First-Principles Enumeration of Persistent Distinguishability in Charged-Lepton PFDs

With a Substrate-Counting Derivation of the Three-Generation Closure of the Charged-Lepton Sector — Record-Current Commitments, Refinement Persistence, Infrastructure Sharing, Primitive-Function Exhaustion, and the Substrate Origin of σ_e, σ_μ, σ_τ in VERSF

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The previous VERSF hierarchy paper proposed that particle masses come largely from what it called the *persistent distinguishability load* — the running cost the substrate pays to keep one particle persistently distinct from another. That paper showed the framework could reproduce the electron, muon, and tau masses if the distinguishability loads took particular values ($\sigma_\mu \approx 3.5$, $\sigma_\tau \approx 4.5$). But it could not derive those values. It chose them. It also took the three-generation structure of the charged-lepton sector as given.

This paper does two things. The headline result, which may be the deeper of the two, is a derivation of *why there are three charged-lepton generations and not four*. The argument is that refinement-persistent charged-lepton structure admits only three distinct loop classes before further loops either gauge-reduce, dissolve, or change sector. The charged-lepton sector is therefore a *finite* topological sequence — three generations, no more — rather than an open ladder. This is the framework's first directly experimentally-testable prediction: if anyone discovers a genuine fourth charged lepton, the closure step of the derivation is wrong and the prediction fails.

The second result is the running-cost calculation itself. With the three-generation structure derived, the paper asks what the substrate is actually counting when it tracks distinguishability, and tries to count it from first principles. The proposal is that the substrate tracks *independent* commitments — features that the substrate must maintain to keep a particle's identity stable, but only counted once each. If two features rely on the same underlying machinery, they share infrastructure and the substrate pays for that machinery once. If a feature is redundant (forced by another commitment, or removable by a gauge choice), it doesn't count at all.

Three classes of commitment emerge for the charged leptons:

1. **Cycle persistence** — one commitment per refinement-persistent loop. Linear in the number of loops.
2. **Phase coherence** — the substrate's ability to maintain phase relations between loops. Established once when the second loop appears; shared from then on.
3. **Refinement infrastructure** — the underlying machinery that makes multi-loop persistent structure possible at all. Established once at the muon; reused by the tau.

The electron is the one-loop baseline. The muon is the *infrastructure-forming* transition: it pays for cycle persistence, phase coherence, and refinement infrastructure all at once. The tau is the *infrastructure-extending* transition: it adds one cycle's worth of persistence, but reuses the phase and refinement infrastructure the muon already paid for.

This naturally produces the non-uniform pattern the previous paper required:

the jump from electron to muon is large; the jump from muon to tau is small.

The derived values land at $\sigma_\mu \in [2.5, 4.5]$ and $\sigma_\tau \in [3.5, 5.5]$, with the toy values from the previous paper sitting at the centre of both bands. The σ values depend on the two infrastructure-sharing parameters only through one combination, so the framework's quantitative reach is really one constraint rather than two. The natural midpoint of that combination reproduces the toy values exactly — though whether the substrate actually realises that midpoint awaits a separate computation.

The paper does *not* derive the lepton masses. What it does is convert the dominant input of the previous hierarchy reconstruction — σ_D, a freely-chosen toy value — into a substrate-counting problem on a sector with a *derived three-generation structure*, governed by three structural rules, two infrastructure-sharing parameters, and explicit falsification conditions. The hierarchy stops being "insert two numbers chosen with the answer in view, applied to a sector whose generation count is taken on faith" and becomes "enumerate three classes of independent commitment on a sector that closes at three generations by construction."

Table of Contents

Abstract

The previous paper, *First Toy Reconstruction of the Charged-Lepton Hierarchy in VERSF*, identified the persistent distinguishability load σ_D as the dominant single contribution to the inter-generational charged-lepton hierarchy within the substrate stiffness framework. The toy reconstruction reproduced the observed lepton ratios using:

$\sigma_e = 0$, $\sigma_\mu \approx 3.5$, $\sigma_\tau \approx 4.5$

with the two non-zero values chosen consistent with the target hierarchy rather than derived. The toy paper additionally carried the n-loop assignments ($n_e = 1$, $n_\mu = 2$, $n_\tau = 3$) as a separate input (Claim 4 of its falsifiability statement), with the three-generation closure of the charged-lepton sector taken as given rather than derived.

The present paper performs two structurally distinct derivations.

The **foundational result** (§4.1) is a derivation of the charged-lepton n-band itself: refinement-persistent charged-lepton PFDs admit exactly the loop counts $n \in \{1, 2, 3\}$, with the lower bound forced by minimal persistence, the generation sequence forced by topological distinction (loop count being the only admissibility-topological invariant available within the charged-lepton sector), and the upper bound forced by closure exhaustion. The hierarchy reconstruction depends on the σ-counting law below, but the more foundational result of the present paper is that the charged-lepton sector closes at three generations as a structural consequence of refinement-

persistent admissibility geometry. Once the admissible refinement-persistent loop classes are exhausted at n = 3, the charged-lepton sector becomes a *finite topological sequence* rather than an open parameter tower. This is the framework's first directly experimentally-testable prediction: discovery of a genuine fourth charged-lepton generation would falsify the closure-exhaustion step.

The **applied result** (§§5–9) is the σ-counting law itself. The central proposal is that σ_D measures the **minimal independent admissibility-fixed informational commitments** required to preserve PFD identity under refinement evolution. The counting is governed by three structural rules:

1. **Refinement-persistence rule.** Only commitments that persist under substrate refinement are counted.
2. **Independence rule.** Only independent commitments are counted; redundant phase relations, gauge-equivalent records, and closure-derivable constraints are quotient-identified.
3. **Infrastructure-sharing rule.** Substrate machinery that supports multiple commitments is paid for once and shared.

A fourth principle — *primitive-function exhaustion* (§3.1) — terminates the counting. A commitment class is primitive-function exhausted at primitive count k when the substrate-physical functions admissible in support of that class at leading order have been enumerated and shown finite. The principle is structurally protective rather than merely numerical: exceeding the exhaustion bound is forbidden not just at the parameter level but at the framework level, because any candidate function beyond the enumerated count either reduces to sub-leading correction, quotient-equivalent structure, or a new commitment class entirely. This principle is the load-bearing tool for both of the present paper's structural ceilings ($\xi_{max} = 2$ and $n_{max} = 3$) and is structurally available for the quark- and neutrino-sector extensions.

Applied to refinement-persistent charged-lepton PFDs, the rules and the closure principle produce three commitment classes:

$\sigma_{cycle}(n) = n - 1$ (linear; one commitment per refinement-persistent loop)

$\sigma_{phase}(n) = \xi \cdot \mathbb{1}_{\{n\geq 2\}}$ (saturating; phase-coherence infrastructure shared)

$\sigma_{refine}(n) = \eta \cdot R \cdot \mathbb{1}_{\{n\geq 2\}}$ (saturating; refinement infrastructure shared)

where R = 2 counts the two substrate-global refinement commitments (refinement identity preservation and admissibility-compatible transport continuity), $\mathbb{1}_{\{n\geq 2\}}$ activates the infrastructure cost only for multi-loop PFDs, and the parameter bands are structurally motivated rather than asserted, with asymmetric epistemic status:

- $\xi \in [1, 2]$ is the phase-coherence infrastructure cost. The upper bound is *structurally pinned* by primitive-function exhaustion (§3.1) applied to the count of primitive phase-coherence functions at leading order (relative phase registration and phase-transport

stability under refinement); a third primitive function would constitute a new commitment class beyond the three identified here, not a higher ξ value. The lower bound is *structurally floored* by the substrate's discrete commitment-counting unit. The ξ band is therefore structurally pinned at both endpoints.

- **η ∈ [¼, ¾]** is the refinement-infrastructure sharing efficiency. The endpoints are *structurally motivated admissible bounds* — maximum-sharing across four effective support channels gives η_min = ¼; minimum-sharing under unavoidable overlap-reduction gives η_max = ¾ — but neither is forced in the way that ξ_max = 2 and ξ_min = 1 are forced. The η band is the weaker of the two band derivations and is honestly flagged as such in §7.2; the working hypothesis "four effective support channels" at η_min and the corresponding ¾ at η_max are not derived from substrate dynamics, only structurally motivated.

Summing yields the minimal quotient-counting law:

$\sigma(n) = (n - 1) + \xi \cdot \mathbb{1}_{\{n \geq 2\}} + \eta \cdot R \cdot \mathbb{1}_{\{n \geq 2\}}$

With the derived n-band $n \in \{1, 2, 3\}$ and the parameter bands above:

$\sigma_e = 0$

$\sigma_\mu = 1 + \xi + 2\eta \in [2.5, 4.5]$ (natural midpoint = 3.5)

$\sigma_\tau = 2 + \xi + 2\eta \in [3.5, 5.5]$ (natural midpoint = 4.5)

The σ values depend on ξ and η only through the combination $\xi + 2\eta \in [1.5, 3.5]$. The framework's quantitative reach is therefore a one-parameter prediction in that combination, not a two-parameter prediction in (ξ, η) separately; the parameter space is genuinely lower-dimensional than the parameter count suggests. The toy values $\sigma_\mu = 3.5$, $\sigma_\tau = 4.5$ sit at the centre of both bands and correspond exactly to the midpoint $\xi + 2\eta = 2.5$. Crucially, the **structural shape** of the hierarchy — $\sigma_\tau - \sigma_\mu < \sigma_\mu - \sigma_e$ — is forced by the saturation of both phase and refinement infrastructure, independent of the specific values of ξ and η within their bands.

The inter-generational gap $\sigma_\tau - \sigma_\mu = 1$ is fixed by the n-loop counting alone and is independent of (ξ, η) entirely. The intra-generational baseline σ_μ depends on $\xi + 2\eta$ — a single combination — which is bounded but not pinned.

The paper does not derive the charged-lepton masses. What it does is convert the dominant input of the previous reconstruction from an opaque toy parameter into a substrate-counting problem governed by three structural rules and two infrastructure-sharing parameters, with the structural shape — but not the precise magnitudes — forced by the counting. The hierarchy problem is thereby reformulated as the enumeration of persistent informational commitments required to preserve refinement-persistent particle identity.

A complementary reading is also developed in §11. The K = 7 closure manifold is strongly non-orthogonal — adjacent Gram-matrix entries $M_{i,i+1} \sim 0.7–0.9$ across the admissible mode-width range, as established in the projected-operator analysis of the closure manifold — and the counting law is therefore most naturally interpreted as counting *effective independent* commitments after overlap reduction on the projected closure basis, rather than as raw combinatorial counting. Whether the counting law is in fact the leading-order approximation to a schematic operator structure $\sigma_{eff} \sim \mathrm{Tr}(M^{-1}L)$, with M the closure-mode Gram matrix and L a refinement-persistence operator, is now a concrete computational question rather than a programmatic one.

---

1. Introduction

The observed charged-lepton hierarchy

$m_e : m_\mu : m_\tau \approx 1 : 206.8 : 3477$

remains one of the deepest unexplained quantitative structures in particle physics. The Standard Model accommodates it through three Yukawa couplings, but offers no internal account of why those couplings span more than three orders of magnitude or why their spacing follows the pattern it does.

The VERSF substrate stiffness programme has approached this problem by reading fermion masses as substrate stabilization costs of Persistent Fold Defects (PFDs), with the hierarchy encoded in four substrate observables: closure-Hessian stiffness, localization compression, persistent distinguishability load, and transport-network complexity. The previous paper, *First Toy Reconstruction of the Charged-Lepton Hierarchy in VERSF*, executed the first explicit reconstruction of this framework on the charged-lepton sector. It identified the persistent distinguishability load σ_D as the dominant single contributor to the inter-generational ratio, supplying roughly 60 % of the log-ratio at both the muon and tau scales.

But the toy reconstruction relied on freely-chosen values $\sigma_\mu \approx 3.5$ and $\sigma_\tau \approx 4.5$. The values were not derived. The toy paper was explicit about this and identified the σ_D derivation as the lead forward-programme task. The §6.4 decomposition proposed a structural model — $\sigma_D = a \cdot \Delta\beta_1 + b \cdot N_{persistent} + c \cdot R_D$, with a saturation hypothesis between μ and τ — but the model fit two totals with six numbers and was therefore underdetermined as a parameter count. What carried genuine falsifiable content was the *shape* of the decomposition: the saturation pattern between μ and τ.

The purpose of the present paper is sharply defined:

Derive σ_D for the charged leptons from substrate-counting principles, recovering the magnitude *and* the structural shape required by the hierarchy reconstruction — and additionally derive the n-loop structure of the charged-lepton sector itself, which the previous paper carried as a separate input.

The central conceptual move is that persistent distinguishability is not arbitrary informational bookkeeping. It is the **minimal independent commitment structure** required to preserve refinement-persistent particle identity. The present paper applies three structural rules (refinement-persistence, independence, and infrastructure-sharing) plus a closure principle (primitive-function exhaustion, §3.1) to derive both the σ-counting law and the n-loop structure. The n-loop assignments ($n_e = 1$, $n_\mu = 2$, $n_\tau = 3$), which the previous paper inherited as Claim 4 of its falsifiability statement, are derived in §4.1 as a structural consequence of the framework. The σ-counting law then operates on a sector with a *derived* three-generation closure rather than on a sector whose generation count is taken as input.

---

2. Persistent Distinguishability as Minimal Commitment Structure

We define persistent distinguishability as follows.

Definition (Persistent distinguishability load). σ_D is the minimal independent admissibility-fixed informational commitment load required to preserve PFD identity under refinement evolution, with redundant commitments quotient-identified and shared infrastructure counted once.

Three structural features are central to this definition:

- Only **refinement-persistent** distinctions count. Distinctions that fade under substrate refinement do not represent ongoing commitments.
- Only **independent** commitments count. Commitments that are derivable from others, removable by gauge choice, or implied by closure consistency are not independent.
- **Shared infrastructure** is identified once. Substrate machinery that supports multiple commitments simultaneously is a single substrate cost, not a per-commitment cost.

These features matter because without them the count grows combinatorially. A naïve count of all distinguishing features of a multi-loop PFD — every pairwise phase, every branching, every orientation — produces numbers in the $\mathcal{O}(n^2)$ range. The observed charged-lepton hierarchy requires σ-values in the $\mathcal{O}$(few) range, with sub-linear growth between generations. The independence rule and the infrastructure-sharing rule are precisely what bring the count down to that range. They are not numerical fudges; they are the structural content of *what counts* as a substrate commitment.

In particular, the substrate need not maintain redundant commitments because it has no mechanism for paying twice for the same constraint. If a phase relation is forced by closure consistency, the substrate does not separately maintain it — closure consistency does the work. If two records are gauge-equivalent, the substrate maintains one and the other is a relabelling. This is not an assumption about parsimony; it is how substrate bookkeeping is structured.

3. Three Structural Rules for Substrate Distinguishability Counting

We formalize the counting rules.

Rule 1 (Refinement-persistence). A commitment C contributes to σ_D only if C persists across substrate refinement: that is, only if there exists no refinement step that removes C without changing PFD identity. Transient commitments — those that emerge and dissolve under refinement without changing what particle is present — do not contribute.

Rule 2 (Independence). A set of commitments $\{C_1, C_2, \ldots, C_k\}$ contributes only via its independent rank. Specifically:

$\mathrm{rank}(\{C_1, \ldots, C_k\}) = k - (\text{number of redundant relations among } C_i)$

Redundant relations arise from three sources: (i) closure consistency (a commitment implied by closure of the refinement structure), (ii) global gauge equivalence (a commitment removable by a substrate-global gauge choice), and (iii) pairwise derivability (a commitment forced by other commitments in the set).

Rule 3 (Infrastructure-sharing). If a class of commitments $\{C_1, C_2, \ldots, C_k\}$ requires a common substrate infrastructure I for its maintenance, the contribution to σ_D from this class is dominated by the infrastructure cost, not by the per-commitment cost:

$\sigma_{class} = \eta \cdot \mathrm{cost}(I) + (\text{sub-leading per-commitment terms})$

with $\eta \in (0, 1]$ capturing the efficiency of infrastructure sharing across the commitments in the class. The infrastructure cost is paid *once* per class, regardless of how many commitments the infrastructure supports.

These three rules are the working content of the framework. They specify *what counts* as a contribution to σ_D. A fourth principle — used implicitly twice in what follows and worth formalising here — specifies *how the count terminates*.

3.1 The closure principle: primitive-function exhaustion

Definition (Primitive-function exhaustion). A commitment class C is *primitive-function exhausted* at primitive count k when the substrate-physical functions admissible in support of C at leading order have been enumerated and shown to be finite in number, with k functions counted as independent (each contributing at most one leading-order commitment under Rules 1–3) and no further admissible functions identifiable at leading order.

The principle has structural rather than merely numerical content. Exceeding the exhaustion bound is *structurally forbidden* in the following strong sense: any additional candidate function is necessarily one of:

- **(a) a sub-leading correction** to an enumerated function — not a new leading-order commitment;
- **(b) quotient-equivalent** to an enumerated function under Rule 2 — eliminable rather than countable;
- **(c) the introduction of a new commitment class entirely** — changing the framework's structure, not its parameter values within the existing class.

Option (c) is what makes the principle load-bearing. It says exhaustion bounds are *structurally protected*: they cannot be relaxed by parameter adjustment within the existing framework, only by introducing a new commitment class that operates at a different level of the substrate's counting structure. A bound set by primitive-function exhaustion therefore differs in kind from a bound set by parameter inference — the former specifies what the framework would have to change to accommodate exceeding the bound, the latter merely specifies what value the parameter is observed to take.

Applications in the present paper. The principle operates in two places.

1. **Phase-coherence class (§6.2.1):** primitive-function exhaustion at $k = 2$. Relative phase registration and phase-transport stability under refinement are the two admissible primitive functions for phase coherence; a third would constitute a new commitment class beyond the three identified in §§5–7, not a higher ξ value.
2. **Charged-lepton sector (§4.1, Step 3):** primitive-function exhaustion at $k = 3$. Three distinct refinement-persistent loop classes are admissible in the charged-lepton sector before further loops either reduce, dissolve, or change sector; a fourth class would belong to a different sector entirely (quark or neutrino), not a fourth charged-lepton generation.

Forward use. The principle is structurally available for sector-extension arguments. Whether the quark sector primitive-function exhausts at $k = 3$ (giving cross-sector consilience with the charged-lepton three-generation closure) or at a different value (indicating the principle's exhaustion counts are sector-local rather than universal) is a forward-task question. The principle applies in both cases; the specific exhaustion count is sector-dependent.

The principle is therefore one of the framework's reusable structural tools, alongside Rules 1–3. Each of the load-bearing structural-ceiling arguments in this paper — $\xi_{max} = 2$ and $n_{max} = 3$ — is an application of it, and the same tool will operate in the quark and neutrino extensions.

A caveat on what the principle does and does not do. The "shown finite in number" clause in the definition is doing substantial work, and it is worth being explicit that the principle is a *tool* whose *application* is itself a structural argument that must be performed sector-by-sector and class-by-class. Showing that the admissible primitive functions for a given commitment class are finite, and enumerating them, is not automatic — it requires (i) identifying the candidate functions admissible for the class at leading order, (ii) demonstrating that the candidate list is

complete, and (iii) walking through each non-enumerated candidate to verify it falls into option (a), (b), or (c) of the principle's exclusion structure. The two applications in the present paper (ξ at k = 2, n at k = 3 for charged leptons) each carry out this enumeration explicitly in §6.2.1 and §4.1 Step 3 respectively. The principle does not relieve the analyst of this work; it provides the structural framework within which the work has a definite shape.

The remaining sections apply Rules 1–3 and the primitive-function exhaustion principle to refinement-persistent charged-lepton PFDs.

---

4. Refinement-Persistent Charged-Lepton PFDs

The toy hierarchy reconstruction inherited the n-loop assignments

$n_e = 1, n_\mu = 2, n_\tau = 3$

as the content of Claim 4 in its falsifiability statement, with the explicit acknowledgement that the derivation was owed to the forward programme. The present paper derives them.

A framing note before the derivation. The hierarchy reconstruction of the previous paper, and the σ-counting law of §§5–9 below, depend on these assignments as inputs. But the more foundational result of the present paper is the derivation itself. Once the admissible refinement-persistent loop classes are exhausted at n = 3, the charged-lepton sector becomes a *finite topological sequence* rather than an open parameter tower. The hierarchy reconstruction is then a calculation *on* a closed sector rather than a fit to an open one. This is a different kind of theoretical object from the σ-counting law: it is a classification result for admissible charged-lepton PFD topologies, with a direct experimental falsifier (the existence of a fourth charged-lepton generation). The σ-counting law inherits its three-generation structure from this classification; the classification does not inherit its content from the σ-counting law.

The argument is given in §4.1 below as a three-step structural derivation; the assignments are then used in §§5–9.

4.1 Derivation of the charged-lepton n-band

The n-band of refinement-persistent charged-lepton PFDs follows from three structural requirements.

Step 1 — Minimal persistence: n ≥ 1. A charged lepton is a refinement-persistent PFD with closure-complete admissibility class C_D = complete (in the notation of the toy paper's invariant tuple). For the PFD to exist as a refinement-persistent object at all, it must contain at least one closed admissibility loop; an open structure dissolves into admissibility-equivalent background under refinement and cannot define a persistent identity. The single closed loop is the minimal carrier of refinement-persistent charged-lepton identity. Therefore n ≥ 1, with the electron saturating the lower bound.

Step 2 — Generation distinction: each successive generation adds exactly one independent refinement-persistent loop. Within the charged-lepton sector, the closure completeness class C_D, hypercharge label h_D, parity assignment π_D, flavour label $\chi\_D = L$, and record-coupling state ρ_D are all common to the three generations (see the toy paper's invariant tuple analysis). Generation differences therefore reside entirely in the hierarchy-relevant invariants (β_1, γ_D, ℓ_D). The argument that exactly one new loop is added per generation proceeds in two parts: first, identifying β_1 as the only *independent* admissibility-topological label available; second, invoking minimality on that label.

Part 2A — β_1 is the only independent topological label. Of the three hierarchy-relevant invariants:

- **γ_D (generation-depth index)** is, in the substrate stiffness framework, not an independent topological invariant of the PFD — it is the *depth of admissible refinement structure*, which is in turn fixed by the refinement-persistent loop content of the PFD. Concretely, $\gamma\_D = n$: the generation-depth index counts the depth at which refinement structure becomes non-trivial, and that depth equals the number of independent refinement-persistent loops the substrate must support. γ_D and β_1 are therefore two labels for the same underlying topological content, not independent generation-distinguishing invariants. A reader who instead prefers γ_D as the primary label can run the entire argument that follows with γ_D in place of β_1; the conclusion is unchanged because the two labels carry the same content.
- **ℓ_D (localization length)** is a continuous geometric quantity, not a topological label. Within the substrate-stiffness framework, ℓ_D is fixed by the substrate's Role-4 localization geometry as a function of the PFD's refinement-persistent loop content — schematically $\ell\_D = \ell(n)$ (see the substrate stiffness paper's localization-compression section). It is therefore a *derived* quantity from β_1 (or equivalently γ_D), not an independent label that could distinguish generations on its own. A continuous geometric quantity also cannot, by itself, supply the discrete generation-distinguishing label the sequence requires — a continuous ℓ_D could in principle take any value at any generation, so it carries no constraint forcing the discrete e/μ/τ sequence.

These two observations leave β_1 as the only independent admissibility-topological label distinguishing charged-lepton generations.

Part 2B — Minimality on β_1. A heavier charged-lepton generation must differ from the lighter generation in at least one independent topological label, since without any topological difference the heavier state would be quotient-equivalent to the lighter one under refinement and would not define a distinct charged-lepton identity. By Part 2A, the only such label available is β_1, so a heavier generation must satisfy $\beta_1(\text{heavier}) > \beta_1(\text{lighter})$. The minimal increment satisfying this constraint is $\Delta\beta_1 = 1$, since β_1 is integer-valued. Adding more than one loop per generation would over-specify the sequence — minimality on the only available distinguishing label fixes:

$n\_e = 1$, $n\_\mu = 2$, $n\_\tau = 3$

The minimality principle applied to *any other* candidate label would not yield this sequence, but Part 2A establishes that β_1 is the only candidate label available. The conclusion therefore rests on (i) the only-label claim of Part 2A plus (ii) integer-valued minimality on that label.

Step 3 — Closure exhaustion: n ≤ 3. The charged-lepton sector supports a finite n-band by **primitive-function exhaustion (§3.1) applied to the charged-lepton loop class**. Beyond three independent refinement-persistent loops, any further loop is necessarily one of the three excluded options identified by the principle:

- **(a) a sub-leading correction** to an already-enumerated loop class (not a new leading-order commitment);
- **(b) quotient-equivalent** to an enumerated loop class under Rule 2 (eliminated by gauge-redundancy with existing loops under the closure-completeness class, or destabilised under refinement as a non-persistent dissolution back into admissibility-equivalent structure);
- **(c) the introduction of a new commitment class** — in the present sector-extension language, *a sector change*: a fourth loop carrying a different hypercharge or admissibility class label belongs to the quark or neutrino sector, not to a fourth charged-lepton generation.

This is the load-bearing step of the n-band derivation. It is the same structural argument used in §6.2.1 for $\xi\_max = 2$ — both are applications of the primitive-function exhaustion principle of §3.1, at different commitment classes and with different exhaustion counts.

The strong-protection content of the principle applies here: a fourth charged-lepton generation cannot exist as a higher n within the framework's existing structure — it would necessarily be a sector change (option (c)), which is to say, *the framework predicts no fourth charged-lepton generation* unless the substrate's commitment-class structure itself is wrong.

Whether closure exhaustion at exactly n = 3 follows from a deeper substrate theorem — for example, from the Catalogue Closure Theorem or from K = 7 closure-manifold dimensionality — is not settled in the present paper. The present paper offers the structural argument (via the §3.1 principle) and the empirical falsifiability hook below; a microscopic derivation of $n\_max = 3$ from corpus-level closure theorems is part of the forward programme (Task 2).

Falsifiability of the n-band. Discovery of a genuine fourth charged-lepton generation would falsify the closure-exhaustion step and force n = 4 into the charged-lepton band. The framework's n-band derivation therefore makes an empirically testable structural prediction at the standard-model level: **the framework predicts that the charged-lepton sector closes at three generations.** This is a structural prediction of the substrate framework, not an absolute mathematical claim — it depends on the closure-exhaustion argument of Step 3, which is itself a structural argument rather than a corpus-level theorem. The prediction is testable in the strong sense that a single experimental discovery would refute it.

4.2 Consequences for the counting law

With the n-band derived, the first Betti increment over the electron reference is:

$\Delta\beta_1(D) = n_D - 1$

giving $\Delta\beta_1(e) = 0$, $\Delta\beta_1(\mu) = 1$, $\Delta\beta_1(\tau) = 2$. These integers feed directly into the cycle-persistence contribution $\sigma_{cycle} = n - 1$ of §5 and into the structural-bound arguments for ξ and η that follow.

With the n-loop structure derived, we identify three classes of commitment that contribute to σ_D for charged-lepton PFDs:

- **Cycle persistence** — commitments associated with each refinement-persistent loop individually.
- **Phase coherence** — commitments associated with the relative phase structure between loops.
- **Refinement infrastructure** — the substrate-global machinery that supports multi-loop refinement-persistent structure at all.

Each class is governed by a distinct combination of the three rules. We analyse them in turn.

---

5. Cycle Persistence Commitments

Each refinement-persistent admissibility loop in the transport graph requires the substrate to maintain a persistence commitment: a constraint that prevents the loop from dissolving into admissibility-equivalent background under refinement.

This commitment is *per-loop* and not naturally shared. Each loop could in principle dissolve independently, and the substrate constraint preventing one loop from dissolving does not prevent any other loop from dissolving. Different loops occupy different regions of the refinement-persistent transport graph, and their persistence constraints are independent in the sense of Rule 2.

The electron's single loop is the reference; we normalise its persistence commitment to zero (a relative count). Each additional loop contributes one independent persistence commitment.

$\sigma_{cycle}(n) = n - 1$

For the charged leptons:

Particle	n	σ_{cycle}
e	1	0
μ	2	1
τ	3	2

This is the only commitment class that grows linearly with n in the resulting count. The substrate-physical reason is straightforward: cycles are spatially distinct features of the transport graph, and their persistence is enforced locally rather than infrastructurally. Each cycle pays its own persistence cost.

---

6. Phase-Coherence Commitments and the Saturation Argument

Multi-loop PFDs carry relative phase structure between loops. A naïve count of pairwise relative phases for an n-loop PFD gives $C(n, 2) = n(n-1)/2$ phase relations. For $n = 2$ this is 1; for $n = 3$ it is 3. Under Rule 2 (independence), this count reduces.

6.1 Independence reduction

Among the $C(n, 2)$ pairwise phases, only $n - 1$ are linearly independent. The remaining $(n - 1)(n - 2)/2$ phases are derivable from the independent set: once $n - 1$ phases are fixed relative to a reference loop, the remaining pairwise phases are determined by the relation $\varphi_{ij} = \varphi_i - \varphi_j$. One further phase is then removable by global gauge choice (substrate-global U(1) phase rotation), leaving $n - 1 - 1 = n - 2$ phases that are independently fixed by the PFD's internal structure.

For $n = 2$ this leaves 0 internally-fixed phases (the one pairwise phase is the global gauge choice); for $n = 3$ it leaves 1; for $n = 4$ it leaves 2.

But this is not yet the full reduction.

6.2 Infrastructure saturation

Rule 3 applies. Maintaining a phase relation between two loops requires the substrate to support a phase-coherence infrastructure: the substrate-level machinery by which two distinct closure cycles share a common phase reference.

This infrastructure is structurally absent for one-loop PFDs, not merely normalised to zero. A single closure cycle has no second cycle to be coherent with; phase coherence is a *binary* relation between distinct loops, and a one-loop PFD has no admissible second argument for that relation. The substrate is not paying for phase-coherence infrastructure at $n = 1$ and then subtracting a reference cost — there is no relation to maintain. This is the substrate-physical content of the $\mathbb{1}_{n\geq 2}$ indicator: it marks a structural transition (the relation becomes well-defined at $n = 2$), not an arbitrary reference choice.

The structural transition is therefore discrete by construction. Smooth turn-on at fractional n is not available because the substrate has no concept of fractional admissibility loops — the n-loop assignment is integer-valued at the level of refinement-persistent transport graphs. A first-

principles enumeration that produced smooth interpolation between n = 1 and n = 2 (for example, via partial-persistence loops or non-integer cycle ranks) would falsify the framework's discrete-activation structure.

Once established at n = 2, the phase-coherence infrastructure supports phase coherence for any number of additional loops without requiring additional infrastructure. The substrate-physical reason is that phase coherence is *relational*. It is a property of how loops share a substrate-time framework, not a property of individual loops. The machinery enabling any phase coherence at all is the machinery enabling all phase coherence at all. Adding a third loop does not create new phase-coherence machinery; it draws on the same machinery already in place.

The saturation is reinforced by an operator-geometric feature of the closure manifold. The K = 7 closure basis is strongly non-orthogonal, with adjacent Gram-matrix entries $M_{i,i+1} \sim 0.7–0.9$ across the admissible mode-width range (established in the projected-operator analysis of the closure manifold; see the K = 7 closure-scale and exact-projection papers in the corpus). Multiple apparent phase commitments therefore correspond to partially overlapping closure support after projection onto the effective independent basis, and the *effective* phase-commitment count rises substantially more slowly than the raw combinatorial count. The two readings — substrate-physical infrastructure sharing and operator-geometric overlap reduction — are reinforcing but not identical; §11.4 distinguishes what each uniquely explains.

Under Rule 3, the contribution from the phase-coherence class is therefore dominated by the infrastructure cost, paid once when the infrastructure is first established:

$$\sigma_{phase}(n) = \xi \cdot \mathbb{1}_{\{n \geq 2\}}$$

where $\xi \in [1, 2]$ captures the cost of phase-coherence infrastructure and $\mathbb{1}_{\{n \geq 2\}}$ activates the infrastructure only for $n \geq 2$. The structural argument for the band endpoints is given in §6.2.1 below.

For the charged leptons:

Particle	n	σ_{phase}
e	1	0
μ	2	ξ
τ	3	ξ

The key structural feature is **saturation**: $\sigma_{phase}(\tau) = \sigma_{phase}(\mu)$. Adding the third loop does not add phase-coherence cost, because the third loop reuses the phase-coherence infrastructure established at the n = 2 transition.

6.2.1 Structural bounds on ξ

The present paper does not compute ξ from the microscopic substrate dynamics. What it does derive is the admissible band, by identifying the minimum and maximum leading-order

independent-commitment content allowed by refinement persistence, quotient-identification, and infrastructure sharing.

Phase coherence at $n \geq 2$ requires two structurally distinct primitive functions:

1. **Relative phase registration** — the substrate-level recording of how two distinct closure cycles are phase-aligned.
2. **Phase transport stability under refinement** — the constraint that the recorded phase relation persists across refinement steps without dissolving.

These are not the same function. Registration is a *static* relation (the phase at a given refinement step); transport stability is a *dynamical* commitment (preservation across refinement). But they share the same underlying phase-coherence infrastructure, since both require the substrate to support a common phase reference between cycles.

Lower bound ($\xi_{min} = 1$). Maximum infrastructure sharing: registration and transport stability are absorbed into a single shared commitment, the substrate paying once for the joint phase-coherence infrastructure. The lower bound is *not zero* because maintaining phase-coherence infrastructure at all requires at least one independent admissibility-fixed commitment — the constraint that a substrate-time framework exist against which loop phases can be referenced. The natural unit of substrate counting (one independent commitment) is the floor: an infrastructure-sharing rule cannot reduce the cost below the cost of having the infrastructure exist at all. A value $\xi < 1$ would correspond to the substrate maintaining phase coherence with less than one commitment of bookkeeping, which is structurally inadmissible by Rule 1 (refinement-persistent commitments are integer-valued at the level of independent admissibility constraints; sub-unit values either fail to persist or are not commitments). The lower bound $\xi_{min} = 1$ is therefore the floor set by the substrate's discrete commitment-counting unit, not an additional asserted parameter.

A clarification on integer vs. fractional values. ξ is integer-valued at its endpoints ($\xi_{min} = 1$, $\xi_{max} = 2$) because it counts admissibility-fixed commitments, and commitments are integer-valued under Rule 1: a partial commitment is either no commitment at all (if it fails to persist) or one commitment expressed at sub-leading order (if it's a correction to an enumerated function). ξ can take continuous values *between* its integer endpoints, but those continuous interior values represent fractional infrastructure sharing between the two integer commitments — they are not themselves commitment counts. By contrast, $\eta \in [¼, ¾]$ in §7 is continuous throughout its range because it measures sharing efficiency (a fractional reduction of a fixed commitment count $R = 2$) rather than counting commitments itself. The integer-valued floor under Rule 1 therefore applies to ξ but not to η: efficiency parameters can take any value in (0, 1], whereas commitment counts cannot fall below 1 without ceasing to exist.

Upper bound ($\xi_{max} = 2$). No infrastructure sharing between the two functions: registration and transport stability remain independent after quotient reduction, each contributing one leading-order commitment. The substrate pays separately for the static relation and for its dynamical stability.

Why exactly two and not more — primitive-function exhaustion. The structural ceiling on ξ is set by primitive-function exhaustion (§3.1) applied to the phase-coherence commitment class. The admissible primitive functions for phase coherence are enumerated above: there are exactly two — registration and transport stability — and a third candidate function falls into one of the three excluded options identified by the principle:

- A "second-order phase derivative across refinement" or similar elaboration is **(a) a sub-leading correction** to transport stability, not a new leading-order commitment.
- A "phase parity" or "phase reflection" relation is **(b) quotient-equivalent** to registration under Rule 2, eliminable by the substrate-global gauge choice.
- A genuinely new phase-coherence primitive would constitute **(c) a new commitment class** beyond the three identified in §§5–7, changing the framework's structure rather than its parameter range.

The framework therefore cannot admit $\xi > 2$ by parameter adjustment alone. $\xi_{max} = 2$ is *structurally protected* in the sense of §3.1: exceeding it requires a change in the framework's commitment-class structure, not merely a higher parameter value.

The actual value of ξ within [1, 2] is set by how strongly registration and transport stability share infrastructure, which in turn depends on substrate thermodynamics. This is the first part of Task 1 in the forward programme.

6.3 Why saturation rather than linear growth

The phase-coherence saturation is the central non-trivial structural prediction of this paper, and it is worth being explicit about why it holds.

The alternative — $\sigma_{phase} = n - 1$ growing linearly with n — would correspond to each additional loop paying its own phase-coherence cost as it joins the PFD. This would treat phase relations as per-pair commitments analogous to cycle persistence. Under this alternative, the substrate would maintain an independent phase commitment for each pair of loops, and the cost would scale with the number of independent pairs.

This is the wrong picture for substrate phase coherence because the substrate does not separately maintain pairwise phase relations. It maintains a global substrate-time framework against which all loop phases are simultaneously referenced. The cost of this framework is the cost of being able to track *any* phase relation, not the cost of tracking each one separately. Once the framework is in place, additional loops draw their phase reference from it for free.

The same structural argument applies, by analogy, to refinement infrastructure in §7. Both phase coherence and refinement infrastructure are *relational* features that scale with the existence of multi-loop structure rather than with the number of loops.

The saturation of σ_{phase} is what allows the framework to reproduce the toy paper's Claim 5 (saturation of $N_{persistent}$ between μ and τ). The toy paper's σ_{record} term — proposed there

with no derivation — is, in the present framework, exactly the saturating phase-coherence infrastructure cost. The structural mapping is made explicit in §11.

---

7. Refinement Infrastructure Commitments

The third commitment class is the most substrate-global. Multi-loop refinement-persistent structure requires two further substrate commitments that exist *at the substrate level* rather than at the level of individual loops or loop-pairs:

1. **Refinement identity preservation.** The substrate must commit to identifying the PFD across refinement steps — that is, recognising that the post-refinement transport graph is the *same* PFD as the pre-refinement transport graph rather than a new one. For a one-loop PFD this is structurally trivial (there is nothing to confuse it with); for $n \geq 2$ it is non-trivial because the substrate must distinguish PFD identity from rearrangements of the loop structure under refinement.
2. **Admissibility-compatible transport continuity.** The substrate must ensure that the transport graph remains admissibility-compatible across refinement steps. Admissibility-violating refinements would either dissolve the PFD or transform it into a different admissibility class. For multi-loop PFDs this is a substrate-global constraint, since admissibility-violation in any loop affects the entire PFD.

Both commitments are infrastructural in the Rule 3 sense: they exist at the substrate level and support all loops simultaneously. Their cost is therefore paid once per multi-loop PFD class, with an efficiency parameter η capturing the degree of sharing.

As with σ_{phase}, the $n = 1$ cost is **structurally zero, not normalised to zero**. A one-loop PFD has no rearrangements of its loop structure that the substrate could confuse with PFD identity loss — there is only one loop, and refinement either preserves it (PFD persists) or dissolves it (PFD is gone). The binary choice does not require an identity-preservation commitment; the commitment becomes meaningful only when multiple loops introduce the possibility of identity-ambiguous rearrangement. Similarly, admissibility-compatible transport continuity becomes a substrate-global constraint only when multi-loop coupling makes admissibility-violation in one loop affect the whole PFD. Both refinement-infrastructure commitments are therefore well-defined only at $n \geq 2$, and the $\mathbb{1}_{n\geq 2}$ indicator is again a structural transition rather than a reference choice.

We count $R = 2$ refinement-infrastructure commitments per multi-loop PFD class, with sharing efficiency η. The contribution is:

$$\boldsymbol{\sigma_{refine}(n) = \eta \cdot R \cdot \mathbb{1}_{n\geq 2} = 2\eta \cdot \mathbb{1}_{n\geq 2}}$$

with $\eta \in [¼, ¾]$. The structural argument for these bounds is given in §7.2 below.

For the charged leptons:

Particle	n	σ_refine
e	1	0
μ	2	2η
τ	3	2η

As with σ_phase, refinement infrastructure saturates between μ and τ.

7.1 Two further saturating commitments?

A natural question is whether there are additional substrate-global commitments not captured by R = 2. Candidate examples include:

- **Substrate-orientation commitments** — preservation of a global orientation across refinement.
- **Substrate-causality commitments** — preservation of refinement causality.

We do not include these in R because they are common to *all* PFDs (including the one-loop electron) and therefore drop out of the inter-generational σ_D differences. They contribute to the absolute scale of σ_D but not to its hierarchy structure, and the present paper concerns hierarchy structure. The absolute scale of σ_D awaits the absolute-mass-scale derivation, which depends on additional substrate-condensate parameters and is deferred to a separate computation.

7.2 Structural bounds on η — and an honesty note

The η band is less tightly derived than the ξ band of §6.2.1, and the paper should be explicit about this. The structural argument bounds η between two limiting cases of how the two refinement commitments share substrate infrastructure, but neither endpoint is forced in the way that ξ_max = 2 is forced by the count of primitive phase-coherence functions.

Lower bound (η_min = ¼). Maximum-sharing hypothesis: each refinement-infrastructure commitment is distributed across four effective support channels — loop identity, loop transport, dual identity, dual transport. Under maximum sharing, only one quarter of each raw commitment remains as independent substrate load after overlap reduction.

Upper bound (η_max = ¾). Minimum-sharing hypothesis: one quarter of each raw commitment is removed by unavoidable overlap-reduction and gauge-quotienting (the irreducible infrastructure-sharing imposed by Rule 3 even in the weakest sharing case), leaving three quarters as independent substrate load.

An honest note about these endpoints. The "four effective support channels" assumption at η_min = ¼ is a working hypothesis about how the substrate distributes refinement infrastructure across multi-loop structure. It is structurally plausible — loop and dual directions are both supported, and identity and transport are distinct aspects — but it is not derived from substrate dynamics. Similarly, the η_max = ¾ figure is the strongest claim about how much of each commitment survives any unavoidable sharing; an alternative analysis could land η_max at ⅔ or

⅘ without much friction. Both endpoints should therefore be read as **structurally motivated admissible bounds**, not as microscopic derivations.

A true derivation of η requires explicit construction of the refinement-support channel space and the overlap-reduction operator on the closure manifold — the same operator-geometric structure introduced conjecturally in §11. This is the second part of Task 1 in the forward programme, and is the less-secure of the two band derivations in the present paper. The ξ band is structurally pinned at the top by the primitive-function ceiling; the η band is not pinned at either endpoint in the same way.

7.3 Combined band

Combining the ξ and η bounds:

$\xi + 2\eta \in [1 + 2(¼), 2 + 2(¾)] = [1.5, 3.5]$

so that:

$\sigma_\mu = 1 + \xi + 2\eta \in [2.5, 4.5]$

$\sigma_\tau = 2 + \xi + 2\eta \in [3.5, 5.5]$

The midpoint $\xi + 2\eta = 2.5$ — reached at the centre of both individual bands — corresponds to $\sigma_\mu = 3.5$, $\sigma_\tau = 4.5$: the toy values from the previous reconstruction.

A note on the relation between this band derivation and Prediction P1 in §14: P1 is the prediction that the substrate-thermodynamic computation will land (ξ, η) near the centre of these bands. The bands themselves are now structurally derived — the previous draft asserted them without argument — but their centre being privileged still rests on the substrate-thermodynamic computation actually landing there. The strengthening from "asserted bands" to "structurally bounded bands" makes the natural-midpoint argument tighter but does not on its own derive the natural-midpoint location.

---

8. The Minimal Quotient-Counting Law

Combining the three classes:

$\sigma(n) = \sigma_{cycle}(n) + \sigma_{phase}(n) + \sigma_{refine}(n)$

$$**= (n - 1) + \xi \cdot \mathbb{1}_{\{n \geq 2\}} + \eta \cdot R \cdot \mathbb{1}_{\{n \geq 2\}}**$$

$$**= (n - 1) + (\xi + 2\eta) \cdot \mathbb{1}_{\{n \geq 2\}}**$$

This is the **minimal quotient-counting law** for refinement-persistent charged-lepton PFDs. Two parameters appear:

- $\xi \in [1, 2]$: phase-coherence infrastructure cost
- $\eta \in [¼, ¾]$: refinement-infrastructure sharing efficiency

The substrate count $R = 2$ is structural and not a free parameter; n is integer and inherited from the transport-graph construction.

Two structural features of this law deserve emphasis.

First: the inter-generational gap is parameter-free. For any $n_a > n_b$, both with $n \geq 2$:

$$\sigma(n_a) - \sigma(n_b) = n_a - n_b$$

The $(\xi + 2\eta)$ term saturates and cancels in the difference. Between μ ($n = 2$) and τ ($n = 3$), the gap is exactly 1, independent of ξ and η. The structural shape of the inter-generational hierarchy is therefore fixed by the n-loop counting alone, not by the parameter values.

Second: the muon baseline is parameter-controlled but bounded. The intra-generational baseline σ_μ depends on $(\xi + 2\eta)$, but this combination is bounded:

$$\xi + 2\eta \in [1 + ½, 2 + 3/2] = [1.5, 3.5]$$

so $\sigma_\mu = 1 + (\xi + 2\eta) \in [2.5, 4.5]$. The baseline cannot be tuned to arbitrary values; it is structurally constrained to a factor-of-≈2 band.

---

9. Charged-Lepton Enumeration

9.1 Electron (n = 1)

$\sigma_e = 0 + 0 + 0 =$ **0**

The electron is the one-loop reference. No multi-loop infrastructure is required; no phase-coherence machinery is required; no cycle-persistence commitments beyond the reference cycle. All hierarchy-contributing commitments are normalised to zero.

9.2 Muon (n = 2)

$\sigma_\mu = 1 + \xi + 2\eta$

```
∈  [1 + 1 + ½,  1 + 2 + 3/2]

=  **[2.5, 4.5]**
```

with natural midpoint at $\xi + 2\eta = 2.5$ (the centre of the [1.5, 3.5] combined band):

$\sigma_\mu \approx 1 + 2.5 =$ **3.5**

The muon is the infrastructure-forming transition. The single cycle-persistence commitment (1) is joined by the establishment of the phase-coherence infrastructure (ξ) and the refinement infrastructure (2η). All three classes activate at once. At the natural midpoint of the combined parameter band, σ_μ reproduces the toy paper's $\sigma_\mu = 3.5$ exactly. The framework's quantitative reach is one-dimensional in $(\xi + 2\eta)$, not two-dimensional in (ξ, η) separately: the same $\sigma_\mu = 3.5$ is reached by $(\xi, \eta) = (1.5, 0.5)$, or $(2.0, 0.25)$, or $(1.0, 0.75)$, and the present framework cannot distinguish among these without an independent determination of one parameter.

9.3 Tau (n = 3)

$\sigma_\tau = 2 + \xi + 2\eta$

```
∈  [2 + 1 + ½,  2 + 2 + 3/2]

=  **[3.5, 5.5]**
```

with natural midpoint at $\xi + 2\eta = 2.5$:

$\sigma_\tau \approx 2 + 2.5 =$ **4.5**

The tau is the infrastructure-extending transition. It adds one cycle-persistence commitment (raising σ_{cycle} from 1 to 2) but draws on the same phase-coherence and refinement infrastructure that the muon already established. The $(\xi + 2\eta)$ contribution is unchanged from the muon.

---

10. The Origin of Non-Uniform Hierarchy Growth

The toy hierarchy reconstruction required the structural feature:

$\sigma_\tau - \sigma_\mu < \sigma_\mu - \sigma_e$

(Claim 5 of the toy paper's falsifiability statement — saturation between μ and τ). The present framework forces this feature without parameter choice:

$\sigma_\tau - \sigma_\mu = (2 + \xi + 2\eta) - (1 + \xi + 2\eta) =$ **1**

$\sigma_\mu - \sigma_e = (1 + \xi + 2\eta) - 0 =$ **$1 + \xi + 2\eta \in [2.5, 4.5]$**

The gap from electron to muon is 2.5 to 4.5 times larger than the gap from muon to tau. This is the substrate-counting origin of the non-uniform hierarchy growth.

The structural reason is now explicit:

- The muon transition ($e \to \mu$) pays for **one cycle**, **the phase-coherence infrastructure**, and **the refinement infrastructure** all at once. Three classes activate simultaneously.
- The tau transition ($\mu \to \tau$) pays for **one additional cycle**, period. The two infrastructure classes saturate and contribute nothing additional.

This is the substrate analogue of the principle that *starting* a new operation has higher overhead than *extending* it. The muon is the "startup cost" of multi-loop refinement-persistent structure; the tau is "one more cycle on top of existing infrastructure." Both transitions add the same integer number of cycles (one), but the muon transition additionally crosses two infrastructure thresholds while the tau crosses zero.

This pattern is a structural prediction of the counting framework, not an input. It would be violated by:

- Phase-coherence machinery that scales linearly with loop number (treated as per-pair commitments rather than as a shared substrate framework).
- Refinement infrastructure that grows with PFD complexity rather than saturating beyond the first multi-loop transition.
- Cycle persistence that, conversely, *saturates* rather than scaling linearly (which would compress the entire hierarchy).

Each of these alternatives would change the shape of $\sigma(n)$ and therefore the shape of the hierarchy. The present framework's prediction is that none of them are realised — that the substrate organises commitments exactly as described in Rules 1–3.

---

11. Operator Interpretation: Distinguishability as Overlap-Corrected Rank

The quotient-counting framework developed in §§5–9 treated refinement-persistent commitments as approximately independent after the explicit redundancy and gauge reductions of Rule 2 and the infrastructure-sharing of Rule 3. This section makes explicit a complementary reading: the same counting law admits a natural operator-geometric interpretation as a count of *effective independent* commitments on a strongly non-orthogonal closure manifold.

The reading is complementary, not foundational. §10 supplied the substrate-physical *why* of saturation (phase coherence and refinement infrastructure are relational substrate features that activate once and are then shared). The present section supplies a parallel *how* in operator language. The two are intended as two descriptions of the same structural fact, not as competing accounts.

11.1 Non-orthogonal closure modes

The projected-operator analysis of the K = 7 closure manifold (see the K = 7 closure-scale paper and the projected-operator paper in the VERSF corpus) establishes that the closure modes ψ_i are strongly non-orthogonal. Adjacent Gram-matrix entries satisfy

$$M_{i,i+1} = \langle \psi_i, \psi_{i+1} \rangle \sim 0.7 – 0.9$$

across the admissible mode-width range. This is structural, not perturbative: the closure-mode overlap is not a small correction to an approximately orthogonal basis, but a leading-order feature of the closure geometry. The figures $M_{i,i+1} \sim 0.7–0.9$ are taken from the projected-operator computation, not assumed here.

In a strongly non-orthogonal basis, raw counting of basis components systematically overcounts independent content. Two basis modes ψ_i, ψ_j with overlap $\langle \psi_i, \psi_j \rangle$ near unity carry mostly the same closure support; their *effective* independent content is much less than their nominal pair count.

11.2 Effective independence via Gram projection

Let M denote the closure-mode Gram matrix and let κ denote a closure configuration whose decomposition onto the mode basis is sought. The exact projection theorem (established elsewhere in the corpus) gives the exact closure amplitudes

$$q = M^{-1}b, \text{ with } b_i = \langle \psi_i, \kappa \rangle$$

so that closure content is not extracted by inner products against the raw basis (which would overcount), but by inner products against the *dual* basis $\tilde{\psi}_i = (M^{-1})_{ij}\, \psi_j$. The dual basis is, by construction, orthogonal to all closure modes other than the one it pairs with, and it is the basis on which independent closure content is genuinely independent.

This gives a precise meaning to "effective independent commitment" on the closure manifold: it is a commitment expressed in the dual basis, not the raw mode basis. Raw commitments that overlap strongly in the mode basis project onto a single direction (or onto a much lower-rank subspace) in the dual basis. The substrate maintains the dual-basis content, not the raw-basis content.

11.3 The conjectured operator structure

These observations motivate the following conjecture about the operator structure underlying the counting law:

Conjecture (Operator structure of σ_D). There exists a refinement-persistence operator L on the closure manifold such that the substrate distinguishability load is given, at leading order, by

$$\sigma_{eff} \sim \mathrm{Tr}(M^{-1} L)$$

where M is the closure-mode Gram matrix and L is the refinement-persistence operator (counting refinement-persistent closure structure). The minimal quotient-counting law $\sigma(n) = (n - 1) + \xi \cdot \mathbb{1}_{n\geq 2} + \eta \cdot R \cdot \mathbb{1}_{n\geq 2}$ is the coarse-grained leading-order approximation to $\mathrm{Tr}(M^{-1}L)$ for refinement-persistent n-loop charged-lepton PFDs.

Two features of this conjecture deserve explicit framing.

The conjecture is not derived in this paper. Validating it requires (i) constructing M for the n-loop charged-lepton PFD closure modes; (ii) constructing the refinement-persistence operator L from the substrate refinement dynamics; (iii) computing $\mathrm{Tr}(M^{-1}L)$ for n = 1, 2, 3; and (iv) verifying that the result reproduces $(n - 1) + (\xi + 2\eta) \cdot \mathbb{1}_{n\geq 2}$ at leading order with (ξ, η) sitting in their structurally-motivated bands. None of these computations is performed here. They are listed as forward tasks in §15.

The conjecture is a working hypothesis about leading-order structure, not a claim of equivalence. Even if the conjecture survives validation, $\sigma(n) = \mathrm{Tr}(M^{-1}L)$ is expected only at leading order; sub-leading operator corrections will appear at the same time as the sub-leading per-pair phase corrections noted in §6.2. The counting law and the operator structure are not the same object; the former is the latter's coarse-grained approximation.

11.4 What each reading uniquely explains

The substrate-physical reading of §10 and the operator-geometric reading of the present section are reinforcing but not identical. They are best understood as addressing different questions about the same structural fact.

What §10 uniquely explains: why infrastructure exists at all once n ≥ 2. The substrate-physical argument says that phase coherence and refinement identity are *relational* features of multi-loop structure — they are well-defined only when there are multiple loops to be coherent across, or multiple loop arrangements to disambiguate. The substrate establishes the infrastructure at n = 2 because at n = 2 the relations first become meaningful. This is a content question (what is the substrate maintaining?), and the operator picture cannot answer it on its own — the operator picture tells you about the rank of the effective independent subspace, not about which substrate features that subspace represents.

What §11 uniquely explains: why the effective rank is reduced regardless of substrate physics. The operator-geometric argument says that on a strongly non-orthogonal closure manifold, raw component counting overcounts effective independent content, by a structural factor controlled by M. This would hold even if the substrate-physical motivation for infrastructure sharing were absent: any counting on a non-orthogonal basis is reduced by overlap. The operator picture supplies a *geometric mechanism* for rank reduction that operates whether or not the substrate-physical reading is correct.

Where the two readings reinforce each other. Both predict saturation of σ_{phase} and σ_{refine} between μ and τ; both predict that infrastructure activates at n = 2; both predict that the effective count is sub-linear in raw commitment count. The reinforcement is non-trivial because the two

arguments arrive at the same predictions from different starting points — one substrate-physical, one operator-geometric. If both arguments are correct, the framework rests on two compatible foundations; if one fails and the other survives, the framework's predictions are sustained by whichever survives.

Where the readings could in principle decouple. If the operator conjecture $\sigma_{eff} \sim \mathrm{Tr}(M^{-1}L)$ is wrong but the substrate-physical reading is right, the counting law still holds via Rule 3 and infrastructure-sharing; the operator structure is then a misleading geometric interpretation. Conversely, if the substrate-physical reading is wrong but the operator structure is right, the counting law holds via geometric rank reduction and the §10 narrative about substrate "not paying twice" would need replacement. The two readings are *consistent* with each other, but the counting framework would survive the failure of either one.

If the operator conjecture survives validation, two structural features of the counting framework gain a deeper reading. Infrastructure sharing (Rule 3) gains a geometric mechanism — closure-mode overlap — as an implementation of the substrate-physical principle that infrastructure is shared. And the saturation pattern $\sigma_\tau - \sigma_\mu < \sigma_\mu - \sigma_e$ acquires a geometric origin: the mode-basis content added at the muon transition spans a much larger effective independent subspace than the content added at the tau, because the latter overlaps strongly with the former in M.

Neither of these is established by the present paper. Both are framed here as the kind of structural reading that would emerge if the conjecture survives — and as part of the case for why the validation task (forward task 3 below) is worth the computational effort it will require.

11.5 Status

The operator interpretation is present in this paper as a connection, not as a foundation. The counting law stands on its three structural rules independently of the operator reading; if Task 3 (operator-framework validation) refutes the $\mathrm{Tr}(M^{-1}L)$ conjecture, the framework's substrate-physical content is unaffected and the counting law continues to derive σ values from infrastructure sharing alone. What the framework would lose is the geometric mechanism by which the substrate-physical sharing is implemented, and the operator-reading sections of the abstract, §11, and §16 would need to be retracted as conjectures that did not survive testing. That is a meaningful loss of structural depth, but it is not a loss of the framework's core counting content. Whether the connection is real — whether $\mathrm{Tr}(M^{-1}L)$ actually reproduces the counting law at leading order — is now a sharply defined computational question.

---

12. Structural Map to the Toy Paper's §6.4 Decomposition

The toy paper's §6.4 decomposition was:

$$\sigma_D = a \cdot \Delta\beta_1 + b \cdot N_{persistent} + c \cdot R_D$$

with toy weights (a, b, c) = (1.0, 0.5, 0.5) and counts that produced $\sigma_\mu = 3.5$, $\sigma_\tau = 4.5$ with saturation of N_persistent and R_D between μ and τ.

The present framework derives the same structural form. The mapping is:

Toy paper §6.4	**Present framework**
$a \cdot \Delta\beta_1$	$\sigma_{cycle} = (n - 1)$
$b \cdot N_{persistent}$	$\sigma_{phase} = \xi \cdot \mathbb{1}_{\{n \geq 2\}}$
$c \cdot R_D$	$\sigma_{refine} = \eta \cdot R \cdot \mathbb{1}_{\{n \geq 2\}}$

In particular:

- The toy paper's σ_{cycle} component with a = 1 is now derived: each additional refinement-persistent cycle contributes one independent persistence commitment under Rules 1 and 2.
- The toy paper's σ_{record} component with $b \cdot N_{persistent} = 1.5$ and saturation between μ and τ is now derived as $\sigma_{phase} = \xi$ with $\xi \approx 1.5$ (one point on the natural-midpoint locus $\xi + 2\eta = 2.5$) and saturation forced by Rule 3 applied to phase-coherence infrastructure.
- The toy paper's σ_{refine} component with $c \cdot R_D = 1$ and saturation between μ and τ is now derived as $\sigma_{refine} = 2\eta$ with $2\eta \approx 1.0$ (correspondingly one point on the natural-midpoint locus) and saturation forced by Rule 3 applied to refinement infrastructure.

The "natural midpoint" of the framework is one-dimensional: it is the locus $\xi + 2\eta = 2.5$, not the point $(\xi, \eta) = (1.5, 0.5)$. The toy paper's specific values $b \cdot N_{persistent} = 1.5$ and $c \cdot R_D = 1.0$ land on this locus at its symmetric centre, but the present framework cannot distinguish that point from any other on the locus without an independent determination of one parameter.

The toy paper's free weights (a, b, c) are now constrained by Rules 1–3 plus the infrastructure-sharing parameter ranges. The toy paper's free substrate counts ($\Delta\beta_1$, N_persistent, R_D) are now structural quantities: $\Delta\beta_1$ from the transport-graph construction, $N_{persistent} \leftrightarrow \xi$ from phase-coherence infrastructure, $R_D \leftrightarrow R = 2$ from refinement-infrastructure enumeration.

The number-absorption is concrete. The toy paper's σ_{record} term was $b \cdot N_{persistent} = 0.5 \cdot 3 = 1.5$ at the muon, with both b and N_persistent freely chosen subject to their product reproducing the toy σ_μ. In the present framework, that entire 1.5 is absorbed into the single parameter $\xi \approx 1.5$: the b and N_persistent split does not exist at the new level of description, because phase-coherence infrastructure is a single substrate cost (Rule 3) rather than a per-record cost multiplied by a record count. Similarly, the toy paper's σ_{refine} term was $c \cdot R_D = 0.5 \cdot 2 = 1.0$, with both c and R_D free; in the present framework, $R_D = 2$ maps directly to the structural count R = 2 (refinement identity preservation + admissibility-compatible transport continuity), and $c = 0.5$ corresponds to $\eta = 0.5$ — a sharing efficiency rather than a free weight.

What the toy paper carried as six freely-chosen numbers (a, b, c, $\Delta\beta_1$, N_persistent, R_D), each contributing independently to fitting two totals, reduces under the present framework to: two bounded parameters (ξ, η); one structural integer (R = 2); and one integer assignment per

generation (n). The §6.4 decomposition's central structural prediction — saturation of two of its three components between μ and τ — is now a *consequence* of the counting framework rather than a hypothesis layered on top of it.

This is the genuine structural advance over the toy paper.

---

13. Comparison with the Toy Hierarchy Reconstruction

The previous reconstruction used $\sigma_\mu = 3.5$ and $\sigma_\tau = 4.5$ to reach the observed lepton ratios under the substrate stiffness identity. The present derivation produces:

$\sigma_\mu \in [2.5, 4.5]$ (midpoint = 3.5)

$\sigma_\tau \in [3.5, 5.5]$ (midpoint = 4.5)

with the inter-generational gap $\sigma_\tau - \sigma_\mu = 1$ fixed by the counting independently of (ξ, η).

The toy values (3.5, 4.5) correspond to $(\xi + 2\eta) = 2.5$, which is the midpoint of the [1.5, 3.5] range of $(\xi + 2\eta)$. This is reached by any point of the natural-midpoint *locus* $\{(\xi, \eta) : \xi + 2\eta = 2.5\}$, which includes $(\xi, \eta) = (1.5, 0.5)$ as its symmetric centre but also includes (2.0, 0.25), (1.0, 0.75), and any other point on the line. The framework cannot distinguish among these without an independent determination of either ξ or η. The toy values therefore correspond to the centre of a structurally-bounded one-dimensional parameter combination, with the natural midpoint of that combination reproducing them exactly.

13.1 Honest reckoning of what's derived

The natural-midpoint reproduction is the strongest single quantitative result in the paper, but it is contingent. A substrate-thermodynamic computation that lands $\xi + 2\eta$ away from 2.5 would shift the σ values without changing the framework's structural shape. Exact reproduction at the natural midpoint is therefore evidence for, but not equivalent to, a derivation of $\xi + 2\eta \approx 2.5$. Distinguishing "the combination lands at the natural midpoint because the structural argument is right" from "the combination lands at the natural midpoint by coincidence" requires the substrate-thermodynamic computation; until that computation is done, the natural-midpoint reproduction is suggestive rather than dispositive.

What the derivation *does* establish:

1. **The σ values land in the few-unit regime.** σ_μ and σ_τ are $\mathcal{O}$(few), not $\mathcal{O}(10)$ or $\mathcal{O}(0.1)$. This is what the toy hierarchy required.
2. **The inter-generational gap $\sigma_\tau - \sigma_\mu = 1$ is forced.** It is not a fit. It is a direct algebraic consequence of the counting law and the n-loop assignments, and is independent of (ξ, η).
3. **The intra-generational baseline $\sigma_\mu = 1 + \xi + 2\eta$ is bounded in [2.5, 4.5].** The toy value 3.5 sits at the centre of this band.

4. **The non-uniform hierarchy shape $\sigma_\tau - \sigma_\mu < \sigma_\mu - \sigma_e$ is forced** for all parameter values in the admissible ranges, not as a tuning choice.
5. **At the natural midpoint of the combined parameter band ($\xi + 2\eta = 2.5$), σ_μ and σ_τ reproduce the toy values exactly.** This is the framework's strongest single quantitative result, but it remains contingent on the substrate-thermodynamic computation actually landing $\xi + 2\eta$ near 2.5.

What the derivation *does not* establish:

1. **The exact values of (ξ, η)** — these are infrastructure-sharing efficiency parameters whose precise values await first-principles substrate-thermodynamic computation.
2. **That the operator conjecture of §11 is correct** — $\mathrm{Tr}(M^{-1}L)$ reproducing the counting law at leading order is offered as a hypothesis, not as a result.

13.2 What changes if (ξ, η) sit elsewhere in their bands

If a first-principles substrate-thermodynamic computation places ξ or η elsewhere in the admissible ranges, the σ values shift but the structural shape is preserved. Two illustrative cases:

- **(ξ, η) = (1.0, 0.25)** — minimal infrastructure cost: $\sigma_\mu = 2.5$, $\sigma_\tau = 3.5$. The hierarchy is somewhat compressed; $m_\mu/m_e \approx 70$, $m_\tau/m_e \approx 1130$. Order of magnitude preserved; quantitative match lost.
- **(ξ, η) = (2.0, 0.75)** — maximal infrastructure cost: $\sigma_\mu = 4.5$, $\sigma_\tau = 5.5$. The hierarchy is somewhat stretched; $m_\mu/m_e \approx 565$, $m_\tau/m_e \approx 9300$. Order of magnitude preserved; quantitative match lost in the other direction.

The framework reproduces the *shape* of the hierarchy across the full parameter range and the *exact values* only at the natural midpoint of the combined band. The natural-midpoint reproduction is the strongest single piece of evidence that $\xi + 2\eta \approx 2.5$ is structurally correct, but it is not yet a derivation of that combination.

---

14. Falsifiability Statement

The framework as instantiated in this paper makes four sharp empirically-testable structural claims (D1–D4) and one quantitative prediction (P1) that has the status of a conditional forecast rather than a peer claim. The distinction matters: D1–D4 are claims about substrate counting and substrate structure that can be tested directly — three by first-principles enumeration and one (D4) by experimental particle-physics data; P1 is a prediction about the outcome of a substrate-thermodynamic computation that has not yet been formulated, let alone performed.

The claims refine the toy paper's Claims 1, 4, and 5.

Claim D1 (σ-counting law). A first-principles enumeration of independent admissibility-fixed informational commitments for refinement-persistent n-loop charged-lepton PFDs yields $\sigma(n) = (n - 1) + \xi \cdot \mathbb{1}_{\{n\geq 2\}} + \eta \cdot R \cdot \mathbb{1}_{\{n\geq 2\}}$ with $\xi \in [1, 2]$, $\eta \in [¼, ¾]$, $R = 2$.

Claim D2 (Phase-coherence saturation). Phase-coherence infrastructure saturates between the muon and tau scales: $\sigma_{phase}(\tau) = \sigma_{phase}(\mu)$ within $\mathcal{O}(10\ \%)$.

Claim D3 (Refinement-infrastructure saturation). Refinement infrastructure saturates between the muon and tau scales: $\sigma_{refine}(\tau) = \sigma_{refine}(\mu)$ within $\mathcal{O}(10\ \%)$.

Claim D4 (Three-generation closure). The framework predicts that the charged-lepton sector closes at three generations: refinement-persistent charged-lepton PFDs admit $n \in \{1, 2, 3\}$, with the upper bound set by closure exhaustion (§4.1, Step 3). This claim is directly testable by experimental search: discovery of a genuine fourth charged-lepton generation falsifies the closure-exhaustion step and forces revision of §4.1.

The quantitative prediction is:

Prediction P1 (Natural-midpoint reproduction). The substrate-thermodynamic computation produces $\xi + 2\eta \approx 2.5$ — the centre of the structurally-derived [1.5, 3.5] band — and the counting law thereby reproduces $\sigma_\mu \approx 3.5$ and $\sigma_\tau \approx 4.5$. Since the counting law depends on (ξ, η) only through $\xi + 2\eta$, P1 is a one-dimensional prediction in the combined parameter, not a two-dimensional prediction in (ξ, η) separately; the framework cannot distinguish (1.5, 0.5) from (2.0, 0.25) or (1.0, 0.75) without an independent determination of one parameter.

If D1–D4 survive their respective tests, the σ_D sector of the charged-lepton hierarchy is corroborated at the structural level. If D1's *form* survives but the parameter bands are wider or shifted, the framework survives with revised σ values. If either saturation claim (D2, D3) fails, the framework's central structural prediction is wrong and the §6.4 decomposition of the toy paper is wrong as currently formulated. If D4 fails — i.e., experimental search reveals a fourth charged-lepton generation — the closure-exhaustion step of §4.1 is wrong but the rest of the framework survives with an expanded n-band, and the counting law's σ values extend naturally to $n = 4$ with the same structural rules. If P1 is refuted (i.e., the substrate-thermodynamic computation produces $\xi + 2\eta$ away from 2.5), the counting law continues to hold but the natural-midpoint reproduction loses its status as evidence for the framework's quantitative reach.

14.1 Failure-mode table

If the first-principles enumeration or experiment finds…	Then…
σ_{phase} grows linearly with n	Claim D2 fails; the saturation hypothesis for phase coherence is wrong; the toy paper's σ_{record} component does not have the interpretation given here.

If the first-principles enumeration or experiment finds…	Then…
σ_refine grows linearly with n	Claim D3 fails; the saturation hypothesis for refinement infrastructure is wrong; σ values for higher generations grow faster than predicted.
ξ or η fall outside their ranges by more than ≈ 50 %	Claim D1's parameter bands are wrong; the counting law's structure may survive but its quantitative predictions revise.
σ_cycle saturates rather than growing linearly	The hierarchy compresses across all generations; framework fails to reproduce the lepton spectrum.
Smooth (non-discrete) interpolation between n = 1 and n = 2 in σ_phase or σ_refine	The 𝟙_{n≥2} indicator is wrong; the discrete-activation structure of the framework fails; framework may need a smooth turn-on form, requiring revision of §§6.2 and 7.
A genuine fourth charged-lepton generation is discovered experimentally	Claim D4 fails; closure exhaustion at n = 3 is wrong; the n-band expands to {1, 2, 3, 4} and the counting law applies to a fourth generation by direct extension. The σ-counting gives σ_4 = 3 + ξ + 2η ≈ 5.5 at the natural midpoint (an increment Δσ = 1 over σ_τ, parallel to σ_τ − σ_μ = 1). Under the full substrate-stiffness identity of the toy paper, the same inter-generational structure that gave m_τ/m_μ ≈ 17 from a unit σ-increment at n = 2 → 3 should give m_4/m_τ ≈ 17 from a unit σ-increment at n = 3 → 4, *if* the localization, transport, and Hessian contributions extend with the same per-generation scaling. The framework's prediction for a fourth-generation mass is therefore m_4 ≈ 17 · m_τ ≈ 30 GeV — which is already ruled out by LEP data, providing a backward-looking sanity check on the closure-exhaustion prediction. A genuine fourth lepton at this mass scale has been experimentally excluded; one at a substantially higher scale would require the framework's per-generation scaling of the other substrate observables to revise, indicating that the n-band extension is not simply a "same machinery, one more generation" extrapolation.
Additional commitment classes appear that we have missed	σ values shift upward; band predictions become lower bounds; framework may need a fourth commitment class.
The same counting law with same (ξ, η) does *not* reproduce the quark and neutrino sectors	The lepton-specific saturation pattern may not generalise; the framework's sector-extension claim fails.

These outcomes are the framework's exposed surface. The toy paper exposed σ_D as a single dominant input; the present paper decomposes that input into three classes governed by three structural rules and two bounded parameters, and identifies which substructure must hold for the framework to survive.

15. Forward Programme

The next stage of the programme follows from the structure exposed here.

1. **Substrate-thermodynamic tightening of (ξ, η)** — derive the phase-coherence infrastructure cost ξ and the refinement-infrastructure sharing efficiency η from substrate thermodynamics, tightening the structurally-bounded bands $\xi \in [1, 2]$ and $\eta \in [¼, ¾]$ to definite values. The task is now better-targeted than the previous draft framed it: $\xi_{max} = 2$ is structurally pinned by the count of primitive phase-coherence functions (§6.2.1) and $\xi_{min} = 1$ by the substrate's commitment-counting unit, so the ξ computation reduces to determining how strongly registration and transport stability share infrastructure within [1, 2]. The η computation is the less-secure of the two, requiring explicit construction of the refinement-support channel space and the overlap-reduction operator (§7.2). The target combination is $\xi + 2\eta \approx 2.5$ — the centre of the structurally-derived [1.5, 3.5] band, with (ξ, η) = (1.5, 0.5) the symmetric centre of the natural-midpoint locus; landing the combination at 2.5 would corroborate Prediction P1 regardless of how (ξ, η) split it.
2. **Closure-exhaustion derivation from corpus theorems** — derive the $n_{max} = 3$ closure-exhaustion step of §4.1 from a deeper substrate theorem. Candidate sources include the Catalogue Closure Theorem, the K = 7 closure-manifold dimensionality, and the no-go theorem for non-simplicial substrates. The §4.1 derivation establishes that $n_{max} = 3$ follows from primitive-function exhaustion (§3.1) applied to the charged-lepton loop class; what is owed is the demonstration that this exhaustion is forced by corpus-level closure theorems rather than asserted at the structural level. This is the analogue of pinning $\xi_{max} = 2$ by the primitive-function exhaustion ceiling.
3. **Operator-framework validation** — construct the closure-mode Gram matrix M and the refinement-persistence operator L for the n-loop charged-lepton PFDs, then compute $\mathrm{Tr}(M^{-1}L)$ for n = 1, 2, 3 and test whether it reproduces the counting law $\sigma(n) = (n - 1) + (\xi + 2\eta) \cdot \mathbb{1}_{\{n \geq 2\}}$ at leading order. This validates (or refutes) the operator conjecture of §11. If validated, the counting framework gains a microscopic operator-geometric foundation; if refuted, the counting law stands on its three structural rules but the operator reading is wrong.
4. **Microscopic admissibility simulation** — simulate the refinement-persistence and admissibility-preservation requirements for n-loop PFDs to test the independence-reduction step of §6.1 ($n - 1$ independent phases from $C(n, 2)$ naïve count) and the infrastructure-saturation step of §6.2.
5. **Closure-operator enumeration** — enumerate the substrate-global closure operators that supply refinement identity preservation and admissibility-compatible transport continuity, confirming R = 2 (or, if not, identifying the additional commitment classes).
6. **Quark-sector extension and the universality of primitive-function exhaustion** — apply the same counting framework to refinement-persistent quark PFDs. The quark sector adds confinement-derived commitment classes and is expected to produce substantially larger σ values; the present framework predicts the *shape* of the quark hierarchy (analogous saturation pattern between generations) but not yet its magnitude.

The strongest consilience test in the programme arises here. The primitive-function exhaustion principle (§3.1) predicts that the quark sector also closes at some finite generation count n_max^quark. If n_max^quark = 3 — the same exhaustion count as the charged-lepton sector — the framework converts an unexplained feature of the Standard Model into a derived consequence of substrate counting. The Standard Model's three-generation structure, identical across leptons and quarks, is at present a *coincidence* of three independent empirical facts (three charged leptons, three quark up-types, three quark down-types) with no theoretical motivation. If primitive-function exhaustion at k = 3 applies universally across these sectors, that coincidence is no longer a coincidence: it is a consequence of the same closure principle operating in each sector, with the universal exhaustion count k = 3 set by some deeper feature of the substrate's admissibility geometry. This would be the framework's strongest single empirical result — the derivation of a Standard Model feature that currently has no theoretical account.

If n_max^quark ≠ 3, the principle's exhaustion counts are sector-local rather than universal, and the framework loses the consilience claim but retains its sector-by-sector predictions. Both outcomes are scientifically substantial: the universal-consilience outcome is the strongest possible result; the sector-local outcome localises a real structural feature of the framework worth understanding on its own.

7. **Neutrino-sector extension** — apply the framework to neutrino PFDs, which carry admissibility-mode commitments rather than full closure commitments. The framework predicts σ_ν substantially suppressed below σ_lepton because admissibility-mode commitments are infrastructurally lighter than closure commitments. The neutrino sector also tests the universality of primitive-function exhaustion: n_max^neutrino = 3 (matching both leptons and quarks) would extend the cross-sector consilience to all three sectors, completing the derivation of Standard Model three-generation structure as a universal closure-principle consequence rather than a triple coincidence.
8. **Cross-sector consistency check** — verify that (ξ, η) take the same substrate-thermodynamic values across leptons, quarks, and neutrinos. Sector-specific (ξ, η) would indicate that the framework's saturation pattern is sector-local rather than universal. Combined with the n_max universality test of Tasks 6 and 7, this sets up a four-dimensional consilience surface: same Rules 1–3, same closure principle, same (ξ, η), same n_max across all three sectors. Survival of all four corroborates the framework strongly; failure of any one localises the failure mode.

These tasks are listed in priority order. Tasks 1 and 3 carry the largest contingent weight, but in different ways.

Task 1 (substrate-thermodynamic derivation of (ξ, η)) would convert the framework's quantitative reach from a parameter-bounded prediction to a parameter-free derivation. If Task 1 lands $\xi + 2\eta \approx 2.5$, Prediction P1 is corroborated and the framework's quantitative reach extends to the toy values without further tuning. If Task 1 lands elsewhere in the [1.5, 3.5] band, the counting law continues to hold but the toy values are no longer reproduced; the framework's structural content (D1–D3) is unaffected.

Task 3 (operator-framework validation) carries no quantitative weight in the present paper — the counting law's σ values are derived from the three structural rules of §3, not from $\mathrm{Tr}(M^{-1}L)$. What Task 3 carries is *interpretive* weight: it determines whether the substrate-physical reading of §10 is implemented by closure-mode overlap geometry (§11), or whether it is implemented by some other substrate mechanism. If Task 3 refutes the $\mathrm{Tr}(M^{-1}L)$ conjecture, the counting law's three rules continue to derive the same σ values, and the framework's predictions D1–D3 and P1 are unaffected; what is retracted is §11 and the operator-reading sentences in the abstract, §6.2, and the conclusion. That is a meaningful loss of interpretive depth but not a loss of the framework's core structural or quantitative content. Conversely, if Task 3 validates the conjecture, the framework gains a microscopic operator-geometric foundation for the substrate-physical infrastructure-sharing principle.

Tasks 2, 4, and 5 carry no contingent-weight asymmetry: they tighten or refute the structural inputs (closure-exhaustion at n = 3, independence reductions, R = 2) that the counting law depends on. Failure of any one falsifies the corresponding piece of the framework directly. Note that the n-loop assignments themselves are no longer a forward task; they are derived in §4.1 of the present paper, with Claim D4 of §14 providing the experimentally-testable falsifiability hook (discovery of a fourth charged-lepton generation).

---

16. Conclusion

This paper performed two structurally distinct derivations. The **foundational result** (§4.1) was a derivation of the charged-lepton n-band {1, 2, 3} from minimal persistence, generation distinction by loop topology, and closure exhaustion at n = 3 — converting the previous toy paper's three-generation input into a structural consequence of refinement-persistent admissibility geometry, and producing the framework's first directly experimentally-testable prediction (Claim D4: discovery of a fourth charged-lepton generation would falsify the closure-exhaustion step). The **applied result** (§§5–9) was the σ-counting law itself: using three structural rules — refinement-persistence, independence, and infrastructure-sharing — applied to three commitment classes — cycle persistence, phase coherence, and refinement infrastructure — the paper derived a minimal quotient-counting law for σ(n) on the closed three-generation sector.

The σ-counting law, operating on the closed three-generation sector derived in §4.1:

- produces σ values whose band $[2.5, 4.5] \times [3.5, 5.5]$ contains the toy values, with exact reproduction at the band midpoint and order-of-magnitude consistency across the band;
- forces the inter-generational gap $\sigma_\tau - \sigma_\mu = 1$ independently of parameter values;
- reproduces the toy paper's $\sigma_\mu = 3.5$, $\sigma_\tau = 4.5$ exactly at the natural midpoint of the framework's two structurally-bounded parameters;
- forces the non-uniform hierarchy shape $\sigma_\tau - \sigma_\mu < \sigma_\mu - \sigma_e$ as a structural consequence of phase-coherence and refinement-infrastructure saturation between μ and τ;
- maps cleanly onto the toy paper's §6.4 three-component decomposition, with the toy paper's six free numbers reducing to two bounded parameters plus a derived integer band.

The central conceptual achievement is not the numerical match. It is the **conversion of σ_D from an opaque toy parameter to a constrained substrate-counting problem**, with the structural inputs of that problem — the n-band, the parameter ranges, the saturation pattern — now derived from substrate-physical principles rather than asserted. Where the toy paper carried σ_μ and σ_τ as freely-chosen values supplying ≈ 60 % of the inter-generational log-ratio at both scales (with the n-loop assignments carried as a separate forward-task input), the present framework derives σ_μ and σ_τ from a counting law governed by three structural rules, on a sector with a *derived* three-generation structure.

The σ_D sector of the charged-lepton hierarchy is no longer described by:

"two free numbers, chosen consistent with the target, applied to a sector whose three-generation structure is taken as input."

It is described by:

"three commitment classes, three structural rules, two structurally-bounded parameters, and a derived three-generation closure — with the inter-generational gap forced by the counting and the intra-generational baseline reproduced at the natural midpoint of the parameter bands."

This is the structural form that a substrate-level theory of fermion mass takes when its dominant input is derived rather than inserted. The framework now exposes four testable claims (D1–D4), with D4 in particular testable by ongoing experimental particle-physics searches for fourth-generation charged leptons. D1 carries the framework's structural content; D2 and D3 carry the saturation predictions; D4 carries the closure-exhaustion prediction; Prediction P1 carries the framework's quantitative reach as a conditional forecast pending the substrate-thermodynamic computation of (ξ, η).

The framework also connects naturally to the projected-operator interpretation of the closure manifold. The K = 7 closure basis is strongly non-orthogonal, and the counting law is most plausibly read as counting *effective independent* commitments after overlap reduction rather than as raw combinatorial counting. Whether the counting law is in fact the leading-order approximation to a schematic operator structure $\sigma_{eff} \sim \mathrm{Tr}(M^{-1}L)$ is now a concrete computational question rather than a programmatic one, and is listed as task 3 of the forward programme. The connection is offered here as a hypothesis to be tested, not as a foundation; the counting law stands on its three structural rules regardless of how that test resolves.

The remaining tasks are now sharply defined. Tighten (ξ, η) from substrate thermodynamics within their structurally-derived bands; derive closure-exhaustion at n = 3 from corpus-level closure theorems; test the operator-framework conjecture by computing $\mathrm{Tr}(M^{-1}L)$ for the n-loop PFDs; and extend the counting framework to the quark and neutrino sectors. If those derivations corroborate the structural form proposed here, the charged-lepton hierarchy is — structurally — a hierarchy of refinement-persistent informational commitments on a sector with a derived three-generation structure. If they fail, the failure mode is now localised to specific substrate-thermodynamic quantities, identifiable saturation patterns, an explicit operator-geometric

conjecture, and an empirically-testable three-generation prediction rather than to a free σ_D parameter.

The hierarchy problem is, in either case, no longer protected by vague emergence language. It now has a derived structural law, a derived n-band, an explicit set of falsifiable substrate-counting predictions, and an empirically-testable three-generation closure — with a candidate operator-geometric foundation made sharp enough to be tested.