more developed: the hierarchy-exponent companion derives  $\kappa \approx \ln(14)$  from the closure-loop count with a first-order additivity theorem, and the Structural Completeness companion proves  $(n_{\text{int}}, n_{\text{role}}) = (3, 4)$  unique,  $\dim(\Gamma_{\text{min}}) = 14$  inevitable, and — crucially — the no-fourth-generation result as a *theorem* (the structural bound  $C_0 \leq 14\kappa'$  gives  $C_3 < C_{\text{crit}}(3)$  for all admissible  $\eta_A$ ), with the  $\kappa = 2.67$  selectivity confirmed numerically in the calibrated regime. The Yukawa-side statement is therefore: the charged-lepton singular-value cascade of Proposition 6.2' is realized by the companions as closure-loop exponential localization ( $\kappa_{\text{clos}} \approx \ln(14)$ , conjecturally the same scale as the focusing value  $8/3$  via  $e^{(8/3)} \approx 14$ , open #8) times a threshold-proximity profile factor at the tau. This is a conditional realization — it carries the companions' fencing (the apparent weight asymmetry  $w \approx 1.055$  is the residual between 14 and  $e^{(8/3)}$ , a rounding artifact only if the two unify, §7.3; the Sector-Curvature Correspondence for the  $8/3$  reading; the tau's threshold proximity  $C_2 \approx 22$  still a calibrated input, not yet derived from Role-4 dynamics) — but it is a *consistent* one, and a worked companion mechanism for the very cascade Proposition 6.2' introduced.

## 7.5 The Closure-Mode Spectrum and the Source of Hierarchy

The realization of §7.4 invites a reframing of the whole flavour question, and a structural reason for the non-uniformity of the hierarchy. Both are conjectural and flagged as such; the second is constructed to *unify* with the companion's threshold-proximity mechanism rather than to compete with it.

**Closure-Mode Spectral Conjecture [conjectural].** Take the spectral decomposition of the mediator from Proposition 6.2',

$$\hat{H}_c = \sum_n \lambda_n |h_n\rangle\langle h_n|, \lambda_n \geq 0,$$

and read the eigenvalues physically: each  $\lambda_n$  is an independent **closure-maintenance cost** — the energetic load of sustaining one independent mode of committed distinguishability against the void. This reading is not arbitrary: since  $\hat{H}_c$  is defined as the positive operator governing closure-stabilization energy, its eigenvalues are precisely the energetic costs of maintaining independent closure modes, so "cost" is what its spectrum already means rather than an interpretation imposed on it. The conjecture is that the observed fermion hierarchy arises primarily from the hierarchy of these eigenvalues, not from arbitrary Yukawa inputs: localization compression (the §7.3 exponential backbone,  $\kappa_{\text{clos}} \approx \ln 14$ ) supplies a smooth geometric envelope, while the spread of the  $\lambda_n$  — persistent distinguishability — supplies the dominant inter-generational separation. This reframes the Standard Model's unanswerable "why are the Yukawas so different?" into the VERSF-native "why are the closure-maintenance costs so different?", which is a question about substrate dynamics rather than about phenomenological constants. It is a conjecture, not a result: it presupposes that the  $\lambda_n$  are computable closure loads and that their spread, rather than the profiles, carries the hierarchy — neither yet derived.

**Conjecture 7.4 (Non-Uniform Hierarchy Growth) [unifies with threshold proximity].** The charged-lepton hierarchy is not generated by repeated identical generation steps. The first transition (electron  $\rightarrow$  muon) must establish new infrastructure — a new refinement-persistent

loop, the phase-coherence machinery to keep it coherent, and the refinement scaffolding that supports it. The second transition (muon  $\rightarrow$  tau) commits only the *additional* loop, reusing the phase-coherence and refinement infrastructure the muon already established. If the per-generation distinguishability load  $\mathcal{S}_g$  (script  $\mathcal{S}$ , to avoid collision with the singular values  $\sigma_n$  of §6.2') charges for newly-established infrastructure, the increment at the second step is therefore smaller than at the first:

$$\mathcal{S}_\tau - \mathcal{S}_\mu < \mathcal{S}_\mu - \mathcal{S}_e,$$

which gives  $m_\tau/m_\mu \ll m_\mu/m_e$  with no fine-tuning. Crucially, this is offered not as a rival to the companion's threshold-proximity account but as a candidate *physical cause* of it: a smaller incremental distinguishability load at the second step is exactly a smaller increment of effective well-capacity, which places the tau sector nearer its binding threshold  $C_{\text{crit}}(2)$  — the suppression of  $\tilde{E}_2$  the companion currently inputs as the calibrated value  $C_2 \approx 22$ . The conjecture is the identification *infrastructure-reuse*  $\Leftrightarrow$  *near-threshold binding*; if it holds, it derives the threshold proximity of open problem #8 from a structural reason (reuse) rather than a fit. It is conjectural precisely there: the identification is plausible but not proven, and the companion's  $C_2 \approx 22$  remains the operative input until it is.

**Hierarchy Dominance Conjecture [conjectural — explicitly not a theorem].** A fuller log-mass decomposition would separate the contributions to  $\ln m_D$  for a defect  $D$ ,

$$\ln m_D \approx \ln \mathcal{D}_D + \frac{1}{2} \ln \lambda_{\max}(H_D) - p \ln \ell_D + \mathcal{S}_D + \ln T_D + \text{const},$$

with a transport amplitude  $\mathcal{D}_D$ , a closure-Hessian stiffness  $\lambda_{\max}(H_D)$ , a localization length  $\ell_D$  (the  $-p \ln \ell_D$  term is the geometric backbone,  $\approx 2\kappa g$  of §7.3), a persistent-distinguishability load  $\mathcal{S}_D$ , and a transport factor  $T_D$ . The conjecture is that  $\mathcal{S}_D$  dominates the inter-generational separation — that the charged-lepton hierarchy is principally an *informational* hierarchy (heavier generations are costly because the substrate must maintain more distinguishable committed structure) rather than a purely geometric one. This is appealing and would tie the mass sector to closure-maintenance complexity, but it is labelled a conjecture, not a theorem, for three honest reasons. (i) The five-term decomposition is not derived here; the companions establish only the two-term form  $\ln M_g = \ln \tilde{E}_g + 2\kappa g + \text{const}$ , and the extra terms ( $\mathcal{D}_D, \lambda_{\max}(H_D), T_D$ ) require independent definition. (ii) The dominance claim currently traces to the log-share of the 13-based reconstruction, which is not a sound basis; it needs an independent derivation. (iii) It must reduce to the companion's two-term form, with  $\mathcal{S}_D$  accounting for the  $\ln \tilde{E}_g$  variation that the companion attributes to threshold proximity — i.e. it must be reconciled with Conjecture 7.4, not stated alongside it. Until (i)–(iii) are met, "the hierarchy is informational, not geometric" is a conjecture about where the dominant term lives, not an established decomposition. What it predicts, stated concretely so the reader has something to hold: if the conjecture is correct, heavier generations correspond to larger persistent distinguishability loads — mass is a measure of how much committed, distinguishable structure the substrate must maintain.

## 7.6 Synthesis: Mass as Closure-Maintenance Cost

The pieces assembled across this paper and its companions are not three separate observations but one chain, and it is worth stating the chain in a single line because the through-line is the actual contribution. Each link is a result already established or conjectured above; this section adds no new claim, only the connection.

**The chain.** Mass is a closure-maintenance cost  $\rightarrow$  Yukawas are overlap operators  $\rightarrow$  the hierarchy is a closure-spectrum hierarchy.

Read link by link, with each pointing to where it lives:

- **Mass  $\leftrightarrow$  overlap (Identification 5.1, §5).** A fermion mass term is the unique substrate object that bridges left- and right-handed sectors — the scalar closure-norm mode (§4) — so the Yukawa matrix element *is* the closure-depth overlap  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$ . This is the established link: where flavour structure resides.
- **Overlap  $\leftrightarrow$  spectrum (Proposition 6.2', §6.2).** Expanding the mediator spectrally,  $Y = \sum_n \lambda_n a_n b_n^\dagger$ , the mass hierarchy is the image of the closure-mode spectrum  $\{\lambda_n\}$ , modulated by the mode profiles. This converts "the masses" into "the spectrum of  $\hat{H}_c$ " — the hierarchy is a closure-spectrum hierarchy. Established as structure; the values await the spectrum.
- **Spectrum  $\leftrightarrow$  maintenance cost (Closure-Mode Spectral Conjecture, §7.5).** Reading each  $\lambda_n$  as an independent closure-maintenance cost — the load of sustaining one mode of committed distinguishability against the void — completes the chain: light fermions are inexpensive closure modes, heavy fermions expensive ones, and the fermion hierarchy becomes a hierarchy of closure-maintenance complexity. This link is the conjectural one, and the paper marks it so.

**What the chain buys, at the rate and the shape.** The companions supply the two quantitative facts that turn the chain from a slogan into a near-realization for the charged leptons, and both are already stated above rather than re-derived here. The *rate* of the hierarchy is inherited, not fitted: the per-generation factor is the closure-loop count,  $\kappa_{\text{clos}} \approx \ln(14)$  with  $e^{(2\kappa_0)} = 196$  from the loop integer, the gap to the observed 206.8 being the 14-vs- $e^{(8/3)}$  residual  $w \approx 1.055$  (a rounding artifact if the two routes unify, open #8; §7.3), and the same  $\dim(\Gamma_{\text{min}}) = 14$  fixes the generation count via the companion no-fourth-generation theorem — so one structural integer controls both how heavy each generation is and how many there are. The *shape* — the strong asymmetry  $m_\tau/m_\mu \ll m_\mu/m_e$  — is the companion's threshold proximity, for which Conjecture 7.4 (§7.5) offers infrastructure-reuse as the candidate structural cause (the first generation transition builds the closure infrastructure; later generations reuse rather than rebuild it, so  $\Delta\mathcal{S}_{\{e \rightarrow \mu\}} > \Delta\mathcal{S}_{\{\mu \rightarrow \tau\}}$ , placing the tau near its binding threshold). The reuse reading is conjectural and is offered as the *cause* of threshold proximity, not a rival to it.

**The reframing this licenses.** The flavour problem is therefore not "why are the Yukawa couplings different?" but "why do closure modes require different stabilization costs?" — a question about substrate dynamics rather than about phenomenological constants. That identifies the closure-spectrum problem — computing  $\{\lambda_n\}$  and the distinguishability loads from substrate dynamics — as the natural next target of the mass programme, and it is the point at which this paper joins the gauge-sector and matter-sector work into a single line: mass is closure-

maintenance cost, Yukawas are its overlap expression, and the hierarchy is the spectrum of that cost. The synthesis is a statement of where the programme is going, fenced at exactly the link (cost-interpretation of  $\{\lambda_n\}$ ) that remains a conjecture; the two earlier links are established, and the rate is inherited.

---

## 8. Relation to Mass Hierarchies and Mixing

This paper does not derive the fermion mass hierarchy or the mixing angles. It locates where they must originate and constrains their form:

- Any future derivation of  $m_e$ ,  $m_\mu$ ,  $m_\tau$ , or the quark masses must proceed through the overlaps  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$ , subject to the assumption-free bound of Theorem 6.1, the rank pattern of Proposition 6.2, the spectral expansion  $Y = \sum_n \lambda_n a_n b_n^\dagger$  of Proposition 6.2', the parameter count of Proposition 6.4, and the mass-squared budget of Proposition 6.7. By Proposition 6.2' the hierarchy is the image of the closure-mode spectrum  $\{\lambda_n\}$ , so the matter-sector hierarchy reduces to the eigenvalue problem for  $\hat{H}_c$  on the depth sectors — the natural subject of the next paper.
- Any future derivation of CKM or PMNS structure must arise from diagonalizing the overlap matrix; the *qualitative* ordering of mixing suppression is already predicted (§7.2) under depth-locality, while the numerical angles belong to the dedicated flavour-mixing programme and are [open] here.
- The hierarchy question is sharply posed: how do the generation-depth sectors and the closure-norm operator determine the overlaps, and is the matrix single-mode-dominated (rank-one leading order) or full-rank? The depth-decay exponent is now fixed ( $\kappa_{\text{clos}} \approx \ln(14)$ , inherited from the companion closure-loop count, convergent with  $8/3$ ), so the open part is the profile factor — in the charged-lepton case, the threshold-proximity suppression at the tau that the companion currently inputs rather than derives (and that §7.5's Conjecture 7.4 conjectures a structural cause for).

This converts vague deferral into a structured set of constraints and open computations.

---

## 9. Why This Matters

The Standard Model treats Yukawa couplings as primitive. VERSF relocates them:

- Fermions = Persistent Fold Defects.
- Generations = refinement-stable closure-depth sectors.
- Higgs field = closure-norm mode (and, by §4, the unique identified substrate scalar mediator).
- Yukawas = closure-depth overlap amplitudes  $Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle$ .

- Mass scale = closure-norm equilibrium; mass structure = overlaps; mass-squared budget = total closure support.
- Gauge representation = PFD class; flavour = depth — orthogonal substrate coordinates.
- CKM / PMNS = diagonalization of the overlap matrix.

The flavour sector is moved from a table of fundamental constants to substrate geometry — and the move now carries content of two clearly-separated grades. With no further assumption it forbids specific textures (Theorem 6.1), fixes the leading-order rank pattern and makes the hierarchy the image of the closure-mode spectrum (Propositions 6.2, 6.2'), explains hierarchical smallness without tuning (Proposition 6.3), fixes the parameter count (Proposition 6.4), separates flavour from gauge structure (Result 6.5, conditional on the dictionary), and fixes the mass-squared budget (Proposition 6.7). Under the separately-flagged depth hypothesis it further predicts the qualitative hierarchy and mixing ordering (§7). The picture is no longer only a relocation; it is a constrained one — theorems where the assumptions allow, conditional results clearly marked as such.

## 10. Open Problems and Inheritances to Confirm

### Open (requiring new computation):

1. Exact generation-depth dynamics and the explicit depth states  $|\gamma_i\rangle$ .
2. Numerical mass-hierarchy derivation from the overlaps.
3. CKM matrix derivation (numerical angles and phase).
4. PMNS matrix derivation.
5. The explicit closure-norm operator  $\hat{H}_c$ , its matrix elements, and its rank (single-mode-dominated or full).
6. The substrate origin of the left/right chirality-sector distinction — the premise on which the mediation rests, and where a substrate account of Yukawas would be most novel.

### Open (inheritances to confirm against prior papers):

7. The closure-norm equilibrium value as the substrate origin of  $v$  (§5).
8. The depth-decay exponent is now fixed —  $\kappa_{\text{clos}} \approx \ln(14)$  from the closure-loop count (convergent with  $8/3$  via  $e^{8/3} \approx 14$ ), inherited from the companion papers — so the free-constant half of this problem is closed. What remains is on the companion side: resolving whether the fundamental structural object is the integer loop count 14 or the focusing value  $e^{8/3} = 14.39$  (their gap is the apparent  $w \approx 1.055$ , i.e. the  $196 \rightarrow 207$  correction is integer-rounding, not a separate parameter to fit); deriving the tau's threshold proximity ( $C_2 \approx 22$  just above  $C_{\text{crit}}(2) = 18.822$ , currently a calibrated input) from Role-4 dynamics, which is what fixes the profile factor modulating the §7.4 cascade; and discharging the companions' own open items (the confinement tension  $T_{\text{conf}}$ , the independent  $\varepsilon_{\text{bit}}$  / absolute-mass-scale determination). Closing these would convert the charged-lepton cascade of Proposition 6.2' from a conditional realization into a derivation.

9. The PFD-dictionary assignment of gauge representation to defect class (Result 6.5).
10. The precise citation for the three-generation count (§3.2) is now the Structural Completeness companion's uniqueness theorem  $((n_{\text{int}}, n_{\text{role}}) = (3, 4), \dim(\Gamma_{\text{min}}) = 14)$  and no-fourth-generation theorem; this item is closed pending only final cross-referencing.

These constitute the next stage of the matter-sector programme.

---

## 11. Conclusion

This paper does not derive the Yukawa matrices. It identifies the substrate structure that corresponds to them and proves what that identification forces.

The closure-norm mode is the unique identified substrate scalar that can bridge left- and right-handed sectors, so the Yukawa operator is identified with the closure-depth overlap amplitude

$$Y_{ij} = \langle \gamma_{L,i} | \hat{H}_c | \gamma_{R,j} \rangle,$$

with the mass scale set by the closure-norm equilibrium and the flavour structure by the overlaps. The reading carries theorems, with the conditional results held separate. From positivity of the mediator alone, the Closure-Overlap Bound forbids large mixing between weakly-coupled sectors and seeds the  $\sqrt{(\text{mass-ratio})}$  texture pattern; its equality case, the Single-Mode Closure Limit, makes the Yukawa matrix rank one when one closure mode dominates, and the general spectral expansion  $Y = \sum_n \lambda_n a_n b_n^\dagger$  makes the mass hierarchy the image of the closure-mode spectrum  $\{\lambda_n\}$  — reducing the hierarchy problem to the eigenvalue problem for  $\hat{H}_c$ ; the framework fixes the Standard Model parameter count, ties the total mass-squared budget to the total closure support via  $\sum m_i^2 \propto \text{Tr}(Y^\dagger Y)$ , explains hierarchical smallness without tuning, and — through the PFD dictionary — separates gauge representation (carried by class) from flavour (carried by depth) as orthogonal substrate structures. These hold with no further assumption. Under a separately and explicitly flagged depth-locality hypothesis, generation separation grades the off-diagonals — giving banding and the observed first-to-third-generation mixing suppression — while the mass hierarchy is the distinct, inhomogeneous statement that there is one rank-one localization block per generation depth, the three summing to  $Y = \sum_g \lambda_g a_g b_g^\dagger$  with  $\sigma_g \sim \exp(-2\kappa a_g)$  — the rank-three spectral form of Proposition 6.2', the split-fermion mechanism, reproducing  $M_g \propto e^{(2\kappa g)}$ .

The numerical masses, angles, and matrix elements remain open, as do two prior-paper inheritances (the closure-norm equilibrium value and the PFD class/gauge assignments); the depth-decay exponent  $\kappa_{\text{clos}} \approx \ln(14)$  is now inherited from the companion closure-loop count (convergent with the  $\text{CP}^2$ -focusing value  $8/3$ ). The significance is structural rather than phenomenological — but it is a *constrained* structure, now carrying theorems: Yukawas are closure-depth overlap operators, that reading forbids textures, fixes the rank pattern and parameter count, separates flavour from gauge structure, predicts the qualitative hierarchy, and

locates both halves of fermion mass — scale and structure — in substrate geometry. The flavour sector is relocated, and the relocation has teeth.

The present work identifies Yukawa structure with closure-depth overlap geometry and conjectures that the dominant source of hierarchy is persistent distinguishability rather than generation depth alone. Combined with the substrate-counting results of the charged-lepton programme — the closure-loop exponent  $\kappa \approx \ln(\dim \Gamma_{\min})$  and the no-fourth-generation theorem — this suggests, if the dominance conjecture of §7.5 can be established, that fermion masses are fundamentally measures of closure-maintenance complexity. The remaining task would then not be to fit Yukawa couplings but to derive the closure-mode spectrum  $\{\lambda_n\}$  and the distinguishability loads directly from substrate dynamics. If that succeeds, the Standard Model flavour sector would be transformed from a collection of phenomenological constants into a computable consequence of closure geometry — and this paper, with its companions, is the bridge from the gauge and matter sectors into that flavour-and-mass programme. The conditional mood is deliberate: the bridge is built and load-bearing where the theorems reach, conjectural where they do not, and the paper keeps the two apart.

---

## References (programme)

The load-bearing inputs are inherited from companion VERSF works: matter as Persistent Fold Defects and the invariant tuple, and the assignment of gauge representation to defect class (the PFD / matter-sector dictionary papers); the three refinement-stable generation-depth sectors and the refinement-stability criterion (the generation-depth paper); the identification of the Higgs with the closure-norm mode and the closure-norm equilibrium value (the closure-norm / electroweak-sector paper); and the Role-4/BCB companion papers: the lepton-mass paper, which derives the localization scaling and the CP<sup>2</sup>-focusing exponent  $\kappa = 8/3$  (conditional on its Sector-Curvature Correspondence) and realizes the charged-lepton cascade of §7.4 via exponential localization and threshold-proximity suppression; and the hierarchy-exponent and Structural Completeness papers, which derive the featured exponent  $\kappa \approx \ln(N_{\text{loop}}) = \ln(14)$  from the closure-loop count (the inter-sector asymmetry  $w \approx 1.055$  being the residual to the focusing value  $e^{(8/3)} = 14.39$  — a rounding artifact if the two routes are the same scale, open #8), prove  $(n_{\text{int}}, n_{\text{role}}) = (3, 4)$  unique and  $\dim(\Gamma_{\min}) = 14$  inevitable, and establish the no-fourth-generation cutoff as a theorem ( $C_0 \leq 14\kappa'$ ). The kinship of §7 is to the wavefunction-localization (split-fermion) flavour literature; the texture relations of §6.1 are of the Gatto–Sartori–Tonin / Fritzsch lineage, with Froggatt–Nielsen relevant there as a texture analogy rather than as the §7 mechanism. These are external comparisons, not VERSF inputs.