First Toy Reconstruction of the Charged-Lepton Hierarchy in VERSF

Closure Graphs, Hessian Spectra, Transport Complexity, and a First Computational Evaluation of the Substrate Stiffness Factors

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

This paper takes the next concrete step in the VERSF programme on where particle masses come from.

The Standard Model contains twelve numbers — the masses of the fundamental matter particles — that have no derivation within it. They are inserted by hand to match what experiments measure. The VERSF programme aims to replace those twelve numbers with a smaller set of computable quantities derived from the substrate it proposes underlies physics. An earlier paper, *The Substrate Stiffness Hierarchy in VERSF*, set up the framework: four substrate observables that together determine the mass of each particle.

The present paper performs the first proof-of-concept calculation. It asks whether the framework can actually be run end to end on a real piece of physics — and if so, what runs the machinery. We chose the simplest case: the electron, muon, and tau, three increasingly heavy versions of the same kind of particle. The toy calculation reproduces the muon at 205 times the electron mass (observed: 207) and the tau at 3 450 times (observed: 3 477).

The numerical fit is not the point. The toy uses inputs that are chosen consistent with the answer we want, so reproducing the answer isn't surprising in itself. What *is* informative is which of the four ingredients carries most of the inter-generational difference. In this toy reconstruction it is the one called *persistent distinguishability load* — the running cost the substrate pays to maintain what makes one particle persistently distinct from another. If that result survives a derivation from substrate first principles (which the framework now sets up as a concrete next step), the spread in particle masses would turn out to be primarily an information-keeping cost — not a geometric or dynamical one. That is a substantial conceptual reframing.

The main contribution of the paper is therefore the framing. We've replaced "insert twelve constants" with "compute four observables, of which one in particular has to come out right — and that one can now be broken down into pieces with separate physical meanings, so that the framework can be tested in pieces rather than only as a whole." The toy reconstruction's value is not that it computes the lepton masses, but that it makes the framework's failure points concrete, so each can be tested independently. That sharp failure surface — the framework can be shown wrong in specific, identifiable places — is what we believe the toy reconstruction earns. The detailed structural and quantitative content lives in the technical sections below.

Table of Contents

Abstract

The previous paper, *The Substrate Stiffness Hierarchy in VERSF*, established the structural framework for fermion mass generation. Fermion masses were interpreted as substrate stabilization costs of Persistent Fold Defects (PFDs), and the hierarchy was encoded in four substrate stiffness factors:

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

That paper stopped at the structural level: it did not demonstrate that the machinery could be instantiated concretely on a real fermion sector.

The present paper performs that instantiation. It executes the first explicit toy computational reconstruction of the framework for the charged leptons. The goal is not a final first-principles derivation of the empirical masses, but a **proof-of-computability** demonstration: that the hierarchy machinery is well-defined on substrate observables rather than on arbitrary Yukawa insertions.

We construct simplified admissibility-preserving transport graphs for the electron, muon, and tau PFD classes; adopt toy closure-Hessian spectra coupled to the transport complexity through a

single toy linkage $k\_D = T(D)$; compute localization-compression scaling from Role-4 geometry with a back-fit localization exponent $p \approx 0.233$; posit persistent distinguishability loads $\sigma\_\mu = 3.5$, $\sigma\_\tau = 4.5$, supported both by a heuristic substrate-counting argument (§6.2) and by an explicit three-component substrate-counting model $\sigma\_D = a \cdot \Delta\beta_1 + b \cdot N\_persistent + c \cdot R\_D$ (§6.4); and calibrate the transport functional by a single muon-scale condition, yielding $\ln T(D) \approx 0.5 \cdot \Delta\beta_1$ to leading order in the graph invariants. These quantities are assembled through the leading-order substrate stiffness identity.

The reconstruction reproduces the observed charged-lepton hierarchy:

$m\_e : m\_\mu : m\_\tau \approx 1 : 205 : 3\,450$

against the observed 1 : 206.8 : 3 477, with all three mass-ratio residuals below 1 %. We are explicit (§8.6) that at this stage the lepton mass ratios cannot falsify the framework. They are reproduced by inputs that include the σ_D values, which were chosen with the target hierarchy in view. Genuine prediction enters only once σ_D (and its substrate-counting components), β_1, and the closure-Hessian/transport linkage are derived independently of the mass spectrum.

We additionally show (§8.5) that once three core structural assumptions are fixed — non-uniform σ_D growth, β_1-dominant transport, and the Hessian/transport linkage — the tau-scale ratio follows without further independent tuning. The structural results that *are* corroborated at the toy level are:

1. **Localization compression alone cannot reproduce the hierarchy.** It supplies an exponential backbone but is sub-dominant at all generations.
2. **Persistent distinguishability dominates the inter-generational ratio, contingent on the toy σ_D values.** It supplies 65.8 % of the muon log-ratio and 55.2 % of the tau log-ratio. The §6.4 substrate-counting model decomposes σ_D into three structurally-interpretable components.
3. **Transport-network complexity supplies the non-uniformity driver** between generations, conditional on the β_1-dominant calibration.
4. **The hierarchy is structurally stable** under perturbation, with mass-ratio uncertainties of order 30–50 % under ±10 % independent input variation.

The significance is structural. The fermion mass problem is transformed from an unconstrained set of phenomenological insertions into a finite computational programme over four well-defined classes of substrate observable. This is the first explicit execution of that programme, and the first to identify a sharp failure surface for the framework — the failure-mode table of §12.3 enumerates the specific outcomes under which the framework would be shown wrong.

The toy reconstruction should not be read as a claim that the lepton masses have been derived. Its value is sharper: it identifies which substrate quantities must take which approximate values for the VERSF hierarchy mechanism to work, decomposes the dominant input (σ_D) into structurally interpretable sub-components, and exposes a specific failure surface (§12.3). The framework is no longer protected by vague emergence language.

1. Introduction

The previous hierarchy paper transformed the fermion mass problem from a list of arbitrary Yukawa couplings into a structured substrate computation problem. The central mass relation was:

$m\_D = D^{\gamma\_D} \cdot S\_H(D) \cdot S\_L(D) \cdot S\_P(D) \cdot S\_I(D) \cdot v$

with:

Quantity	Interpretation
D^{γ_D}	generation-depth operator
v	closure-condensate vacuum scale
$S\_H(D)$	closure-Hessian stiffness
$S\_L(D)$	localization compression
$S\_P(D)$	persistent distinguishability load
$S\_I(D)$	transport-network complexity

The present paper has a single, sharply defined purpose:

perform the first explicit toy reconstruction of the framework.

This is not a final exact derivation, and it is not a fully microscopic substrate computation. It is a demonstration that the hierarchy machinery actually runs end to end on a real fermion sector — and that, when it runs on plausible toy inputs, the output is recognisable as the observed mass spectrum.

We construct toy PFD transport graphs, evaluate toy Hessian spectra, compute localization scaling, posit distinguishability loads, calculate graph invariants, and assemble the resulting hierarchy. The sector chosen is the charged leptons — the simplest non-trivial generation sequence — comprising the electron, muon, and tau.

Throughout the paper we distinguish carefully between three classes of statement:

- **Toy inputs.** Values assigned to substrate observables for illustrative purposes. These are not derived in this paper. The σ_D values, the localization exponent p, the transport coefficient α_1, and the closure-Hessian/transport linkage $k\_D = T(D)$ all fall in this class. The σ_D values are the most consequential: §10 shows that they carry most of the inter-generational log-ratio, §6.4 provides a structurally-interpretable decomposition of them, and §11 explicitly identifies them as the framework's principal vulnerability.
- **Consistency outcomes.** Quantities computed from the toy inputs via the substrate stiffness identity — most notably the reconstructed mass ratios. These demonstrate that

the machinery is internally consistent and that the inputs reach the observed ratios; they do not constitute predictions, because the inputs were chosen with the ratios in view.

- **Structural outputs.** Qualitative conclusions independent of the precise toy values: which substrate observables dominate, perturbative stability against input variation, computational executability of the framework as a whole. These are the principal scientific content of the paper.

The forward programme (§12) makes explicit which inputs the framework will live or die by under microscopic derivation. The toy reconstruction's role is to make those forward-programme questions sharply formulable — not to settle them.

---

2. Toy Charged-Lepton PFD Models

We begin by constructing simplified toy invariant tuples for the charged leptons. The full PFD invariant tuple, as fixed in the previous paper, is:

$I(D) = (C\_D, \beta_1, h\_D, \pi\_D, \chi\_D, \gamma\_D, \ell\_D, \rho\_D)$

with the components carrying their previously assigned interpretations:

Component	Meaning
C_D	closure completeness class
β_1	first Betti number of the refinement-persistent transport graph
h_D	hypercharge label
π_D	parity assignment
χ_D	charge-parity flavour label
γ_D	generation-depth index
ℓ_D	localization length
ρ_D	persistent record-coupling state

For the toy computation we retain only the hierarchy-relevant components. Within the charged-lepton sector, C_D = complete, the hypercharge label, the parity assignment (+1), the flavour label (L), and the record-coupling state (free) are common to all three generations and cancel in generation-level ratios. The hierarchy-relevant variation lives in (β_1, γ_D, ℓ_D), together with the derived quantities σ_D, $T(D)$, and λ_max(H_D).

A notational caution before the per-particle tables: $T(D)$ and λ_max(H_D) are taken equal at the toy level under the closure-stiffness/transport-complexity linkage adopted in §5. They are *not* independent inputs. We list them as separate columns for transparency, but each muon and tau table entry below carries the same value in both rows by construction.

2.1 Electron

$I(e)$ = (complete, 1, h_1, +1, L, 1, ℓ_e, free)

Observable	Value
γ_D	1
σ_D	0
$T(D)$	1 (reference)
λ_max(H_D)	1 (= $T(D)$ by linkage; reference)

The electron is the reference PFD: minimal generation depth, minimal closure bookkeeping, minimal refinement-stable transport structure, and minimal closure curvature. All log-ratios are normalised against this reference.

2.2 Muon

$I(\mu)$ = (complete, 2, h_2, +1, L, 2, ℓ_μ, free)

Observable	Value
γ_D	2
σ_D	3.5 (toy input; structural decomposition in §6.4)
$T(D)$	e^0.5 ≈ 1.65 (set by §7.3 calibration)
λ_max(H_D)	e^0.5 ≈ 1.65 (= $T(D)$ by linkage)

2.3 Tau

$I(\tau)$ = (complete, 3, h_3, +1, L, 3, ℓ_τ, free)

Observable	Value
γ_D	3
σ_D	4.5 (toy input; structural decomposition in §6.4)
$T(D)$	e^1.0 ≈ 2.72 (consequence of §7.4)
λ_max(H_D)	e^1.0 ≈ 2.72 (= $T(D)$ by linkage)

The toy values reflect the qualitative expectation that higher-generation PFDs carry more closure bookkeeping (higher σ_D), more refinement-stable transport structure (higher $T(D)$), and consequently stiffer closure curvature (higher λ_max). The per-generation increments are deliberately modest — fractional in the exponents, not exponential — so that the resulting hierarchy is driven by the structure of the substrate stiffness identity rather than by hand-tuned input gaps. The σ_D values in particular are chosen consistent with the target hierarchy; §6.2 supplies a heuristic substrate-counting argument and §6.4 a three-component structural decomposition that together place the values within a structurally interpretable model.

---

3. Generation-Depth Contribution

The generation operator is inherited from the flavour hierarchy programme of the substrate stiffness paper:

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

so that:

Particle	D^{γ_D}	$\ln D^{\gamma_D}$
e	1	0
μ	2	0.693
τ	4	1.386

As established previously, this contribution alone is many orders of magnitude too small to account for the observed hierarchy. Its role is to provide generation ordering — a discrete index counting the depth of admissible refinement — not the dominant amplification.

---

4. Localization Compression

Localization scaling follows the Role-4 structure of the substrate refinement geometry. We use the canonical form:

$\ell_g = \ell_0 \cdot e^{-\kappa g}$, $\kappa = \sqrt{8/3} \approx 1.633$, $g = \gamma_D - 1$

and the corresponding stiffness contribution:

$S_L(D) = \ell_D^{-p}$

We adopt the toy value

$p = 0.380 / \kappa \approx 0.233$

so that $\ln S_L = \kappa \cdot p \cdot g = 0.380 \cdot g$. This gives:

Particle	g	$\ln S_L$
e	0	0.00
μ	1	0.380
τ	2	0.760

The choice of p deserves an honest framing. The numerical value adopted here is **not derived**. It is a back-fit: p is chosen so that ln S_L per generation lands at 0.380, which positions localization as a moderate but sub-dominant contributor in the log-share decomposition of §10. A first-principles derivation of p from Role-4 substrate refinement geometry is deferred to the forward programme. Stated more bluntly: the choice of p reflects the framework's assignment of *which* contribution drives the hierarchy, not an independent geometric measurement. If a Role-4 derivation produces p markedly larger or smaller than 0.233, the log-share decomposition of §10 will need to be redone accordingly, and localization could move from sub-dominant to comparable with σ_D.

What is generic, independent of the value of p, is the *form* of localization scaling: an exponential backbone with constant per-generation increment, fixed entirely by κ and p. Localization scaling is purely geometric and carries no generation-specific bookkeeping. As a result, it contributes an exponential backbone of structure across generations but cannot, on its own, supply the inter-generational gradients required to reach $m_\tau / m_e \approx 3\,477$.

---

5. Closure-Hessian Toy Spectrum

5.1 Toy curvatures and the linkage

We construct a toy closure Hessian from the simplified local closure functional:

$F[\rho] = \frac{1}{2} k_D \rho^2 + \frac{1}{4} \lambda \rho^4$

so that, evaluated at the closure equilibrium,

$H_D = \delta^2 F / \delta\rho^2 = k_D + 3\lambda\rho^2$

The toy effective curvatures we adopt are:

Particle	k_D
e	1.00
μ	1.65
τ	2.72

We have adopted a **toy-level closure-stiffness/transport-complexity linkage**:

$k_D = T(D)$

This is a working assumption being tested for consistency, not a relation derived in this paper. Physically, it captures the expectation that PFDs supporting more elaborate refinement-stable transport must also be stiffer closure structures, because the same internal architecture that supports richer transport modes also supplies more constraint against closure dispersal.

The Hessian contribution to the stiffness identity is:

$S\_H(D) = \sqrt{\lambda\_max(H\_D)}$

giving the logarithmic contributions:

Particle	½ · ln λ_max(H_D)
e	0.00
μ	0.250
τ	0.500

5.2 Sensitivity to relaxing exact equality

The previous draft adopted exact equality $k\_D = T(D)$. Claim 3 of the falsifiability statement (§12.2) only requires the linkage to hold within ≈ 20 %. The gap matters: exact equality silently halves the effective parameter count (the Hessian sector ceases to be an independent input) and contributes 0.25 / 0.50 log-units to the muon / tau totals respectively. If the linkage is genuinely approximate rather than exact, the toy reconstruction's accuracy depends on how much slack the linkage actually has.

We therefore test the sensitivity of the reconstruction to relaxing exact equality. Let

$k\_D = \alpha \cdot T(D)^\beta$

with α, β perturbed by O(1) around the baseline (α, β) = (1, 1). Three variant cases:

Variant	(α, β)	½ ln k_μ	½ ln k_τ	m_μ/m_e	m_τ/m_e	m_τ/m_μ
Baseline	(1.0, 1.0)	0.25	0.50	205	3 450	16.83
A	(1.0, 1.2)	0.30	0.60	215	3 812	17.73
B	(0.8, 1.0)	0.14	0.39	184	3 081	16.74
C	(1.2, 0.8)	0.29	0.49	214	3 419	15.97

Two observations:

1. **An overall rescaling α ≠ 1 with β = 1 shifts both Hessian contributions equally and therefore preserves m_τ / m_μ exactly.** This is a generic structural feature: the inter-generational ratio is insensitive to the overall normalisation of the linkage. Only β ≠ 1 perturbs the ratio.
2. **O(1) variation of (α, β) around (1, 1) shifts m_μ / m_e and m_τ / m_e by 5–15 %, and m_τ / m_μ by up to 5 %.** This is well within the ±20 % tolerance set by Claim 3.

The honest reading is that the toy adopts the exact-equality limit of a relation that, structurally, only needs to hold within O(20 %) for the reconstruction to land. The Hessian sector is therefore doing *some* concealed parameter-count reduction work — exact equality is stronger than the framework requires — but the work is bounded: relaxing equality to a 20-%-band linkage costs at most ~15 % of mass-ratio accuracy. We retain exact equality for the body of the reconstruction in the interest of clean bookkeeping, with the understanding that the linkage carries the toy-level slack documented here.

A further note worth flagging for the eventual quark extension: variant A above ($\beta = 1.2$) already overshoots the tau by roughly 10 %, which is borderline for Claim 3. Any β meaningfully greater than unity will progressively break the linkage at higher generations, where Hessian curvatures grow further. This propagates into a quark-scale failure mode that is not visible at the lepton scale, and constrains the form a derived $k_D = \alpha \cdot T(D)^\beta$ relation may take if the framework is to extend coherently beyond the charged leptons.

---

6. Persistent Distinguishability

This section warrants more space than the others because, on the toy inputs adopted here, σ_D is the largest single contribution to the inter-generational hierarchy. The robustness of every quantitative conclusion in this paper turns on whether the $\sigma_\mu \approx 3.5$ and $\sigma_\tau \approx 4.5$ values are roughly correct under a first-principles enumeration.

6.1 The distinguishability factor

The distinguishability factor is:

$$S_P(D) = e^{(\sigma_D)}$$

with the toy distinguishability loads:

Particle	σ_D
e	0 (reference)
μ	3.5 (toy input)
τ	4.5 (toy input)

The physical interpretation is direct: higher-generation leptons carry more committed closure bookkeeping — more persistent records of distinguishability that the substrate must continuously maintain against re-coherence. Each such record is a persistent record-current commitment in the sense of the record-current sector papers. The cost of maintaining those records is paid in stabilization energy, and therefore in mass.

6.2 Heuristic substrate-counting estimate

The previous draft did not say why one might expect σ_μ near 3–4 and σ_τ near 4–5 from substrate considerations. The argument below is a heuristic, not a derivation, but it positions the toy σ_D values within an order of magnitude of what a structurally honest count would produce.

In substrate terms, a refinement-persistent PFD must maintain bookkeeping for each persistent record-current commitment that distinguishes it from a featureless reference. The simplest counting argument runs as follows. For each refinement-persistent admissibility loop in the transport graph, the substrate must maintain:

1. **The loop's persistence.** One commitment per loop, ensuring the loop does not dissolve under substrate refinement.
2. **The loop's orientation/handedness** in admissibility space. One commitment per loop.
3. **The loop's relative phase** against other loops in the same PFD. One commitment per pair of loops.
4. **The branch-point connections** at vertices where loops meet. Roughly one commitment per branch point.

For the **one-loop reference** (electron): 1 persistence + 1 orientation + 0 phases + 0 branch points = 2 raw commitments. We take this as the baseline; $\sigma_e = 0$ in the reconstruction is a *relative* zero.

For the **two-loop PFD** (muon): 2 persistence + 2 orientation + 1 pairwise phase + ~2 branch points ≈ 7 raw commitments. Net over the electron reference: ≈ 5 commitments.

For the **three-loop PFD** (tau): 3 persistence + 3 orientation + 3 pairwise phases + ~3 branch points ≈ 12 raw commitments. Net over the electron reference: ≈ 10 commitments.

The substrate-thermodynamic load σ_D is *not* the raw commitment count: it is the cost weighted by the persistence cost of each commitment, with sub-unit weights expected because substrate bookkeeping infrastructure is shared between commitments. This sub-linear behaviour — the same way a database has higher startup cost per table than per row — is made structurally explicit in §6.4 below.

A note on how this heuristic relates to the structural model of §6.4: the heuristic reaches $\sigma_\mu \approx 3.5$ with effective weight ≈ 0.7 per raw commitment and $\sigma_\tau \approx 4.5$ with effective weight ≈ 0.45 per raw commitment — that is, the heuristic implements the sub-linear growth through a weight that decreases between generations. The §6.4 structural model recovers the same sub-linear growth from the opposite direction: it holds the per-commitment weight constant and lets the substrate counts $N\_persistent$ and R_D saturate between μ and τ. The two pictures are not in conflict. They are two formulations of the same physical assumption — that substrate bookkeeping infrastructure is shared between commitments rather than duplicated — with §6.4 making the assumption sharper, discrete, and locally testable as Claim 5 in §12.2.

6.3 What the heuristic does and does not establish

The heuristic does *not* derive $\sigma_\mu = 3.5$ and $\sigma_\tau = 4.5$. It establishes only that σ values in the "few" range, with sub-linear growth between generations, are structurally plausible for refinement-persistent PFDs of the relevant graph complexities. A genuine first-principles enumeration could land σ_μ anywhere in the range ≈ 2–6 and σ_τ anywhere in the range ≈ 3–7 without surprising the heuristic. The toy values sit comfortably inside that range.

What the heuristic does establish is the order-of-magnitude consistency: $\sigma_D \sim \mathcal{O}(\text{few})$ per additional refinement-persistent loop is the natural substrate scale, and $\sigma \sim \mathcal{O}(50)$ or $\sigma \sim \mathcal{O}(0.1)$ would be structurally surprising. This is the level of motivation that the toy values warrant at the present stage of the programme; sharper motivation is provided by the explicit three-component model below.

6.4 Minimal structural model for σ_D

Rather than treating σ_μ and σ_τ as free numbers, we introduce a substrate-counting decomposition that exposes their internal structure. The proposed form is:

$\sigma_D = \sigma_{base} + \sigma_{cycle}(D) + \sigma_{record}(D) + \sigma_{refine}(D)$

with three substrate-counting components and a (currently zero) base offset:

Component	Form	Substrate interpretation
σ_{base}	constant	reference offset; set to 0 in this normalisation
$\sigma_{cycle}(D)$	$a \cdot \Delta\beta_1(D)$	cost of maintaining additional admissibility cycles
$\sigma_{record}(D)$	$b \cdot N_{persistent}(D)$	cost of maintaining persistent record commitments
$\sigma_{refine}(D)$	$c \cdot R_D$	cost of maintaining refinement infrastructure

A parameter-count caveat before going further: this decomposition fits two totals (σ_μ, σ_τ) with six numbers (three weights a, b, c plus three substrate counts at the muon scale, with one further count at the tau scale). As a pure parameter count it is therefore massively underdetermined — many (a, b, c) and count combinations reproduce (3.5, 4.5). What carries falsifiable content is not the specific weight values but the **saturation structure** between μ and τ ($N_{persistent}$ and R_D roughly constant across this transition), isolated explicitly as Claim 5 in §12.2. The specific (a, b, c) values below should be read as one structurally-interpretable instantiation of the decomposition, not as the decomposition itself.

Taking the toy weights $a = 1.0$, $b = 0.5$, $c = 0.5$ and the toy substrate counts below:

Particle	$\Delta\beta_1$	$N_{persistent}$	R_D	σ_{cycle}	σ_{record}	σ_{refine}	σ_D
e	0	0	0	0	0	0	0
μ	1	3	2	1.0	1.5	1.0	3.5
τ	2	3	2	2.0	1.5	1.0	4.5

The structural reading is now explicit:

- **σ_cycle grows linearly** with the number of additional admissibility cycles. This is the only component that distinguishes the tau from the muon at this toy level.
- **σ_record saturates at the muon scale** (N_persistent = 3 for both μ and τ). The interpretation is that the persistent record-current commitments established for the muon's two-loop structure are reused, not duplicated, when a third loop is added.
- **σ_refine saturates at the muon scale** (R_D = 2 for both μ and τ). The refinement infrastructure required to support multi-loop refinement-persistent structures is established at the first transition (electron → muon) and shared with the tau.

The sub-linear growth from σ_μ to σ_τ (3.5 → 4.5, only 1.0 increment, against the σ_e → σ_μ jump of 3.5) is no longer a numerical coincidence: it is the direct quantitative signature of σ_record and σ_refine saturation. The muon transition is the "infrastructure-establishment" step; the tau transition is the "cycle-addition" step over already-existing infrastructure.

This decomposition narrows the falsifiability target substantially. Instead of asking only "does the enumeration return σ_μ ≈ 3.5 and σ_τ ≈ 4.5?", the framework now demands:

1. that the substrate counts $\Delta\beta_1$, N_persistent, and R_D admit a substrate-level definition;
2. that the muon-to-tau saturation pattern (N_persistent and R_D roughly constant between μ and τ) emerges from refinement-persistence;
3. that the weights a, b, c are derivable from substrate thermodynamics and land in the O(0.5–1) band.

If a first-principles enumeration reproduces σ_μ ≈ 3.5 and σ_τ ≈ 4.5 by a different decomposition (e.g. with σ_record growing linearly and σ_cycle small), the totals would be the same but the structural story would be different, and §6.4 would need revision even though the §10 dominance result might survive intact. This is exactly the level of decomposable structure the toy framework needs to expose in order to be testable in pieces, rather than as a single opaque parameter.

6.5 σ_D is the framework's principal vulnerability

Because σ_D supplies 65.8 % of the muon log-ratio and 55.2 % of the tau log-ratio (§10), the framework's quantitative reconstruction is dominated by this single input. The fermion mass hierarchy is, on the toy inputs adopted, primarily an **informational cost** — a result whose implications are substantial and whose contingency on σ_D (and on the §6.4 decomposition) is correspondingly substantial. If a first-principles enumeration produces σ_μ and σ_τ within roughly a factor of two of the toy values, and admits a sub-linear-growth decomposition along the lines of §6.4, the structural reading of the hierarchy as informationally-driven is corroborated. If the enumeration produces values markedly outside that range, or grows uniformly rather than saturating, the framework as currently structured is wrong — and the failure mode is now sharply identifiable rather than buried under a degeneracy of parameter choices.

This is, in our view, the single most consequential framework-level claim the paper makes, and §12 elevates the σ_D enumeration to the lead forward-programme task.

---

7. Transport-Graph Construction

We now make the transport-complexity input concrete by constructing toy admissibility-preserving transport graphs for each charged-lepton PFD class.

7.1 Transport functional and calibration

The general transport functional aggregates four graph invariants:

$T(D) \approx \exp[\ \alpha_1 \beta_1 + \alpha_2 \lambda_2^{-1} + \alpha_3 \chi + \alpha_4 \langle \deg \rangle\ ]$

with:

- β_1 — first Betti number (number of independent admissibility cycles)
- λ_2^{-1} — inverse spectral gap (slowest refinement mode)
- χ — branching factor (refinement-stable bifurcation density)
- $\langle \deg \rangle$ — average degree (local connectivity)

The toy coefficients $\alpha_1 \ldots \alpha_4$ are constrained by a single calibration condition: that the muon graph, relative to the electron reference, yield $\ln T(\mu) \approx 0.5$. This condition is **under-determining**: with four coefficients and one condition, infinitely many $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ tuples are admissible.

We adopt the simplest calibration that respects the structural intuition that the first Betti number is the dominant graph invariant for refinement-persistent transport:

$\alpha_1 = 0.5,\ \alpha_2 = \alpha_3 = \alpha_4 = 0$ (leading order)

Under this calibration the transport functional reduces, to leading order in the graph invariants, to

$\ln T(D) \approx 0.5 \cdot \Delta\beta_1$

where $\Delta\beta_1$ is the difference in first Betti number from the electron reference. The other invariants are present in the full functional but their coefficients are sub-leading at this toy level and their derivation from substrate refinement geometry is left to the forward programme.

7.2 Electron graph

The minimal refinement-persistent graph $G\_e = (V\_e, E\_e)$, with one independent admissibility cycle:

Invariant	Value
β_1	1
λ_2^{-1}	0.2

Invariant	Value
χ	0.1
$\langle deg \rangle$	0.1

Under the β_1-dominant calibration, $T(e) = 1$ by reference normalisation.

7.3 Muon graph

Two coupled loops with refinement-stable branching:

Invariant	Value
β_1	2
λ_2^{-1}	0.6
χ	0.4
$\langle deg \rangle$	0.3

With the calibrated coefficients, $\ln T(\mu) = 0.5 \cdot (2 - 1) = 0.5$, hence $T(\mu) \approx 1.65$.

7.4 Tau graph

Three loops with a richer refinement structure:

Invariant	Value
β_1	3
λ_2^{-1}	1.2
χ	0.9
$\langle deg \rangle$	0.9

With the same coefficients, $\ln T(\tau) = 0.5 \cdot (3 - 1) = 1.0$, hence $T(\tau) \approx 2.72$.

The relation $\ln T(\tau) = 2 \cdot \ln T(\mu)$ is, under the β_1-dominant calibration, a direct linear consequence of $\beta_1(\tau) - \beta_1(e)$ being twice $\beta_1(\mu) - \beta_1(e)$. It is not an independent dynamical prediction; it is a structural consequence of the calibration choice combined with the integer cycle counts. As a corollary, the tau-scale residual in the assembled reconstruction (§8) is *not* independent evidence that the framework reaches the second generation correctly — see §8.6 for an explicit statement.

The graph β_1 values themselves are assignments, not derivations: that the muon graph has $\beta_1 = 2$ and the tau graph has $\beta_1 = 3$ is the natural minimal extension consistent with generation depth, but a substrate-level argument for these integer counts (and against, say, $\beta_1(\mu) = 3$ or $\beta_1(\tau) = 4$) is owed to the forward programme.

The phrase "natural minimal extension" carries a structural prior worth naming: that the generation ladder corresponds to adding *exactly one* independent admissibility cycle per generation step. This is not a derivation. It is the content of Claim 4 in the falsifiability statement (§12.2), and it is one of the framework's sharpest claims awaiting microscopic validation. A refinement-persistence calculation that produced, say, $\beta_1(\mu) = 2$ but $\beta_1(\tau) = 4$ (skipping $\beta_1 = 3$) would falsify this prior and require a substantial revision to §7 and §8.

Transport complexity therefore acts as the **non-uniformity driver** across generations: it supplies the inter-generational gradient — driven by integer cycle counting — that distinguishes μ from τ in the correct proportion, conditional on the calibration and on the β_1 assignments.

---

8. Full Hierarchy Reconstruction

We now assemble the four substrate contributions.

8.1 Reference: the electron

$\ln m_e = 0$ (normalization)

8.2 Muon

$\ln (m_\mu / m_e) = \ln 2 + \ln S_L(\mu) + \sigma_\mu + \ln T(\mu) + ½ \ln \lambda_{max}(H_\mu)$

```
=  0.693  +  0.380  +  3.500  +  0.500  +  0.250
=  5.323
```

so that:

$m_\mu / m_e = e^{5.323} \approx 205$

Observed: 206.8.

The muon reconstruction absorbs the single muon-scale calibration of the transport functional and is therefore not even a within-toy prediction.

8.3 Tau

$\ln (m_\tau / m_e) = \ln 4 + \ln S_L(\tau) + \sigma_\tau + \ln T(\tau) + ½ \ln \lambda_{max}(H_\tau)$

```
=  1.386  +  0.760  +  4.500  +  1.000  +  0.500
=  8.146
```

so that:

$m_\tau / m_e = e^{8.146} \approx 3\,450$

Observed: 3 477.

8.4 Reconstruction summary

Ratio	Reconstructed	Observed	Residual	Status
m_μ / m_e	205	206.8	−0.9 %	single muon-scale calibration absorbed
m_τ / m_e	3 450	3 477	−0.8 %	structurally constrained, but σ_τ is a free input
m_τ / m_μ	16.83	16.81	+0.1 %	follows from the two ratios above

8.5 Reduced-parameter reconstruction

The toy reconstruction uses a large number of visible quantities — closure curvatures, four graph invariants per generation, an exponent p, a transport functional coefficient α_1, a linkage relation, distinguishability loads. It is fair to ask how many of those quantities are genuinely independent inputs, as opposed to derived consequences of a small number of structural assumptions.

The working content of the reconstruction reduces to **three structural assumptions**:

1. **Non-uniform σ_D growth.** $\sigma_e = 0$, $\sigma_\mu = 3.5$, $\sigma_\tau = 4.5$. The non-uniformity (sub-linear growth between μ and τ) is the key structural feature, decomposed in §6.4 into a saturating record/refinement sector plus a linear cycle sector.
2. **Linear cycle-rank transport.** $\ln T(D) = 0.5 \cdot (\beta_1(D) - 1)$. Calibrated by a single muon-scale condition, with the β_1-dominant form following from the structural intuition that refinement-persistent cycles are the leading transport invariant.
3. **Linked Hessian stiffness.** $\lambda_{max}(H_D) \approx T(D)$. Taken as exact equality at the toy level (§5.1), with a 20-%-band tolerance (§5.2).

Once these three assumptions are fixed, the tau-scale ratio is **not independently tuned**. Given the integer Betti counts ($\beta_1(\mu) = 2$, $\beta_1(\tau) = 3$), the localization scaling of §4, and the generation-depth operator of §3, the tau-scale log-ratio is the algebraic consequence:

$$\ln (m_\tau / m_e) = \ln 4 + 0.380 \cdot 2 + \sigma_\tau + 0.5 \cdot 2 + 0.5 \cdot 0.5 \cdot 2$$

There is no further free parameter at the tau scale beyond the three structural assumptions above. The §6.4 decomposition then narrows assumption 1 further by replacing the two toy σ values with a three-parameter substrate-counting model (a, b, c) plus three substrate counts ($\Delta\beta_1$, $N_{persistent}$, R_D); this widens the apparent count of toy inputs but narrows the falsifiability target. An erroneous σ-structure is now detectable component-by-component, not only by failure to reach the totals.

This is the appropriate frame in which to read the sub-percent residuals of §8.4: not as a triple-coincidence agreement, but as a propagation of the σ_D choice — decomposed in §6.4 into a saturation structure plus three weight numbers — through three structural assumptions.

8.6 What the toy reconstruction can and cannot falsify

The lepton mass ratios, at this stage of the programme, **cannot falsify the framework**. The reasons are direct:

1. The two σ_D values (σ_μ, σ_τ) are freely chosen toy inputs. Between them they fit roughly 60 % of the inter-generational log-ratio at both scales. Any "agreement with observation" at the σ_D-dominated scales is therefore an agreement between the data and a deliberate input fit.
2. The localization exponent p is a toy back-fit (§4), supplying a further 7–9 % of the log-ratio at each scale.
3. The transport coefficient α_1 is calibrated against the muon scale (§7.1). It supplies the muon transport contribution by construction, and the tau transport contribution by integer-Betti propagation — which, under the β_1-dominant calibration, is automatic.
4. The closure-Hessian/transport linkage is taken as exact equality (§5.1), removing the Hessian sector as an independent input.

What this means concretely: the structural feature "tau-scale residual no larger than muon-scale residual" is **not independent evidence** that the framework reaches the second generation correctly. Given the integer Betti counts, the β_1-dominant calibration, and a freely chosen σ_τ, that property is structurally guaranteed at the toy level.

The toy reconstruction therefore demonstrates **consistency**, not **prediction**. The machinery is internally well-defined, the inputs propagate without producing pathologies, and the resulting numbers sit in a recognisable neighbourhood of the observed ratios. Genuine prediction enters only when σ_D (with its §6.4 decomposition), the β_1 assignments, the localization exponent p, and the closure-Hessian/transport linkage are derived **independently of the mass spectrum**. The forward programme of §12, and especially the failure-mode table of §12.3, sets up exactly that derivation.

---

9. Sensitivity Analysis

A central question is whether the hierarchy is fragile. We perturb the four primary inputs by ±10 % and examine the propagation to mass ratios. Asterisks (*) mark rows whose perturbation is not independent of another row under the working linkage of §5.

9.1 Single-input perturbations

Perturbation (±10 %)	$\Delta(m_\mu / m_e)$	$\Delta(m_\tau / m_e)$	$\Delta(m_\tau / m_\mu)$
σ_D (distinguishability)	± 23 %	± 24 %	± 15 %
$T(D)$ (transport)	± 5 %	± 10 %	± 11 %
p (localization)	± 4 %	± 8 %	± 8 %
k_D (Hessian)*	± 3 %	± 5 %	± 5 %

* Under the linkage k_D = $T(D)$ (§5), the Hessian row is not independent of the transport row. Reported separately for diagnostic purposes only.

The σ_D contribution is by far the most sensitive single input, consistent with its identification as the dominant hierarchy driver in §10.

9.2 Independent random perturbations (quadrature)

Treating the σ_D, transport, and localization perturbations as statistically independent and combining in quadrature on the log-ratio:

Ratio	Δ ln(ratio)	Δ(ratio)
m_μ / m_e	± 0.36	± 36 %
m_τ / m_e	± 0.47	± 47 %

9.3 Correlated worst-case perturbations

If σ_D, transport, and localization perturb simultaneously by the same sign:

Ratio	Δ ln(ratio), all +10 %	Δ(ratio)
m_μ / m_e	+ 0.46	+ 60 %
m_τ / m_e	+ 0.68	+ 97 %

9.4 σ_D-only large-perturbation test

We also test how much σ_D can be wrong while leaving the order-of-magnitude reconstruction intact. We perturb σ_μ and σ_τ independently by ±50 % and ×2:

Perturbation	New m_μ / m_e	New m_τ / m_e	Comment
σ_D × 0.5	35	365	hierarchy too compressed; e < μ < τ ordering preserved
σ_D × 1.5	1 192	31 175	hierarchy too stretched; ordering preserved
σ_D × 2.0	6 904	313 500	ordering preserved, scales misplaced

The framework constrains the *shape* of the hierarchy (robust under input perturbation) more strongly than its *exact magnitudes* (contingent on σ_D landing within roughly a factor of two of the toy values).

9.5 Interpretation

Four observations follow:

1. The σ_D contribution is the most sensitive single input, consistent with its dominant log-share role.
2. Small (±10 %) perturbations propagate to 30–50 % mass-ratio variation in quadrature, and 60–100 % under correlated worst-case. Both are well within an order of magnitude.
3. Large (±50 % – ×2) σ_D perturbations break the quantitative fit but preserve the qualitative hierarchy ordering. The framework's *shape* claim is robust; its *magnitude* claim is contingent on σ_D landing within roughly a factor of two of the toy values.
4. The hierarchy is structurally stable to small input variation. No single input must be tuned to a fraction of a percent.

We deliberately avoid the stronger phrasing "explained rather than engineered" used in some drafts: the σ_D values were chosen with the target in view, and the toy reconstruction is not an explanation in the first-principles sense. What the sensitivity analysis establishes is that the framework's *shape* claim does not require fine-tuning, and that the *magnitude* claim is sensitive to σ_D within a wide but bounded band.

---

10. Dominant Mechanisms: Log-Share Decomposition

The qualitative role of each contribution can now be made quantitative. We decompose the log-ratios into their additive components and report each component's share of the total. Under the toy linkage $k_D = T(D)$ (§5), the Hessian contribution is not an independent input; we therefore lead with the **linkage-aware four-factor decomposition** as the structurally honest accounting, and provide the naive five-factor decomposition for comparison.

10.1 Linkage-aware four-factor decomposition (primary)

Muon (total log-ratio 5.323)

Contribution	ln-share	% of total
Generation depth (ln 2)	0.693	13.0 %
Localization (S_L)	0.380	7.1 %
Distinguishability (S_P)	3.500	65.8 %
Transport + Hessian (linked)	0.750	14.1 %

Tau (total log-ratio 8.146)

Contribution	ln-share	% of total
Generation depth (ln 4)	1.386	17.0 %
Localization (S_L)	0.760	9.3 %
Distinguishability (S_P)	4.500	55.2 %
Transport + Hessian (linked)	1.500	18.4 %

10.2 Structural reading

Contribution	Role
D^(γ_D)	generation ordering only
S_L (localization)	exponential backbone, sub-dominant
S_P (distinguishability)	**dominant hierarchy driver (contingent on σ_D toy values and §6.4 structure)**
S_I × S_H (transport+Hessian, linked)	non-uniformity driver across generations

Three structural observations:

1. **Distinguishability is the dominant single contributor at both generation scales.** σ_D accounts for 65.8 % of the muon log-ratio and 55.2 % of the tau log-ratio. The "informational" reading of the hierarchy is the natural reading of the toy reconstruction.
2. **The dominance claim is contingent on σ_D and on the §6.4 substrate-counting decomposition.** The σ_D values are toy inputs, motivated by the heuristic of §6.2 and the structural model of §6.4 but not derived. If a first-principles enumeration produces σ values markedly different from 3.5 and 4.5, or grows uniformly with generation rather than saturating, the dominance structure changes.
3. **Localization is sub-dominant at 7–9 % across both scales — but this is contingent on the back-fit value of p.** A substantially different Role-4 derivation could move localization into the dominant band.

The framework's toy reconstruction points squarely at the *informational* reading — closure bookkeeping plus transport structure — as the dominant mechanism, with localization and dynamical stiffness playing supporting roles. If this reading is corroborated by the microscopic σ_D enumeration (Claim 1 of §12.2), the fermion mass hierarchy is, structurally, a hierarchy of *informational cost*. That reframing is the most consequential outcome of the present paper, well beyond the numerical fit, but it is honestly a hypothesis being staked rather than a conclusion being drawn.

10.3 Naive five-factor decomposition (for comparison)

Without applying the linkage:

Muon

Contribution	ln-share	% of total
Generation depth (ln 2)	0.693	13.0 %
Localization (S_L)	0.380	7.1 %
Distinguishability (S_P)	3.500	65.8 %
Transport (S_I)	0.500	9.4 %
Hessian (S_H)	0.250	4.7 %

Tau

Contribution	ln-share	% of total
Generation depth (ln 4)	1.386	17.0 %
Localization (S_L)	0.760	9.3 %
Distinguishability (S_P)	4.500	55.2 %
Transport (S_I)	1.000	12.3 %
Hessian (S_H)	0.500	6.1 %

This naive bookkeeping is what a reader would assemble from the abstract identity m_D = D^(γ_D) · S_H · S_L · S_P · S_I · v without applying any linkage between the factors. It is included for completeness, but it overstates the number of independent contributions by treating S_H and S_I as separate inputs when, at the toy level, they are not.

---

11. Scope and Limitations

This paper demonstrates that:

- the hierarchy framework is computationally executable end to end,
- its observables are well-defined and have explicit numerical values,
- a recognisable hierarchy emerges from modest substrate-complexity increments, and
- the resulting hierarchy is perturbatively stable.

It does **not** claim:

- a final first-principles derivation,
- exact substrate graphs for the charged leptons,
- exact distinguishability counts,
- exact closure Hessian operators,
- a complete Standard Model reconstruction, or
- a derivation of the absolute charged-lepton mass scale. Only ratios are computed here; the absolute scale requires v (derived in the closure-condensate sector of the programme; see

the master action and matter coupling papers) and the absolute prefactor in the stiffness identity, both of which are deferred to a separate computation.

The PFD graphs, Hessian curvatures, distinguishability loads, and graph invariants used here are illustrative toy inputs, chosen consistent with the qualitative substrate picture and with the target hierarchy. Their numerical values are not derived in this paper.

Counting the toy freedoms honestly:

- **σ_μ and σ_τ** are toy inputs, but decompose under §6.4 into substrate counts ($\Delta\beta_1$, N_persistent, R_D) and weights (a, b, c). The decomposition does not derive the σ values but reduces the falsifiability target to a structured set of sub-claims.
- **The localization exponent p** is a back-fit (§4).
- **The transport coefficient α_1** is calibrated to the muon scale (§7.1).
- **The linkage $k\_D = T(D)$** is taken as exact equality, removing the Hessian sector as an independent input (§5).
- **The integer β_1 assignments** are topological but still assigned by hand (§7.4).

§8.5 shows that the working content of the reconstruction reduces to three structural assumptions plus the σ values; the tau ratio is not independently tuned given those assumptions. §8.6 then states explicitly that the lepton mass ratios at this stage **cannot falsify the framework** — they can only fail to be reachable, and they are reachable. Prediction enters only once σ_D (and its §6.4 components), β_1, and the linkage are derived independently of the mass spectrum.

The contribution of this paper is the **demonstration that the framework operates correctly when realistic-shape inputs are supplied**, the identification of which substrate observables would dominate the hierarchy conditional on those inputs, the structural decomposition of those inputs into testable sub-components, and the precise specification of which forward-derivation outcomes would corroborate or falsify the framework's quantitative reach.

---

12. Forward Programme and Falsifiability

The next stage of the programme follows directly from the structure exposed here. We list it in priority order, with the σ_D enumeration first because it carries the largest contingent weight in the present reconstruction.

12.1 Forward computations

1. **Distinguishability enumeration (highest priority).** Derive σ_D from a direct count of persistent record-current commitments in each PFD class, replacing the toy σ_D values. The toy reconstruction places $\sigma\_\mu \approx 3.5$ and $\sigma\_\tau \approx 4.5$, and the §6.4 decomposition further requires the saturation pattern (N_persistent and R_D roughly constant between μ and τ) to emerge from refinement-persistence. The enumeration must validate both the totals and the saturation structure.

2. **Transport invariant computation.** Construct the full refinement-persistent transport graphs for each PFD class and compute their invariants β_1, λ_2^{-1}, χ, $\langle deg \rangle$ from substrate geometry. Test whether β_1 is the dominant invariant, and whether the integer assignments $\beta_1(e) = 1$, $\beta_1(\mu) = 2$, $\beta_1(\tau) = 3$ are forced by refinement-persistence.
3. **Closure-Hessian derivation.** Compute $\lambda\_max(H\_D)$ for the charged-lepton PFD classes from the substrate closure functional. Test whether the linkage $k\_D \approx T(D)$ emerges, or whether the Hessian and transport sectors are genuinely independent.
4. **Localization derivation.** Derive the exponent p from Role-4 substrate refinement geometry, and test whether the derived value sits near the back-fit $p \approx 0.233$.
5. **σ-decomposition weights.** Derive the toy weights a, b, c of §6.4 from substrate thermodynamics. Test whether they land in the $\mathcal{O}(0.5–1)$ band.
6. **Sector extension.** Repeat the programme for the quark sector (with the additional confinement contribution) and the neutrino sector (with the additional admissibility-mode contribution).

12.2 Falsifiability statement

The framework as instantiated in this paper makes five sharp empirically-testable structural claims:

Claim 1 (σ_D landing zone). A first-principles enumeration of persistent record-current commitments yields $\sigma\_\mu \in [2.0, 6.0]$ and $\sigma\_\tau \in [3.0, 7.0]$. (Heuristic justification: §6.2; structural model: §6.4.)

Claim 2 (Betti dominance). Among the four graph invariants, the first Betti number β_1 is the dominant contributor to $\ln T(D)$, with the other invariants contributing at most $\lesssim 30$ % of the inter-generational gradient.

Claim 3 (Hessian/transport linkage). A first-principles computation of $\lambda\_max(H\_D)$ yields values within ≈ 20 % of $T(D)$ for each charged-lepton PFD class.

Claim 4 (Integer β₁ assignments). The refinement-persistence requirement forces $\beta_1(e) = 1$, $\beta_1(\mu) = 2$, $\beta_1(\tau) = 3$.

Claim 5 (σ-decomposition saturation). A first-principles enumeration of $N\_persistent(D)$ and R_D produces approximate saturation between the muon and tau scales (i.e., $N\_persistent(\tau) \approx N\_persistent(\mu)$ and $R\_\tau \approx R\_\mu$ to within ~30 %), with the cycle-rank component σ_cycle alone supplying the inter-generational increment.

If all five claims survive their respective microscopic derivations, the toy reconstruction is corroborated and the framework gains genuine predictive content. If any single claim fails, the framework as currently structured is wrong on that dimension; the specific failure mode then identifies which sector of the framework needs revision.

12.3 Failure-mode table

The framework's failure surface can now be stated explicitly. The table below maps possible outcomes of the forward computations to their consequences for the framework as currently structured.

If the next computation finds…	Then…
σ_μ, σ_τ within ≈ ±50 % of 3.5, 4.5	hierarchy framework strongly supported at the lepton scale
σ grows uniformly with generation (no μ↔τ saturation)	current σ-decomposition (§6.4) fails; sub-linear growth was a key structural prediction
β_1 is not dominant in ln $T(D)$	transport sector must be revised; β_1-dominant calibration is wrong
β_1 is leading but other invariants contribute > 30 % (Claim 2 soft fail)	framework survives; transport sector partially revised; β_1-only calibration replaced by a multi-invariant form
$\lambda_{max}(H_D)$ does not track $T(D)$ within 20 %	linkage removed; framework operates with four independent factors and §10 dominance ordering needs revisiting
p from Role-4 geometry is markedly different from 0.233	localization may move into dominant band; §10 reading revised
§6.4 weights (a, b, c) cannot be assigned consistent $\mathcal{O}(0.5–1)$ values	structural σ-decomposition is wrong; σ_D may not have a substrate-counting reading
quark σ and T are not substantially larger than lepton values	quark hierarchy explanation fails by extension
quark σ is substantially larger than lepton σ but with a different internal decomposition than §6.4	framework survives at the quark scale; §6.4 structural model needs revision; the saturation hypothesis may be sector-specific to the leptons
neutrino σ is not suppressed below lepton σ	neutrino mass smallness explanation fails

This table is the framework's *skin in the game*. It enumerates the specific, identifiable outcomes under which the framework would be shown wrong — not in vague terms, but at the level of which substrate quantity has misbehaved and what part of the toy reconstruction needs revision. The toy reconstruction's principal scientific contribution is precisely that this table can now be written.

12.4 Subsidiary objective

A subsidiary computational objective — the explicit derivation of the Hessian/transport linkage $k_D \approx T(D)$ — would simultaneously discharge the most significant toy-level assumption of the present paper and reduce the four-factor framework to a tighter three-factor one. This is captured by Claim 3 above.

---

13. Conclusion

This paper performed the first explicit toy reconstruction of the VERSF substrate stiffness hierarchy framework. Using toy closure graphs, toy Hessian spectra, localization scaling, distinguishability counts, and graph transport complexity, the framework reproduced the observed charged-lepton hierarchy to leading order, with sub-percent residuals at all three mass ratios.

The central result is not the numerical fit, which depends on toy inputs chosen with the target in view. It is this:

the hierarchy machinery can actually run, its dominant inputs are now identifiable, quantifiable, and falsifiable, and the σ_D sector — the most consequential single input — has been given a structurally interpretable internal decomposition that further narrows the falsifiability target.

The fermion mass problem is no longer:

"insert arbitrary Yukawa numbers."

It becomes:

"compute four substrate observables — of which one, the persistent distinguishability load σ_D, decomposes into a saturating record/refinement sector plus a linear cycle sector, and does most of the work."

The log-share decomposition (§10) sharpens this: among the four observables, the informational contribution — σ_D — carries roughly 60 % of the inter-generational ratio in the toy reconstruction, with transport and localization providing the supporting structure. If this pattern survives a microscopic derivation of σ_D (Claim 1) with the saturation structure required by §6.4 (Claim 5), fermion masses are, structurally, **measurable expressions of substrate informational cost**, not arbitrary constants inserted into nature. If it does not survive, the failure-mode table of §12.3 localises the failure to a specific sector of the framework rather than diffusing it across twelve free parameters.

The toy reconstruction should therefore not be read as a claim that the lepton masses have already been derived. Its value is sharper: it makes the framework's failure surface concrete and identifiable. Where the previous structural paper could be ambiguous about which substrate observables matter and by how much, the present reconstruction states the answer in numbers and ties each number to a forward derivation that will either land it or refute it. That is the form in which a substrate-level theory of fermion mass can be done — not as a fitting exercise, but as a sequence of structurally exposed, locally testable claims.

Either outcome — corroboration or failure — is scientifically productive. The next step is well-defined: perform the σ_D enumeration (with the §6.4 decomposition validation), the β_1 computation, and the closure-Hessian derivation for the charged-lepton PFD classes — and then, if those land, extend to quarks, neutrinos, and the flavour-mixing sector.