The Substrate Stiffness Hierarchy in VERSF

Closure-Hessian Geometry, Localization Compression, Persistent Distinguishability, Interface Transport, and the Emergence of the Fermion Mass Spectrum

Keith Taylor — VERSF Theoretical Physics Programme

General reader summary

One of the strangest facts in physics is that particles that look almost identical can have wildly different masses. The electron is light. The muon is two hundred times heavier. The tau is twenty times heavier still. Quarks span an even wider range — the top quark outweighs the electron by hundreds of thousands of times. The Standard Model has names for these masses (the Yukawa couplings) and a mechanism by which the Higgs field hands them out, but it does not explain why nature chose these particular numbers rather than any others. The numbers are inserted by hand.

This paper attempts to move beyond that.

Earlier VERSF papers proposed that particles are not fundamental point objects but stable closure structures — Persistent Fold Defects, or PFDs — embedded in a deeper informational substrate. Those papers established how gauge structure, confinement, electroweak symmetry breaking, and the existence of three generations all emerge from closure geometry. The substrate Higgs paper showed that fermion masses take the form $m_D = \mathcal{D} \cdot S(D) \cdot v$, where $\mathcal{D}$ is the generation-depth operator (delivering generation count and ordering), v is the closure-condensate vacuum scale (the Higgs-like background), and $S(D)$ is a substrate stiffness factor (still to be computed) that determines the actual mass magnitudes. That same paper identified the computation of $S(D)$ as the central open problem of the programme.

The present paper opens that problem. The central proposal is that fermion mass is not a fundamental property but a measure of *stabilization complexity*: how energetically difficult it is for the substrate to maintain a given closure structure against admissibility decay and dispersive distinguishability flow. A heavier particle is not "made of more stuff" — it is a closure geometry that is harder for the substrate to hold coherently. Some closures are tightly compressed, strongly localized, and carry large persistent-distinguishability loads. Others are diffuse and easy to sustain. The energetic cost of preserving these structures against dispersal is what we observe macroscopically as mass.

The paper develops this idea by decomposing the substrate stiffness factor $S(D)$ into four physically distinct contributions: a closure-Hessian curvature term S_H (how steep the local closure free-energy minimum is at the PFD's support), a localization compression term S_L (how tightly the PFD must be squeezed against substrate dispersal), a persistent-distinguishability term S_P (how much committed informational bookkeeping the closure carries), and an interface

transport complexity term *S_I* (how many admissibility-preserving transport channels the closure structure must support to maintain refinement persistence). Each factor is constrained by a distinct substrate mechanism and depends on specific components of the PFD invariant tuple.

The combined scaling law turns the fermion mass hierarchy from a list of arbitrary inputs into a calculational problem in closure geometry. The mass hierarchy becomes the logarithmic sum of these four substrate-complexity contributions — exponential in additive structural quantities, in the same way that mass gaps in lattice gauge theory are exponential in coupling, or Boltzmann factors in energy. Quark/lepton separation, neutrino suppression, and the relative sizes of CKM and PMNS mixing all fall out of differences in how the four factors scale across PFD classes.

The paper does not compute exact fermion masses from substrate primitives. What it does is convert the hierarchy problem from "twelve mysterious phenomenological numbers" into "four constrained sub-factors with specific substrate mechanisms and specific dependences on the PFD invariant tuple, computed independently and combined." That conversion is the structural step the programme has been pointing toward since the matter paper specified the four-contribution mass scaffold without computing the contributions; the dictionary classified the PFDs without assigning mass values; and the Higgs paper supplied the energetic mechanism without explicit calculation.

What remains is the calculation itself, and the calculation now has a shape.

Contents

Abstract

We develop the substrate-level account of the fermion mass hierarchy within the Void Energy-Regulated Space Framework (VERSF), opening the central open problem identified by the closure-norm condensation paper: the computation of the substrate stiffness factor *S(D)* from substrate primitives. The closure-norm condensation paper established the mass formula

$m_D = \mathcal{D} \cdot S(D) \cdot v$

where the generation operator $\mathcal{D} = diag(1, 2, 4)$ delivers generation count and ordering, v is the closure-condensate vacuum scale (the empirical electroweak scale $v_{SM} \approx 246\ GeV$), and *S(D)* is the substrate stiffness factor that determines the empirical mass magnitudes. The present paper supplies the structural decomposition of *S(D)* and the scaling architecture from which the observed hierarchy emerges.

The central thesis is that *fermion mass is the substrate energetic cost of maintaining a stable closure geometry against admissibility decay and dispersive distinguishability flow*. The substrate stiffness factor decomposes into four physically distinct contributions inherited from the PFD invariant tuple:

$S(D) = S_H(D) \cdot S_L(D) \cdot S_P(D) \cdot S_I(D)$

where S_H is the closure-Hessian stiffness (local energetic curvature at the PFD's support), S_L is the localization compression (energetic cost of maintaining the PFD's localization length against substrate dispersal), S_P is the persistent distinguishability load (Landauer-type cost of holding the PFD's committed folds in their admissibility-fixed state), and S_I is the interface transport complexity (cost of maintaining admissibility-preserving transport channels supporting refinement persistence). Each factor is structurally constrained by a distinct substrate mechanism with a specific functional dependence on components of the invariant tuple $\mathcal{I}(D)$.

The paper establishes seven principal structural results:

1. **The four-factor decomposition is canonical.** The substrate stiffness factor decomposes uniquely into four contributions corresponding to the four substrate mechanisms identified in the matter paper §6.5 mass scaffold and the closure-norm condensation paper §11.2a, with no residual ambiguity in the assignment of stabilization costs to factor categories.
2. **The localization compression contribution carries exponential structure.** From the Role-4 lepton programme's localization scaling $\ell_g = \ell_0 \cdot exp(-\kappa g)$ combined with the

Bessel localization mode normalization, $S\_L(D) \sim exp(p \cdot \kappa \cdot \gamma\_D)$ with $p > 0$ — naturally generating exponential generation hierarchy without parameter insertion.

3. **The persistent distinguishability contribution carries exponential structure forced by BC1.** From bit conservation (BC1) combined with commitment persistence and the substrate's natural admissibility-relaxation flow, $S\_P(D) \sim exp(\sigma\_D)$ — with Landauer's principle recovered as the thermodynamic limit, not assumed as a separate input. This provides a second exponential channel structurally distinct from localization, derived from substrate primitives rather than imported from conventional thermodynamics.
4. **The interface transport complexity contribution is graph-theoretic with concrete invariants.** The factor $S\_I(D) = \alpha_1\beta_1(G\_D) + \alpha_2\lambda_2^{-1}(G\_D) + \alpha_3\chi(G\_D) + \alpha_4 deg\_avg(G\_D)$ is determined by four computable invariants of the admissibility-preserving transport graph G_D — cycle complexity, spectral gap, branching number, and average connectivity. Each invariant has a definite substrate interpretation, with the four substrate-derived coefficients constituting concrete computational targets rather than schematic placeholders.
5. **The combined scaling law is additive in log-space as a leading-order EFT.** Taking logarithms of the four-factor decomposition gives $\ln m\_D \approx \ln \mathcal{D}\_{\gamma\_D} + (1/2) \ln \lambda\_max(H\_D) - p \cdot \ln \ell\_D + \sigma\_D + \ln T(D) + const$ (Theorem 9.1), with each independent substrate stabilization free-energy contribution adding linearly in log-space. The factorization itself is the leading-order substrate stiffness hierarchy theory; inter-sector corrections $\Delta\_\{ab\}$ are expected at higher order and are structurally computable from substrate primitives without introducing new free parameters. This frames the framework as a substrate EFT analogous in structure to perturbative expansions elsewhere in physics.
6. **Quark/lepton separation is structurally forced.** Quark PFDs carry *additional* contributions to S_I (multi-channel confinement transport) and S_P (colour-charge committed distinguishability) that lepton PFDs do not. The mass amplification factor for quarks relative to leptons of the same generation depth is therefore structurally fixed, even though its numerical value depends on the open calculation.
7. **Neutrino suppression is structurally forced.** Neutrino PFDs have $\ell\_D = 0$ (ledger charge), removing the dominant ledger-charge contributions to S_L, S_P, and S_I — producing tiny $S(\nu)$ and therefore tiny neutrino masses without fine-tuning.

Epistemic status. Results 1, 3, 4, 6, 7 are proven from inherited substrate primitives — Result 3 with Landauer derived from BC1 + commitment persistence, Result 4 with concrete graph-theoretic invariants. Result 5 is conditional on the leading-order EFT factorization (§4.1), under which inter-sector corrections are sub-leading and computable. Result 2 is structurally forced once Role-4 localization scaling is taken as input; the exponent p is structurally constrained positive but its precise value is open. The paper does *not* claim to derive: exact numerical mass values from substrate primitives (the §13.8 toy reproduces lepton ratios under hand-chosen substrate-complexity assignments, demonstrating consistency at the few-tenths-of-a-percent level pending substrate-derivation of the parameters), the precise functional form of inter-sector corrections $\Delta\_\{ab\}$, the running of stiffness factors with substrate scale, or the second-quantised reconstruction of substrate stiffness dynamics.

The paper completes the structural framework anticipated by the closure-norm condensation paper §17.3, reduces the fermion mass hierarchy problem from "twelve mysterious phenomenological numbers" to "four constrained sub-factors with substrate-derivable functional forms," and supplies a concrete seven-step computational roadmap (§13.6) for the explicit evaluation. A worked toy calculation (§13.8) with hand-chosen but *modest* substrate-complexity increments reproduces $m_\mu/m_e \approx 206.4$ (observed 206.8, 0.2% agreement) and $m_\tau/m_e \approx 3491$ (observed 3477, 0.4% agreement); the toy demonstrates feasibility, not derivation, but establishes that the empirical charged-lepton hierarchy is consistent with the four-factor framework under bit-count differentials of order a few committed bits per generation. The remaining computation — performing the seven steps for substrate-derived values — is the natural target of the next computational paper. The framework is further conjectured (§17) to be *universal* across all stable PFD sectors and across closure-substrate frameworks supporting committed-distinguishability defects, positioning it as a candidate universality class of substrate-derived fermion mass theories.

---

1. Introduction

The Standard Model contains the largest unexplained hierarchy in fundamental physics. The fermion masses span more than twelve orders of magnitude:

$$m_\nu \ll m_e \approx 0.511\ \text{MeV} < m_\mu \approx 106\ \text{MeV} < m_\tau \approx 1.78\ \text{GeV} < m_t \approx 173\ \text{GeV}$$

and the corresponding Yukawa couplings range from $y_e \approx 3 \times 10^{-6}$ for the electron to $y_t \approx 1$ for the top quark. Within the Standard Model, these values are phenomenological inputs. They are measured, not derived. No principle in the Standard Model selects them, predicts them, or relates them to other observables.

The hierarchy presents three distinct puzzles:

1. *Why does the hierarchy exist at all?* The Higgs mechanism provides a universal mass scale $v \approx 246\ \text{GeV}$. There is no obvious reason that different fermions should couple to that scale with strengths spanning twelve orders of magnitude rather than being clustered around the same value.
2. *Why does the hierarchy take its observed shape?* The mass ratios are not random. They exhibit systematic structure across generations and between sectors: charged leptons cluster below their corresponding up-type quarks, neutrinos are anomalously light, quark mass ratios are larger than charged-lepton mass ratios, and the heaviest member of each sector dominates by orders of magnitude.
3. *Why do mixing angles (CKM, PMNS) exhibit complementary hierarchical structure?* The off-diagonal generation mixing tracks the mass hierarchy in specific ways — small CKM angles for the quark sector, large PMNS angles for the lepton sector, suppression of non-adjacent generation mixing.

The VERSF programme has progressively addressed several adjacent questions. The matter paper established Persistent Fold Defects (PFDs) as the substrate ontology of matter and identified four contributions to PFD mass (commitment-density loading, closure-Hessian stiffness, confinement/localization cost, persistent distinguishability content) in its §6.5 mass scaffold, but did not compute the contributions. The PFD–Standard Model Dictionary classified PFD invariant tuples onto Standard Model representation classes, with §16 specifying three generations from refinement persistence (P3) and §19.3 reformulating mass as closure cost, but did not assign mass values. The closure-norm condensation paper established the mass formula $m_D = \mathcal{D} \cdot S(D) \cdot v$ with the closure condensate providing the universal scale v via the Higgs mechanism, decomposed $S(D)$ into the four substrate contributions in its §11.2a, and identified the explicit computation of $S(D)$ as *the* central open problem of the programme (its §17.3).

The present paper takes that problem up.

The central thesis is:

Fermion mass is the substrate energetic cost of maintaining a stable closure geometry against admissibility decay and dispersive distinguishability flow. The substrate stiffness factor S(D) decomposes into four structurally distinct contributions — closure-Hessian curvature, localization compression, persistent distinguishability load, and interface transport complexity — each constrained by a specific substrate mechanism and depending on specific components of the PFD invariant tuple.

The paper is companion to three previously established works. The *closure-norm condensation paper* established the mass formula and identified the four substrate contributions whose product is $S(D)$; this paper develops each contribution. The *matter paper* established the four-contribution mass scaffold whose substrate mechanisms the present paper makes structurally precise. The *Role-4 lepton programme* established the exponential localization scaling $\ell_g = \ell_0 \cdot \exp(-\kappa g)$ with $\kappa \approx 3/8$ from CP^2 geometric structure and confinement dynamics; this localization scaling enters as the substrate input for the S_L term developed here in §6.

The paper is organised as follows. §2 collects inherited substrate architecture. §3 states the central thesis precisely. §4 introduces the four-factor decomposition of $S(D)$. §5–§8 develop each factor in turn, identifying its substrate mechanism, its dependence on components of $\mathcal{I}(D)$, and its functional form. §9 assembles the combined scaling law and states the central structural theorem. §10–§12 work through the structural consequences for quark–lepton separation, neutrino suppression, and flavour mixing. §13 provides an empirical consistency check against the observed lepton hierarchy. §14 places the result within the broader VERSF programme. §15–§16 collect falsification criteria, epistemic status, and open problems. §17 concludes.

---

2. Inherited substrate architecture

The substrate stiffness hierarchy developed here builds on a stack of inherited results. This section summarises only what is needed; full derivations are in the upstream programme.

2.1 Persistent Fold Defects and the invariant tuple

Matter sectors are interpreted as stable closure-defect topologies in the committed distinguishability graph $\Sigma = (F, E)$, where F is the set of committed folds and E is the set of admissible committed relations among them. A Persistent Fold Defect (PFD) is a connected closure structure $D \subset \Sigma$ satisfying:

(P1) nontrivial closure topology, $\beta_1(D) \geq 1$, with at least one admissibility-fixed homology class; (P2) nontrivial closure holonomy on at least one non-contractible loop; (P3) refinement persistence; (P4) positive closure stability, $\delta^2 S_{int}[D] \succ 0$.

Its internal structure is encoded by the invariant tuple

$$\mathcal{I}(D) = (C_D, \beta_1(D), h_D, \pi_D, \chi_D, \gamma_D, \ell_D, \rho_D)$$

with components as specified in the dictionary §2.1 and the closure-norm condensation paper §2.1: C_D = closure completeness, $\beta_1(D)$ = first Betti number, h_D = holonomy class, π_D = orientation parity, χ_D = chirality, γ_D = generation depth $\in \{1, 2, 3\}$, ℓ_D = ledger charge vector, ρ_D = confinement requirement. The present paper takes the invariant-tuple structure as fixed input and develops $S(D)$ as a function of these components.

2.2 The generation operator $\mathcal{D}$

The flavour-mixing programme derives the generation-space stiffness operator

$$\mathcal{D} = diag(1, 2, 4)$$

acting on the three-dimensional generation space. $\mathcal{D}$ fixes two structural facts: there are exactly three generations ($\mathcal{D}$ has three eigenvalues, with $\gamma_D \leq 3$ from refinement persistence), and the generation ordering is $m_{gen\ 1} < m_{gen\ 2} < m_{gen\ 3}$ (the eigenvalues are strictly increasing). What $\mathcal{D}$ does *not* deliver are the empirical mass-ratio magnitudes; the linear-in-$\mathcal{D}$ contribution alone would give $m_2/m_1 = 2$ and $m_3/m_1 = 4$, whereas the empirical charged-lepton ratios are 207 and 3477. The empirical magnitudes require the substrate stiffness factor $S(D)$ developed below.

2.3 The closure-norm condensate and the vacuum scale v

The closure-norm condensation paper established that the substrate vacuum carries a closure condensate $\langle\rho\rangle = v \approx v_{SM} = 246$ GeV (in standard convention), with the Higgs boson identified as the radial closure fluctuation around this condensed background. Fermion masses arise from PFD-class-specific coupling to the condensate:

$$m_D = g_D \cdot v$$

with $g_D = S(D)$, the substrate stiffness factor. The condensate scale v provides the universal multiplicative factor in m_D; all PFD-class differentiation lives in $S(D)$. The present paper develops $S(D)$.

2.4 The mass formula in expanded form

Combining §2.2 and §2.3 with the closure-norm condensation paper §11.2:

$$m_D = \mathcal{D}_{\gamma_D} \cdot S(D) \cdot v$$

This formula is the upstream architectural input to the present paper. $\mathcal{D}$ delivers generation count and ordering; v delivers the universal scale; $S(D)$ delivers the empirical magnitudes. The paper develops $S(D)$.

2.5 Bessel localization modes

The closure-field programme established that spatially localized closure-field excitations satisfy modified Bessel structures

$$f(r) = K_n(\kappa r)$$

with characteristic decay rate κ and winding number n. The energy density of such a localized mode increases as κ^{-1} (the coherence radius) decreases. This is the substrate input that supplies the energetic cost of localization compression developed in §6.

2.6 Role-4 localization scaling

The Role-4 lepton programme established that PFD localization lengths scale exponentially with generation depth:

$$\ell_g = \ell_0 \cdot \exp(-\kappa g)$$

with $\kappa \approx 3/8$ derived from CP^2 geometric structure and Role-4 confinement dynamics (not an adjustable parameter). Higher-generation PFDs are progressively more deeply localized — the muon's closure structure is compressed relative to the electron's, the tau's relative to the muon's. This exponential localization scaling is the substrate input that gives the localization compression factor its exponential structure in §6.

2.7 Position relative to the closure-norm condensation paper

The closure-norm condensation paper §11.2a introduced the four-factor decomposition

$$S(D) = S_H(D) \cdot S_L(D) \cdot S_P(D) \cdot S_I(D)$$

with schematic functional forms for each factor and an explicit acknowledgement that the explicit substrate-level derivation of each factor — and its evaluation for specific PFD classes —

was deferred to a companion paper. The present paper is that companion paper. The structural decomposition is inherited here; the substrate-mechanism derivation of each factor is developed in §5–§8; the assembly of the combined scaling law is the substantive contribution of §9.

The closure-norm condensation paper §11.2b also stated a *scaling conjecture*:

$$\ln m_D \sim a \cdot \gamma_D + b \cdot \ell_D^{-1} + c \cdot \sigma_D + d \cdot T(D) + \text{const.}$$

This conjecture is taken up here as a structural prediction of the four-factor framework and developed into a theorem (§9) under stated substrate assumptions. The §11.2b form is *qualitative* — it captures the structural reading that smaller ℓ_D and larger $T(D)$ increase the mass without claiming to be the proper log-space expansion of the four-factor decomposition. The proper expansion, supplied by Theorem 9.1, is $\ln m_D \approx \ln \mathcal{D}_{\gamma_D} + (1/2) \ln \lambda_{max}(H_D) - p \cdot \ln \ell_D + \sigma_D + \ln T(D) + \text{const}$ — additive in log-space across the four substrate stabilization free-energy contributions. The two forms agree qualitatively but differ in the precise functional dependence on ℓ_D (logarithmic, not inverse-power) and $T(D)$ (logarithmic, not linear); §9 takes Theorem 9.1's form as the substantive result.

---

3. The central thesis

Before developing the four factors, we state the central thesis precisely. In the closure-condensed substrate phase, every PFD class D experiences the same universal closure-density background v. What differs across PFD classes is the *substrate energetic cost* of holding the PFD's specific closure geometry in its admissibility-fixed configuration against perturbative decay.

This stabilization cost is what we identify with mass.

Central thesis. The fermion mass m_D of a PFD class D is the substrate free-energy cost per unit time of maintaining the PFD's closure structure against admissibility decay and dispersive distinguishability flow, evaluated in the closure-condensed phase. This cost is $m_D = \mathcal{D}_{\gamma_D} \cdot S(D) \cdot v$, with v the universal condensate scale, $\mathcal{D}$ the generation-depth amplifier, and $S(D)$ the PFD-class-specific stabilization complexity factor.

This identification is not an ad hoc reframing. It is the natural reading of the matter paper §6.5 mass scaffold ("mass is the energetic cost of maintaining persistent closure") combined with the closure-norm condensation paper §10.4 equivalence between the closure-condensate-induced mass $m_D = g_D \cdot v$ and the dictionary §19.3 closure-cost mass $m_D \sim E_{closure}(D)$. The two formulations are the same substrate statement viewed from different angles: closure-condensate coupling and substrate stabilization cost are the same quantity.

The four-factor decomposition of $S(D)$ developed in §4–§8 is the structural mechanism by which this stabilization cost varies across PFD classes. The combined scaling law of §9 is the resulting hierarchy structure.

4. Structural decomposition of the substrate stiffness factor

We decompose the substrate stiffness factor into four physically distinct contributions, each corresponding to a distinct substrate mechanism by which a PFD's closure structure incurs stabilization cost:

$S(D) = S_H(D) \cdot S_L(D) \cdot S_P(D) \cdot S_I(D)$

where:

— $S_H(D)$: **closure-Hessian stiffness** — the local energetic curvature of the substrate closure free-energy at the PFD's support, measuring how steeply the substrate resists perturbations of the PFD's closure-norm profile;

— $S_L(D)$: **localization compression** — the energetic cost of maintaining the PFD's localization length ℓ_D against substrate dispersal, measuring how tightly the PFD must be squeezed against the condensate background;

— $S_P(D)$: **persistent distinguishability load** — the Landauer-type cost of holding the PFD's committed folds in their admissibility-fixed state, measuring the informational bookkeeping required to maintain the PFD's distinguishability content;

— $S_I(D)$: **interface transport complexity** — the cost of maintaining admissibility-preserving transport channels supporting refinement persistence (P3), measuring the topological complexity of the closure transport network the PFD must support.

Each factor maps to one of the four matter-paper §6.5 mass contributions and to one of the four closure-norm condensation paper §11.2a factors. The assignment is canonical:

Matter paper §6.5 contribution	Closure-norm paper §11.2a factor	Present paper §4 factor
(i) Commitment-density loading	(absorbed into v)	(absorbed into v)
(ii) Closure-Hessian stiffness	$S_H(D)$	$S_H(D)$
(iii) Confinement/localization cost	$S_L(D)$, $S_I(D)$	$S_L(D)$, $S_I(D)$
(iv) Persistent distinguishability content	$S_P(D)$	$S_P(D)$

Contribution (i) (commitment-density loading) is *not* part of $S(D)$ — it is absorbed into the universal condensate scale v via the closure-norm condensation paper §10.4 Identification (a). Contribution (iii) (confinement/localization cost) splits naturally into two distinct substrate mechanisms — localization compression (S_L) and interface transport complexity (S_I) — that the matter paper treated as a single bullet point but that have structurally different functional forms and dependences on $\mathcal{I}(D)$. The four-factor decomposition is therefore not a refactoring of

the matter paper's scaffold but a sharpened version of it that separates compositionally distinct effects.

The decomposition is multiplicative because each factor $S_X(D)$ is a *dimensionless mass amplification* contributed by mechanism X — the effective Yukawa-coupling contribution from that mechanism, in the sense that $m_D = \mathcal{D} \cdot (S_H \cdot S_L \cdot S_P \cdot S_I) \cdot v$ with each factor amplifying the bare condensate scale v by a multiplicative substrate-derived factor. Independent multiplicative amplifications compound multiplicatively for the same reason that successive scattering amplitudes multiply: each represents a separate substrate response, and the total response is the product of the individual responses.

Equivalently — and this is the more illuminating framing — identify $\ln S(D)$ with the *substrate stabilization free-energy contribution* required to maintain the PFD's closure structure against admissibility decay. Then

$$\ln S(D) = \ln S_H(D) + \ln S_L(D) + \ln S_P(D) + \ln S_I(D)$$

expresses the additivity of *independent* free-energy contributions from the four mechanisms. The multiplicative decomposition of $S(D)$ and the additive decomposition of $\ln S(D)$ are the same statement viewed from different sides; Theorem 9.1's additive log-space scaling law is the direct expression of this free-energy additivity. This framing also tightens the link to §4.1's EFT structure: inter-sector corrections Δ_{ab} appear as cross-coupling terms in the substrate free-energy, expected at higher order in substrate-coupling parameters.

A clarification of scope: free-energy additivity explains *why* the decomposition is additive in log-space — it follows directly from independent stabilization mechanisms contributing additively to the substrate free-energy budget. But free-energy additivity by itself does *not* determine the *specific functional dependence* of each factor on the PFD invariant-tuple components — whether S_L scales as ℓ_D^{-p} with some specific p, whether S_P scales exponentially in σ_D, whether S_I depends graph-theoretically on $T(D)$. Those functional forms come from the substrate-mechanism derivations of §5–§8, each of which exhibits the specific substrate process responsible for the corresponding stabilization cost.

The four factors are developed in turn in §5–§8.

4.1 Leading-order factorization and higher-order coupling corrections

The multiplicative decomposition $S(D) = S_H \cdot S_L \cdot S_P \cdot S_I$ should be understood as the *leading-order* factorization of the substrate stiffness hierarchy, not as an exact substrate identity. At leading order, the four stabilization mechanisms correspond to distinct substrate sectors (local closure curvature, localization compression, persistent distinguishability bookkeeping, transport-network maintenance) whose energetic costs add independently to the PFD's stabilization budget. The leading-order approximation neglects cross-coupling terms between sectors.

More generally, the full stiffness functional may contain inter-sector interaction corrections:

$S\_full(D) = S\_H(D) \cdot S\_L(D) \cdot S\_P(D) \cdot S\_I(D) \cdot \exp(\Delta\_\{HL\} + \Delta\_\{HP\} + \Delta\_\{HI\} + \Delta\_\{LP\} + \Delta\_\{LI\} + \Delta\_\{PI\} + ...)$

where $\Delta\_\{ab\}$ represents inter-sector coupling corrections capturing how each pair of stabilization mechanisms back-reacts on the other. For example, $\Delta\_\{LP\}$ would capture the coupling between localization compression and persistent distinguishability — a more deeply localized PFD with large committed-bit count might pay an *additional* cost beyond the sum of its localization and distinguishability costs because the two structures interact through closure curvature.

The present paper studies the zeroth-order factorized structure. The existence of higher-order corrections is expected and does *not* invalidate the decomposition; it places the current framework in the role of the *leading-order effective substrate hierarchy theory*, analogous to how Standard Model perturbation theory is structured as an expansion in coupling constants around the free-field point. Inter-sector corrections $\Delta\_\{ab\}$ would contribute at higher order in substrate-coupling parameters that are themselves structurally constrained and computable from substrate primitives — they are not new free parameters.

This EFT-style framing has two important consequences. First, it explains why the four-factor decomposition is *useful* even though it is not exact: the leading-order structure captures the dominant hierarchy mechanisms, and corrections appear at sub-leading order with computable structure. Second, it identifies a natural calculational path forward: compute the leading-order factors first (the target of the next computational paper), then evaluate inter-sector corrections to assess where the leading-order approximation breaks down. The empirical 12-order-of-magnitude span of the fermion mass spectrum, if reproduced at leading order to within (say) 30%, would constitute strong evidence that the leading-order EFT captures the dominant substrate physics and that $\Delta\_\{ab\}$ corrections are sub-dominant.

(Structural; the leading-order factorization is the substantive assumption of the present paper. The existence and structure of inter-sector corrections $\Delta\_\{ab\}$ is open and is a target for the second-quantised stiffness reconstruction companion paper.)

---

5. Closure-Hessian stiffness $S\_H(D)$

5.1 Substrate mechanism

The closure free-energy functional $F[\rho]$ developed in the closure-norm condensation paper §4 defines, for each spatial point and each PFD's local closure-norm profile, a Hessian operator

$H\_D = (\delta^2 F / \delta \rho^2)\_D$

evaluated at the PFD's closure support. The eigenvalues of H_D measure the local energetic curvature: how steeply the substrate free-energy increases as the PFD's closure-norm profile is perturbed away from its admissibility-fixed configuration. A PFD with steeper local curvature is

more deeply locked into its closure minimum and requires more substrate stabilization energy to maintain.

5.2 Functional form

The closure-Hessian stiffness factor is

$S_H(D) \sim \sqrt{\lambda}_{max}(H_D)$

where λ_{max} denotes the largest eigenvalue of H_D restricted to the PFD's local neighbourhood. The square-root structure arises because stabilization energy scales as the substrate's natural frequency, which is $\sqrt{(stiffness/inertia)}$ for a quadratic minimum — exactly as harmonic-oscillator energy scales as $\omega = \sqrt{(k/m)}$ with k the stiffness.

For the structural framework here, what matters is that $S_H(D)$ is a positive scalar function of the PFD invariant tuple via the dependence of H_D on closure topology and holonomy structure:

— C_D (closure completeness): complete-closure PFDs sit in deeper free-energy minima than partial-closure PFDs; — h_D (holonomy class): nontrivial holonomy increases the local curvature by tethering the PFD's closure-norm profile to specific admissibility-fixed angular configurations; — $\beta_1(D)$ (Betti number): higher topological genus contributes additional stiffness from the additional admissibility-fixed homology classes the PFD must maintain.

5.3 Dependence on generation depth

A key question is how $S_H(D)$ depends on generation depth γ_D. The Role-4 lepton programme established that higher-generation PFDs sit in deeper closure minima — $\gamma_D = 2$ PFDs in deeper minima than $\gamma_D = 1$, $\gamma_D = 3$ deeper still. The closure-Hessian curvature $\lambda_{max}(H_D)$ should therefore *increase* monotonically with γ_D:

$\lambda_{max}(H_D)|\{\gamma_D = 2\} > \lambda_{max}(H_D)|\{\gamma_D = 1\}$ $\lambda_{max}(H_D)|\{\gamma_D = 3\} > \lambda_{max}(H_D)|\{\gamma_D = 2\}$

The functional dependence — whether linear, polynomial, or exponential in γ_D — is the open computation. A natural starting hypothesis, given that $\mathcal{D} = diag(1, 2, 4)$ already encodes a doubling of generation-space stiffness per generation, is

$\lambda_{max}(H_D) \sim \mathcal{D}_{\gamma_D}^2$

with the square coming from the Hessian being a quadratic functional of the closure-norm. This would give

$S_H(D) \sim \mathcal{D}_{\gamma_D} = (1, 2, 4)$ for $\gamma_D = (1, 2, 3)$

This contributes a factor of 2 between generations 1 and 2, and a factor of 4 between generations 1 and 3 — substantially smaller than the empirical magnitudes (207, 3477) and therefore a *sub-dominant* contribution to the hierarchy.

(Structural; the explicit dependence of $\lambda\_max(H\_D)$ on γ_D requires evaluation of the closure free-energy Hessian for specific PFD classes, which is part of the open computation.)

5.4 Epistemic status

The functional form $S\_H(D) \sim \sqrt{\lambda\_max(H\_D)}$ is structurally forced by the substrate free-energy framework. The dependence on closure topology (C_D, h_D, $\beta_1(D)$) is structurally constrained but the exact functional form is open. The dependence on generation depth γ_D is conjectural — the $\sim \mathcal{D}\_\{\gamma\_D\}$ hypothesis is natural but unproven, and resolving it requires explicit computation of the closure free-energy Hessian for specific PFD classes.

---

6. Localization compression $S\_L(D)$

6.1 Substrate mechanism

A PFD with localization length ℓ_D must be sustained against substrate dispersal — the natural tendency of any localized closure excitation to spread out into the condensate background. The energetic cost of localization compression is the work done against this dispersive tendency, evaluated in the closure-condensed phase.

For a Bessel-mode-localized closure excitation $f(r) = K\_n(\kappa r)$, the energy density goes as κ^2, and the total energy of the localized configuration goes as $\kappa^2 \cdot \ell\_D^3$ for a three-dimensional configuration of characteristic size ℓ_D. The coherence radius κ^{-1} and the localization length ℓ_D are inversely related: a more deeply localized PFD has larger κ. So tighter localization (smaller ℓ_D) costs more energy.

6.2 Functional form

The localization compression factor is

$$S\_L(D) \sim \ell\_D^{-p}$$

with $p > 0$ a positive scaling exponent. The exponent p is constrained by the dimensional structure of the Bessel-mode normalization integral and the substrate free-energy density's dependence on closure compression. A natural starting hypothesis, from the energy-density scaling argument above, is $p = 1$ (energy $\sim \kappa^{-1} \cdot (\kappa^2 \cdot \ell\_D^3) \cdot \ell\_D^{-1} \approx \kappa \sim \ell\_D^{-1}$ for the dominant cost mode), but the precise value depends on the substrate's elastic response to compression and is part of the open computation.

6.3 Exponential generation hierarchy from Role-4 scaling

The substantive prediction of this paper enters here. The Role-4 lepton programme established

$$\ell_g = \ell_0 \cdot \exp(-\kappa \cdot g)$$

with $\kappa \approx 3/8$ derived from CP^2 geometric structure and confinement dynamics. Substituting into S_L:

$$S_L(D) \sim (\ell_0 \cdot \exp(-\kappa \cdot \gamma_D))^{-p} = \ell_0^{-p} \cdot \exp(p \cdot \kappa \cdot \gamma_D)$$

The localization compression factor therefore scales *exponentially* in generation depth:

$$S_L(\gamma_D = 1) \sim \exp(p \cdot \kappa)\ S_L(\gamma_D = 2) \sim \exp(2 \cdot p \cdot \kappa)\ S_L(\gamma_D = 3) \sim \exp(3 \cdot p \cdot \kappa)$$

The mass ratio across generations from localization compression alone is

$$S_L(\gamma_D = 2) / S_L(\gamma_D = 1) = \exp(p \cdot \kappa)\ S_L(\gamma_D = 3) / S_L(\gamma_D = 1) = \exp(2 \cdot p \cdot \kappa)$$

With $\kappa \approx 3/8$: for $p = 1$, the per-generation factor is $\exp(3/8) \approx 1.45$; for $p = 2$, it is $\exp(3/4) \approx 2.12$; for $p = 5$, it is $\exp(15/8) \approx 6.5$. These are all substantially smaller than the empirical lepton ratios (207 between consecutive generations would require $p \cdot \kappa \approx 5.33$, i.e., $p \approx 14$).

Localization compression alone is therefore *not sufficient* to generate the observed empirical hierarchy at moderate values of p. The empirical hierarchy requires either large p (which would be unnatural) or substantial contribution from the other three factors S_H, S_P, S_I. We will see in §13 that the natural reading is that S_L contributes the *exponential structure* of the hierarchy while S_P, S_I contribute the dominant magnitudes.

6.4 Epistemic status

The functional form $S_L(D) \sim \ell_D^{-p}$ is structurally forced by the substrate dispersal mechanism and the Bessel-mode energy scaling. The exponent p is structurally constrained to be positive but the precise value is open — it requires explicit evaluation of the substrate elastic response to compression for specific PFD classes. The exponential dependence on γ_D is *forced* once Role-4 localization scaling $\ell_g = \ell_0 \cdot \exp(-\kappa \cdot g)$ is taken as input; this is the structural origin of exponential generation hierarchy.

---

7. Persistent distinguishability load $S_P(D)$

7.1 Substrate mechanism

A PFD carries committed-fold distinguishability content — the subset of admissibility-fixed folds whose specific committed configuration must be maintained for the PFD to remain a coherent closure structure. Let σ_D denote this persistent distinguishability load, measured in committed bits per natural admissibility cycle.

Each committed bit carries a *minimum stabilization cost* per cycle, by direct application of Landauer's principle: maintaining a bit of distinguishability against thermalization requires $k_B T \cdot \ln 2$ of energy per bit per natural relaxation time, where T is the substrate's effective temperature. For a PFD carrying σ_D bits of persistent distinguishability load, the stabilization cost is therefore

$$\text{(stabilization cost from distinguishability load)} \sim \sigma_D \cdot (k_B T \cdot \ln 2)$$

per natural cycle. Converting to a *stabilization-cost factor* (the mass amplification associated with this stabilization mechanism):

$$S_P(D) \sim \exp(\sigma_D)$$

The exponential form arises because S_P measures *multiplicative* mass amplification, and exponentials are the natural way to convert additive bit-counts to multiplicative cost factors. (Compare: free-energy differences from informational sources produce $\exp(\Delta S / k_B)$ Boltzmann factors via standard statistical mechanics.)

7.1a Why persistent distinguishability necessarily costs energy

The application of Landauer-type reasoning to persistent closure structure is not imported arbitrarily from conventional computation theory. It follows naturally from the substrate ontology itself, via the BCB primitives.

Under BC1 (bit conservation), distinguishability is *conserved* across substrate evolution — it is neither created nor destroyed by admissible substrate dynamics. But conservation is not the same as gratuitous availability: maintaining a *committed* fold against dispersive substrate evolution requires continuous suppression of admissible state relaxation. The substrate's natural evolution under unconstrained admissibility transport drives committed configurations toward maximal local distinguishability dispersal (the substrate analogue of thermalization). A committed fold therefore represents:

— a *constrained* local distinguishability configuration, — *maintained away from* maximal substrate entropy, — against the substrate's *natural relaxation flow*.

Such maintenance necessarily carries energetic cost.

The argument is structural, not phenomenological. Let σ_D count the persistent committed distinguishability associated with a PFD class. Maintaining that structure requires continuous stabilization work against three substrate processes:

(i) *closure relaxation* — the natural decay of committed-fold configurations toward unrestricted admissibility states; (ii) *dispersive transport* — the spreading of localized distinguishability content into the surrounding substrate; (iii) *local entropy increase* — the BC1-conserving but rearranging flow that, while preserving total distinguishability, redistributes it away from the PFD's specific committed configuration.

Each of these processes is admissible under BC1 and would proceed spontaneously absent stabilization work. The energetic contribution required to suppress them scales monotonically with the amount of committed distinguishability being held against relaxation — i.e., with σ_D — independently of any specific microscopic realization of the substrate.

Landauer's principle provides the *low-energy thermodynamic limit* of this broader substrate requirement. In the conventional Landauer setting, bit erasure costs $k_B T \cdot \ln 2$ energy because erasing a bit reduces local information content while conservation requires that the corresponding entropy be exported to the environment. In the present substrate setting, the analogous statement is that *holding* a committed bit against substrate relaxation costs energy because the natural substrate evolution would dissipate that bit absent stabilization work. The two settings are dual: conventional Landauer measures the cost of *creating* an information differential against entropic equilibrium; substrate Landauer measures the cost of *maintaining* an information differential against admissibility flow.

(The substrate "temperature" T in the expression $k_B T \cdot \ln 2$ should be understood as a *dimensional placeholder* for the substrate's natural admissibility-relaxation rate scale, not as a thermodynamic temperature in the conventional sense. The substrate-derived argument of §7.1a does not require a defined substrate temperature to go through — it derives the exponential cost structure $S_P(D) \sim \exp(\sigma_D)$ directly from BC1 + commitment persistence + admissibility-relaxation flow. The Landauer expression appears only when one takes the thermodynamic limit, where the substrate's relaxation rate acquires the dimensional structure of $k_B T$. A first-principles construction of the substrate temperature from substrate dynamics is a target for the second-quantised closure-condensate QFT companion paper.)

Structural interpretation. In conventional computation, *bit erasure costs energy*. In the present framework, *persistent admissibility-fixed distinguishability costs energy*. The substrate stiffness contribution $S_P(D)$ is therefore the closure-theoretic generalization of Landauer stabilization cost — not an arbitrary import from thermodynamics but the natural substrate-level expression of the same underlying principle that BC1 plus committed persistence demands.

This identification has a programmatic implication: Landauer's principle is not a *separate input* to the framework but a *consequence* of BCB combined with PFD commitment persistence. The exponential dependence $S_P(D) \sim \exp(\sigma_D)$ is therefore not an imported Landauer assumption but a substrate-derivable result of which Landauer is the thermodynamic limit. This strengthens the framework's epistemic position substantially: nothing in §7's argument requires external thermodynamic axioms beyond what BCB already supplies.

(Forced by BC1 + commitment persistence + admissibility relaxation. The exponential functional form is the substrate-level statement of which conventional Landauer is the thermodynamic limit.)

7.2 Dependence on PFD invariants

The persistent distinguishability load σ_D depends on multiple components of the PFD invariant tuple:

— $\beta_1(D)$ (Betti number): higher topological genus carries more committed bits because more admissibility-fixed homology classes must be maintained; — χ_D (chirality): chiral PFDs carry an additional committed bit per chirality (the L/R distinguishability); — ℓ_D (ledger charge): ledger-charge-bearing PFDs carry additional committed bits from the charge-distinguishability bookkeeping; — C_D (closure completeness): complete-closure PFDs carry more bits than partial-closure PFDs.

Specifically, the persistent distinguishability load of a PFD class D satisfies a structural lower bound

$$\sigma_D \geq \beta_1(D) + |\chi_D| + |\ell_D| + I(C_D)$$

where the notation is: $\beta_1(D)$ is the first Betti number (counting independent admissibility-fixed homology classes that must be maintained); $|\chi_D| \in \{0, 1\}$ is the *chirality bit-count* (1 if the PFD is chiral, 0 if chirality-trivial); $|\ell_D|$ is the *ledger-charge bit-count* — the number of *nonzero components* of the ledger-charge vector ℓ_D, counting how many independent ledger-charge distinctions the PFD must maintain; and $I(C_D) = 1$ for complete closure, 0 for partial closure. The bound is saturated for PFDs with minimal closure structure and is conservative for PFDs with additional internal complexity (e.g., extended closure cycles, multi-channel ledger structures, holonomy-class refinement).

7.3 Exponential structure naturally generates hierarchies

The exponential dependence $S_P(D) \sim exp(\sigma_D)$ is structurally important: it provides a *second* exponential channel for hierarchy generation, structurally distinct from the localization-compression exponential of §6. The combination of two exponential mechanisms — one in localization scaling (§6) and one in distinguishability load (§7) — together with the smaller polynomial contributions from S_H and S_I, generates the empirical magnitudes naturally.

For lepton PFDs, the persistent distinguishability loads can be ordered structurally: $\sigma_e < \sigma_\mu < \sigma_\tau$, because higher-generation lepton PFDs carry additional refinement-stable closure structure. The specific values are open (they require explicit enumeration of the committed-fold structure for each PFD class), but the ordering is forced.

7.4 Epistemic status

The functional form $S_P(D) \sim exp(\sigma_D)$ is structurally motivated by Landauer's principle applied to substrate informational bookkeeping. The structural lower bound on σ_D is forced by the PFD invariant tuple structure. The explicit value of σ_D for each PFD class is open and is part of the open computation. The exponential dependence is the natural functional form for stabilization mechanisms whose costs are additive in distinguishability bits, by the same logic that generates Boltzmann factors in statistical mechanics.

---

8. Interface transport complexity *S_I(D)*

8.1 Substrate mechanism

A PFD's closure structure is maintained against admissibility decay through transport along admissibility-preserving channels in the closure transport network. The number of such channels — and the topological complexity of the network — determines how robustly the PFD's closure structure can be sustained against perturbative decay. A PFD that admits more independent admissibility-preserving transport paths is more stable (the structure can repair itself via multiple routes); a PFD that depends on a single transport path is more fragile.

The interface transport complexity *T(D)* measures this. Concretely, *T(D)* is the topological complexity of the *admissibility-preserving transport graph* of the PFD — the directed graph whose nodes are the PFD's closure cells and whose edges are admissibility-fixed transport channels between them. Natural measures include:

— *T(D) = (number of independent admissibility-preserving paths supporting refinement persistence)*; — or, equivalently, the *minimum cut* of the transport graph that disconnects the PFD into admissibility-incomplete sub-structures.

8.2 Functional form

The interface transport complexity factor is

$S_I(D) \sim T(D)$

with *T(D)* a positive integer (counting admissibility-preserving paths) or a positive real (when fractional path-weighting from substrate occupancies is incorporated). The linear dependence reflects that each independent transport path contributes additively to the maintenance cost.

The closure-norm condensation paper §11.2a noted that this factor corresponds to matter-paper contribution (iii) — confinement/localization cost — and that its substrate origin lies in the Role-4 stabilization machinery. The present paper sharpens that to: *T(D)* is the topological complexity of the PFD's admissibility-preserving transport graph.

8.2a Graph-theoretic realization of transport complexity

Let $G_D = (V_D, E_D)$ be the admissibility-preserving transport graph associated with PFD class D, where vertices V_D correspond to admissibility-fixed closure cells of the PFD's support and edges E_D correspond to admissibility-preserving transport channels between them. G_D is *directed*, with edge orientations inherited from the directional structure of admissibility-preserving transport (each transport channel has a definite admissibility-preferred direction at the substrate level). Refinement persistence (P3) requires G_D to be connected; closure stability (P4) requires the substrate transport operator restricted to G_D to have a strictly positive spectral gap.

We define the transport complexity functional concretely as

$$T(D) = \alpha_1 \cdot \beta_1(G_D) + \alpha_2 \cdot \lambda_2^{-1}(G_D) + \alpha_3 \cdot \chi(G_D) + \alpha_4 \cdot deg\_avg(G_D)$$

where:

— $\beta_1(G_D)$: the first Betti number of the (undirected) underlying graph — its *cycle rank*, counting the number of independent closed transport loops the PFD must support; — $\lambda_2(G_D)$: the second-smallest eigenvalue of the (symmetrised) graph Laplacian — the *spectral gap*, with λ_2^{-1} measuring transport bottleneck severity (small spectral gap = severe bottleneck = high stabilization cost); — $\chi(G_D)$: the *transport branching number*, defined as the average *out-degree* per vertex (a measure of forward-channel multiplicity); — $deg\_avg(G_D)$: the *average total connectivity*, the mean total degree (in-degree + out-degree) per vertex of the directed graph.

The coefficients $\alpha_1, \ldots, \alpha_4$ are substrate-level constants of order unity, set by the relative weight of cycle, bottleneck, branching, and connectivity contributions to substrate stabilization cost. They are not free parameters — they are computable from the substrate transport operator's response to perturbations along each of the four structural directions.

Independence of the invariants. For directed graphs the four invariants are formally independent: β_1 counts cycles regardless of graph size; λ_2 depends on the global spectral structure of the Laplacian; χ (out-degree) and deg_avg (total degree) coincide only in undirected graphs and decouple for directed graphs with non-trivial admissibility orientations. The standard algebraic identity $\beta_1 = |E| - |V| + 1$ relates β_1 to graph size but not to spectral or directional properties, so β_1 and λ_2^{-1} remain independently variable. In specific PFD-class realizations the four invariants will correlate (a confined quark transport graph is *both* densely connected *and* spectrally bottlenecked, for example), but their formal independence justifies treating the four contributions separately and assigning each its own coefficient α_i.

Interpretation. The four graph invariants capture structurally distinct contributions to stabilization difficulty. Larger $\beta_1(G_D)$ means more independent transport cycles must be maintained against decay. Smaller $\lambda_2(G_D)$ (larger λ_2^{-1}) means the transport network has weaker spectral robustness — perturbations equilibrate slowly, requiring sustained stabilization work. Larger $\chi(G_D)$ means more parallel transport paths must each be kept admissibility-fixed. Larger $deg\_avg(G_D)$ means more local connections to maintain. Each invariant contributes additively to $T(D)$; their combination determines the total transport complexity.

Structural expectation for the quark/lepton split. Quark PFDs are structurally expected to possess:

— *larger* β_1: confinement transport requires multi-cycle colour-singlet bookkeeping (counting independent baryon and meson formation loops); — *smaller* λ_2 (larger λ_2^{-1}): the colour confinement bottleneck is severe — perturbations cannot escape the confinement scale through admissibility-preserving channels, giving small spectral gap; — *larger* χ: multi-channel colour-flavour branching gives high transport branching number; — *larger deg_avg*: confinement networks have dense local connectivity across colour channels.

All four contributions push $T(D)|_{quark} \gg T(D)|_{lepton}$ for PFDs of the same generation depth. This makes the structural quark/lepton separation of §10 concrete and graph-theoretically computable.

Computational implication. With this definition, evaluating $T(D)$ for a specific PFD class is a *graph-theoretic computation* on a finite admissibility-fixed graph: count the cycles, compute the Laplacian spectrum, evaluate branching and degree statistics, combine with substrate-derived coefficients. Each step is mechanical given the closure transport graph as input. The open computation is the construction of G_D for specific PFD classes (electron, muon, tau, up, down, etc.) from the substrate transport operator — a target for the next computational paper.

(Forced by substrate transport architecture; the four graph invariants and their substrate-derived coefficients constitute a concrete computational target rather than a schematic placeholder.)

8.3 Dependence on PFD invariants

The transport graph complexity depends on:

— ρ_D (confinement requirement): confined PFDs (quarks) require multi-channel transport networks supporting colour singlet formation, dramatically increasing $T(D)$; — C_D (closure completeness): complete-closure PFDs require denser transport networks than partial-closure; — γ_D (generation depth): higher-generation PFDs require additional refinement-persistent transport paths beyond those of lower generations.

The structural prediction is

$T(D)|_{quark} \gg T(D)|_{lepton}$

for PFDs of the same generation depth, with the order-of-magnitude difference reflecting the multi-channel colour confinement transport required by quark PFDs. This is the structural origin of the quark/lepton mass separation developed in §10.

8.4 Epistemic status

The graph-theoretic definition of $T(D)$ is structurally forced by the substrate transport architecture. The dependence on ρ_D, C_D, γ_D is structurally constrained. The explicit value of $T(D)$ for each PFD class is open and requires graph-theoretic computation on the closure transport network — a target for the next computational paper.

---

9. The substrate stiffness hierarchy theorem

We now state the central structural result. Combining the four factors of §5–§8:

Theorem 9.1 (Substrate Stiffness Hierarchy). *Let D be a PFD class with invariant tuple $\mathcal{I}(D) = (C_D, \beta_1(D), h_D, \pi_D, \chi_D, \gamma_D, \ell_D, \rho_D)$, and suppose the four-factor decomposition of §4 holds. Then the fermion mass m_D is*

$$m_D = \mathcal{D}_{\{\gamma_D\}} \cdot S_H(D) \cdot S_L(D) \cdot S_P(D) \cdot S_I(D) \cdot v$$

where v is the closure-condensate vacuum scale, $\mathcal{D} = \mathrm{diag}(1, 2, 4)$ is the generation operator, and the four substrate stiffness factors satisfy the functional forms

$S_H(D) \sim \sqrt{\lambda_{\max}(H_D)}$ — from substrate free-energy Hessian (§5);

$S_L(D) \sim \ell_D^{-p}$ with $p > 0$ — from substrate dispersal cost (§6);

$S_P(D) \sim \exp(\sigma_D)$ — from Landauer cost of persistent distinguishability bookkeeping (§7);

$S_I(D) \sim T(D)$ — from admissibility-preserving transport graph complexity (§8).

Taking logarithms gives the additive scaling law

$$\ln m_D \approx \ln \mathcal{D}_{\{\gamma_D\}} + (1/2) \ln \lambda_{\max}(H_D) - p \cdot \ln \ell_D + \sigma_D + \ln T(D) + \ln v.$$

Substituting the Role-4 scaling $\ell_D = \ell_0 \cdot \exp(-\kappa \cdot \gamma_D)$ and collecting constants:

$$\ln m_D \approx \ln \mathcal{D}_{\{\gamma_D\}} + (1/2) \ln \lambda_{\max}(H_D) + p \cdot \kappa \cdot \gamma_D + \sigma_D + \ln T(D) + \text{const.}$$

The fermion mass hierarchy is therefore the sum of (i) a sub-linear-in-γ_D contribution from $\mathcal{D}$, (ii) a Hessian-curvature contribution from S_H, (iii) an exponential-in-γ_D contribution from S_L driven by Role-4 localization scaling, (iv) an exponential-in-σ_D contribution from S_P driven by Landauer informational cost, and (v) a transport-graph-complexity contribution from S_I.

(Conditional on the four-factor decomposition holding at leading order, the Role-4 scaling for ℓ_D, and Landauer's principle applied to substrate informational bookkeeping.)

9.1 Structural interpretation

The theorem converts the fermion mass hierarchy from a phenomenological insertion into a structured calculational programme. The key observations are:

— *Two exponential channels are active.* The factor S_L contributes exponential structure in γ_D via Role-4 localization scaling; the factor S_P contributes exponential structure in σ_D via Landauer's principle. The combination of two distinct exponential mechanisms is what naturally generates the empirical twelve-order-of-magnitude hierarchy.

— *The factors are logically independent.* Each factor measures a distinct substrate energetic cost from a distinct substrate mechanism. There is no double-counting; a PFD with both deep localization and large distinguishability load pays both costs.

— *Quark/lepton separation lives in S_I and S_P.* The factors that distinguish quark PFDs from lepton PFDs of the same γ_D are S_I (multi-channel confinement transport) and S_P (colour-charge distinguishability bookkeeping). This naturally generates the empirical quark/lepton mass separation, as developed in §10.

— *Neutrino suppression lives in $\ell_D = 0$.* Neutrino PFDs have $\ell_D = 0$, removing the dominant ledger-charge contributions to S_L, S_P, S_I. This produces tiny $S(\nu)$ and therefore tiny neutrino masses, as developed in §11.

9.2 What the theorem delivers and what remains open

The theorem delivers the *structural framework* — the functional form of $S(D)$ as a product of four substrate-mechanism-derived factors, and the resulting additive-in-log-space scaling law. What it does not deliver are the *explicit numerical values* of the four factors for specific PFD classes, which require:

(a) evaluation of $\lambda_{max}(H_D)$ from the substrate free-energy Hessian for each PFD class; (b) computation of the exponent p from the substrate dispersal response; (c) enumeration of the committed-fold distinguishability content σ_D for each PFD class; (d) graph-theoretic evaluation of $T(D)$ from the closure transport network.

Each of these is a well-posed computational problem on substrate primitives. The combination is the next computational paper.

---

10. Quark–lepton separation

A central structural feature of the empirical mass spectrum is the systematic separation between quark and lepton masses. Up-type quarks of generation g are typically heavier than charged leptons of the same generation by factors ranging from roughly 4 (electron vs up) to 1.5 (tau vs

top, though the top dominates by another order). The structural origin of this separation in the four-factor framework is straightforward.

10.1 The three sources of quark/lepton separation

Quark PFDs differ from lepton PFDs of the same generation depth in three key invariants:

— C_D = *partial* for quarks (colour-triplet, partial closure under SU(3)_C) vs C_D = *complete* for charged leptons (colour-singlet). This affects S_H and S_P directly.

— ρ_D = *confined* for quarks vs ρ_D = *free* for charged leptons. This affects S_I dramatically: quark PFDs require multi-channel colour-confinement transport networks supporting colour-singlet baryon and meson formation.

— ℓ_D = *(electric + colour)* for quarks vs ℓ_D = *electric only* for charged leptons. This affects S_P through the additional colour-charge distinguishability bits.

10.2 Structural prediction

Combining the three effects gives the structural prediction

$$S(D)|quark = \alpha\{q/l\} \cdot S(D)|_{lepton}$$

with the quark/lepton amplification factor

$$\alpha_{q/l} = (S_H|_q / S_H|_l) \cdot (S_L|_q / S_L|_l) \cdot \exp(\sigma_q - \sigma_l) \cdot (T_q / T_l)$$

The dominant contributions are from S_P (exponential in additional colour-distinguishability bits) and S_I (linear in additional multi-channel transport paths). The S_L contribution is structurally similar between quarks and leptons of the same generation (both have similar localization lengths ℓ_D set primarily by γ_D rather than by colour quantum numbers). The S_H contribution can go either way depending on whether colour-triplet PFDs sit in deeper or shallower free-energy minima than colour-singlet PFDs at the same γ_D.

The structural prediction is therefore $\alpha_{q/l} > 1$: quark PFDs are heavier than charged-lepton PFDs of the same generation. The precise numerical value of $\alpha_{q/l}$ requires the open computations.

10.3 Why quark hierarchy is steeper than lepton hierarchy

Empirically, the quark mass hierarchy is steeper than the lepton hierarchy: $m_c/m_u \approx 600$ and $m_t/m_u \approx 75{,}000$, versus $m_\mu/m_e \approx 207$ and $m_\tau/m_e \approx 3477$. The four-factor framework predicts this structurally:

— S_L contributions: similar between quarks and leptons (both follow Role-4 scaling); — S_P contributions: steeper for quarks (additional colour-distinguishability bits compound

exponentially across generations); — S_I contributions: dramatically steeper for quarks (higher-generation quark PFDs require even more complex confinement transport).

The combination predicts that quark/lepton separation increases with generation depth, which matches the empirical observation.

10.4 Epistemic status

The structural prediction $\alpha_{q/l} > 1$ is forced by the four-factor decomposition plus the PFD invariant differences between quark and lepton classes. The numerical value of $\alpha_{q/l}$ is open and depends on the explicit evaluations enumerated in §9.2. The structural prediction that quark hierarchy is steeper than lepton hierarchy is also forced, with the open numerics determining the precise steepness ratio.

---

11. Neutrino suppression

The neutrino mass scale is anomalously small — at most an eV, ten orders of magnitude below the electron. The substrate origin of this suppression follows directly from the four-factor framework.

11.1 The structural role of ℓ_D = 0

Neutrino PFDs have:

— C_D = *complete* (singlet under SU(3)_C); — $\ell_D = 0$ (no ledger charge — electromagnetically neutral); — $\pi_D = +1$ (orientation-trivial); — $\gamma_D \in \{1, 2, 3\}$.

The invariant $\ell_D = 0$ is structurally consequential. Three of the four substrate stiffness factors depend on ledger charge:

— $S_L(D)$ depends on ℓ_D through the localization length ℓ_D (the localization compression mechanism is partly driven by ledger-charge concentration); — $S_P(D)$ depends on ℓ_D through the ledger-charge contribution to persistent distinguishability load σ_D; — $S_I(D)$ depends on ℓ_D through the ledger-charge transport channels required for refinement persistence.

A neutrino PFD with $\ell_D = 0$ has the ledger-charge contributions to each of S_L, S_P, S_I set to zero (or to the minimal value supported by the remaining non-ledger contributions). The result is that $S(\nu)$ is dramatically smaller than $S(e)$ for the corresponding charged lepton of the same generation:

$$S(\nu) \ll S(e)$$

and therefore

$m_\nu = \mathcal{D}_{\gamma_D} \cdot S(\nu) \cdot \nu \ll m_e = \mathcal{D}_{\gamma_D} \cdot S(e) \cdot \nu.$

11.2 Estimating the suppression scale

If the ledger-charge contributions to S_P and S_I are comparable in magnitude to the non-ledger contributions, the neutrino-to-electron mass ratio scales as

$m_\nu / m_e \sim \exp(\sigma_\nu - \sigma_e) \cdot (T_\nu / T_e)$

with $\sigma_\nu - \sigma_e$ being negative (neutrinos carry fewer committed-distinguishability bits) and T_ν / T_e being less than one (neutrino transport networks lack ledger-charge channels). Both contributions push m_ν / m_e far below unity.

Empirical anchors. Neutrino oscillation data fixes the mass-squared splittings $\sqrt{|\Delta m^2_{32}|} \approx 50$ meV and $\sqrt{|\Delta m^2_{21}|} \approx 8.6$ meV, so at least one mass eigenstate must satisfy $m_{\nu,\text{heaviest}} \gtrsim 0.05$ eV. Cosmological constraints from CMB and large-scale structure give an upper bound on the sum, $\Sigma m_\nu \lesssim 0.12$ eV (Planck), bounding the heaviest eigenstate below roughly 0.1 eV. The framework must therefore accommodate a *spectrum* of three eigenstate masses spanning roughly 1 meV to 0.1 eV, not a single suppression scale.

Predicted spectrum, not single value. The four-factor framework predicts a *split spectrum* across generations: the diagonal stiffness factor $S(\nu)$ takes three distinct values $S(\nu_1), S(\nu_2), S(\nu_3)$ corresponding to $\gamma_D = 1, 2, 3$, with the generation amplifications from $\mathcal{D} = \text{diag}(1, 2, 4)$ and from S_L, S_P at higher γ_D producing a hierarchy $m_{\nu,1} < m_{\nu,2} < m_{\nu,3}$. The 50 μeV figure cited below refers to the *lightest* eigenstate; the heavier eigenstates are amplified by the same generation-depth mechanisms that produce the charged-lepton hierarchy, scaled by the overall $S(\nu)/S(e)$ ratio.

Bit-count calibration. A back-of-envelope estimate with ledger-charge bit-count $\Delta\sigma \approx 15$ gives $m_{\nu,\text{lightest}} / m_e \sim \exp(-15) \approx 3 \times 10^{-7}$, corresponding to $m_{\nu,\text{lightest}} \approx 1.5$ meV — within the regime allowed by normal-hierarchy oscillation data. With $m_{\nu,1} \approx 1.5$ meV and $\sqrt{|\Delta m^2_{21}|} \approx 8.6$ meV, normal hierarchy requires $m_{\nu,2} \approx \sqrt{(m_{\nu,1}^2 + |\Delta m^2_{21}|)} \approx 8.7$ meV — i.e., the framework would need to predict a *factor-of-six* amplification from $m_{\nu,1}$ to $m_{\nu,2}$ via $\mathcal{D}$, S_L, S_P, S_I at $\gamma_D = 2$. This is structurally similar to the charged-lepton case (where the gen 1 → gen 2 amplification is 207) but smaller in magnitude; explicit substrate computation of the neutrino $S(\nu_1), S(\nu_2), S(\nu_3)$ and their combination with $\mathcal{D} = (1, 2, 4)$ is a target for the next computational paper. (A smaller-Δσ scenario, $\Delta\sigma \approx 10$, would give $m_{\nu,\text{lightest}} \approx 25$ meV, in which case all three eigenstates would be closer in mass and dominated by the splittings — giving a quasi-degenerate spectrum $m_1 \approx 25$, $m_2 \approx 26$, $m_3 \approx 56$ meV summing to roughly 107 meV, approaching the cosmological $\Sigma m_\nu \lesssim 120$ meV Planck bound but not yet excluded; a larger-Δσ scenario, $\Delta\sigma \gtrsim 20$, would give $m_{\nu,\text{lightest}}$ well below the splitting scale, with $m_{\nu,2}$ and $m_{\nu,3}$ set entirely by $\sqrt{|\Delta m^2_{21}|}$ and $\sqrt{|\Delta m^2_{32}|}$ rather than by $S(\nu_1)$ amplification.) The exact value of $\Delta\sigma$ is open and depends on explicit enumeration of the committed-fold structure of the ledger-charge sector. The structural prediction is that the empirical neutrino mass scale is *consistent with substrate primitives at the bit-count level* once $\Delta\sigma$ is computed — the

framework is empirically calibrated, with the derivation reduced to the substrate computation of $\Delta\sigma$ and the gen 1 → gen 2 amplification ratio.

11.3 Why the suppression is structural, not fine-tuned

In the Standard Model, the small neutrino mass requires either a very small Dirac Yukawa coupling (fine-tuned without explanation) or a seesaw mechanism with a heavy right-handed neutrino at an inserted high scale (model-dependent). In the present framework, the suppression is structural: it follows from the neutrino PFD having a specific value of one invariant ($\ell_D = 0$), which removes one of the dominant contributions to the stiffness factor without any parameter tuning. The empirical neutrino mass scale becomes a *consequence* of the substrate invariant structure rather than an unexplained fine-tuning.

11.4 Epistemic status

The structural prediction $S(\nu) \ll S(e)$ is forced by the four-factor decomposition plus $\ell_D = 0$ for neutrinos. The estimate $m_\nu / m_e \sim \exp(\sigma_\nu - \sigma_e) \cdot (T_\nu / T_e)$ is conditional on the relative weights of the four factors. The explicit numerical magnitude of the suppression is open and depends on the explicit committed-fold enumeration.

---

12. CKM and PMNS mixing structure

The closure-norm condensation paper §11.4 introduced the structural formula for inter-generation mixing:

$$\sin\theta_{ij} \sim S_{ij}(D) / |\mathcal{D}_{\gamma_i} - \mathcal{D}_{\gamma_j}|$$

with $S_{ij}(D)$ the off-diagonal stiffness factor between generations i and j (the closure-condensate-induced coupling between PFD classes of different generation depths). The present paper's four-factor decomposition extends naturally to off-diagonal stiffness factors.

12.1 Off-diagonal stiffness from the four-factor structure

The off-diagonal stiffness factor decomposes the same way as the diagonal:

$$S_{ij}(D) = S_{H,ij}(D) \cdot S_{L,ij}(D) \cdot S_{P,ij}(D) \cdot S_{I,ij}(D)$$

with each *off-diagonal* factor measuring the substrate-mediated coupling between PFD classes of different generation depths via the corresponding mechanism. The diagonal cases recover the §4 factors when $i = j$: $S_{ii}(D) \equiv S(D)$.

The off-diagonal factors are structurally smaller than the diagonal factors because inter-generation coupling involves admissibility-preserving transport between distinct closure-depth sectors, which is constrained by refinement persistence.

12.2 Quark CKM vs lepton PMNS structure

The structural prediction is that *CKM mixing is small* and *PMNS mixing is large*, set by the ratio of off-diagonal to diagonal stiffness factors:

$|U_{CKM}|_{ij} \sim S_{ij}^q / S^q$ for the quark sector, $|U_{PMNS}|_{ij} \sim S_{ij}^\nu / S^\nu$ for the neutrino sector.

For quarks, both S^q (diagonal) and S_{ij}^q (off-diagonal) are large in absolute terms, but their *ratio* is small because the diagonal grows faster across generations than the off-diagonal. For neutrinos, both are small in absolute terms, but the ratio is large because the diagonal S^ν is anomalously suppressed (by $\ell_D = 0$) while the off-diagonal mixing structure is not similarly suppressed.

The structural prediction $|U_{CKM}| \ll |U_{PMNS}|$ is therefore forced by the same $\ell_D = 0$ mechanism that produces neutrino mass suppression. Neutrino mass smallness and large PMNS mixing are two manifestations of the same substrate fact.

12.3 Suppression of non-adjacent generation mixing

The factor $|\mathcal{D}_{\gamma_i} - \mathcal{D}_{\gamma_j}|$ in the denominator of the mixing formula provides the suppression of non-adjacent generation mixing:

— Between adjacent generations: $|\mathcal{D}_{\gamma_1} - \mathcal{D}_{\gamma_2}| = 1$, $|\mathcal{D}_{\gamma_2} - \mathcal{D}_{\gamma_3}| = 2$; — Between non-adjacent generations: $|\mathcal{D}_{\gamma_1} - \mathcal{D}_{\gamma_3}| = 3$.

This naturally gives $\sin\theta_{13} \ll \sin\theta_{12}$ in the CKM matrix, matching the empirical hierarchy $|V_{us}| \approx 0.22 \gg |V_{ub}| \approx 0.004$.

12.4 Epistemic status

The structural prediction $|U_{CKM}| \ll |U_{PMNS}|$ is forced by the four-factor framework plus $\ell_D = 0$ for neutrinos. The suppression of non-adjacent generation mixing is forced by the structure of the $\mathcal{D}$ eigenvalue gap. The explicit numerical values of CKM and PMNS matrix elements are open and depend on the off-diagonal stiffness factors S_{ij} — the same open computation as the diagonal $S(D)$.

---

13. Empirical consistency check

The four-factor framework does not, in its current form, compute exact fermion masses. What it can do is provide a *consistency check*: do the four factors plausibly combine to produce the observed hierarchy under reasonable substrate parameter values?

We check the charged lepton hierarchy $m_e : m_\mu : m_\tau \approx 1 : 207 : 3477$ in log-space.

13.1 Setup

Taking logarithms of the mass formula at each generation $g \in \{1, 2, 3\}$:

$$\ln m_g = \ln \mathcal{D}_g + (1/2) \ln \lambda_{max}(H_g) + p \cdot \kappa \cdot g + \sigma_g + \ln T_g + const.$$

The *empirical* log-ratios for charged leptons are:

$$\ln(m_\mu / m_e) = \ln 207 \approx 5.33 \quad \ln(m_\tau / m_e) = \ln 3477 \approx 8.15 \quad \ln(m_\tau / m_\mu) = \ln(3477/207) = \ln(16.8) \approx 2.82$$

These ratios are *not* uniformly exponential in generation depth: the gap between g=1 and g=2 is much wider in log-space than the gap between g=2 and g=3. Any framework claiming to derive the hierarchy must accommodate this non-uniform structure.

13.2 Contribution from $\mathcal{D}$ alone

The generation operator contributes:

$$\ln(\mathcal{D}_2 / \mathcal{D}_1) = \ln 2 \approx 0.69 \quad \ln(\mathcal{D}_3 / \mathcal{D}_1) = \ln 4 \approx 1.39$$

— substantially smaller than the empirical ratios. $\mathcal{D}$ alone accounts for roughly 13% of the e→μ log-gap and roughly 17% of the e→τ log-gap. The bulk of the hierarchy comes from *S(D)*.

13.3 Contribution from S_L alone

With Role-4 scaling $\ell_g = \ell_0 \cdot \exp(-\kappa \cdot g)$ and $\kappa = 3/8$:

$$\ln(S_L(g=2) / S_L(g=1)) = p \cdot \kappa = 3p/8 \quad \ln(S_L(g=3) / S_L(g=1)) = 2p \cdot \kappa = 3p/4$$

For this to account for the empirical e→μ gap of 5.33 − 0.69 = 4.64 (subtracting the $\mathcal{D}$ contribution only), we would need $3p/8 \approx 4.64$, giving $p \approx 12$. For the e→τ gap of 8.15 − 1.39 = 6.76, we would need $3p/4 \approx 6.76$, giving $p \approx 9$. (Under the §5.3 hypothesis $S_H \sim \sqrt{\lambda_{max}(H_D)} \sim \mathcal{D}_{\gamma_D}$, the $\mathcal{D}$ + S_H contributions would together subtract approximately $(3/2) \cdot \ln \mathcal{D}$ rather than $\ln \mathcal{D}$ alone, shifting the inferred p values down by roughly one unit each; the qualitative conclusion is unchanged.)

These two values of p are inconsistent. S_L alone, with constant p, cannot account for the empirical hierarchy. The non-uniformity of the hierarchy across generations requires *additional* contributions from S_H, S_P, S_I that vary differently across generations.

13.4 The natural reading

The structural reading is that *S_L* contributes the *exponential structure* of the hierarchy (per-generation amplification roughly *exp(p·κ)*) while *S_P* (exponential in distinguishability bits) and *S_I* (linear in transport complexity) contribute the dominant magnitudes and the non-uniformity across generations.

Specifically: if *S_P(g=2) − S_P(g=1)* in committed bits is larger than *S_P(g=3) − S_P(g=2)* (a structurally plausible scenario in which most of the additional distinguishability content of higher-generation closures is added between g=1 and g=2, with diminishing additions at higher g), the resulting hierarchy is non-uniform in the empirically-observed direction. The transport complexity *T(D)* similarly can grow non-linearly with γ_D — the additional transport network needed to support g=2 closures might be a substantial step-up from g=1, with g=3 representing only an incremental further step.

Structural prediction (P-1): non-uniform σ_D and T(D) growth. *The empirical hierarchy is consistent with the four-factor framework if and only if the dominant non-uniformity across generations lives in σ_D (committed distinguishability bit-count) and T(D) (transport complexity), with the per-generation gap σ_2 − σ_1 substantially larger than σ_3 − σ_2, and similarly for T_2 − T_1 vs T_3 − T_2. The next computational paper must find this non-uniform growth pattern when σ_D and T(D) are evaluated explicitly from substrate primitives for the electron, muon, and tau PFD classes. If the substrate computations instead yield uniform per-generation growth of σ_D and T(D), the framework is falsified at the consistency-check level: the empirical hierarchy ratio m_τ/m_μ ≈ 16.8 cannot be reconciled with m_μ/m_e ≈ 207 under any uniform-growth scenario.*

This is a sharp falsifiable prediction at the level of substrate observables: not "the framework fits" or "the framework doesn't fit," but a specific structural statement about how the bit-counts and transport graph sizes must scale across generations. It is formally distinct from the falsification criteria of §15 in that it constrains the *internal substrate computation* rather than the framework's overall agreement with experiment. The two are complementary: §15 falsifies the framework against external data; the P-1 prediction falsifies it against the substrate computations the framework itself enables.

13.5 Estimated factor magnitudes

A natural starting assignment, consistent with the empirical lepton hierarchy:

Generation	ln $\mathcal{D}$	ln S_L	σ	ln T	(1/2) ln λ_max(H_D)	total	predicted ratio to m_e	observed
g = 1 (e)	0	0	0	0	0	0	1	1
g = 2 (μ)	0.69	0.38	3.5	0.5	0.26	5.33	206	207
g = 3 (τ)	1.39	0.76	4.5	1.0	0.51	8.16	3491	3477

(All quantities in natural log. The $\mathcal{D}$ contribution uses 1, 2, 4. The S_L contribution uses $p \cdot \kappa \cdot g$ with $p \cdot \kappa = 0.38$. The σ values are illustrative committed-bit counts. The ln T values assume

linear growth in transport-graph size with generation. The (1/2) ln λ_max(H_D) column uses *hand-chosen increments of order 0.25 per generation*; the §5.3 hypothesis $S_H \sim \sqrt{\lambda_{max}(H_D)} \sim \mathcal{D}_{\{\gamma_D\}}$ would give differentials (1/2) ln 2 ≈ 0.347, not the 0.26 used here, so the toy's S_H values do not strictly implement that hypothesis. The illustrative values are chosen for arithmetic convenience to make the toy reproduce the empirical ratios under the P-1 constraint; if the substrate-derived S_H values instead match the §5.3 hypothesis at 0.347 per generation, the budget for σ and $\ln T$ across generations would absorb the difference. The consistency check therefore demonstrates that the four-factor framework *can* accommodate the empirical hierarchy under reasonable substrate parameter assignments, not that it *must*. The next computational paper will replace these illustrative values with substrate-derived ones.)

The assignment is internally consistent and reproduces the empirical hierarchy without parameter fitting beyond fixing the four substrate quantities (σ, T, λ_max, p·κ) at each generation. The *predictivity* of the framework comes from these four substrate quantities being independently computable from the PFD invariant tuple — they are not free parameters but specific substrate observables. Computing them constitutes the next computational paper.

13.6 Computational roadmap and closure observables

The present paper establishes the structural decomposition of the substrate stiffness hierarchy but does not yet evaluate the four factors explicitly for individual PFD classes. The substantive claim of this subsection is that the hierarchy problem has now been reduced to a *finite computational programme on substrate observables* rather than an unconstrained phenomenological insertion. The four stiffness contributions correspond to four concrete computational targets, each with a specific operator-level definition:

Factor	Required computation	Observable substrate quantity
$S_{H(D)}$	Hessian eigenvalue analysis on the closure free-energy functional	$\lambda_{max}(H_D)$
$S_{L(D)}$	Bessel localization normalization on the closure transport network	coherence radius ℓ_D
$S_{P(D)}$	Committed-fold enumeration on the admissibility-fixed structure	persistent distinguishability count σ_D
$S_{I(D)}$	Transport-network analysis on the admissibility graph	graph complexity $T(D) = \alpha_1\beta_1 + \alpha_2\lambda_2^{-1} + \alpha_3\chi + \alpha_4 deg\_avg$

The fermion hierarchy problem therefore reduces to *four substrate computations*, each with explicit operator-level inputs and outputs. None is residual; none is a free fit.

Explicit algorithmic pathway. A first-pass computational programme proceeds as follows:

1. *Construct the admissibility-fixed closure graph* for a chosen PFD class, using the dictionary's invariant-tuple specification of which folds are committed and which transport channels are admissibility-preserving.

2. *Compute the closure Hessian spectrum* from the local free-energy functional $F[\rho]$ of the closure-norm condensation paper §4, evaluated at the PFD's closure support. Extract $\lambda_{max}(H_D)$.
3. *Evaluate the localization profile* via the Bessel-mode closure solution $K_n(\kappa_D \cdot r)$, identifying the coherence radius $\ell_D = \kappa_D^{-1}$ from the PFD's confinement geometry.
4. *Enumerate committed distinguishability content* by counting admissibility-fixed folds in the PFD's structure that are constrained to remain in committed configurations (the lower bound $\sigma_D \geq \beta_1 + |\chi_D| + |\ell_D| + I(C_D)$ from §7.2 gives a starting count; extended closure cycles and multi-channel ledger structures add further contributions).
5. *Construct the admissibility-preserving transport graph* G_D and evaluate its four invariants — first Betti number $\beta_1(G_D)$, second Laplacian eigenvalue $\lambda_2(G_D)$, branching number $\chi(G_D)$, average degree $deg_{avg}(G_D)$ — combining into the transport complexity $T(D)$ via §8.2a.
6. *Assemble the four-factor stiffness product*

$$S(D) = \sqrt{\lambda_{max}(H_D)} \cdot \ell_D^{-p} \cdot \exp(\sigma_D) \cdot T(D)$$

and combine with the generation operator and condensate scale to compute $m_D = \mathcal{D}_{\{\gamma_D\}} \cdot S(D) \cdot v$.

7. *Compare resulting hierarchy ratios* against the observed fermion masses (lepton ratios first, then quark ratios), and assess whether the leading-order four-factor decomposition reproduces the empirical spectrum to within the few-percent envelope set by inter-sector corrections (§4.1).

Each step is mechanical given the upstream substrate-operator framework. The combination is the next computational paper.

Programmatic implication. The fermion hierarchy problem is no longer an unconstrained phenomenological insertion of twelve mysterious numbers. It is a definite computational programme on substrate primitives, with seven explicit calculational steps, each with operator-level inputs and outputs from existing programme content. The framework's testability is now sharp: either the seven-step computation, performed for the electron, muon, tau, then quarks, reproduces the empirical spectrum to leading order — in which case substrate stabilization complexity is the correct origin of the hierarchy — or it does not, in which case Theorem 9.1 is falsified at the level identified in §15.

(Forced framework; the computational roadmap is now algorithmically defined rather than gestured at. Its execution is the explicit target of the next computational paper.)

13.7 Epistemic status

The empirical consistency check demonstrates that the four-factor framework is *consistent* with the observed lepton hierarchy under reasonable substrate parameter assignments. It does not constitute a *derivation* of the hierarchy from substrate primitives — that requires the explicit substrate evaluations enumerated in §13.6. The check does, however, sharpen the structural

prediction: the empirical non-uniformity of the hierarchy requires the dominant contributions to live in S_P and S_I rather than S_L, a falsifiable structural feature.

13.8 Toy calculation: charged-lepton hierarchy

To make the consistency check fully explicit, we work the leading-order formula end-to-end with the §13.5 parameter assignment and compare the resulting mass ratios against observation. The leading-order mass relation is

$$m_g \propto \mathcal{D}_g \cdot S_H(g) \cdot S_L(g) \cdot S_P(g) \cdot S_I(g)$$

with the toy parameter values:

$\mathcal{D} = (1, 2, 4)$ — generation operator (forced by the flavour-mixing programme); $S_L(g) = \exp(0.38 \cdot (g - 1))$ — localization compression from Role-4 scaling $\kappa = 3/8$ with $p = 1$; $S_P(g) = \exp(\sigma_g)$ with $\sigma = (0, 3.5, 4.5)$ — persistent distinguishability bit-count; $S_I(g) = \exp(\tau_g)$ with $\tau = (0, 0.5, 1.0)$ — log-transport-complexity; $S_H(g) = \exp(h_g)$ with $h = (0, 0.26, 0.51)$ — closure-Hessian; hand-chosen of order 0.25 per generation. The §5.3 hypothesis $S_H \sim \sqrt{\lambda_{max}(H_D)} \sim \mathcal{D}_{\gamma_D}$ would predict $(1/2) \ln \mathcal{D} = (0, 0.347, 0.693)$ for these increments — close in magnitude but not identical to the toy values. The toy treats S_H as an independent hand-chosen contribution rather than committing the §5.3 hypothesis at this stage.

All factors are normalised to $g = 1$ (electron) so that the toy directly computes mass ratios m_g / m_e.

Muon (g = 2). Summing the log-contributions:

$$\ln(m_\mu / m_e) = \ln 2 + 0.38 + 3.5 + 0.5 + 0.26 = 0.69 + 0.38 + 3.5 + 0.5 + 0.26 = 5.33$$

Therefore:

$$m_\mu / m_e = \exp(5.33) \approx 206.4$$

Observed value: $m_\mu / m_e \approx 206.77$. Agreement to *0.18%*.

Tau (g = 3). Summing the log-contributions:

$$\ln(m_\tau / m_e) = \ln 4 + 0.76 + 4.5 + 1.0 + 0.51 = 1.39 + 0.76 + 4.5 + 1.0 + 0.51 = 8.16$$

Therefore:

$$m_\tau / m_e = \exp(8.16) \approx 3491$$

Observed value: $m_\tau / m_e \approx 3477$. Agreement to *0.40%*.

Decomposing the empirical non-uniformity. The empirical lepton hierarchy is non-uniform: the gap $\ln(m_\mu / m_e) = 5.33$ is substantially wider than the gap $\ln(m_\tau / m_\mu) = 2.83$. The toy decomposes this non-uniformity contribution-by-contribution:

Contribution	g = 1 → g = 2	g = 2 → g = 3
$\Delta(\ln \mathcal{D})$	0.69	0.69
$\Delta(\ln S_L)$	0.38	0.38
$\Delta\,\sigma$	3.50	1.00
$\Delta(\ln T)$	0.50	0.50
$\Delta((1/2)\ln \lambda_{max})$	0.26	0.25
total	5.33	2.82

Four of the five contributions are essentially uniform across the two transitions ($\Delta(\ln \mathcal{D})$, $\Delta(\ln S_L)$, $\Delta(\ln T)$, $\Delta((1/2)\ln \lambda_{max})$); the *single* source of empirical non-uniformity is $\Delta\sigma$, which drops from 3.5 bits between gen 1 and gen 2 to 1.0 bits between gen 2 and gen 3. This is exactly the structural prediction **P-1** of §13.4: the dominant non-uniformity lives in σ_D (the persistent distinguishability bit-count), not in S_L, S_H, or S_I.

What the toy demonstrates. Three things are worth stating clearly.

First, the four-factor framework *can* reproduce the empirical charged-lepton hierarchy to under half a percent with *modest* substrate-complexity increments — bit-count differentials of order a few committed bits between generations, transport-graph log-complexity differentials of order half a unit. The hierarchy does not require any single large arbitrary number: it emerges from the *sum* of five small log-contributions, each tied to a distinct substrate observable.

Second, the agreement is a *feasibility check*, not a derivation. The toy has effectively four hand-chosen parameters per heavier generation ($\sigma, \tau, h, p\cdot\kappa$); these are tuned to reproduce the empirical ratios subject to the structural constraint that the dominant non-uniformity live in σ (P-1, extracted from the empirical non-uniformity in §13.4). What the toy demonstrates is the *existence* of substrate-parameter assignments that simultaneously (a) fit the empirical ratios and (b) exhibit the P-1 structural pattern — i.e., the P-1 prediction is not in conflict with the empirical hierarchy. Whether substrate primitives actually produce values close to these toy values is the empirical question the next paper must settle. A scenario in which the empirical ratios were $m_\mu/m_e \approx 1000$ and $m_\tau/m_\mu \approx 1000$ (uniform exponential) would require $\Delta\sigma$ uniformly large across generations — itself a substrate-structural prediction subject to falsification by direct enumeration. The non-uniformity of $\Delta\sigma$ between generations is therefore a *substrate observable* that the next computational paper will fix, not a constraint imported from elsewhere.

Third, the toy is one explicit rehearsal of the §13.6 seven-step computational roadmap: with the four substrate quantities ($\lambda_{max}(H_g)$, ℓ_g, σ_g, T_g) supplied as *toy inputs*, the framework produces explicit numerical mass predictions. The next computational paper replaces the toy inputs with values computed from substrate primitives via the seven-step procedure. The toy demonstrates that *if* the substrate computations yield numbers of the indicated magnitudes (a few

bits, half a unit of log-transport, etc.), the framework reproduces the empirical hierarchy. The empirical question is whether the substrate computations *do* yield those numbers — and that is what the next paper must show.

Epistemic status. The toy *demonstrates feasibility*, not *derivation*. The five substrate parameters per generation are hand-chosen to reproduce the empirical ratios; they are not derived from substrate primitives in the present paper. What the toy *does* show is that (a) the empirical hierarchy is *consistent* with the four-factor framework under modest substrate-complexity assignments, (b) the dominant non-uniformity lives precisely where P-1 predicts it should (in σ, not elsewhere), and (c) the framework's predictivity comes from the four substrate quantities being independently computable from the PFD invariant tuple via §13.6's roadmap — not from any free fitting freedom at the present stage. The toy therefore serves as a *waypoint*: a specific, falsifiable target that the substrate computations of the next paper must reach if Theorem 9.1 is to be empirically supported.

(Illustrative; the toy parameters are hand-chosen to reproduce empirical ratios but are constrained by P-1 to exhibit the empirically-observed non-uniform structure in σ. The substantive prediction the next computational paper must verify is that the substrate-derived values of $\sigma_e, \sigma_\mu, \sigma_\tau$ and T_e, T_μ, T_τ exhibit the non-uniform growth pattern $\Delta\sigma_{1\to2} \gg \Delta\sigma_{2\to3}$ that the toy and P-1 both require.)

---

14. Cross-programme synthesis

The substrate stiffness hierarchy developed here sits within the broader VERSF programme. The synthesis with five other parts of the programme is worth making explicit.

14.1 With the matter paper

The matter paper §6.5 mass scaffold enumerated four contributions to PFD mass — commitment-density loading, closure-Hessian stiffness, confinement/localization cost, persistent distinguishability content — without computing any of them. The present paper assigns each to a specific factor in the four-factor decomposition, derives the functional form of each from a distinct substrate mechanism, and identifies what the open computation requires for each factor. The matter paper's scaffold is the structural ancestor of the present paper's framework; this paper makes the scaffold concrete enough to compute on.

14.2 With the PFD–Standard Model Dictionary

The dictionary classified PFD invariant tuples onto Standard Model representation classes via the invariant tuple $\mathcal{I}(\mathrm{D})$. The present paper takes this classification as fixed input and asks: for each classified PFD class, what is the substrate stiffness factor? The decomposition $S(D) = S_H \cdot S_L \cdot S_P \cdot S_I$ converts the dictionary's invariant-tuple classification into a mass formula by assigning each substrate stiffness factor a specific dependence on specific components of $\mathcal{I}(\mathrm{D})$.

The dictionary classifies which PFDs exist; this paper classifies how energetically expensive each is for the substrate to maintain.

14.3 With the closure-norm condensation paper

The closure-norm condensation paper established the mass formula $m_D = \mathcal{D} \cdot S(D) \cdot v$ and decomposed $S(D)$ into four factors in its §11.2a, while explicitly deferring the substrate-mechanism derivation of each factor and the assembly of the combined scaling law to a companion paper. The present paper is that companion paper. The four-factor decomposition is inherited; each factor's substrate mechanism is derived here; the combined scaling theorem (Theorem 9.1) assembles the result. Together, the closure-norm condensation paper and the present paper provide the complete substrate-Higgs + substrate-stiffness mechanism for fermion mass generation.

14.4 With the Role-4 lepton programme

The Role-4 lepton programme derived the exponential localization scaling $\ell_g = \ell_0 \cdot \exp(-\kappa g)$ with $\kappa \approx 3/8$ from CP^2 geometric structure and confinement dynamics. This is the substrate input that drives the exponential structure of $S_L(D)$ via §6.3. The Role-4 programme's localization scaling, applied to $S_L(D) \sim \ell_D^{-p}$, generates the exponential-in-γ_D structure that is one of the two exponential channels in the hierarchy. The present paper therefore extends the Role-4 programme from charged leptons to the full PFD spectrum, with the S_L factor providing the bridge.

14.5 With the flavour-mixing programme

The flavour-mixing programme provided the generation operator $\mathcal{D} = diag(1, 2, 4)$ and the structural framework for CKM and PMNS mixing. The present paper extends this in two ways: it sharpens the role of $\mathcal{D}$ (delivering generation count and ordering but *not* the empirical magnitudes), and it provides the substrate mechanism for the off-diagonal stiffness factors $S_{ij}(D)$ that determine the mixing-angle magnitudes via §12. The flavour-mixing programme provides the basis structure; the present paper provides the substrate-mechanism explanation for why CKM is small and PMNS is large.

14.6 Cross-vocabulary identifications

Concept	Matter paper §6.5	Dictionary	Closure-norm paper	Present paper
Mass formula	$m_D \sim E_{closure}(D)$ (scaffold)	$m_D \sim E_{closure}(D)$ (§19.3)	$m_D = g_D\, v$ (§10)	$m_D = \mathcal{D} \cdot S(D) \cdot v$ (Theorem 9.1)
Hierarchy origin	(open)	(open)	(deferred to present paper)	substrate stabilization complexity

Concept	Matter paper §6.5	Dictionary	Closure-norm paper	Present paper
Closure-Hessian stiffness	contribution (ii)	(implicit in §19.3)	S_H (§11.2a)	$S_H \sim \sqrt{\lambda}_{max}(H_D)$ (§5)
Localization compression	contribution (iii)	(implicit in §19.3)	$S_L \sim \ell^{-p}$ (§11.2a)	$S_L \sim exp(p \cdot \kappa \cdot \gamma_D)$ via Role-4 (§6)
Persistent distinguishability	contribution (iv)	(implicit in §19.3)	$S_P \sim exp(\sigma)$ (§11.2a)	$S_P \sim exp(\sigma_D)$ via Landauer (§7)
Interface transport	contribution (iii)	(implicit in §19.3)	$S_I \sim T$ (§11.2a)	$S_I \sim T(D)$ via transport graph (§8)

The cross-vocabulary table makes explicit that the four-factor framework consolidates and sharpens results that the upstream papers had introduced piecewise.

---

15. Falsification criteria

The framework is structurally falsifiable. It would fail under any of the following observations:

1. **Discovery that fermion masses do not depend on closure topology.** If precision tests revealed that fermion masses were independent of the topological invariants of the PFD invariant tuple (β_1, h_D, C_D), the four-factor framework would be falsified — *S_H, S_P, S_I* all depend on these invariants.
2. **Discovery that the lepton hierarchy is uniformly exponential in generation depth.** §13 established that the empirical lepton hierarchy is *non-uniform* across generations, with a much wider log-gap between e and μ than between μ and τ. If precision measurements revealed the hierarchy was actually uniform — $m_\mu/m_e = m_\tau/m_\mu$ — the structural prediction that *S_P* and *S_I* contribute the dominant non-uniformity would be falsified.
3. **Discovery of a fourth fermion generation.** §11.3 of the closure-norm condensation paper, inherited here, derived $\gamma_D \leq 3$ from refinement persistence. A fourth refinement-stable generation with standard-charge fermion content would refute the bound and the present mass-hierarchy mechanism.
4. **Discovery that quark and lepton masses follow the same hierarchy structure.** The framework predicts structural separation between quark and lepton masses driven by *S_I* (multi-channel confinement) and *S_P* (colour distinguishability). If precision tests revealed quark and lepton masses followed identical hierarchy patterns once normalized to the same electric charge, the structural quark/lepton separation prediction would be falsified.
5. **Discovery that neutrino masses do not arise from ledger-charge suppression.** The framework predicts $m_\nu \ll m_e$ from $\ell_D = 0$ removing dominant contributions to *S_L, S_P, S_I*. If the empirical neutrino mass mechanism were established (e.g., via cosmological constraints or $0\nu\beta\beta$ measurements) to require a distinct substrate origin not reducible to $\ell_D = 0$, the structural prediction would be falsified.

6. **Failure of additive log-space scaling.** Theorem 9.1 predicts $\ln m_D \approx \ln \mathcal{D}_{\gamma_D} + (1/2) \ln \lambda_{max}(H_D) - p \cdot \ln \ell_D + \sigma_D + \ln T(D) + const$. If, in the next computational paper, the substrate evaluations of the four factors did *not* assemble into this additive log-space form — for example, if there were strong cross-couplings between factors that violated the multiplicative decomposition of §4 — the framework would be falsified.
7. **Discovery of substrate non-isotropy that breaks the four-factor decomposition.** The four-factor decomposition assumes that each factor's substrate mechanism operates independently of the others (multiplicative decomposition). If detailed substrate calculations revealed strong inter-factor cross-couplings that prevented this decomposition, the framework's structural assumption would be falsified — though the underlying mass-as-stabilization-cost identification would survive in a more complex form.
8. **Substrate computations yield uniform σ_D growth across generations (P-1 falsification).** Structural prediction P-1 (§13.4, §13.8) predicts that the dominant empirical non-uniformity in the lepton hierarchy must live in the substrate-derived σ_D values, with $\Delta\sigma_{1\to2}$ substantially larger than $\Delta\sigma_{2\to3}$. If the next computational paper's explicit enumeration of $\sigma_e, \sigma_\mu, \sigma_\tau$ from substrate primitives instead yields uniform per-generation growth, the framework cannot reconcile the empirical ratios $m_\mu/m_e \approx 207$ and $m_\tau/m_\mu \approx 16.8$ under any consistent assignment of the other three factors. This is the framework's sharpest internal falsification target: P-1 is structurally embedded in the four-factor framework and is directly testable against the substrate enumeration that §13.6's computational roadmap will produce.

---

16. Epistemic status and open problems

16.1 Epistemic colour-coding

Forced by substrate primitives (proven from inherited results):

— Mass formula $m_D = \mathcal{D} \cdot S(D) \cdot v$ (inherited from closure-norm condensation paper §11.2). — Four-factor decomposition $S(D) = S_H \cdot S_L \cdot S_P \cdot S_I$ as the canonical assignment of substrate mechanisms to mass contributions (§4). — *Persistent distinguishability cost from BC1 + commitment persistence* (§7.1a): the exponential $S_P(D) \sim exp(\sigma_D)$ follows from BC1's conservation of distinguishability combined with the substrate's natural admissibility-relaxation flow. This is a substrate-derived result, with Landauer's principle recovered as the thermodynamic limit, not an imported assumption. — *Graph-theoretic transport complexity* (§8.2a): $T(D) = \alpha_1\beta_1(G_D) + \alpha_2\lambda_2^{-1}(G_D) + \alpha_3\chi(G_D) + \alpha_4 deg\_avg(G_D)$ is a concrete graph-theoretic functional on the admissibility-preserving transport graph, computable from substrate primitives. — Quark/lepton structural separation $S(D)|_q > S(D)|_l$ for same generation depth, driven by PFD-invariant differences in C_D, ρ_D, ℓ_D (§10). — Neutrino structural suppression $S(\nu) \ll S(e)$, driven by $\ell_D = 0$ (§11). — Suppression of non-adjacent generation mixing from $|\mathcal{D}_{\gamma_i} - \mathcal{D}_{\gamma_j}|$ gap structure (§12.3). — Structural prediction $|U_{CKM}| \ll |U_{PMNS}|$ from $\ell_D = 0$ for neutrinos (§12.2). — Computational roadmap (§13.6): the hierarchy problem

reduces to a finite seven-step computational programme on substrate observables, with explicit operator-level inputs at each step.

Leading-order EFT framing (introduced in this paper):

— *Leading-order factorization* (§4.1): the multiplicative decomposition $S(D) = S_H \cdot S_L \cdot S_P \cdot S_I$ is the zeroth-order substrate stiffness hierarchy, with inter-sector corrections Δ_{ab} expected at higher order. The factorization is not exact substrate dogma but the leading-order effective substrate hierarchy theory, analogous in status to leading-order EFT structures elsewhere in physics. Inter-sector corrections are computable from substrate primitives, not new free parameters.

Conditional theorems (under stated assumptions):

— *Substrate Stiffness Hierarchy Theorem (Theorem 9.1)* — conditional on the leading-order factorization (§4.1), Role-4 localization scaling for ℓ_D, and the BC1-derived persistent-distinguishability cost (§7.1a). Gives the additive-in-log-space scaling law for m_D with inter-sector corrections appearing at higher order via §4.1. — *Localization compression exponential* $S_L \sim \exp(p \cdot \kappa \cdot \gamma_D)$ — conditional on Role-4 scaling $\ell_g = \ell_0 \cdot \exp(-\kappa g)$ with $\kappa \approx 3/8$ inherited from the Role-4 lepton programme. — *Two-exponential hierarchy generation* — conditional on Theorem 9.1 plus structural validity of the S_L and S_P exponential forms.

Conditional on prior programmes (with explicit dependencies):

— Mass formula $m_D = \mathcal{D} \cdot S(D) \cdot v$ — conditional on closure-norm condensation paper Theorem 6.1 (Higgs mass bound) plus matter paper §6.5 mass scaffold. — Generation count and ordering — conditional on the flavour-mixing programme's $\mathcal{D} = \mathrm{diag}(1, 2, 4)$. — Bessel-mode-based localization energy scaling — conditional on the closure-field programme's Bessel solutions $K_n(\kappa r)$.

Schematic / structural but not numerically derived:

— Functional form $S_H(D) \sim \sqrt{\lambda_{max}(H_D)}$ (§5.2) — substrate origin clear, but explicit evaluation of $\lambda_{max}(H_D)$ for specific PFD classes is open. — Exponent p in $S_L(D) \sim \ell_D^{-p}$ (§6.2) — structurally constrained positive, but precise value requires explicit substrate-dispersal evaluation. — Specific committed-bit count σ_D for each PFD class (§7.2) — structural lower bound forced, but explicit enumeration is open. — Specific transport-graph complexity $T(D)$ for each PFD class (§8.3) — graph-theoretic definition forced, but explicit graph evaluation is open. — *Toy calculation reproducing lepton ratios* (§13.8) — $m_\mu/m_e \approx 206.4$ and $m_\tau/m_e \approx 3491$ under hand-chosen substrate-complexity assignments; the agreement at 0.2–0.4% level demonstrates *feasibility* but the substrate parameters are not yet derived. The toy exhibits the P-1 non-uniform pattern in σ (consistent with §13.4). — Numerical mass values for specific PFD classes — the entire computational programme remains open; this paper supplies the framework and a feasibility demonstration.

Conjectural at programme level:

— *Specific claim that the dominant hierarchy non-uniformity lives in σ_D alone, with T(D) growth approximately uniform across generations* (§13.8 toy pattern) — this is a *sharper* variant of P-1 than the §13.4 statement requires. §13.4's P-1 says the non-uniformity lives in *S_P* and *S_I* combined; the §13.8 toy realises this with the non-uniformity entirely in *σ_D* and *T(D)* uniform across generations. The toy's sharper pattern is structurally suggested but not derived from substrate primitives; the next computational paper will determine whether *T(D)* grows uniformly or non-uniformly across generations, refining or revising the toy assignment. — Specific numerical relation between *Δσ* across generations and substrate distinguishability content (§7.3) — natural under the BC1-derived persistent-distinguishability cost but unproven from substrate operators. — *Universality hypothesis* (§17, "Universality hypothesis"): the four-factor hierarchy structure is universal across all stable PFD sectors and across all closure-substrate frameworks supporting committed-distinguishability defects. Structurally supported by the genericity of the four substrate mechanisms but not derived from substrate-independent reasoning within the present paper. — Inter-sector corrections *Δ_{ab}* (§4.1): the existence and structure of higher-order corrections to the leading-order factorization is expected, but the specific functional form of each *Δ_{ab}* is open and is a target for the second-quantised stiffness reconstruction companion paper.

16.2 Open problems

★ Explicit numerical evaluation of the four factors [primary open problem]. The structural framework developed here transforms the fermion mass hierarchy problem from "twelve mysterious phenomenological numbers" into "four constrained sub-factors with specific substrate mechanisms and specific dependences on the PFD invariant tuple." Computing each of *S_H, S_L, S_P, S_I* explicitly for specific PFD classes — and verifying that the empirical lepton ratios $m_e : m_\mu : m_\tau \approx 1 : 207 : 3477$ follow from substrate primitives alone — is the next computational target. The problem is now structured rather than residual: it is a calculational programme on substrate operators, not an unconstrained fit.

Localization compression exponent. The exponent *p* in $S_L(D) \sim \ell_D^{-p}$ requires evaluation of the substrate elastic response to compression. Whether $p = 1$ (dimensional natural value) or some larger value depends on the substrate's stress-strain behaviour under closure compression and is part of the open computation.

Persistent-distinguishability bit-count. The committed-bit count *σ_D* for each PFD class requires explicit enumeration of the committed-fold structure. For lepton PFDs the structural lower bound is moderate; for quark PFDs with colour-distinguishability bookkeeping the bound is larger. Establishing the *exact* committed-bit count for each PFD class is the next-level enumeration problem.

Transport-graph complexity. The transport-graph complexity *T(D)* requires evaluating the admissibility-preserving transport network for each PFD class — a graph-theoretic problem on the closure transport architecture inherited from the closure-field programme. For specific PFD classes (electron, muon, tau, up, down, etc.) this evaluation can be set up concretely; computing it is the next step.

Off-diagonal stiffness factors. The off-diagonal stiffness factors $S_{ij}(D)$ governing CKM and PMNS matrix elements are open in the same way as the diagonal $S(D)$. Computing them requires extending the four-factor decomposition to inter-generation closure couplings — a target for the flavour-mixing extension companion paper.

Quark hierarchy verification. The structural prediction that quark mass hierarchy is steeper than lepton mass hierarchy is forced by S_I and S_P differences. Verifying this numerically requires substrate evaluation of S_I^q vs S_I^l and σ_q vs σ_l — extending the lepton consistency check of §13 to the quark sector.

Second-quantised stiffness reconstruction. The present treatment is at the level of mean-field / classical field theory. A full second-quantised treatment of substrate stiffness dynamics — accounting for quantum corrections to the four-factor framework and to specific factor evaluations — remains a programme target.

RG running of stiffness factors. Whether the four factors S_H, S_L, S_P, S_I exhibit RG running between the substrate scale and the empirical electroweak scale is open. If they do, the values relevant for mass determination would be the IR-running values, not the bare substrate values. This is part of the open computational programme.

16.3 Programme positioning

This paper opens the primary open problem identified in the closure-norm condensation paper §17.3 — the explicit substrate-level account of the fermion mass hierarchy. Combined with the closure-norm condensation paper (substrate Higgs mechanism) and the upstream programme content (PFDs, gauge structure, generation operator), it provides the structural framework in which fermion masses become substrate-derivable quantities rather than phenomenological inputs.

The single most pressing remaining target is the *explicit numerical evaluation* of the four factors S_H, S_L, S_P, S_I for specific PFD classes, leading to predicted numerical values for m_e, m_μ, m_τ (and then quarks) compared against the empirical spectrum. This is a calculational programme on substrate primitives — not an unconstrained fit — and is the natural target of the next computational paper.

The remaining major programme targets, beyond this paper and the explicit numerical computation, are: the second-quantised closure-condensate stiffness reconstruction (closing the QFT layer above the substrate stiffness framework), the ν_R PFD existence companion paper (closing dictionary Question A), the flavour-mixing extension to off-diagonal stiffness factors (closing CKM and PMNS computations), and a unified treatment of generation-depth and stabilization-depth dynamics (closing dictionary §14.4).

---

17. Conclusion

The VERSF programme has progressively reconstructed the foundations of gauge structure, geometry, gravity, matter ontology, electroweak symmetry breaking, and now — with the present paper — the structural origin of the fermion mass hierarchy.

The central result of the present paper is structural:

Fermion mass is the substrate energetic cost of maintaining a stable closure geometry against admissibility decay and dispersive distinguishability flow. The substrate stiffness factor S(D) decomposes into four physically distinct contributions — closure-Hessian curvature, localization compression, persistent distinguishability load, and interface transport complexity — each constrained by a specific substrate mechanism and depending on specific components of the PFD invariant tuple.

This identification produces a unified picture in which:

— Mass is *not* a fundamental property of particles but a measure of substrate *stabilization complexity*: how energetically difficult it is for the substrate to maintain a given closure geometry.

— The substrate stiffness factor $S(D)$ decomposes canonically into four factors $S_H \cdot S_L \cdot S_P \cdot S_I$, each tied to a distinct substrate mechanism and a distinct subset of PFD invariant-tuple components.

— Generation hierarchy arises from *two exponential channels* — localization compression $S_L \sim \exp(p \cdot \kappa \cdot \gamma_D)$ from Role-4 localization scaling, and persistent distinguishability load $S_P \sim \exp(\sigma_D)$ from Landauer cost — together with polynomial contributions from S_H and S_I.

— Quark/lepton mass separation is *structurally forced* by PFD invariant differences (C_D, ρ_D, ℓ_D), driven primarily by additional contributions to S_I (multi-channel confinement) and S_P (colour-charge distinguishability) for quarks.

— Neutrino mass suppression is *structurally forced* by $\ell_D = 0$, removing dominant contributions to three of the four factors and producing $S(\nu) \ll S(e)$ without parameter tuning.

— CKM and PMNS mixing structure arises from off-diagonal stiffness factors $S_{ij}(D)$ assembled from the same four-factor mechanism applied to inter-generation closure couplings, with quark CKM ≪ lepton PMNS structurally forced by the same $\ell_D = 0$ mechanism that suppresses neutrino mass.

Quantitatively, the framework establishes:

— *Structural prediction.* Additive log-space scaling: $\ln m_D \approx \ln \mathcal{D}_{\gamma_D} + (1/2) \ln \lambda_{max}(H_D) - p \cdot \ln \ell_D + \sigma_D + \ln T(D) + const$ (Theorem 9.1) — the log-mass is the sum of independent substrate stabilization free-energy contributions. — *Structural prediction.* Two-exponential hierarchy generation, naturally accommodating the empirical twelve-order-of-magnitude span. — *Structural prediction (P-1, §13.4).* The dominant non-uniformity across

generations lives in σ_D and $T(D)$, with $\Delta\sigma_{1\to2} \gg \Delta\sigma_{2\to3}$. The next computational paper must find this pattern when σ_D is computed from substrate primitives; otherwise the framework is falsified at §15 criterion 8. — *Structural prediction.* Quark mass hierarchy steeper than lepton mass hierarchy, with hierarchy steepness increasing with generation depth. — *Structural prediction.* Split neutrino spectrum $m_{\nu,1} < m_{\nu,2} < m_{\nu,3}$ set by ledger-charge bit count $\Delta\sigma$ and transport ratio T_ν/T_e — both pushing the spectrum below the charged-lepton scale, with the lightest eigenstate at $m_{\nu,\text{lightest}} \sim 1\text{–}50\ meV$ depending on the substrate-derived value of $\Delta\sigma$. Under the normal-hierarchy scenario $m_{\nu,1} \approx 1.5\ meV$, the framework requires a *structurally specific gen 1 → gen 2 amplification of order six* in $S(\nu)$ via $\mathcal{D}$, S_L, S_P, S_I at γ_D = 2 — a sharper computational target alongside the charged-lepton P-1. — *Feasibility demonstration (§13.8).* Charged-lepton hierarchy reproduced under hand-chosen but modest substrate-complexity increments. The toy gives $m_\mu/m_e \approx 206.4$ vs observed 206.8 (0.2% agreement) and $m_\tau/m_e \approx 3491$ vs observed 3477 (0.4% agreement). The toy *fits* rather than *predicts* the ratios — the substrate quantities are hand-chosen — but establishes that the empirical hierarchy is consistent with the four-factor framework under bit-count differentials of order a few committed bits per generation.

Qualitatively, the framework converts the fermion mass hierarchy from twelve arbitrary phenomenological inputs into a structured calculational programme on substrate primitives. Four substrate quantities (closure-Hessian eigenvalues, localization-compression exponent, persistent-distinguishability bit counts, transport-graph complexities) — each independently computable from the PFD invariant tuple — together produce the entire hierarchy.

Universality hypothesis

The four-factor hierarchy structure is conjectured to be *universal* across all stable PFD sectors, independent of the detailed microscopic realization of the closure substrate. The substrate-level mechanisms that generate the four factors — free-energy curvature, dispersal cost, BC1-derived informational maintenance, and admissibility-preserving transport — are not specific to the particular geometric realization of the closure architecture (hexagonal tiling, K = 7, etc.) but follow from the general structure of *any* substrate that supports committed-distinguishability defects against admissibility relaxation. If the universality hypothesis holds, the observed fermion hierarchy reflects a *universality class* of admissibility-preserving closure stabilization rather than a model-dependent parameterization tied to specific substrate microstructure.

Three structural features support the universality conjecture:

— The four-factor decomposition arises from logically distinct stabilization mechanisms, not from features of any particular substrate geometry. Any substrate that supports persistent localized closure structures will exhibit (i) a free-energy Hessian at the defect's support, (ii) an energetic dispersal cost for localization, (iii) a Landauer-derived informational maintenance cost, and (iv) a transport-network complexity. The four factors are *generic substrate observables* of any admissibility-preserving closure system.

— The exponential structures of S_L and S_P arise from generic mechanisms (energy density scaling with coherence-radius inverse-power; Boltzmann factors from additive bit costs) that do

not depend on substrate-specific geometric details. They would appear in any substrate exhibiting Bessel-like localization modes and BC1-like distinguishability conservation.

— The graph-theoretic structure of S_I depends on the *topological* properties of the admissibility-preserving transport network, not on its specific embedding in any particular substrate geometry. Different microscopic substrate realizations could share the same transport-graph topology (and therefore the same $T(D)$) while differing in geometric details.

If the universality conjecture is correct, the framework's predictive content is broader than the specific VERSF realization: the fermion mass hierarchy would emerge from substrate stabilization complexity in *any* closure-based substrate framework, with the same four-factor scaling structure. This positions the framework as a candidate *universality class* of substrate-derived fermion mass theories — a structural status that places it alongside other universality classes in physics (the Ising universality class for second-order phase transitions, the Wilson-Fisher fixed point for critical phenomena, etc.) where the macroscopic behaviour is determined by symmetry and dimensionality rather than microscopic details.

(Conjectural at programme level; the universality hypothesis is a structural prediction of the framework but is not derived from substrate primitives within the present paper. Testing it requires either explicit comparison with alternative substrate frameworks sharing the same admissibility/closure structure, or a direct substrate-independent proof that the four-factor mechanism follows from generic stabilization-cost reasoning.)

The broader implication of this picture:

Particle masses are not arbitrary inputs to nature. They are measurable expressions of how difficult different closure structures are for the substrate to maintain coherently. Mass, flavour, confinement, and generation hierarchy all become manifestations of one deeper principle — the energetic cost of preserving committed distinguishability against substrate dispersal.

In this view, the fermion mass spectrum is a *substrate readout*: each measured mass is a window into the closure topology, localization length, distinguishability content, and transport complexity of the corresponding PFD class. The Standard Model's twelve mysterious Yukawa numbers become twelve substrate-structural quantities, each independently computable from inherited substrate primitives. What remains is the computation — which now has a shape.

Appendix A. Dependency structure

The present paper depends on the following inherited substrate results. The dependency map is intended as a navigation aid and a check on non-circularity: every claim traces back to upstream programme content, and no claim circularly depends on downstream phenomenological inputs.

Result	Substrate-level dependencies
Mass formula $m_D = \mathcal{D} \cdot S(D) \cdot v$ (§2.4)	closure-norm condensation paper §11.2 + Higgs Theorem 6.1
Four-factor decomposition (§4)	closure-norm condensation paper §11.2a + matter paper §6.5 mass scaffold
Leading-order EFT factorization (§4.1)	independent-mechanism structure of §4 + inter-sector correction expansion
Closure-Hessian stiffness $S_H \sim \sqrt{\lambda_{max}(H_D)}$ (§5)	closure-norm condensation paper §4 free-energy functional + matter paper §6.5 contribution (ii)
Localization compression $S_L \sim \ell_D^{-p}$ (§6)	Bessel localization modes (closure-field programme) + Role-4 localization scaling
Exponential $S_L \sim exp(p \cdot \kappa \cdot \gamma_D)$ (§6.3)	Role-4 lepton programme $\ell_g = \ell_0 \cdot exp(-\kappa g)$ with $\kappa \approx 3/8$
BC1-derived persistent distinguishability cost (§7.1a)	BC1 (bit conservation) + commitment persistence + admissibility-relaxation flow; Landauer's principle as thermodynamic limit
Persistent distinguishability $S_P \sim exp(\sigma_D)$ (§7)	BC1-derived cost (§7.1a) + matter paper §6.5 contribution (iv)
Lower bound *σ_D ≥ β₁ +	χ_D
Transport complexity $S_I \sim T(D)$ (§8)	closure transport network (closure-field programme) + matter paper §6.5 contribution (iii)
Graph-theoretic $T(D) = \alpha_1\beta_1 + \alpha_2\lambda_2^{-1} + \alpha_3\chi + \alpha_4 deg\_avg$ (§8.2a)	admissibility-preserving transport graph G_D + substrate-derived coefficients
Substrate Stiffness Hierarchy Theorem (§9, Theorem 9.1)	All four §5–§8 derivations + leading-order factorization (§4.1)
Quark/lepton separation (§10)	PFD invariant differences in C_D, ρ_D, ℓ_D (dictionary)
Neutrino suppression (§11)	Neutrino PFD invariants $\ell_D = 0$ (dictionary §12.2)
Off-diagonal stiffness S_{ij} (§12)	closure-norm condensation paper §11.4 + four-factor decomposition
*	U_CKM
Empirical consistency check (§13)	Theorem 9.1 + empirical lepton mass values
Toy calculation reproducing lepton ratios (§13.8)	Theorem 9.1 + hand-chosen substrate-complexity parameters reproducing empirical lepton ratios under the P-1 structural constraint
Computational roadmap (§13.6)	Theorem 9.1 + operator-level definitions of λ_{max}, ℓ_D, σ_D, $T(D)$

Result	Substrate-level dependencies
Falsifiability criteria (§15)	Theorem 9.1 + PFD invariant-tuple structure
Universality hypothesis (§17)	Generic substrate stabilization mechanisms underlying each of S_H, S_L, S_P, S_I
Cross-vocabulary identifications (§14.6)	matter paper §6.5 + dictionary §19.3 + closure-norm condensation paper §11.2a + present paper §4

The present paper introduces *no new free continuous parameters* beyond those already inherited from upstream programme content. The four substrate quantities ($\lambda_{max}(H_D), p, \sigma_D, T(D)$) that determine the stiffness factor are structurally constrained by the PFD invariant tuple and are independently computable from substrate primitives — they are not free fits to empirical data. The Role-4 scaling parameter $\kappa \approx 3/8$ is inherited from the Role-4 lepton programme, not introduced here. The empirical electroweak scale $v \approx 246\ GeV$ is inherited from the closure-norm condensation paper's empirical anchor. No additional empirical inputs are required to set up the present paper's framework.

The present paper's contribution is therefore *purely structural*: it converts the fermion mass hierarchy from an unconstrained collection of phenomenological inputs into a structured calculational programme on substrate primitives. The next computational paper will perform the explicit evaluations and produce numerical predictions.