Orbit Counts and the Refinement Structure of Completion Density

Return Counts, Commitment Counts, the Completion-Density Order, and a Reuse-Monotonicity Theorem for Finite Closure Dynamics

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The hierarchy programme consumes three quantities — a return length, a commitment count, and their ratio, the completion density — and asks one structural question of them: why do the deeper refinement levels carry higher density? In the programme's prior structure these quantities are named but not built, and the structural question has no machinery in which to be posed.

This paper does two things, in two parts of different strength.

Part I builds the quantities. For any finite closure-state machine of the right general type — a deterministic one-tick update on finitely many closure states, with a ledger that records commitment and never reverses it — the return count, the commitment count, and the completion density are well-defined, parameter-free, mass-blind properties of each self-sustaining cycle. This is proven for the whole class of such machines, not one chosen machine, and it carries no assumption about how many levels exist or how they are ordered. It banks the machinery.

Part II asks the structural question. It introduces a *refinement operation* — a relation between orbit classes, one depth to the next — and asks when refining raises the completion density. The answer is a sharp condition: refinement raises density exactly when it adds commitments **faster than** it lengthens the return — when each added commitment costs less return-length than the orbit's current average. This is the *reuse* condition: a refinement that threads existing structure (reusing channels) buys commitment cheaply in return-length and raises density; one that merely extends the cycle buys it dearly and lowers density. This first-order condition governs whether density *ascends* with depth. A second, independent condition governs whether the *rate* of increase compresses or expands — and the two come apart in the real world: the down-type quarks ascend (density rising with depth) but expand (the jumps widen rather than shrink). The paper keeps these two statements rigorously separate, and neither implies the other. The ordering theorem is clean and is not stressed by the down sector; the compression condition is the

falsifiable one, built to fail exactly where the measured table expands — which is the honest place to put the empirical stress.

What the paper does **not** do: it does not exhibit the physical machine, does not select the count of three, and does not claim refinement raises density unconditionally. It states the structural condition under which it does, and freezes that condition before any physical operator is written — so the construction paper that follows can be judged against a standard it did not get to author.

Table of Contents

0. Scope, Strength, and the Two Frozen Surfaces

This paper has two parts at two grades.

Part I (§§2–5) discharges obligations **D1–D4** of the operator specification — the existence and well-definedness of K_c, p_v, and C as orbit counts — at grade **[Proven]**, for the entire class of finite closure-state machines.

Part II (§§6–8) adds refinement structure: the **completion-density order** on orbit classes [Proven, unconditional] and the **Reuse Monotonicity Theorem** [Proven under a stated, frozen hypothesis]. Part II does not assume, and does not establish, that refinement raises density in general — that statement is false for the class and would be the density clause (α′) inserted rather than derived. It establishes the structural condition under which the increase holds, and names that condition in advance.

Neither part exhibits a machine, selects the depth census, or produces a physical number. The depth census **D5** remains [Open], deferred to the construction paper and governed by the Selection Audit frozen in the operator specification.

The frozen hypotheses (Part II). Stated here, at the front, before any theorem invokes them, so their later use is audit rather than choice. There are two, at two orders, and the paper keeps them strictly distinct:

First order — hierarchy-positivity (governs ordering). A refinement operation $\mathscr{R}$ is *hierarchy-positive* at an orbit class exactly when the marginal density of the added structure exceeds the orbit's current completion density:

Δp_v / ΔK_c > p_v / K_c.

By the mediant identity (Lemma 7.2) this holds iff C strictly increases under the refinement, $C(\mathscr{R}\mathcal{O}) > C(\mathcal{O})$. It is the condition on whether density *ascends* with depth — the density clause (α′).

Second order — compression-positivity (governs the profile). A refinement chain is *compression-positive* when its successive density ratios contract — the density log-gaps shrink, $G_2 < G_1$ (the symbol G is reserved throughout for density-ratio log-gaps; measured mass log-gaps are written Γ, and the two are kept distinct — §8). This is a condition on the realized density sequence directly, not on the marginals: a chain can ascend at every step yet have its ratios widen. Compression and ascent are **independent** — neither implies the other.

These conditions can each hold or fail, and they fail independently. A sector may be hierarchy-positive (densities ascending, (α′) holding) yet compression-negative (ratios widening with depth) — and the measured fermion table contains exactly such a sector: the down-type quarks ascend but expand (§8). The falsifiable edge of Part II therefore attaches to *compression-positivity*, the second-order condition, not to the ordering theorem: Theorem 7.3 keeps a clean iff for ascent, and the down sector stresses the compression condition, which is where the world already shows the condition is not universal. A theorem that could not fail would test nothing; the compression condition is built to fail where the down sector expands, and Theorem 7.3 is honest precisely by *not* being billed as the thing under that stress.

1. Inherited Structure and Markers

Layer A. [Inherited] Tick-prior substrate; $\tau \in \mathbb{N}$ the sole ordering primitive; all quantities here are dimensionless Layer-A objects defined before any calibration to time.

Closure states. [Inherited] The objects acted on carry the closure invariants (k, w, n). The proofs treat the state set abstractly as a finite set $\boldsymbol{S}$; the invariant labelling enters only in the remarks connecting to the programme.

Commitment ledger. [Inherited] Commitment advances an irreversible ledger by a non-negative quantum per tick; committed facts are not reversed (ANCH-IRR). This is the structural reason the ledger is a separate monotone accumulator, not a state coordinate.

Refinement. [Inherited, as a relation to be made operational] The programme speaks of refinement depth n and of holding a mark to pass from depth n to n+1. Part II makes this an explicit operation between orbit classes; what the *physical* refinement operation is remains the construction paper's burden.

Markers. [Proven] established here; [Proven under hypothesis] established contingent on a stated assumption; [Hypothesis] a carried assumption; [Inherited] imported; [Inherited-empirical] imported from measured data at the protocol's scoring layer, not derived here; [Open] deferred. *These descriptive labels map to the corpus's canonical scheme as follows: [Proven] = Proven; [Proven under hypothesis] = Conditional (on the named hypothesis); [Hypothesis] = the antecedent of such a Conditional; [Inherited] and [Inherited-empirical] = Inherited (the latter flagging an empirical rather than theorem-grade import); [Open] = Open. The descriptive forms are kept because Part II's central results are Conditional in a specific, named way it is clearer to label at point of use; the deviation from the bare five-label scheme is deliberate and carries this cost.*

PART I — THE ORBIT COUNTS

2. The Finite Closure-State Machine

Definition 2.1 (Closure-state machine). A *finite closure-state machine* is a triple $\boldsymbol{\mathcal{M}} = (\boldsymbol{S}, \mathfrak{A}, \ell)$ where $\boldsymbol{S}$ is a finite non-empty set of closure states; $\mathfrak{A} : \boldsymbol{S} \to \boldsymbol{S}$ is the **anchoring update**, total and deterministic, one application being one tick; and $\ell : \boldsymbol{S} \times \boldsymbol{S} \to \mathbb{N}$ is the **commitment-increment rule**, a non-negative integer on each ordered transition $(s, \mathfrak{A}(s))$. The ledger is the monotone accumulator of ℓ along a trajectory — not a coordinate of $\boldsymbol{S}$.

Remark 2.2 (Why the ledger is separate). Were the ledger a coordinate, a closed structural orbit would have to return it to its start value — to un-commit over one return. ANCH-IRR forbids this. Modelling ℓ as a monotone accumulator encodes irreversibility: the closure state is periodic, the ledger is monotone, and they need not share a period.

Definition 2.3 (Trajectory). From s ∈ 𝑺: the sequence s, $\mathfrak{A}(s)$, $\mathfrak{A}^2(s)$, … with $\mathfrak{A}^0$ = identity.

Definition 2.4 (Recurrent state; admissible mode). A state s is **recurrent** iff $\mathfrak{A}^k(s) = s$ for some $k > 0$. An **admissible mode** is a recurrent state, identified with the cycle it generates. Non-recurrent states are **transient**. *Throughout this paper "admissible" carries only this bare dynamical sense — lying on a cycle of 𝔄 — and never the programme's physics-selection sense (which mode is a physically realized particle). The latter is the depth census D5, governed by closure balance, and is deferred entirely; nothing here turns on it.*

Remark 2.5 (Eventual capture). Every trajectory in a finite deterministic system is eventually periodic: a transient tail of length $\mu(s) \geq 0$ feeds a unique cycle of length $\lambda(s) \geq 1$. A state is recurrent iff $\mu(s) = 0$. The theorems concern recurrent states; which states are recurrent — the physics of admissibility — is the construction paper's subject.

---

3. The Orbit Count Theorem

Notation. $\mathfrak{A}^t$ denotes t-fold composition. Where the exponent is a bare lowercase variable it is set as a Unicode superscript ($\mathfrak{A}^t$, $\mathfrak{A}^i$, $\mathfrak{A}^j$); where it is a compound or capital expression it is set with a caret (𝔄^K, 𝔄^(t+1), 𝔄^(K−1)). The two forms denote the same operation.

Theorem 3.1 (Orbit Count Theorem). Let 𝓜 = (𝑺, 𝔄, ℓ) be a finite closure-state machine. Every admissible mode 𝒪 has:

1. a finite return count K_c(𝒪) ∈ $\mathbb{Z}_+$, well-defined as a start-independent cycle invariant;
2. a finite commitment count p_v(𝒪) ∈ ℕ, the ledger increment over one return, equal at every starting point;
3. a completion density C(𝒪) = p_v(𝒪)/K_c(𝒪), a dimensionless rational cycle invariant with no free parameter and no mass-derived input;

and, under the non-degeneracy hypothesis H+ (Def. 3.6), p_v(𝒪) > 0 and C(𝒪) > 0.

Lemma 3.2 (Recurrence — the generic step). If s is recurrent, $\{ k > 0 : \mathfrak{A}^k(s) = s \}$ is non-empty with a least element.

Proof. Non-empty by Def. 2.4; least element by well-ordering of $\mathbb{Z}_+$. ■

This uses only determinism on a finite set and would hold for any such map. It gives existence of K_c, not its value, and not invariance under choice of starting state — those are Lemmas 3.4–3.5.

Definition 3.3 (Return count; orbit). For recurrent s, $K_c(s) = \min\{ k > 0 : \mathfrak{A}^k(s) = s \}$, and $\mathcal{O}(s) = \{ s, \mathfrak{A}(s), \ldots, \mathfrak{A}^{K_c(s)-1}(s) \}$.

Lemma 3.4 (Start-independence of the return count). If $s' \in \mathcal{O}(s)$, then s' is recurrent, $\mathcal{O}(s') = \mathcal{O}(s)$, and $K_c(s') = K_c(s)$.

Proof. Let $K = K_c(s)$ and $s' = \mathfrak{A}^i(s)$, $0 \le i < K$. Then $\mathfrak{A}^K(s') = \mathfrak{A}^i(\mathfrak{A}^K(s)) = \mathfrak{A}^i(s) = s'$ (powers commute), so s' is recurrent with $K_c(s') \le K$. For the reverse: the iterates $s, \mathfrak{A}s, \ldots, \mathfrak{A}^{K-1}s$ are pairwise distinct, for if $\mathfrak{A}^i s = \mathfrak{A}^j s$ with $i < j < K$, then $\mathfrak{A}^{K-(j-i)}(s) = s$ with $0 < K-(j-i) < K$, contradicting the minimality of K. Hence $\mathfrak{A}$ permutes this K-element orbit cyclically and is injective on it, so $\mathfrak{A}^i$ is left-cancellable there. Now if $K_c(s') = K' < K$, then $\mathfrak{A}^{i+K'}(s) = \mathfrak{A}^i(s)$; cancelling $\mathfrak{A}^i$ gives $\mathfrak{A}^{K'}(s) = s$ with $0 < K' < K$, the same contradiction. So $K_c(s') = K$ and the cyclic sets coincide. ■

First non-generic content: "a return length from here" becomes "the mode's return length." Determinism is used essentially — the in-orbit bijectivity. Without it, K_c and p_v would be path-dependent, not mode-invariant.

Lemma 3.5 (Well-definedness of the commitment count). For recurrent s with $K = K_c(s)$, let $p_v(s) = \Sigma_{t=0}^{K-1} \ell(\mathfrak{A}^t(s), \mathfrak{A}^{t+1}(s))$. Then $p_v(s') = p_v(s)$ for all $s' \in \mathcal{O}(s)$.

Proof. The orbit is the cyclic edge sequence $e_0, \ldots, e_{K-1}$ with $e_t = (\mathfrak{A}^t s, \mathfrak{A}^{t+1} s)$, and $p_v(s)$ sums ℓ over all K edges. A different start $s' = \mathfrak{A}^i(s)$ (same orbit, same K, Lemma 3.4) re-indexes the same K edges by a cyclic rotation; a finite sum is invariant under rotation of its terms. Hence $p_v(s') = p_v(s)$. ■

Second non-generic content, and the reason the separate-accumulator model is clean: ℓ is an edge-function on a fixed cyclic edge set, so p_v is a sum over the cycle's edges — start-independent — and because the ledger never reverses, p_v is a sum of non-negative increments, not a signed quantity that could cancel.

Definition 3.6 (Non-degeneracy hypothesis H+). An admissible mode is **non-degenerate** iff $\ell > 0$ on at least one of its cycle edges. A mode with $\ell \equiv 0$ around its cycle is **degenerate** — structurally periodic, committing nothing (a maintenance-silent pattern).

H+ is a hypothesis, not a derivation. Whether self-sustenance forces commitment is closure physics, deferred. The degenerate zero-commitment orbit is named, not hidden — it is the limiting object the Distinction Cost lean points at [Inherited; the lean is to be cited to its source paper at submission], and whether the physical dynamics admits it is [Open].

Lemma 3.7 (Positivity and the density). Under H+, p_v(𝒪) ≥ 1 > 0, and C(𝒪) = p_v/K_c is a strictly positive, dimensionless, rational cycle invariant determined entirely by (𝔄, ℓ) and the orbit, with no free parameter.

Proof. H+ gives one positive summand, all non-negative, so p_v ≥ 1; K_c ≥ 1; both invariant (3.4, 3.5); the ratio is a dimensionless rational invariant; no quantity beyond 𝔄, ℓ, the orbit — in particular no mass, no constant — enters. ■

Proof of Theorem 3.1. (1)–(3) from Lemmas 3.2–3.5 with the ratio of Lemma 3.7 (well-defined and parameter-free with or without H+); (4) from Lemma 3.7 under H+. ■

---

4. What Part I Establishes, and Does Not

Banked. K_c and p_v are orbit counts, not parameters — computed by following a cycle and counting, invariant under indexing (3.4, 3.5). C carries no knob (3.7). No mass enters. The result holds for the whole class, so the physical machine's counts inherit well-definedness without separate argument; the open question of *which* machine is physical cannot threaten well-definedness, only values.

Not established. No machine is exhibited; no number produced; no census selected; positivity not derived (H+ is a hypothesis); and nothing is ordered yet — Part I makes the ordering question *expressible*, Part II asks it.

---

5. Worked Illustration (Non-Physical, Machinery Only)

S = {a, b, c, d}; 𝔄: a→b→c→a (a 3-cycle), d→a (transient). ℓ = 1 on (a,b) and (c,a); ℓ = 0 on (b,c) and (d,a).

Recurrent: a, b, c (return in 3); transient: d (μ=1). Mode 𝒪 = {a,b,c}: K_c = 3. Ledger over one return from a: 1+0+1 = 2 = p_v; from b: 0+1+1 = 2 (same, by Lemma 3.5). H+ holds, so C = 2/3. The transient d is not a mode; its counts belong to the cycle that captures it.

This computes the counts the only permitted way — follow the cycle, sum the edge-rule — and exhibits start-independence. It selects nothing and means nothing physically.

PART II — THE REFINEMENT STRUCTURE

6. Refinement Operations and the Completion-Density Order

Part I gives, for one machine, a finite set of admissible modes, each with a density. Part II adds the structure the generation problem actually concerns: a relation *between* modes, taking one depth to the next.

Definition 6.1 (Realized density set; the density order). For a machine $\mathcal{M}$, let $\mathcal{C}(\mathcal{M}) = \{ C(\mathcal{O}) : \mathcal{O}$ an admissible non-degenerate mode $\} \subseteq \mathbb{Q}_+$ be the set of realized completion densities. As a finite subset of $\mathbb{Q}$, $\mathcal{C}(\mathcal{M})$ is totally ordered by <. Ordering classes by their density value yields a **total preorder** on classes — every pair is comparable — which fails antisymmetry, and hence fails to be a partial order, exactly when distinct classes share a density.

Proposition 6.2 (The density order is free and parameter-free). $\mathcal{C}(\mathcal{M})$ is a finite totally ordered set, determined entirely by $(\mathfrak{A}, \ell)$. [Proven — immediate from Theorem 3.1: each C is a well-defined rational invariant, and a finite set of rationals is totally ordered.]

*This is honest structure but thin, and its limits must be stated. It orders classes **by their density**; it therefore cannot, by itself, **explain** an ordering, since ordering-by-C is circular if C is the thing whose ordering is to be derived. The content of Part II is not this order but the* operation *that moves along it, and the condition under which that operation raises C — which is §7.*

Definition 6.3 (Refinement operation). A **refinement operation** is a partial map $\mathscr{R}$ on orbit classes, taking a class $\mathcal{O}$ at depth n to a class $\mathscr{R}(\mathcal{O})$ at depth n+1, with finite increments

$$\Delta K_c = K_c(\mathscr{R}\mathcal{O}) - K_c(\mathcal{O}), \quad \Delta p_v = p_v(\mathscr{R}\mathcal{O}) - p_v(\mathcal{O}).$$

$\mathscr{R}$ is **expansive** if $\Delta K_c > 0$ (the refined orbit is longer) and **commitment-adding** if $\Delta p_v > 0$. The physical refinement — holding a mark to pass from depth n to n+1 — is one such operation; its explicit form is the construction paper's burden. Part II proves what *any* $\mathscr{R}$ does to density under a stated condition, exhibiting no particular $\mathscr{R}$.

Definition 6.3 carries no claim that refinement raises density. It defines the increments and leaves their sign to the operation. This is the fence: $\mathscr{R}$ is not defined so as to raise C; the condition under which it does is stated separately, in §7, and can fail.

---

7. The Reuse Monotonicity Theorem

Definition 7.1 (Hierarchy-positive refinement — the frozen hypothesis). A refinement operation $\mathscr{R}$ is **hierarchy-positive** at a class $\mathcal{O}$ (with K_c, p_v, and increments $\Delta K\_c > 0$, $\Delta p\_v \geq 0$ as in Def. 6.3) iff

$$\Delta p\_v / \Delta K\_c > p\_v / K\_c.$$

Equivalently, the marginal density of the added structure exceeds the orbit's current density. We call this the **reuse condition**: refinement that threads existing structure adds commitment at sublinear return-length cost, so its marginal density beats the running average.

This is the §0-frozen first-order hypothesis of Part II. It is stated before Theorem 7.3 invokes it, it is a property an operation has or lacks, and it governs whether density ascends — the density clause (α'). It is not the condition the down sector falsifies: a sector may satisfy Def. 7.1 at every step (densities ascending) and still widen its gaps. That second-order behaviour is Def. 7.1b.

Definition 7.1b (Compression-positive refinement — the second-order condition). A chain of expansive refinements $\mathcal{O} \rightarrow_{\mathscr{R}} \mathcal{O}' \rightarrow_{\mathscr{R}} \mathcal{O}''$ is **compression-positive** at the second step when the successive density ratios contract:

$$C(\mathcal{O}'')/C(\mathcal{O}') < C(\mathcal{O}')/C(\mathcal{O}),$$

equivalently when the density-ratio log-gaps shrink, $G_2 < G_1$, where $G_i \coloneqq \ln[\, C(\text{of the i-th refined class}) / C(\text{of its predecessor})\,]$ is the log of the i-th **density** ratio — an abstract orbit-count quantity, defined purely on the C-sequence and carrying no mass. Def. 7.1b is stated **directly on the realized C-sequence** and carries no marginal-density characterization: unlike ascent, compression is *not* controlled by the marginals. (The second difference of a mediant sequence is not fixed by the monotonicity of the marginal densities — a chain can have rising marginals at both steps while its density ratios widen, so "marginals rise" and "ratios compress" are inequivalent.) Compression and ascent are **independent**: a chain may ascend at every step (Def. 7.1 holding throughout) while its ratios widen (Def. 7.1b failing), and neither condition implies the other. Theorem 7.3 governs ascent; the down-type sector (§8) is the empirical witness to the independence — ascent holding while compression fails.

Lemma 7.2 (The mediant inequality). For positive integers (or positive reals) p, K and non-negative increments Δp, ΔK with $\Delta K > 0$,

$$(p + \Delta p)/(K + \Delta K) > p/K \iff \Delta p/\Delta K > p/K.$$

Proof. $(p+\Delta p)/(K+\Delta K) > p/K \iff K(p+\Delta p) > p(K+\Delta K)$ [cross-multiplying, denominators positive] $\iff K\Delta p > p\Delta K \iff \Delta p/\Delta K > p/K$ [dividing by $K\Delta K > 0$]. All steps are equivalences since $K, \Delta K > 0$. ■

The refined density is the ***mediant*** *of the old density p/K and the marginal density Δp/ΔK. A mediant lies strictly between its two constituents (when they differ); so the refined density exceeds the old one iff the marginal density does. This is the arithmetic content: the reuse condition is exactly the statement that the marginal beats the average, and the mediant identity makes "raises the average" and "marginal exceeds average" the same inequality.*

Theorem 7.3 (Reuse Monotonicity). Let ℛ be a refinement operation, expansive at a non-degenerate class 𝒪 (ΔK_c > 0). Then

C(ℛ𝒪) > C(𝒪) if and only if ℛ is hierarchy-positive at 𝒪 (Def. 7.1).

Consequently, along any chain 𝒪 →ℛ 𝒪' →ℛ 𝒪'' of refinements each of which is hierarchy-positive at the class it acts on, the completion density is strictly increasing:

C(𝒪) < C(𝒪') < C(𝒪'').

Proof. C(ℛ𝒪) = (p_v + Δp_v)/(K_c + ΔK_c) and C(𝒪) = p_v/K_c; by Lemma 7.2 (with ΔK_c > 0), C(ℛ𝒪) > C(𝒪) ⇔ Δp_v/ΔK_c > p_v/K_c, which is Def. 7.1. The chain statement is iterated application, the strict increase at each step holding under hierarchy-positivity at that step. ■

Corollary 7.4 (The non-expansive cases, for completeness). Throughout, the class is non-degenerate (p_v > 0, H+). If ΔK_c = 0 and Δp_v > 0 (commitment added at no return-length cost), C strictly increases. If ΔK_c < 0 (return compressed) with Δp_v ≥ 0, C strictly increases — using p_v > 0 in the Δp_v = 0 sub-case, where the same numerator over a smaller denominator raises the ratio. In both, "increases" holds *without invoking Def. 7.1*, still within non-degeneracy. The reuse condition (Def. 7.1) is the criterion precisely in the substantive case where **both** counts grow (ΔK_c > 0, Δp_v > 0) — where it is not automatic that refinement helps, and the question is whether the added commitment outpaces the added length. [Proven]

Corollary 7.4 places the theorem: the two "easy" directions raise density for free, and the reuse condition is the real criterion exactly where the route's situation lives — both p_v and K_c growing, the increase contingent on the ratio. This is why the hypothesis is neither thin nor circular: it bites in the one regime that matters.

---

8. The Failure Case, Owned — Where Compression Fails

A theorem that cannot fail tests nothing. The first-order condition (Def. 7.1) governs ascent and is not the thing under empirical stress; the second-order condition (Def. 7.1b) governs compression and *can* fail, and the measured fermion table shows a sector where it does.

The down-type column ascends but expands — on the measured table. The measured quantities are the natural-log gaps of the *measured* down-type quark masses: $\Gamma_1 = \ln(m\_s/m\_d) \approx 2.99$ and $\Gamma_2 = \ln(m\_b/m\_s) \approx 3.80$, the second step the larger. The symbol Γ is reserved for these **mass** log-gaps and is kept distinct from the **density** log-gaps G of Def. 7.1b; the two are not the same quantity, and the paper does not let one stand for the other.

Both the column's ascent and its expansion are readings of the measured table, and **both are [Inherited-empirical]** — neither is derived here, since Part I's machinery is mass-blind, produces no densities, and identifies no sector. Translating these mass readings into statements about the abstract density sequence requires the bridge ANCH-COMP, and the two translations stand at different strengths, which the paper marks rather than blurs:

- **Ascent transfers under the order bridge.** The masses ascend ($m\_d < m\_s < m\_b$); under ANCH-COMP's order-preserving relation between mass and density, the densities ascend too, so Def. 7.1 holds at every step of the down column. This uses only that the bridge is monotone — the form of ANCH-COMP invoked everywhere in the corpus.
- **Expansion transfers only under a stronger, named bridge.** The mass log-gaps widen ($\Gamma_2 > \Gamma_1$). For this to imply that the *density* log-gaps widen ($G_2 > G_1$, the failure of Def. 7.1b) requires that mass-ratio spacing tracks density-ratio spacing — a second-order, log-linear form of ANCH-COMP, strictly stronger than order-preservation. **This is carried as [Hypothesis — the log-linear bridge], not asserted:** the second-order transfer Γ-spacing → G-spacing is exactly the bridge the construction paper and the Eigenmode Decision's mass-aware layer must establish, and it is named here so the down-column conclusion is visibly conditional on it rather than smuggled through a shared letter.

On the inherited reading, then — ascent under the order bridge, expansion under the named log-linear bridge — the down column would be a clean separation of the two conditions: Def. 7.1 holding (ascent) while Def. 7.1b fails (G-gaps widening). The down sector enters here only as this inherited-empirical constraint the future $\mathscr{R}$ must reproduce, conditional on the bridge made explicit above; it is not a result of this paper.

The Γ-quantities are empirical mass log-gaps and live at the Eigenmode Decision's mass-aware scoring layer, on the far side of the mass firewall. This paper does not equate Γ with G, and does not assert the second-order bridge between them — it states it as the [Hypothesis] under which the density-spacing reading follows. The order bridge (ascent) it does use, since that is the corpus-standard monotone form of ANCH-COMP.

This is the theorem working, not failing. On the inherited reading the down densities ascend, so Def. 7.1 holds there and Theorem 7.3 is *not* contradicted — it is left intact, conditional on that reading, exactly as it should be. The condition the down sector falsifies is compression-positivity (Def. 7.1b), and falsifying it is the point. Two consequences are bound here, in advance:

- **Ordering and compression are independent, and the down sector is the empirical witness.** A column can be hierarchy-positive at every step (densities ascending, (α') holding) while compression-negative (density ratios widening). Theorem 7.3 concerns ordering through the mediant; the reuse-*compression* profile — the secondary readout the

Eigenmode Decision logs at firewall discipline (a readout it records but does not let feed any verdict) — concerns Def. 7.1b. On the inherited reading, and under the named log-linear bridge for the second-order step, the down sector withholds compression consonance while satisfying (α') — the mixed-per-column behaviour the Eigenmode Decision's profile readout anticipates.

- **The physical refinement operation $\mathscr{R}$ must be compression-positive where columns compress and compression-negative in the down sector. This is a constraint on $\mathscr{R}$, frozen here:** a construction whose $\mathscr{R}$ forced uniform compression would contradict the measured table and is thereby excluded in advance. The down-sector expansion is compression-positivity's falsifiable edge — the place the world already shows the second-order condition is not universal.

Remark 8.1 (What this hands the construction paper). Theorem 7.3 converts the density clause (α') from an assertion into a first-order structural condition: *(α') holds along a refinement chain iff each step's marginal density exceeds the running average* (Def. 7.1). Separately, the reuse-compression profile is a second-order condition (Def. 7.1b) the construction must reproduce as *mixed*: compression-positive in the ascending-and-compressing columns, compression-negative in the down sector. The construction paper, having exhibited $\mathscr{R}$, must therefore show two things, sector by sector — that $\mathscr{R}$'s increments (Δp_v, ΔK_c) satisfy Def. 7.1 wherever the table ascends (which is everywhere, all three charged columns ascending), and that they reproduce the *compression pattern* of Def. 7.1b, including the down-sector expansion. The theorem proves neither for the physical $\mathscr{R}$; it states exactly what proving them on computed counts would require, and freezes the down sector as the test of honesty — the one place a construction that merely forced compression everywhere would be caught.

One asymmetry, noted so it is not mistaken for a debt. The [Hypothesis: log-linear bridge] of §8 is this paper's import device, not the construction's obligation. This paper has no densities, so to read the down-sector compression off the world at all it must borrow the measured mass gaps Γ and assume they transfer to the density gaps G — hence the bridge. The construction paper computes the G-sequence directly from its $\mathscr{R}$ and tests Def. 7.1b on those computed density gaps natively; it never routes through Γ, so the log-linear bridge does not appear in its obligations. The hypothesis is load-bearing here and nowhere downstream.

---

9. What Part II Establishes, and Does Not

Established. The realized densities form a finite ordered set [Prop. 6.2, Proven]. Refinement is an operation on classes with definable increments [Def. 6.3]. Under the first-order reuse condition (Def. 7.1), refinement raises density, by the mediant identity, exactly in the substantive both-counts-growing regime [Thm. 7.3 + Cor. 7.4, Proven under hypothesis]. The two conditions are distinguished and shown independent: ordering (Def. 7.1) and compression (Def. 7.1b) can diverge, and the down-type sector is a measured case where ordering holds and compression fails [§8, the down-sector reading itself Inherited-empirical].

Not established. No physical $\mathscr{R}$ is exhibited; the reuse condition is not shown to hold for any actual refinement (that is the construction paper's, sector by sector); (α') is not proven, only reduced to the reuse condition on computed counts; and nothing about the *count* of admissible depths is touched — D5 remains [Open].

---

10. Connection to the Programme's Frozen Instruments

To the operator specification. Part I discharges the K- and P-clauses in their well-definedness part; the value clauses (exhibit the orbit, do not assert the count) remain owed by the construction. Part II adds two new frozen surfaces — the first-order reuse condition (Def. 7.1, a structural form of (α')) and the second-order compression condition (Def. 7.1b) — that the construction's $\mathscr{R}$ must meet sector by sector: Def. 7.1 wherever the table ascends, and the *pattern* of Def. 7.1b including the down-type expansion as the audit case.

To the Eigenmode Decision. When the construction supplies counts, the Eigenmode Decision's T1 (ordering) tests (α') on the computed densities. Theorem 7.3 tells the construction what T1 will be testing structurally: not that densities ascend by fiat, but that each refinement step's marginal density beat the running average (Def. 7.1). The Eigenmode Decision's reuse-*profile* readout — logged at firewall discipline, expected mixed per column, the down sector withholding consonance — is the Def. 7.1b picture, now given its mechanism: compression holds where columns compress and fails where they expand, while ordering (the thing T1 actually scores) holds across all three charged columns regardless.

---

11. The Open Problem, Narrowed Again

The programme's burden is now:

Exhibit the physical finite closure-state machine $(\mathfrak{A}, \ell)$ and its refinement operation $\mathscr{R}$ from the closure algebra; (i) compute K_c, p_v, C as orbit counts (values owed); (ii) check the first-order reuse condition (Def. 7.1) sector by sector, reproducing ascent everywhere the table ascends, and reproduce the second-order compression *pattern* (Def. 7.1b), including the down-sector expansion; and (iii) determine whether the admissible (recurrent) set, organized by (k, w, n), selects exactly the depths $n \in \{0, 1, 2\}$.

Parts (i)–(ii) are the magnitude content the Eigenmode Decision's counting bound showed the charged sector cannot adjudicate by fitting — now reduced to a frozen structural condition the construction must meet, not a curve it may draw. Part (iii) is the depth census **D5**, governed by the Selection Audit, and remains [Open]. Theorem 3.1 guarantees the counts are well-defined whatever the machine; Theorem 7.3 guarantees the ordering follows from the reuse condition whatever the $\mathscr{R}$; neither can be undone by the still-open identity of the physical machine.

---

12. Position in the Programme

```
Operator Specification            Eigenmode Decision
  A/B-requirements, S1–S5            R1–R7, T1/T2, verdicts,
  frozen first                      frozen first
        │                               │
        └───────────────┬───────────────┘
                        │
        THIS PAPER
   PART I — Orbit Count Theorem
     K_c return count   [Proven, invariant]
     p_v ledger increment [Proven, invariant]
     C = p_v/K_c        [Proven, parameter-free]
     p_v > 0 under H+    [Proven under hypothesis]
     → D1–D4 banked for the whole class
   PART II — Refinement Structure
     density order        [Proven, thin]
     Def 7.1  reuse / ordering  Δp_v/ΔK_c > p_v/K_c
     Def 7.1b compression  G₂ < G₁ (density gaps; second order)
     Reuse Monotonicity (ordering iff Def 7.1)
                          [Proven under hypothesis = Conditional]
     down sector (inherited): masses ascend → Def 7.1 holds
       (order bridge); mass gaps Γ₂>Γ₁ → Def 7.1b fails
       only under [Hypothesis: log-linear bridge] —
       the falsifiable edge, identification [Inherited-empirical]
     → (α′) reduced to Def 7.1; compression profile = Def 7.1b
                        │
        THE OPERATOR CONSTRUCTION  [next paper]
        exhibit (𝔄, ℓ, ℛ) from the closure algebra;
        compute counts (values owed);
        meet Def 7.1 where the table ascends (everywhere);
        reproduce the Def 7.1b compression pattern,
          including the down-sector expansion;
        attempt the census D5 against S1–S5;
        submit to the protocol
```

The principal achievement, without inflation. Part I proves the return count, commitment count, and completion density are well-defined, invariant, parameter-free, mass-blind properties of self-sustaining cycles in any finite closure-state machine. Part II proves that refinement raises completion density exactly when it adds commitment faster than it lengthens the return — the mediant criterion — and separates that first-order ordering condition from a second-order compression condition, identifying the measured down-type sector as one that ascends (ordering intact) while expanding (compression failing). Together they convert the hierarchy programme's central structural claim, (α'), from an assertion into a sharp first-order condition on computed orbit counts, with the compression profile a second frozen condition, both fixed before the operator exists. No physical number is produced and the count of three is untouched. What the paper banks is the machinery *and* its first real structure — and it leaves the construction a single, isolated, fenced problem.

13. Conclusion

The programme had named its quantities and asserted their structure. This paper builds the quantities and proves the structure — as far as the structure can be proven without the physical machine, and not one step further.

The return count is a cycle's length; the commitment count is the irreversible ledger's increment over one return; their ratio is the completion density, dimensionless and parameter-free and blind to mass. These are now objects, not assumptions, for the whole class of finite closure machines. And the structural question the hierarchy turns on — why density rises with depth — now has a precise answer in the abstract: density rises along refinement exactly when refinement reuses structure, adding commitment faster than it lengthens the return, so that the marginal density beats the average. That first-order condition governs ascent and is clean. A second, independent condition governs whether the rate compresses — and here the world already complicates the picture: the down-type quarks ascend (density rising) but expand (their mass jumps widening, the density reading following under the bridge §8 names). The paper keeps the two apart — neither implies the other — attaches its falsifiable edge to the compression condition the down sector actually stresses, and leaves the ordering theorem standing where it belongs — uncontradicted on the inherited reading, because the down densities ascend.

What remains is the physical machine: its update, its commitment rule, its refinement operation, drawn from the closure algebra, and the question of whether its admissible set terminates the census at three. Those are owed, and they are owed against two frozen instruments now — the Eigenmode Decision's protocol and this paper's reuse condition — neither of which the construction will get to rewrite. The machinery is built and the structure it carries is proven; the count is the one thing left to earn.