The Matter Skeleton Theorem

A Partial Structural Derivation of Standard Model Matter from VERSF

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

Ordinary matter shows a handful of stubborn regularities that any deep theory eventually has to account for. Quarks carry charges in thirds. Those thirds are never seen on their own — a lone one-third charge does not walk around free. Matter comes in three near-copies, the generations, identical in charge but climbing in mass: electron, muon, tau, and the same pattern twice more among the quarks. And within each such family the later members are heavier than the earlier ones.

This paper asks how much of that short list follows from the VERSF closure picture, and it is careful to claim only what genuinely follows. It does not try to derive the whole Standard Model — not the full force structure, not the precise masses, not the mixing between families. It proves one thing and proves it cleanly: that a small set of structural assumptions the programme has established elsewhere, taken together, *force* a connected piece of the list above. The result is called the **Matter Skeleton Theorem**, because what it delivers is a skeleton — the shape of matter, not yet its full flesh.

The picture is built in stages, and the paper's discipline is to mark each stage with how sure it is. Some results are *proven*: that members read fractional shares of a class integer, so thirds appear; that a detached member cannot carry its fraction away, so free fractional charge does not occur; that matter repeats, and repeats only finitely. Some results are *conditional*: that the repetition lands on exactly the three-level pattern, and that mass climbs with depth — each true given a stated extra assumption, named in the open. And two things are left frankly *open*, behind two clearly separated gates. One gate is about counting: does the framework's own bookkeeping stop the family count at three? The other is about matching: are the particles we actually measure the framework's objects, sitting in the levels the theory describes, rather than look-alikes? The paper keeps these two questions apart on purpose, because they can fail independently — the count could be right while the match is wrong, or the reverse.

The strength of the paper is not that it settles everything. It is that it draws the line between settled and unsettled exactly, and then checks its own line. For each claimed result, the paper tests, against deliberately altered model situations, that the result really does rest on the assumptions it names and on no others — so that a reader can see, for instance, that "matter repeats finitely" survives even when the assumption fixing the family count is switched off, which is why that result earns the strong grade and the count does not. These checks are logical: they show what each result depends on. The simplest of them stand on their own; the more elaborate test-situations would, to be fully concrete, need the same physical machine the paper leaves to be built — so the checking is honest about resting, in part, on work still owed.

What remains, stated plainly, is a physical machine to be built from the closure rules, the matching of the framework to the measured particle table, and the closure of the family count. Those are owed. The skeleton, the paper argues, is in hand; the two open gates are named; and the whole is laid out so that nothing is claimed more firmly than it has been earned.

Table of Contents

Abstract

This paper proves a single composition theorem — the **Matter Skeleton Theorem** — internal to the VERSF closure programme. The theorem does not re-derive the programme's established results; it proves that those results, taken as hypotheses, *entail* a connected subset of Standard Model matter structure, and it derives the epistemic grade of each entailment from the grades of the hypotheses it uses. Its content is the **logical structure of the chain**: which inputs force which outputs, and at what grade.

The theorem composes three things it does **not** itself establish, and the honesty of the result rests on saying so plainly: **refinement existence** (that distinct refinement classes exist — supplied by the Generation Theorem, not derived here); the **world-to-framework identification** (that the observed particles are the admissible classes — an occupancy map, O1, gated not derived); and the **census closure** (that the admissible count terminates — the admissible census D5, over capacity and depth, open). The theorem derives only the consequences of these, never the things themselves. It does not derive refinement, and it does not derive the world; it derives what follows once refinement exists, once the identification holds, and once the census closes.

From six structural hypotheses — finite distinguishability, closure (bath) transport, class carriers, realization, saturation, and refinement — with one named additional clause for the mass conclusion and the occupancy map O1 for the generation identification, the theorem establishes: (1) fractional member-grain charge appears; (2) free fractional charge does not; (3) matter repeats finitely; (4) the repetition count is bounded by a refinement census, in **five** sub-clauses of decreasing grade that separate count-individuation, the two-mark architecture, the world-identification, and the census; and (5) realized mass is ordered by refinement depth within each column.

The proof is graded throughout, and the grading is **certified modulo witness-model realizability** (§4, Remark 4.0): a tightness check proves by independence models that no conclusion uses a hypothesis it does not cite and that each cited hypothesis is necessary — sound as logical-independence arguments, with the realization of the structured witnesses as concrete closure models owed to the operator construction. Two certifications carry weight — finite repetition provably routes around the conditional refinement architecture (hence [Proven]), and the mass ordering provably does *not* follow from the six structural hypotheses but requires the named closure-burden clause. A **two-gate Collapse Proposition** (§6) then separates the framework's two distinct exposures: one witness model varies the census (testing D5), a second misidentifies the occupancy (testing O1), and the census-independent skeleton survives both. D5 closes the framework's census; O1 maps the world into the framework; they are orthogonal gates, and the skeleton stands without either.

Status Table — Read First

The strongest claim a speculative programme can make is an accurate one. This table states, before any argument, what the paper establishes and at what grade. There are four grades, ordered **Proven ≻ Conditional ≻ Conjectural ≻ Open** (strongest to weakest); "[Inherited]" marks an imported result carrying its source grade. Where a grade is **Open**, a gate tag names the undischarged hypothesis it waits on — **Gate: D5** (the framework's count) or **Gate: O1** (the world-to-framework map). The tag is presentation, not a fifth grade (see §1.1). Every grade is certified in §4, modulo witness-model realizability (Remark 4.0).

Claim	Grade	Gating dependency
Fractional member-grain charge appears	**[Proven]**	carrier + realization
Free fractional charge does not appear	**[Proven]**	saturation
Residue-family *structure* (partition + conjugation)	**[Proven]**	realization
Matter repeats finitely (a finite census exists)	**[Proven]**	finite distinguishability
Refinement levels are count-individuated	**[Proven]**	Exchangeability Lemma
Refinement levels linearly ordered by count (bare order, no mass)	**[Proven]**	Exchangeability Lemma
Within-column mass ordering	**[Conditional]**	closure-burden clause
Two-mark architecture gives R_0, R_1, R_2	**[Conditional]**	two-mark architecture
Dimension-two refinement allocation	**[Conjectural]**	Fold-interface candidate
Observed generations occupy the refinement levels	**[Open · Gate: O1]**	occupancy map O1
Residue-family *catalogue* (the six families)	**[Open · Gate: D5]**	admissible census D5 (capacity)
Repetition count is *exactly three*	**[Open · Gate: D5]**	admissible census D5 (depth exhaustiveness)
Density ascent $C(R_0)<C(R_1)<C(R_2)$, *transferred* from mass ordering	**[Inherited]**	order bridge
Density-ratio expansion (gaps grow), *transferred*	**[Conditional]**	log-linear bridge
Independently *computed* densities $C = p\_v/K\_c$	**[Open]**	Eigenmode Decision, not yet run
Gauge group, hypercharge, chirality, CKM, PMNS, exact spectrum	**[Open]**	not addressed here

1. Hypotheses and Marker Calculus

1.1 The marker calculus

Every claim carries one of four grades: [Proven], [Conditional] (contingent on a stated, carried hypothesis), [Conjectural] (a candidate discharge), [Open] (no current discharge), ordered Proven ≻ Conditional ≻ Conjectural ≻ Open.

Grade-inheritance rule. A conclusion inherits the **meet** (the weakest) of the grades on its derivation path. The proof (§3) applies this; the tightness check (§4) certifies the cited path is exact.

Gate tags on the Open grade. There are exactly four grades; **Open is one grade, not several.** But the programme has *two distinct* undischarged gates, and a bare [Open] would not say which one a conclusion waits on — information the rest of the paper depends on. So an Open grade carries a **tag** naming its gate. The tag is presentation, not a grade:

- **[Open · Gate: D5]** — waits on **D5**, the closure of the framework's own **admissible census over both capacity and depth** — the (k, w, n)-organized recurrent set: whether the admissible capacities terminate (gating the residue catalogue) and whether the admissible refinement levels terminate at the three mark-levels (gating 4e). This is the joint scope of D5; the Orbit Count Theorem states D5 over the same (k, w, n) organization, and the catalogue (a capacity enumeration) and 4e (a depth-exhaustiveness claim) are its two faces. Discharging it would lift the gated conclusion to the grade it inherits from the rest of its path. Failure changes *how many* admissible classes the framework contains.
- **[Open · Gate: O1]** — waits on **O1**, the map from the framework to the observed world. A conclusion here is *conditional on O1* and, while O1 is undischarged, formally Open; discharging O1 would lift it to [Conditional] (the grade it inherits from its structural path). Failure changes *whether the observed particles are* the admissible classes, leaving the framework's count untouched.

The two gates are orthogonal: O1 can hold with D5 open (the identification is right but the count is unproven), and D5 can close with O1 false (the count is proven three, but the world does not occupy the levels as claimed). A conclusion gated by one carries that one tag; §6 proves the orthogonality by separate witness models. Throughout, read "[Open · Gate: X]" as "grade Open, gate X" — never as a distinct grade.

1.2 The hypotheses, restated in one line each

The six structural inputs are theorems of prior papers, restated and used, not re-proven.

- **H1 — Finite distinguishability.** [Inherited, Proven] At bounded refinement, finitely many distinguishable closure states.
- **H2 — Closure (bath) transport.** [Inherited] Member identity is not standing structure; only the class survives transport.

- **H3 — Class carriers (Carrier Theorem).** [Inherited, Proven] Under H2, the carrier of quantized response is the class; the winding integer *q* belongs to the class.
- **H4 — Realization (Realization Theorem).** [Inherited, Proven] A member-grain registration of class winding *q* in a capacity-*k* class reads *q*/*k*, under the registration premises.
- **H5 — Saturation (Saturation Theorem).** [Inherited, Proven] A detached member does not carry the fraction; admissible free configurations are transport-complete, with integral total response.
- **H6 — Refinement (Generation Theorem).** [Inherited] Factored into four sub-claims of distinct grade, because the proof and tightness check use them separately:
  - **H6e (existence)** [Proven]: distinct admissible refinement classes over a shared winding tuple exist.
  - **H6x (count-individuation, Exchangeability Lemma)** [Proven, on the Exchangeability Lemma]: distinct refinement levels are individuated by their activation count — two levels coincide iff their counts coincide.
  - **H6a (architecture)** [Conditional]: the minimal refinement marks are **binary** (each held at most once), exchangeable, and two in number. H6a constrains the *mark-activation count*; it is **not** assumed exhaustive of admissible refinement — whether the two marks generate every admissible level, or whether the recurrent set contains a level no mark-combination produces, is D5 (conclusion 4e), a separate question.
  - **H6d (dimension)** [Conjectural]: the refinement allocation has dimension two (Fold-interface candidate).

Two further inputs are **not** structural closure hypotheses and are named separately so the theorem cannot appear to derive them.

- **H7 — Closure-burden clause (Mass Hierarchy Theorem).** [Inherited, Conditional] Realized mass manifests closure-maintenance burden, monotone in refinement depth. Used only for conclusion (5).
- **O1 — Occupancy map (world-to-framework identification).** [Inherited, gated] The observed Standard Model entries are read as admissible VERSF closure classes occupying the catalogue's columns: (e, μ, τ) as three distinct admissible classes in one charged-lepton column; (u, c, t) in one up-type column; (d, s, b) in one down-type column; neutrinos, if included, in the corresponding neutral column; and — the load-bearing clause — the observed entries correspond to the refinement levels R_0, R_1, R_2 rather than being unrelated duplicates or misidentified catalogue entries. O1 maps the world into the framework; it is used only at conclusion (4d), and it is not discharged here.

1.3 What the theorem does not derive (stated before it is assumed)

To pre-empt the circularity reading directly: the theorem **assumes** refinement existence (H6e) and **derives its consequences**; it does not derive that refinement classes exist. The refinement programme (which establishes H6) and the matter-skeleton programme (this paper, which composes H6 with the rest) are distinct, and the join is explicit. Symmetrically, the theorem does not derive the world: O1 is assumed, not proven, and conclusion (4d) is the consequence of the identification, not the identification itself. The theorem's content is the composition of supplied existence, gated identification, and open census — none of the three established here.

---

2. The Matter Skeleton Theorem

Theorem 2.1 (Matter Skeleton Theorem). Assume H1–H6, and invoke neither an occupancy identification (O1) nor a census closure (D5). Then conclusions 1–4c hold and are **framework-internal** — derived from the closure structure alone, with no appeal to the observed particle table and no claim that the marks exhaust the admissible census (the count of *mark-generated* levels is fixed; whether further non-mark levels exist is left to D5). Conclusion 4d additionally requires O1; conclusion 4e additionally requires D5; conclusion 5 additionally requires H7. Explicitly:

Framework-internal (H1–H6 only): conclusions 1, 2, 3, 4a, 4b, 4c. **Requires the world-map O1:** conclusion 4d. **Requires the census closure D5:** conclusion 4e. **Requires the closure-burden clause H7:** conclusion 5.

The conclusions:

1. **(Fractions appear)** Member-grain registrations take values q/k; at $k = 3$ with q not divisible by 3, a member registers a third. **[Proven]**
2. **(Free fractions do not appear)** No isolated member is a free standing carrier of fractional charge; every admissible free configuration carries integral total response. **[Proven]**
3. **(Finite repetition)** The admissible classes over a fixed winding tuple form a finite, non-empty set of distinct refinement classes — matter repeats, finitely. **[Proven]**
4. **(Bounded repetition count)** The count is bounded by a refinement census, in five sub-clauses separating four logical joints:
    - **(4a)** a finite census exists **[Proven]**;
    - **(4b)** refinement levels are count-individuated **[Proven]** (Exchangeability Lemma H6x), and are thereby linearly ordered by activation count **[Proven]** (4b′) — a bare structural order carrying no mass content;
    - **(4c)** under the two-mark architecture H6a the mark-activation counts are exactly {0, 1, 2}, individuating three **mark-levels** R_0, R_1, R_2 **[Conditional]** — a statement about what the marks generate, *not* a claim that no further admissible level exists; the dimension-two allocation H6d **[Conjectural]** is a logically separate rider, not used for the count;
    - **(4d)** assuming O1, the observed Standard Model generations occupy these refinement levels **[Open · Gate: O1]** (conditional on O1; would lift to [Conditional] once O1 is discharged);
    - **(4e)** the two-mark architecture is *exhaustive* — no admissible level beyond the three mark-levels — equivalently, the census closes at exactly three **[Open · Gate: D5]**; this is D5, and it is not entailed by H6a, which constrains only what the marks generate.

5. **(Mass ordering)** Assuming additionally H7, realized mass is strictly ordered by refinement depth within each populated column: $m(R_0) < m(R_1) < m(R_2)$. **[Conditional]**

The four joints of (4) are deliberately separated: **4b** is structural (counts individuate levels), proved by exchangeability; **4c** is where the *number* of marks enters, fixing the count of *mark-generated* levels; **4d** is where the *world* enters, through O1; **4e** is where the *census closes*, through D5 — the distinct claim that the marks are exhaustive of admissible refinement. They are four different dependencies and four different grades; fusing them would conceal exactly the joins a referee should inspect. In particular, 4c and 4e are not the same "three": 4c says the marks generate three levels, 4e says there are no others.

---

3. Proof

Each path is annotated; §4 proves the annotations complete.

Conclusion 1 — Fractions appear. [Proven]

By H3 the integer q is class-owned. By H4 a member-grain registration reads it at member grain, fixed by permutation symmetry and additivity at

q_member = q / k.

At $k = 3$ with q not divisible by 3 this is a third; the reading is forced under H4's premises. Path **{H3, H4}**, meet [Proven]. ■

Conclusion 2 — Free fractions do not appear. [Proven]

By H5 the fraction is a reading of class support, not portable with a detached member; the free sector is restricted to transport-complete structures,

free admissible = saturated closures + conjugate completions + their compositions,

with integral total response. Path **{H5}**, meet [Proven]. ■

Conclusion 3 — Finite repetition. [Proven]

By H6e distinct refinement classes over a shared tuple exist (non-emptiness). By H1 they are finite in number. Matter repeats, finitely. Path **{H1, H6e}** — neither the architecture H6a nor the count enters. Meet [Proven]. ■

The sharpest grade boundary: finite repetition routes around the conditional architecture. Certified in §4.

Conclusion 4 — Bounded repetition count.

(4a) A finite census exists. [Proven]. By H1 the admissible refinement classes are finite, so the enumeration of levels is a finite object. Path **{H1}**. ■

(4b) Refinement levels are count-individuated. [Proven]. By the Exchangeability Lemma H6x, two refinement levels coincide iff their activation counts coincide — the level is individuated by its count, giving a bijection level ↦ count (set-individuation). This is structural: it fixes *how levels are told apart*, prior to and independent of *how many marks there are*. (The natural order on counts then induces an order on levels; that induced order is invoked where it is used, in conclusion 5, not asserted here.) Path **{H6x}**, [Proven on the Exchangeability Lemma]. ■

(4c) Two marks give three mark-levels R_0, R_1, R_2. [Conditional]; dimension-two [Conjectural]. By H6a the marks are binary, exchangeable, and two in number; with two binary exchangeable marks the mark-activation counts are exactly 0, 1, 2, which by 4b individuate three distinct mark-levels R_0, R_1, R_2. This is a statement about what the two marks *generate* — it does **not** assert that no admissible level lies outside the mark-image; that exhaustiveness claim is 4e (D5). The two-mark input is H6a ([Conditional]). The dimension-two allocation H6d ([Conjectural]) is a logically separate rider — it concerns *why* there are two marks, and is **not** on the path that fixes the count of mark-levels — so the count's meet is [Conditional], with H6d's [Conjectural] confined to the separate dimension claim (status table). Path **{H6x, H6a}**; meet [Conditional]. ■

(4d) Observed generations occupy the levels. [Open · Gate: O1]. By O1 the observed entries (e,μ,τ), (u,c,t), (d,s,b) are admissible classes in their columns and correspond to R_0, R_1, R_2 rather than being duplicates or misidentifications. Given O1 and 4c, the observed three-fold generation pattern *is* the three-level refinement structure. This step uses O1 essentially: without it the framework's levels and the world's generations are formally unrelated. The structural part of the path, {H6x, H6a}, is [Conditional]; O1 is undischarged, so by the meet rule the conclusion is formally Open, tagged **Gate: O1** to mark that it waits on the world-map and not the census. Were O1 discharged, 4d would lift to [Conditional] (inheriting 4c). Path **{H6x, H6a, O1}**. ■

(4e) The marks are exhaustive — exactly three and no more. [Open · Gate: D5]. That the two-mark architecture *exhausts* admissible refinement — that the recurrent set contains no level beyond the three mark-levels of 4c — is not entailed by H1–H6 or O1. H6a fixes what the marks generate (4c); it is silent on whether anything else is admissible. The exhaustiveness question is D5, governed by the frozen Selection Audit over the (k, w, n) census. Path undischarged; Open, tagged **Gate: D5**. ■

4d and 4e are different opens. 4d can fail while 4e holds (the count is three but the observed table does not occupy the levels as O1 asserts); 4e can fail while 4d holds (the observed three occupy three levels, but the framework admits a fourth the world has not shown). The grades name which gate each waits on.

Conclusion 5 — Mass ordering. [Conditional]

The content of conclusion 5 is a *rule*: within each populated column, the later refinement level is heavier. It needs two ingredients and no more. The **level-order** is delivered by 4b′ (H6x) — the count-individuation bijection composed with the natural order on counts linearly orders the levels. The **burden law** is **H7**, the closure-burden clause inherited self-standing from the Mass Hierarchy Theorem: deeper refinement carries more closure structure, more maintenance burden, and realized mass manifests burden monotonically. Composing, within each populated column the later level is heavier — a rule that holds for *whatever* ordered levels exist. Instantiated at the three mark-levels of 4c, it reads

$m(R_0) < m(R_1) < m(R_2)$.

Path of the rule: {H6x, H7} — H6x for the order, H7 for the mass content; meet [Conditional] (H6x [Proven], H7 [Conditional]); **within-column only**. The rule does not use the two-mark cardinality: it makes no claim about *how many* levels there are, only that heavier tracks deeper among those present, so — like C3 and 4b — it **routes around H6a** (certified in §4). H6a enters only when the rule is *instantiated* at R_0, R_1, R_2: those names are the mark-levels of 4c, so the displayed inequality carries H6a as a label on its subscripts, doing no ordering or mass work. H7 is a standing inherited hypothesis, not a result awaiting the Eigenmode computation: the completion-density ascent of §7 is an *independent* preregistered corroboration of the same closure-burden physics, not a premise the rule rests on, so the rule does not inherit the [Open] grade of that not-yet-run test. This cardinality-free path is what lets §6 carry the mass rule through a census change.

This completes the proof of Theorem 2.1. ■

---

4. Tightness — The Cited Paths Are Exact

The meet-rule grade of §3 is correct only if each cited path is exact: a missing hypothesis would lower the true grade, a superfluous one would overstate the dependency and mislead §6. This section certifies both, by **independence models** — to show C does not use H, exhibit a model on C's path with H toggled where C still holds; to show H necessary, a model on the rest of the path with H false where C fails.

Witness models. A *closure model* fixes a census (capacities $\mathcal{K}$, depths $\mathcal{N}$), a transport type, an occupancy map, and the status of each hypothesis. The intended model is

M_0: $\mathcal{K} = \{1,2,3\}$, $\mathcal{N} = \{0,1,2\}$, bath transport, O1 holding, H1–H7 holding.

Toggled models:

- **M_census:** $\mathcal{K} = \{1,2,3,4\}$, $\mathcal{N} = \{0,1,2,3\}$, where the additional refinement level ($n = 3$) is **non-mark-generated** — admissible but not produced by the two marks, and the additional capacity $k = 4$ likewise admitted. H1–H7 and O1's structural clauses still hold: the two marks still generate exactly three mark-levels (H6a, hence 4c, intact), while a further admissible non-mark level exists (so exhaustiveness, 4e, fails). Each surviving hypothesis is about *whatever* census is admitted. Tests D5-sensitivity (the exhaustiveness/enumeration gate).

- **M_misid:** as M_0 but O1's correspondence clause false — the observed entries are admissible classes that do *not* occupy R_0,R_1,R_2 as claimed (e.g. duplicates within one level). Census unchanged. Tests O1-sensitivity. (O1 is a conjunction — class-membership, column assignment, and the load-bearing $R_0/R_1/R_2$ correspondence; M_misid negates only the correspondence clause, so the independence test exercises that clause, and the [Open · Gate: O1] grade covers the whole stack.)
- **M_∞:** infinitely many distinguishable states (H1 false). Tests necessity of H1.
- **M_ledger:** ledger transport (H2 false, H3 false — member-owned integer). Tests necessity of H3.
- **M_¬sat:** saturation absent (H5 false). Tests necessity of H5.
- **M_¬card:** the marks remain binary and exchangeable but are **not** two in number (only the cardinality clause of H6a is negated), so H1, H6e, and H6x all survive (count-individuation rides on exchangeability, which is retained). Tests whether the cardinality clause is used by C3, 4b, and the C5 rule, and whether it is necessary for 4c. (Negating exchangeability as well would break H6x, so the witness toggles cardinality alone — the separable clause 4c actually needs.)
- **M_¬indiv:** exchangeability fails, so distinct refinement levels can share an activation count — count-individuation (H6x) is false — while H1, H6e, and mark-existence hold. Tests necessity of H6x for 4b, the induced level-order, and the C5 rule.
- **M_¬burden:** closure-burden absent (H7 false), H6 holding. Tests necessity of H7 for C5.

Remark 4.0 (Status of the witnesses — what "certified" means here). A witness certifies a grade only if it is a consistent closure model in which the named hypothesis takes the stated truth value. The pure single-hypothesis toggles — M_∞, M_ledger, M_¬sat, M_¬card, M_¬indiv, M_¬burden — are manifestly consistent: each negates one abstract property of an otherwise-finite closure structure, and a finite state set with or without that property exists trivially, so these certifications are unconditional. The two *structured* witnesses, M_census and M_misid, specify finite combinatorial objects (a census with a flagged non-mark level; an occupancy that breaks the correspondence clause); these are logically consistent as specifications — they contain no internal contradiction — and the tightness method needs only that logical consistency, not a concrete dynamics. What it does **not** supply is a realizing closure machine ($\mathfrak{A}$, ℓ, $\mathscr{R}$) exhibiting any such specification as a genuine model of the closure algebra; that construction is owed to the operator paper (§9). The certifications in §4 therefore hold **modulo the realizability of the witness specifications as closure models** — sound as logical-independence arguments, with the realization owed. This modulo is orthogonal to D5: the construction discharges *which census is physical*; it does not underwrite the independence method, which concerns only the logical relations among the abstract hypotheses. Read every "certified" below as carrying this qualification.

4.1 Certification, conclusion by conclusion

C1 (path {H3, H4}). *Necessity:* M_ledger (H3 false) → member owns an integer, no *q*/*k* → C1 fails; dropping H4 removes the reading. *No-smuggling:* C1 holds in M_∞, M_¬sat, M_¬card, M_¬burden, M_misid (the reading is indifferent to the world-map). Bath content enters only via H3. **[Proven] certified.**

C2 (path {H5}). *Necessity:* M_¬sat → free sector unrestricted → C2 fails. *No-smuggling:* holds in M_∞, M_ledger, M_¬card, M_¬burden, M_misid, and with H4 false (integrality survives). **[Proven] certified.**

C3 (path {H1, H6e}). *Necessity:* M_∞ → repetition not finite; ¬H6e → no repetition. *No-smuggling — decisive:* in **M_¬card** H1, H6e, H6x hold (exchangeability retained, only cardinality-two negated) but H6a as a whole is false; classes are still finite and non-empty → **C3 holds without the two-mark architecture**. Also holds in M_misid (repetition is a framework fact, not a world fact). **[Proven] certified — the proven/conditional boundary is real.**

4a (path {H1}). *Necessity:* M_∞. *No-smuggling:* holds in M_¬card, M_misid. **[Proven] certified.**

4b (path {H6x}). *Necessity:* in **M_¬indiv** exchangeability fails, two levels share a count, and the level ↦ count map is not a bijection → 4b fails, so H6x is necessary. *No-smuggling — the second decisive case:* in **M_¬card** exchangeability holds (so H6x holds) but the marks are not two in number → levels are still individuated by their counts even though the cardinality is not two → **4b holds without the cardinality clause**, confirming count-individuation rides on exchangeability and is prior to and independent of the number of marks. Holds in M_misid (individuation is framework-internal). **[Proven on the Exchangeability Lemma] certified.**

4b′ (induced level-order, path {H6x}). [Proven]. The level ↦ count bijection of 4b composes with the natural order on activation counts to linearly order the levels — higher count, higher level — a relation fixed by H6x alone, independent of *how many* levels there are or *which* counts occur (so independent of the cardinality clause: it survives M_¬card). *Necessity:* in M_¬indiv the levels are not individuated by count and carry no such induced order → 4b′ fails. This is the bare structural ordering, distinct from the mass ordering of C5: it carries no H7 and no claim about mass. Path **{H6x}**; **[Proven] certified**, routing around the two-mark architecture exactly as 4b does.

4c (path {H6x, H6a}). *Necessity:* M_¬card (cardinality-two false) → the activation counts need not be {0,1,2} → the three-mark-level count fails; dropping H6x removes individuation. *No-smuggling:* 4c does not use H7 (M_¬burden), H5, or O1 (M_misid leaves the framework's three mark-levels intact); and it does not use the census value — in **M_census** the two marks still generate exactly three mark-levels (the fourth level there is non-mark), so 4c *survives* the census change, confirming 4c is the mark-level count and not an exhaustiveness claim. *H6d-independence:* toggling H6d (the dimension-two allocation) off while retaining two binary exchangeable marks leaves the activation counts {0,1,2} and the three mark-levels intact → the count does **not** use H6d, so H6d's [Conjectural] grade is confined to the separate dimension claim and does not enter the count's meet. **[Conditional] certified; dimension-two [Conjectural] separately, off the count path.**

4d (path {H6x, H6a, O1}). *Necessity — the world-gate:* in **M_misid** the framework still has R_0,R_1,R_2 (H6x, H6a hold) but O1's correspondence is false, so the observed generations do *not* occupy the levels → 4d fails. **The world-identification therefore provably requires O1 and is not delivered by the framework alone.** *No-smuggling:* 4d does not depend on the census *value* — in M_census the observed three still occupy R_0,R_1,R_2 (O1 unchanged), so 4d is insensitive to D5. **This is the certification that 4d's gate is O1, not D5** — it waits on the world-map, not the census. **Grade [Open · Gate: O1] certified.**

4e (path: D5). *Sensitivity:* in **M_census** the census is four levels → 4e fails. *Orthogonality to O1:* 4e fails in M_census regardless of O1's status, and holds-or-fails independently of M_misid. **Grade [Open · Gate: D5] certified, orthogonal to 4d.**

C5 mass-rule (path {H6x, H7}). *Necessity — the structural guard:* in **M_¬burden** H1–H6 hold, H7 false → levels ordered but no mass content → the rule fails. **The mass ordering provably does not follow from the six; H7 is necessary and correctly outside the six.** *H6x-necessity:* in **M_¬indiv** the levels are not individuated by count and carry no induced order (4b′ fails), so there is no "deeper" for "heavier" to track and the rule fails — H6x is on the path. *No-smuggling — parallel to C3 and 4b:* in **M_¬card** (exchangeability retained, cardinality ≠ 2) the rule holds for the levels that exist → it does **not** use the two-mark cardinality; it is independent of O1 (M_misid), of the census value (in **M_census** the rule survives the census change — the survivor fact §6 relies on), and of the not-yet-run density test (H7 standing-inherited). The displayed instance $m(R_0)<m(R_1)<m(R_2)$ carries H6a only through the names R_0,R_1,R_2. **[Conditional] certified.**

4.2 Summary

Conclusion	Path	Necessity witness	Decisive no-smuggling	Grade
C1	{H3,H4}	M_ledger	M_∞, M_misid	[Proven]
C2	{H5}	M_¬sat	M_∞	[Proven]
C3	{H1,H6e}	M_∞	**M_¬card (no two-mark arch)**	[Proven]
4a	{H1}	M_∞	M_¬card	[Proven]
4b	{H6x}	M_¬indiv	**M_¬card (exch. retained)**	[Proven]
4b′ order	{H6x}	M_¬indiv	**M_¬card**	[Proven]
4c	{H6x,H6a}	M_¬card	**M_census (3 mark-levels survive)**, M_¬burden, M_misid	[Conditional]
4d	{H6x,H6a,O1}	**M_misid (O1)**	M_census (not D5)	[Open · Gate: O1]
4e	D5	**M_census**	orthogonal to M_misid	[Open · Gate: D5]

Conclusion	Path	Necessity witness	Decisive no-smuggling	Grade
C5 mass-rule	{H6x,H7}	**M_¬burden (H7)**; ¬H6x	**M_¬card (no cardinality)**, M_census, M_misid	[Conditional]

Three certifications carry the structure: **C3, 4b, and the C5 mass-rule all route around the two-mark architecture** — finite repetition needs only finiteness and existence, count-individuation needs only exchangeability, and the within-column mass rule needs only the induced order and the burden law; the cardinality-two clause is off all three paths, certified by M_¬card, so all three are established prior to the architecture. **C5 still requires H7** (mass ordering does not follow from the structural six — M_¬burden), and **4c, 4d, 4e separate** — 4c survives M_census (the mark-level count), 4d fails under M_misid (the world-gate), 4e fails under M_census (the census/exhaustiveness gate) — so the mark-count, the world-gate, and the census-gate are three distinct dependencies. §6 builds the collapse on exactly this separation.

---

5. The Residue-Family Structure the Theorem Supports

For capacity *k*, classes sort by *q* mod *k*. The assignment structure is the residue-family partition with conjugation:

F(k, r) = { capacity-k classes with q ≡ r (mod k) }, F(k, r) ↦ F(k, −r).

[Proven] The *partition and conjugation* follow from H4 (the *q*/*k* reading fixes the residue label) and, by the census-independence certified for C1 in §4, do not depend on which capacities are admitted.

[Open · Gate: D5] The *catalogue* — F(1,0), F(2,0), F(2,1), F(3,0), F(3,1), F(3,2), with F(3,1)/F(3,2) a conjugate pair at *k*=3 — ranges over the census {1,2,3} and is conditional on **D5**. The structure is proven; the enumeration waits on the census gate.

---

6. The Two-Gate Collapse — What Each Gate Actually Controls

The programme has two distinct undischarged gates, and they fail in different directions. This section proves their separation with the two witness models built for it: **M_census** (D5) and **M_misid** (O1). The census-independent, identification-independent skeleton survives both.

Proposition 6.1 (Two-Gate Collapse).

(a) Census gate (D5). Evaluate in **M_census** (census {1,2,3,4}/{0,1,2,3}, the fourth level non-mark-generated, all hypotheses and O1's correspondence holding). The items with no census assumption on their certified path — C1, C2, C3, 4a, 4b, **4c** (the two marks still generate exactly three mark-levels R_0,R_1,R_2; the fourth level is non-mark), C5-as-rule (the within-column ordering held independent of which columns are populated), the residue *structure*, and **4d** (the observed three still occupy R_0,R_1,R_2; O1 unchanged) — hold unchanged. The items asserting exhaustiveness or ranging over the census value — the six-family *catalogue*, **4e** (a fourth, non-mark level now exists, so the marks are not exhaustive), and census *completeness* — are falsified. **D5 controls the exhaustiveness and enumeration of the census, not the mark-architecture, the structure, or the world-occupancy.**

(b) Identification gate (O1). Evaluate in **M_misid** (census {1,2,3}, but O1's correspondence false). The framework-internal items — C1, C2, C3, 4a, 4b, 4c (R_0,R_1,R_2 still individuated and counted), 4e's census status, the residue structure and catalogue, C5-as-rule — hold unchanged, because none asserts that the *observed* particles are these classes. The single item that fails is **4d**: the observed generations no longer occupy the levels. **O1 controls the world-occupancy, nothing else.**

Proof. Each item's certified path (§4) either contains a census assumption, contains O1, or contains neither. Items with neither survive both models (the toggled feature is off their path). Items with the census assumption fail in M_census; the unique item with O1 fails in M_misid. The partition is read directly off the §4.2 table. ■

Corollary 6.2 (The two gates, named). D5 closes the framework's census; O1 maps the world into the framework. They are orthogonal: the matter *skeleton* (C1–C3, 4a, 4b, the within-column mass rule on path {H6x, H7}, the residue structure) is both census-independent and identification-independent, and stands whatever D5 and O1 turn out to be. The mass rule's cardinality-free path is what makes it a genuine survivor here — it does not assume there are exactly three levels, so a census change cannot touch it. What waits on D5 is the *count* (catalogue, exactly-three, completeness). What waits on O1 is the *world-occupancy* (4d alone). No single failure collapses both, and neither touches the skeleton.

This is the precise centrality of the two gates: the skeleton is derived and stands alone; the count and the world-identification are the two distinct things still owed, and the paper says which is which.

---

7. The Computational Decision, and Its Owned Failure Edge

Conclusion 5 establishes within-column mass ordering on the closure-burden clause H7, inherited self-standing from the Mass Hierarchy Theorem; it does not await the computation below. A *separate and independent* line of corroboration is preregistered rather than assumed: the Eigenmode Decision tests whether completion density rises with refinement depth — a computation from closure structure and anchoring dynamics with masses excluded, outputting

$C = p_v / K_c$,

with K_c the return count and p_v the commitment count over one return (parameter-free, mass-blind orbit counts from the Orbit Count Theorem). The clause under test is, within each column,

$C(R_0) < C(R_1) < C(R_2)$.

[Inherited] Two facts are carried in so this is not misread as expected-clean.

First, the ascent is governed by the **reuse condition**: density rises at a step exactly when its marginal density exceeds the running average,

$\Delta p_v / \Delta K_c > p_v / K_c$,

the mediant criterion — refinement raises density when it adds commitment faster than it lengthens the return.

Second, **the test is not expected to pass uniformly**, and two distinct inheritances must be kept apart — exactly as the Reuse Monotonicity paper keeps them, on pain of the corpus contradicting itself across the two papers. The down-type column's adjacent *mass* log-gaps run $\Gamma_1 \approx 2.99$ then $\Gamma_2 \approx 3.80$ [Inherited, empirical] ($\ln(m_s/m_d)$, $\ln(m_b/m_s)$; these are scheme- and scale-dependent and carry whatever scheme the home paper fixes), the second the larger. Here Γ denotes the *mass*-gap, reserved distinctly from G, the *density-ratio* log-gap; the two are not interchangeable.

The **ascent** — that the down-type completion densities rise with depth, $C(R_0) < C(R_1) < C(R_2)$ — transfers from the mass ordering under the order bridge and is carried clean [Inherited]. The **expansion** — that the density-ratio gaps themselves grow ($G_2 > G_1$, the C-sequence spreading rather than compressing) — does **not** follow from the mass-gap expansion ($\Gamma_2 > \Gamma_1$) by the order bridge alone. It requires the named **log-linear bridge**, a second-order strengthening of ANCH-COMP strictly stronger than the order bridge, carried here as a hypothesis (the companion's [Hypothesis — log-linear bridge]); the expansion claim is therefore [Conditional on the log-linear bridge], not [Inherited]. Drawing the density-expansion conclusion directly from the mass numbers, with no bridge carried, is the second-order firewall crossing the companion paper forbids, and is avoided here.

The reading is column-by-column: densities ascend where the table ascends (clean, under the order bridge); whether a column's density-ratios *compress or expand* is the second-order question the log-linear bridge governs, and the down-type column is the owned case where, under that bridge, they expand. A construction forcing uniform compression would contradict the measured table and is excluded in advance.

---

8. Why This Is Not Yet a Full Standard Model Derivation

Stated more sharply than a referee would.

Derived (at the certified grades of the Status Table): class carriers [Proven]; fractional readings [Proven]; saturation / free-sector integrality [Proven]; residue structure [Proven]; finite repetition [Proven]; count-individuation and the induced level-order [Proven]; the {0,1,2} mark-count [Conditional]; within-column mass ordering [Conditional]; the density *ascent* transferred from the mass ordering [Inherited, under the order bridge].

Gated, not derived: the world-occupancy of the levels [Open · Gate: O1]; the catalogue and exactly-three [Open · Gate: D5]; the density-ratio *expansion* [Conditional on the log-linear bridge].

Not derived at all: SU(3) × SU(2) × U(1) as a complete product gauge group with full representation content; the hypercharge table; weak chirality; electroweak symmetry breaking in full; CKM and PMNS mixing; exact Yukawa magnitudes; absolute masses; and the **independently computed** completion densities C = p_v/K_c (the Eigenmode computation is preregistered, not run — this is the unrun *computation*, distinct from the *transferred prediction* above, which is inherited).

The derivation reaches the matter *skeleton* — the structure that survives both gates (Proposition 6.1) — not the whole Standard Model, not the closed count, and not the world-identification. The honest final claim is graded: a proven fractional-charge-and-saturation core, a conditional generation-and-mass reach, and two distinct open gates, D5 and O1. The theorem assumes refinement existence and derives its consequences; it does not derive refinement, and it does not derive the world.

---

9. Position in the Programme

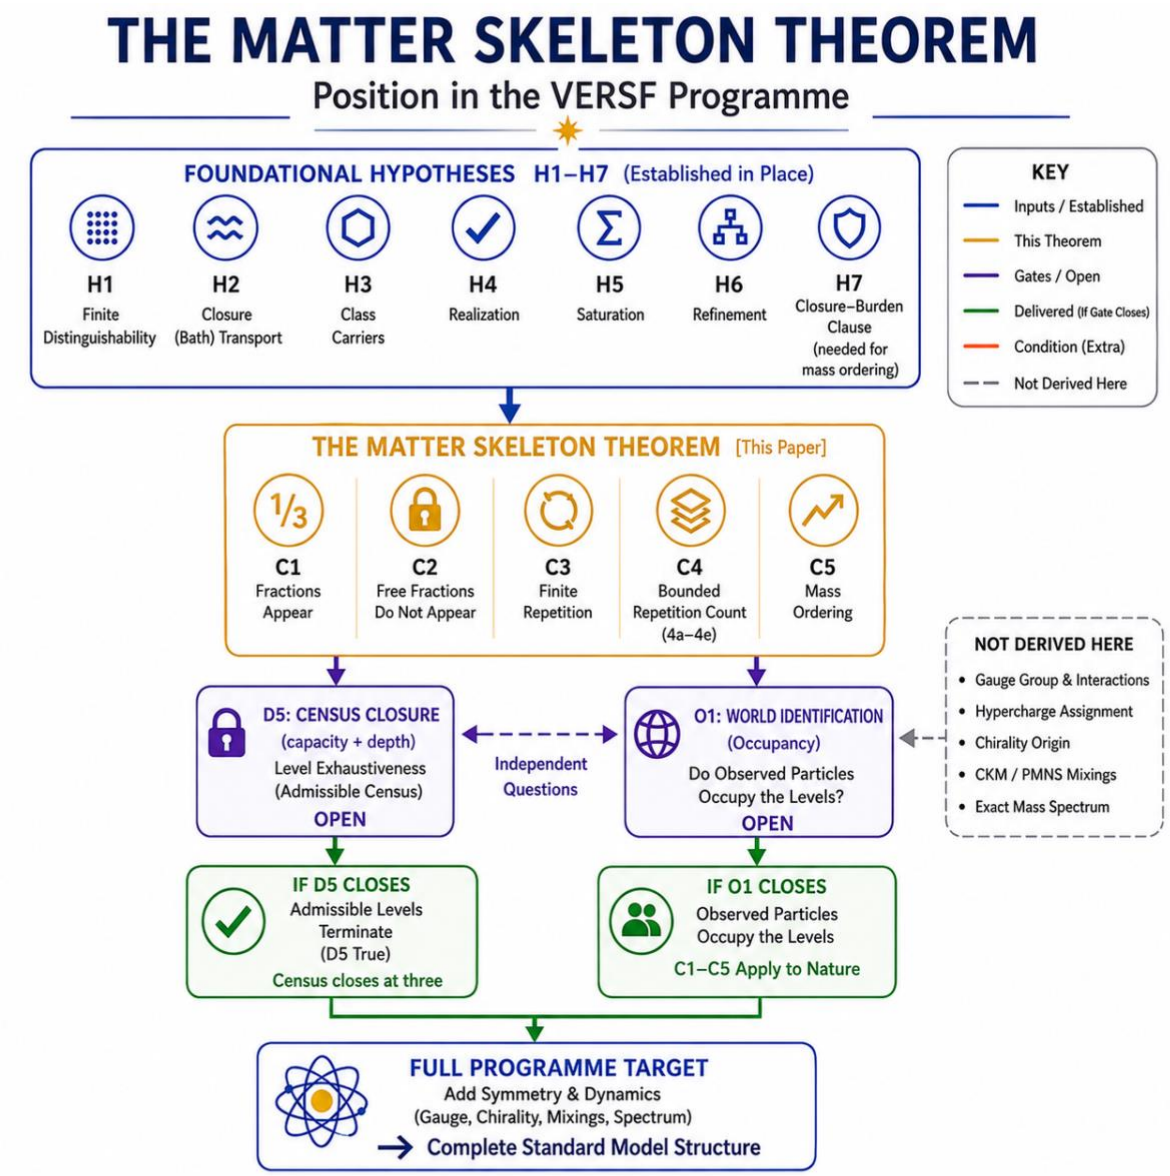

The principal achievement, without inflation. One theorem — six structural hypotheses entail a graded matter skeleton — with its grades certified against independence models (modulo their realizability, Remark 4.0), and one proposition — the framework's two gates, the census D5 and the world-map O1, are orthogonal, and the skeleton survives both. The proven core is census-independent and identification-independent; the conditional reach is named at its certified grade; the two opens are distinguished as the count gate and the world gate. No physical number is produced and the count of three is not closed. The theorem assumes refinement existence and derives its consequences — it derives neither refinement nor the world.

10. Conclusion

VERSF can now state, as a graded composition rather than a survey — its grades certified modulo witness realizability (Remark 4.0) — that its closure architecture entails a connected part of Standard Model matter structure, and can say exactly which part, at which grade, behind which gate.

The five-way split of the count conclusion is where the structure became honest: a finite census exists [Proven]; levels are count-individuated [Proven, by exchangeability]; two marks give three levels [Conditional]; the observed generations occupy them [Open · Gate: O1]; the census closes at three [Open · Gate: D5]. These are four different logical joints behind one phenomenon, and naming them separately is what lets the theorem compose H6 with the rest without appearing to derive H6 — and without appearing to derive the world. The theorem assumes refinement existence and derives its consequences; it assumes the occupancy map and derives what the world then is; it derives neither.

The tightness check makes the grades trustworthy rather than asserted, and the two-gate collapse locates the two distinct exposures precisely: O1 maps the world into the framework, D5 closes the framework's census, and the skeleton stands without either. What remains is the physical machine drawn from the closure algebra, the discharge of O1 against the observed table, and the closure of D5 against the frozen Selection Audit. Those are owed, against instruments the construction will not get to rewrite. The skeleton is derived, and its grades certified modulo the realizability of the witness models (Remark 4.0). The count and the identification are the two things left to earn.