The Yukawa Hierarchy Theorem

Hypercharge Completion, Species–Family Identification, and the Emergence of Mass Ordering from Refinement Depth

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

Matter comes in three "generations." The electron, the muon, and the tau are, as far as every known force can tell, the same particle three times over — same electric charge, same interactions — except for one thing: the muon is about two hundred times heavier than the electron, and the tau heavier still. The same threefold repetition runs through the quarks. The Standard Model records this but does not explain it. It hands each particle a private number, the Yukawa coupling, that sets its mass, and it offers no reason why those numbers come in three repeated families or why the heaviest is so much heavier than the lightest.

This paper asks the sharp version of that question: if the three generations are identical under the forces, what makes them differ in mass at all? The answer proposed here is that two different properties of a particle are read off two different things. *What a particle is* under the forces — its charges — is fixed by one kind of structure. *How heavy it is* is fixed by another, which the programme calls completion density. The first kind of structure is blind to the second. That single fact does most of the work: because charge cannot "see" the property that distinguishes the generations, the three families come out identical under the forces; because mass can see it, they come out different in mass.

A useful consequence falls straight out of this and has already been tested. Since the mass-generating part of the theory cannot touch charge, the muon is forced to carry exactly the same electric charge as the electron — and that equality is one of the most precisely confirmed facts in physics. A theory built the other way would have predicted the muon and electron carrying different charges, which is plainly false. Here, charge sameness across generations is not a coincidence to be absorbed; it is a requirement of the construction.

That leaves the masses. The claim is that the Yukawa numbers are not fundamental constants but readouts of how strongly each family is anchored into the underlying structure — heavier means more strongly anchored. The ordering then has two parts, and they carry different levels of confidence, which the paper is careful to keep separate. The secure part: *if* the three families are ranked by anchoring strength in the natural order, their masses come out in that same order —

this follows directly, because mass is proportional to anchoring strength. The less secure part: *whether* they are ranked that way in the first place. Here the paper makes partial progress rather than a full answer. The families fill a kind of register of fixed size in doubling steps — one, then two, then four — and the third generation is the one that exactly *fills* it: it adds no new force-charge, it completes the count. This "saturation" picture is a solid, almost arithmetic fact about how the families pack into the register. What it does *not* by itself deliver is the mass ranking. Getting from "the third family completes the register" to "the third family is heaviest" needs two further assumptions — that this packing fraction tracks the anchoring strength, and that anchoring strength tracks mass — and the paper marks both as unproven bridges rather than results. There is also a subtlety the paper is explicit about: the register fact gives a single ranking by family, whereas the masses split by type (the up-type, down-type, and charged-lepton ladders behave somewhat differently), and tying the one ranking to the several is itself left open. So the honest reading is: the saturation picture supplies a clean structural *reason to expect* the third family to sit at the top, but it does not on its own *prove* that it does.

None of this replaces the Higgs mechanism. The Higgs still turns these couplings into masses in the usual way; the proposal here is only about where the couplings come from. And the paper does not claim to compute the actual mass values — that is left for later work. What it does is prior and necessary: explain why identical-charge families can differ in mass at all, why the mass order should track the family order, and why the Yukawa numbers are better read as emergent quantities than as an unexplained list. The precise statements, with their exact assumptions and grades, are in the status table and the sections that follow; this summary is the plain-language version of them.

Contents

Status Table — Read First

Claim	Grade	Dependency
HC1 — Species bundle contains the five SM Weyl multiplet slots {Q, u^c, d^c, L, e^c}	Conditional	inherited seat/capacity catalogue (T1a roster)
HC2 — Hypercharge assignment matches the SM values	Conditional	inherited winding/seat charge lattice (T1b charges)
HC3 — The five anomaly conditions cancel for the completed bundle	Proven (arithmetic)	direct check; a necessary consistency confirmation, not a certification
HC4 — A durable species is a complete anomaly-free SM family	Conditional	HC1 + HC2 directly (exact roster, exact charges); HC3 confirms consistency but does not establish identity
R1 — Species replication is gauge-identical	Inherited Proven	Replication Theorem
F1 — Charge-side factorization: χ is independent of refinement class c, so gauge identity and refinement depth are separable coordinates and charge cannot read depth	Proven (structural / by construction)	definitions of χ; the mass consequence is not part of F1
O1 — The amplification operator $\hat{Y}$ is positive, monotone, refinement-sensitive, and charge-blind, with $[\hat{Y}, \chi] = 0$	Proven (structural) given the charge maps	F1
Y1 — Charge identity and mass difference are mutually consistent because charge and mass read different structural coordinates (a non-contradiction result, not an existence result)	Proven on factorization	F1
Y2 — Completion-density order supplies the generation mass ordinal	Conditional	ANCH-COMP bridge
Y2′ — Mass ordering, downstream: given the completion-density ordering $C_1<C_2<C_3$, the per-charged-sector masses are ordered	Proven, modulo (anchoring relation $m\propto C$, itself Conditional on $P\hbar$) ∧ (ε below the gaps)	$m \propto C$ is order-preserving; monotonicity automatic, not assumed
Y2″ — The completion-density ordering itself, $C_1<C_2<C_3$ (per	Open, but no longer bare: the depth ordinal is derived (Y7); the per-	eigenmode programme (anchoring §16); ECB (Corollary 5.8)

Claim	Grade	Dependency
charged sector), from interface eigenmodes	sector triples and ECB remain open	
Y7 — Refinement-efficiency ordinal $E_1<E_2<E_3$ from the dyadic loads (1,2,4) in the K=7 register; third generation is the saturation generation ($E_3=1$)	Proven (arithmetic), conditional on inherited loads and K=7	refinement/census arc
ECB — Efficiency–Completion Bridge: completion density is a positive monotone function of refinement efficiency, so E-ordering induces C-ordering	Conditional	the named upstream hinge (Corollary 5.8)
Y6 — Sectoral sensitivity λ_a is forced sector-dependent (shared-K determination)	Proven within the exponential schema ∧ ε below the gaps	measured ratios + c a single shared index
Y3 — Yukawa couplings are effective completion-density readouts	Conjectural / Conditional	depends on explicit mass-map construction
Y4 — Exact Yukawa magnitudes are derived	Open	not done here
Y5 — CKM/PMNS mixing entries are derived	Open	downstream

(Token note: HC1–HC4 retoken what earlier drafts called H1–H4, to avoid collision with the distinguishability/finiteness budget H1 used elsewhere in the corpus. The amplification operator is written Ŷ, not $\mathcal{A}$, to avoid collision with the state-space and anchoring-update glyphs of the operator paper.)

The paper's strongest new move is not a numerical mass prediction. It is the structural separation:

hypercharge is fixed by winding and seat; generation mass hierarchy is fixed by refinement depth through completion density.

A note on what hypercharge completion is and is not. The anomaly cancellation (HC3) is a *consistency* result, not a *derivation* and not a *certification*. It does not establish that the bundle is the Standard Model family — anomaly-free chiral sets beyond the SM exist, so cancellation refutes on failure but cannot verify on passage. What establishes the identification is HC1 and HC2 directly: the roster is exactly the five SM representations, and the charges are exactly the SM hypercharges, both by inherited assignment. HC3 then confirms those inherited values are mutually compatible. The Replication Theorem already drew this necessary-not-sufficient line for anomaly cancellation, and this paper holds to it.

That separation is what allows replicated families to carry identical charges yet unequal masses. Y1 records only that the two facts do not contradict. The mass-ordering claim now splits in two, at different grades: *given* the completion-density ordering, the masses are ordered (Y2′, Proven modulo the anchoring relation and ε, §6a); whether the species are *actually* ordered by completion density is Open, the output of the eigenmode programme (Y2″, §6a, Input D). The separation itself is the charge-side Factorization Theorem (F1, §4a) — which asserts only that charge cannot read refinement depth — given operator form (O1, §5a); the operator's commutator with charge forbids generation-dependent charge, a test already passed (§5a). The measured ratios force one positive structural fact: the sectoral sensitivity must be sector-dependent (Y6, §7), because the completion-depth spacing is shared across sectors and so cannot carry the cross-sector divergence. What none of these supply is the count of refinement classes or the magnitudes — those remain inherited and open respectively, as the ledger in §13a records.

---

1. The Problem: Identical Charges, Unequal Masses

The Standard Model presents two facts side by side without deriving their common origin.

First, the matter fields repeat. There are three generations. Each generation contains the same pattern of quarks and leptons. The electron, muon, and tau share the same electric charge and the same gauge interactions. The up, charm, and top quarks share the same charge and force assignments. The down, strange, and bottom quarks likewise repeat the same charge structure.

Second, the masses are wildly unequal. The electron is light, the muon heavier, the tau heavier still. The same holds, more dramatically, in the quark sector, where the top quark is enormously heavier than the up quark. The Standard Model encodes these differences through Yukawa couplings. For a fermion f,

$m_f = y_f v / \sqrt{2}$,

where v is the Higgs vacuum expectation value and y_f is the fermion's Yukawa coupling. But the y_f are not explained by the Standard Model. They are fitted inputs.

The question is therefore not merely:

Why are there three generations?

That question was addressed by the replication and species-count arc.

The sharper question is:

Why are the three generations gauge-identical but mass-distinct?

A correct answer must preserve both halves. If the structural coordinate that distinguishes generations also changed charge, replication would fail. If it touched no mass-relevant property, hierarchy would fail. The solution must therefore separate the charge coordinate from the mass coordinate.

VERSF has exactly such a separation. The charge map reads winding and seat structure, and the refinement coordinate does not enter it. That is why families replicate gauge-identically. But refinement depth does enter completion density, return stability, and anchoring strength. That is where the mass hierarchy can enter. This paper makes that separation explicit.

---

2. Inherited Inputs

The following results are inherited.

Input A — Species replication. The Replication Theorem established that the species/refinement coordinate is charge-blind. The charge map depends on winding and seat data, not on refinement class. Changing species therefore changes the refinement class but leaves gauge charge unchanged.

Input B — Species bundle. A durable species σ carries a bundle $\mathfrak{B}(\sigma)$: the set of charge-carrying modes sharing σ's refinement class while ranging over admissible winding and seat data.

Input C — Standard Model target family. In left-handed Weyl convention, one Standard Model generation consists of:

$Q = (3, 2)_{+1/6}$, $u^c = (\bar{3}, 1)_{-2/3}$, $d^c = (\bar{3}, 1)_{+1/3}$, $L = (1, 2)_{-1/2}$, $e^c = (1, 1)_{+1}$.

Input D — Completion-density order (assumed, not derived). The spectral arc *assigns* durable species an order by completion density $C = p\_v/K\_c$ (the flip-completion density of §5c). This order is taken here as the species ordinal. It is essential to be exact about its status: this paper *assumes* the ordering $C_1 < C_2 < C_3$; it does not derive it. Whether the VERSF interface dynamics actually *force* the three stable replicated species to satisfy

$p_{v,1}/K_{c,1} < p_{v,2}/K_{c,2} < p_{v,3}/K_{c,3}$

is the output of the eigenmode programme (anchoring paper §16, explicitly future work) and is Open. It is no longer *bare*, however: §5d derives a refinement-efficiency ordinal $E_1 < E_2 < E_3$ from the inherited dyadic loads (1, 2, 4) in the K = 7 closure register, and the Efficiency–Completion Bridge (ECB, Corollary 5.8) carries that to $C_1 < C_2 < C_3$. So the *depth* ordinal has a structural origin; what remains open is ECB and the relation of the single depth ordinal to the per-sector triples below. Moreover, because the sector-dependent sensitivity forced by Result 7.1 (§7) must enter somewhere, and the natural carrier is a sector-dependent $p_{v,a}(c)$, the completion density is in general itself sector-dependent, $C_a(c) = p_{v,a}(c)/K_c(c)$, with $K_c(c)$ shared across sectors (a refinement-class property) and $p_{v,a}$ carrying the sectoral structure. The upstream open problem is therefore not "compute three numbers" but "compute three numbers per charged sector and show each triple is ordered." Completion density C is a spectral quantity; it is not yet a mass.

Input E — Anchor–completion bridge (ANCH-COMP). The mass hierarchy arc carries the relation that completion density maps to mass. Under the anchoring foundation (§5c), this map is the near-identity $m \propto C$ ($m = \eta\hbar/(c^2\Delta t) \cdot C$, positive constants), so the *bridge from C to m is monotone automatically* rather than by separate hypothesis — multiplication by a positive constant preserves order. What remains conditional in Input E is therefore not the monotonicity of $C \mapsto m$ but the anchoring relation itself (Conditional on the action postulate P$\hbar$ and the Δt calibration, §5c). This is a sharper placement of the conditionality than earlier drafts, which treated the monotonicity as the assumption; the monotonicity is free once the anchoring relation holds, and the anchoring relation is the thing that is Conditional.

---

3. Hypercharge Completion

Before the Yukawa hierarchy can be addressed, the species bundle must be completed. A replicated object is not yet a Standard Model generation merely because it repeats. It must repeat the right thing.

This section is a **consistency theorem, not a derivation theorem.** It does not derive the hypercharges +1⁄6, −2⁄3, +1⁄3, −1⁄2, +1 from below; their origin is the inherited winding/seat charge arc, upstream of this paper. What this section does is twofold: it states that the inherited roster and charges *are* the Standard Model family (a matter of reading off the inherited assignment), and it checks that those inherited values pass the anomaly conditions (a matter of arithmetic). The identification rests on the first; the second is a necessary confirmation that the inherited values are mutually compatible.

The logical direction matters, and it is the direction the Replication Theorem was careful about. Anomaly cancellation is *necessary* for a consistent chiral set but not *sufficient* to single out the Standard Model family — anomaly-free chiral sets beyond the SM exist, so cancellation refutes on failure but cannot verify on passage. Therefore the identification "this bundle is the SM family" must rest on the roster being exactly the five SM representations (HC1) and the charges being exactly the SM hypercharges (HC2), *not* on the fact that they cancel (HC3). The cancellation is the test those inherited values must pass, not the reason the bundle is the SM family.

The Standard Model family is not an arbitrary collection of fields, and its hypercharges cancel gauge anomalies — a severe consistency test that a missing field, an extra exotic field, or a wrong hypercharge generically breaks. The species bundle must therefore pass the following check:

Do the inherited roster and charges cancel all five anomaly conditions?

If yes, the inherited Standard Model assignment is confirmed consistent. If no, the programme fails at this point. Passing does not by itself certify the family; the certification is HC1 + HC2.

3.1 The bundle slots

The inherited seat structure supplies the five multiplet slots:

Q, u^c, d^c, L, e^c.

These correspond to:

- Q: colored weak doublet;
- u^c: colored weak singlet, up-conjugate slot;
- d^c: colored weak singlet, down-conjugate slot;
- L: colorless weak doublet;
- e^c: colorless weak singlet.

The non-abelian representation structure is therefore:

$Q = (3, 2)$, $u^c = (\bar{3}, 1)$, $d^c = (\bar{3}, 1)$, $L = (1, 2)$, $e^c = (1, 1)$.

This completes the non-abelian half of the family identification.

3.2 The hypercharge assignment

The inherited winding/seat charge lattice assigns the following hypercharges:

$Y(Q) = +1/6$, $Y(u^c) = -2/3$, $Y(d^c) = +1/3$, $Y(L) = -1/2$, $Y(e^c) = +1$.

So the completed bundle is:

$Q = (3, 2)_{+1/6}$, $u^c = (\bar{3}, 1)_{-2/3}$, $d^c = (\bar{3}, 1)_{+1/3}$, $L = (1, 2)_{-1/2}$, $e^c = (1, 1)_{+1}$.

This is the Standard Model family in left-handed convention. The electric charges follow from $Q_{em} = T_3 + Y$ and reproduce the observed values: the Q doublet yields $+2/3$ and $-1/3$, the L doublet yields 0 and -1, and the conjugate singlets supply the remaining entries.

3.3 Anomaly cancellation

The bundle must now pass the anomaly tests. Multiplicities below count colour and weak-isospin degrees of freedom.

Linear U(1)–gravity anomaly. Sum over all left-handed Weyl degrees of freedom:

$\Sigma Y = 6(1/6) + 3(-2/3) + 3(1/3) + 2(-1/2) + 1(1) = 1 - 2 + 1 - 1 + 1 = 0$.

Cubic U(1)³ anomaly. Over a common denominator of 36:

$\Sigma Y^3 = 6(1/6)^3 + 3(-2/3)^3 + 3(1/3)^3 + 2(-1/2)^3 + 1(1)^3 = 1/36 - 8/9 + 1/9 - 1/4 + 1 = (1 - 32 + 4 - 9 + 36)/36 = 0$.

$[SU(3)]^2U(1)$ anomaly. The colored fields contribute (weak multiplicity included for Q):

$2(1/6) + (-2/3) + (1/3) = 1/3 - 2/3 + 1/3 = 0.$

$[SU(2)]^2U(1)$ anomaly. The weak doublets contribute (colour multiplicity included for Q):

$3(1/6) + (-1/2) = 1/2 - 1/2 = 0.$

Witten SU(2) global anomaly. The number of SU(2) doublets is:

3 colored Q doublets + 1 lepton L doublet = 4.

Four is even, so the Witten anomaly is absent.

3.4 Hypercharge completion theorem

Theorem 3.1 — Hypercharge Completion (consistency). If the inherited seat catalogue supplies exactly the five Standard Model multiplet slots Q, u^c, d^c, L, e^c (HC1) and the inherited winding/seat lattice assigns them exactly the hypercharges

$+1/6, -2/3, +1/3, -1/2, +1$ (HC2),

then the durable species bundle *is* one Standard Model family by inherited assignment, and that family is anomaly-free, since the five anomaly conditions vanish (HC3).

Grade: Conditional on the inherited roster and charges (HC1, HC2); the anomaly cancellation (HC3) is a Proven arithmetic confirmation that those inherited values are mutually consistent. Cancellation is necessary, not the certifier of identity.

Proof. By HC1 the slots are exactly the five Standard Model non-abelian representations. By HC2 the hypercharges are exactly the Standard Model left-handed Weyl hypercharges. Roster and charges being the Standard Model ones by inherited assignment, the bundle is one Standard Model family — this is the identification, and it rests on HC1 and HC2, not on what follows. Separately, the five anomaly conditions vanish as computed in §3.3, so the inherited assignment is consistent; this confirms the values are mutually compatible but does not, by itself, single out the family, since anomaly-free chiral sets beyond the Standard Model exist. The identification and the consistency check are therefore distinct, and the family identity is carried by the former.

3.5 Result of Part I

The Replication Theorem gave:

species differ by refinement class, and refinement class is charge-blind.

The Hypercharge Completion Theorem now gives:

each species bundle is, by its inherited roster and charges, one Standard Model family — and that family passes the anomaly check.

Together they yield:

durable species = replicated Standard Model generation.

This is the necessary starting point for the Yukawa hierarchy. There are now three charge-identical families. The question becomes why their masses differ.

---

4. The Separation Principle

The central structural principle of the paper is this:

charge cannot see refinement class; mass can.

Charge reads winding and seat data alone. Mass reads winding and seat data together with refinement depth, the latter entering through completion density and anchoring strength. The asymmetry is not that charge and mass read disjoint coordinates — both read winding and seat. The asymmetry is in one coordinate only: refinement class c is invisible to charge and visible to mass.

Let a mode be indexed as:

(w, s, c),

where:

- w is winding data;
- s is seat/capacity data;
- c is refinement class.

The charge map has the form:

$\chi(w, s)$,

with no dependence on c. So c is invisible to charge.

The mass-readout map has the form:

M(w, s, c), more precisely M_{w,s}(C(c)),

where C(c) is the completion-density measure associated with refinement class c, and the (w, s) dependence carries the sectoral scale. So c is visible to mass, and so is the sector.

This gives the programme exactly what the Standard Model requires:

- same charges across generations, because χ does not depend on c;
- different masses across generations, because M does.

The Standard Model encodes this through arbitrary Yukawa constants. VERSF reads the constants as effective readouts of refinement-depth structure modulated by sector.

4a. The Charge-Side Factorization Theorem

The separation principle, stated as a theorem, is the deepest structural claim of the paper. It concerns one thing only: that charge cannot read refinement depth. That much is true by construction. The further claim that mass *does* read depth — that the hierarchy actually tracks completion density — is not part of this theorem; it is the conditional content of Y2/Y2′, and it must not be smuggled in under the structural grade.

Theorem 4.1 — Standard Model Family Factorization (charge side). On the space of charge-carrying modes indexed by (w, s, c), the charge map factorizes off the refinement coordinate:

$\chi = \chi(w, s)$, independent of c.

Gauge identity and refinement depth are therefore separable coordinates, and charge cannot read depth: for fixed (w, s), χ is constant as c varies, and c ranges independently of (w, s). The refinement coordinate is *available* to a mass map, but it is invisible to charge.

Grade: Proven (structural / by construction). This covers the charge half only. Whether mass actually varies with c — whether the available coordinate is used to generate a hierarchy — is Y2/Y2′, Conditional on ANCH-COMP, and is not asserted here.

Proof. The charge map is $\chi(w, s)$, a function of winding and seat data alone (Input A). It does not take c as an argument, so for fixed (w, s) it is constant as c varies: the gauge factor is independent of refinement depth. The refinement class c ranges over the admissible catalogue independently of (w, s) (Input B), so the depth coordinate is unconstrained by gauge identity. The index (w, s, c) therefore separates into a gauge-identity coordinate read by χ and a refinement-depth coordinate that χ cannot read. The mass map M_{w,s}(C(c)) *may* read the second coordinate — that availability is what makes a hierarchy possible — but whether it does, and with what ordering, is supplied by ANCH-COMP, not by this factorization.

What this proves and what it does not. The charge-side factorization is true the moment χ is written as a function of (w, s) alone — it is a decomposition theorem, not a dynamical discovery. Its force is precise and limited: it is the exact statement of why replicated families can be gauge-identical, and it is exactly what the operator relation $[\hat{Y}, \chi] = 0$ (§5a) requires. It does *not* establish that masses differ, nor that they differ in completion-density order; those are the conditional theorems Y2/Y2′. Read correctly, F1 is the conceptual spine for *charge invariance* *under replication*, and the mass hierarchy hangs off it only through the separate, conditional bridge. The replication, charge-invariance, and (conditionally) hierarchy results are all corollaries of this one separation: replication because χ is blind to the coordinate that distinguishes species; charge-invariance because that blindness is exact; hierarchy because the coordinate χ cannot read is nonetheless available to mass — the last step requiring ANCH-COMP.

---

5. Yukawa Couplings as Effective Readouts

In the Standard Model, a fermion mass is written as:

$m_f = y_f v / \sqrt{2}$.

The Higgs vacuum expectation value v is common. The mystery is y_f.

For each fermion species f the Standard Model simply assigns a number y_f. These numbers span a huge range. The top quark's Yukawa coupling is of order one. The electron's is tiny. The neutrino sector is smaller still, depending on the mass mechanism. Nothing in the Standard Model explains the pattern.

VERSF proposes that y_f is not fundamental. It is an effective parameter summarising how a charge-carrying mode couples to the electroweak mass-generating channel through its completion structure.

The proposed form is:

$y_{(a,c)} = Y_a \cdot A(c)$,

where:

- a labels the within-family sector: up-type, down-type, charged lepton;
- c labels refinement class / generation;
- Y_a is the sectoral anchor for that type of fermion, read from (w, s) data;
- A(c) is the refinement-depth amplification factor, the image of the completion-density ordinal C(c) under the anchor–completion bridge.

This separates two things the Standard Model blends together:

1. **Sector difference:** why up-type, down-type, and charged leptons sit at different mass scales. Carried by Y_a.
2. **Generation difference:** why generations 1, 2, 3 form an ascending hierarchy. Carried by A(c).

This is not yet a numerical derivation. It is the structural decomposition required before numerical derivation becomes possible. Note that A is a mass-side quantity: it is C(c) pushed through ANCH-COMP, not C(c) itself.

5a. The Refinement-Depth Amplification Operator

The factor A(c) is written above as a scalar function. Its proper home is an operator — but an operator introduced by fiat, with eigenvalues A(c) declared, would be a relabeled diagonal and would do no work. The stronger construction makes the amplification operator *emerge* from structure the programme already has, rather than positing it independently.

The programme already carries a completion-density operator $\mathcal{C}$, diagonal in refinement class, returning the inherited completion-density ordinal on each class:

$\mathcal{C}\,|w, s, c\rangle = C(c)\,|w, s, c\rangle$.

This is not new to this paper; C(c) is Input D. The amplification operator is then *defined as a function of it*:

$\hat{Y} = F(\mathcal{C})$,

where F is the ANCH-COMP map — positive and monotone — carrying the completion-density ordinal to the mass-amplification factor. Acting on the basis,

$\hat{Y}\,|w, s, c\rangle = F(C(c))\,|w, s, c\rangle = A(c)\,|w, s, c\rangle$,

so the eigenvalue A(c) is not a new postulate but the image F(C(c)) of an already-existing spectral quantity. ANCH-COMP *is* the statement that F is positive and monotone; the operator inherits exactly that and nothing more.

Its properties now follow from those of $\mathcal{C}$ and F rather than being asserted:

- **Positive:** $F > 0$, so $A(c) = F(C(c)) > 0$ and $\hat{Y}$ is a positive operator. Mass amplification cannot be zero or negative.
- **Monotone:** F is monotone and the spectrum of $\mathcal{C}$ is ordered, so $C(\sigma_i) < C(\sigma_j) \Longrightarrow A(\sigma_i) < A(\sigma_j)$. The eigenvalues of $\hat{Y}$ inherit the order of the eigenvalues of $\mathcal{C}$.
- **Refinement-sensitive:** C(c) is non-constant across classes and F is order-preserving, so $\hat{Y}$ is not proportional to the identity.
- **Charge-blind:** because $\hat{Y}$ is a function of $\mathcal{C}$, and $\mathcal{C}$ is diagonal in c while every charge operator χ is a function of (w, s) alone, $\hat{Y}$ commutes with every χ:

  $[\hat{Y}, \chi] = [F(\mathcal{C}), \chi] = 0$.

The commutator is not decoration. It is the operator form of the entire thesis: the operator that generates the mass hierarchy is a function of completion density, the operators that read charge are functions of (w, s), and the two coordinate sets are disjoint — so refinement depth can be amplified without disturbing any charge. The charge-side Factorization Theorem (4.1) is exactly the assertion that χ is a function of (w, s) alone; $[\hat{Y}, \chi] = 0$ is its immediate operator consequence once $\hat{Y}$ is built from $\mathcal{C}$. The scalar decomposition $y_{(a,c)} = Y\_a \cdot A(c)$ is the matrix element of $\hat{Y}$, with the sectoral anchor Y_a the (w, s)-side coefficient $\hat{Y}$ does not touch.

The commutator forbids a specific failed theory — and that is a test already passed. $[\hat{Y}, \chi] = 0$ is not merely a restatement; it has generative content, because the negation is a concrete, experimentally excluded theory. Any model in which the generation-distinguishing operator *failed* to commute with charge — $[\hat{Y}, \chi] \neq 0$ — would have $\hat{Y}$ and χ non-simultaneously-diagonalizable, so amplifying refinement depth would shift charge eigenvalues. Such a model predicts **generation-dependent charge**: the muon would carry a different electric charge than the electron, the tau different again. That is the single most obvious wrong answer a theory of generations can give, and it is excluded to the limits of measurement — the electron and muon (and tau) carry identical electric charge as far as any experiment has probed. So $[\hat{Y}, \chi] = 0$ is the structural reason VERSF *cannot* produce that wrong answer: charge universality across generations is, in this framework, a **theorem rather than a coincidence**, and any non-commuting amplification is immediately falsified by the measured equality of electron and muon charge. This is the strongest claim form available to a foundations paper — a prediction (charge universality across generations) that is both forced by the structure and already confirmed — and it costs nothing in grade, since the commutator is structural (Theorem 5.1) and the empirical fact is among the most precisely established in physics.

Theorem 5.1 — Charge-Blindness of Amplification. Let $\hat{Y} = F(\mathcal{C})$ with $\mathcal{C}$ the completion-density operator (diagonal in refinement class) and F positive and monotone (ANCH-COMP). Then $\hat{Y}$ is positive, monotone, and refinement-sensitive, and $[\hat{Y}, \chi] = 0$ for every charge operator χ of the gauge factor.

Grade: Proven (structural), given the charge maps of Input A and the existence of $\mathcal{C}$ (Input D). The *form* of F and the *spacing* of the eigenvalues are not fixed here.

Proof. $\mathcal{C}$ is diagonal in c (Input D); any function $F(\mathcal{C})$ is therefore also diagonal in c with eigenvalues F(C(c)). Each χ is a function of (w, s) alone (Input A), acting within fixed (w, s) without moving c, while $F(\mathcal{C})$ acts within fixed c without moving (w, s). On the basis $|w, s, c\rangle$ both are diagonal, hence $F(\mathcal{C})\chi = \chi F(\mathcal{C})$, i.e. $[\hat{Y}, \chi] = 0$. Positivity, monotonicity, and refinement-sensitivity transfer from F (positive, monotone, by ANCH-COMP) composed with the ordered, non-constant spectrum of $\mathcal{C}$.

Because $\hat{Y}$ is built from $\mathcal{C}$ rather than posited, its *existence and charge-blindness* are structural — they need only that $\mathcal{C}$ exists and that F is some positive monotone function. What remains conditional is the *content* of F: that it is monotone is ANCH-COMP (Conditional), and its precise form, hence the eigenvalue spacing, is open numerically (Y4).

5b. Relation to the Higgs Mechanism

A clarification forecloses a likely misreading. VERSF does not replace the Higgs mechanism, and it does not compete with electroweak symmetry breaking as the origin of mass. The mass formula $m_f = y_f v / \sqrt{2}$ stands unaltered: the Higgs vacuum expectation value v supplies the scale, and electroweak symmetry breaking supplies the mechanism by which the Yukawa coupling becomes a mass.

What VERSF addresses is the one ingredient the Higgs mechanism leaves undetermined — the coupling y_f itself. In the Standard Model the Yukawa couplings are free parameters that the Higgs sector requires but does not predict. VERSF reads them as the matrix elements $Y_a \cdot A(c)$ of the amplification operator $\hat{Y} = F(\mathcal{C})$: the Higgs mechanism converts those couplings into masses, and VERSF supplies the structural origin of the couplings.

This is a stronger and more defensible position than displacing the Higgs sector would be. The claim is not that the Higgs mechanism is wrong but that its free Yukawa inputs are not fundamental — they are effective readouts of refinement depth, fed into an otherwise unchanged electroweak mass-generation channel.

5c. Relation to the Anchoring Foundation, and a Vocabulary Correction

The amplification operator $\hat{Y} = F(\mathcal{C})$ is built here on the completion-density operator $\mathcal{C}$ (Input D) and a positive monotone bridge F (ANCH-COMP, Input E). Two companion papers in the programme attempt to supply F from below — the void-anchoring mass-scale paper and the phase-coherence paper. This subsection records what they provide, at honest grade, and corrects a vocabulary collision that would otherwise make this paper appear to contradict them.

The vocabulary correction (the word "deep"). The anchoring paper's mass-tracking quantity is the flip-completion density

$$C = p_v / K_c,$$

with mass $m \propto C$, where p_v is the per-tick coupling probability and K_c is the anchoring depth (micro-events per committed bit-flip). In that paper's vocabulary, large K_c means *deep* anchoring — many ticks per flip — and therefore *lower* mass; the anchoring paper states this in bold as a caution. Heavier means *higher completion density C*, i.e. *shallower* K_c.

This paper's ordering coordinate is that same completion density: the $\mathcal{C}$ of $\hat{Y} = F(\mathcal{C})$ is the flip-completion density p_v/K_c, heavier = higher C. The collision is on the word "deep." Where §6, §7, and §7a of this paper call a heavier generation "deeper," the intended meaning is *higher completion density* — shallower in the anchoring paper's K_c sense, more flips per tick — and **not** larger anchoring depth K_c. Read with the anchoring paper's K_c convention, "deepest =

heaviest" would invert that paper's central caution. The physics reconciles cleanly under the identification $C \equiv p\_v/K\_c$; only the word "deep" must be read as "high completion density," and is so read throughout this paper.

What F is, under this identification. With $C = p\_v/K\_c$, the anchoring paper's central conditional result is $m = \eta\hbar/(c^2\Delta t) \cdot C$ — mass proportional to the completion-density coordinate up to imported constants. The coherence paper supplies the remaining substructure, $p\_v = g \cdot p\_align(R)$, with R a phase-coherence order parameter and g a chirality-class gate. Composing, F is essentially the map $C \mapsto m$, which is the near-identity up to the universal prefactor $\eta\hbar/(c^2\Delta t)$. This is good for ANCH-COMP's monotonicity — a near-identity is trivially positive and monotone — but it is honest to note the cost: if A(c) is proportional to the completion-density eigenvalue C(c), then $\hat{Y} = F(\mathcal{C})$ is close to F = identity dressed by constants, and the operator's content reduces to "mass is proportional to the completion-density eigenvalue." The operator framing buys the clean commutator $[\hat{Y}, \chi] = 0$ (§5a) and nothing dynamical beyond it; the dynamical content lives in what sets C, which is exactly what the companion papers defer to their unexecuted eigenmode programs.

Grade — Conditional, and doubly so. Citing these papers does not upgrade ANCH-COMP toward Proven. The anchoring paper is explicit that its mass scale is conditional on the action postulate $P\hbar$ ("the primary load-bearing assumption … not derived"), on the calibration $\Delta t = t\_P$ ("a calibration assumption, not a derived result"), and on imported relations ($\hbar$, c, $E = mc^2$); and that no particle mass is predicted from first principles. The coherence paper is likewise conditional on a multiplicative-alignment postulate and an unexecuted eigenmode computation. So the most this paper may claim is: *a candidate dynamical realization of the completion-density coordinate exists — $m \propto p\_v/K\_c$, itself Conditional on $P\hbar$, the Layer B bridge, and the Δt calibration — under which ANCH-COMP's monotone map is the near-identity $m \propto C$.* Input E thereby moves from a bare bridge to a bridge with a candidate mechanism whose own grade is Conditional-on-a-postulate. It does not move to Proven, and Result 7.1 below does not depend on the mechanism being correct.

Two cross-references worth flagging. First, the neutrino sector held back in §6 acquires a consistent candidate across both companion papers: ultra-light neutrino mass as high coherence ($R \approx 1$) with strong chirality gating ($g \ll 1$), so p_v/K_c is suppressed by g rather than by loss of coherence. This is the same mechanism in both, and it is why the neutrino is excluded from the charged-sector ordering of §6 rather than assumed to follow it. Second, the two companion papers do **not** yet agree on the functional form of the hierarchy: the coherence paper gives $m \propto g \cdot q^{N\_eff(K\_c)}/K\_c$ (exponential-of-a-power), while the anchoring paper's scale-invariance argument gives a pure power law $m \propto K\_c^{-(\alpha+1)}$. Reconciling these two forms is an open problem for the foundation papers, not for this one. Result 7.1 (§7) survives either way, because what it establishes — that the sectoral sensitivity must be sector-dependent, since the shared completion-depth spacing cannot carry the cross-sector divergence — constrains both forms equally and excludes the single-λ geometric law that both already reject.

Token hygiene. The coherence paper uses ε for its phase-alignment window; this paper uses $\varepsilon_{(a,c)}$ for the §9 off-diagonal correction. These are unrelated quantities and must not be conflated across papers — a future cross-reference should retoken one of them. The symbol K_c is shared and consistent across all three papers (anchoring depth, micro-events per flip), as is the completion density C = p_v/K_c. The class gate g of the companion papers has no separate token here; it is absorbed into the sectoral anchor Y_a together with the universal prefactor, with A(c) carrying the completion-density factor.

5d. Refinement Efficiency and the Origin of the Generation Ordinal

The preceding sections treat completion density as the quantity read by the mass side of the theory. A remaining question is why completion density should be ordered by generation depth rather than inserted as a separate ordinal assumption. This is precisely the upstream gap left Open in Input D and §6a: the ordinal $C_1 < C_2 < C_3$ has so far been assumed, not derived. This section narrows that gap by supplying an intermediate structural quantity — refinement efficiency — whose ordering *is* derived, from the dyadic refinement architecture the corpus already carries.

The refinement papers establish that depth is not a geometric direction but a record of realised admissible distinction. Under binary refinement and uniform single-Fold individuation, the admissible generation loads are dyadic:

1, 2, 4.

Within the closure register of capacity

K = 7,

the cumulative occupancies are therefore

1, 3, 7.

Define the cumulative refinement efficiency of generation depth c by

$E(c) = L_{(\leq c)} / K$,

where $L_{(\leq c)}$ is the cumulative realised load up to depth c, and K is the available closure capacity. For the dyadic sequence inside the K = 7 closure register,

$E_1 = 1/7$, $E_2 = 3/7$, $E_3 = 7/7 = 1$.

Hence

$E_1 < E_2 < E_3$.

The third generation is therefore the *saturation generation*: it does not introduce a new gauge identity, but it completes the admissible refinement register. The saturation reading is exact — $E_3 = 1$ — rather than a quantity exceeding unity or diverging; the third generation fills the register, it does not overrun it.

Theorem 5.7 — Refinement-Efficiency Ordinal

Assume the inherited dyadic loading structure (1, 2, 4) and the inherited closure capacity $K = 7$. Then cumulative refinement efficiency increases strictly across the three admissible generation depths:

$$E_1 < E_2 < E_3.$$

Grade: Proven (arithmetic), conditional on the inherited dyadic loads and closure capacity. The loads (1, 2, 4) and the capacity $K = 7$ are not derived here; they are inherited from the refinement/census arc and are cited, not re-proven. Given them, the ordinal is arithmetic.

Proof. The cumulative loads are 1, $1 + 2 = 3$, and $1 + 2 + 4 = 7$. Dividing by the fixed closure capacity $K = 7$ gives

$$1/7 < 3/7 < 7/7.$$

Thus cumulative refinement efficiency is strictly ordered across the three admissible depths.

Corollary 5.8 — From Efficiency Ordinal to Completion-Density Ordinal

Efficiency–Completion Bridge (ECB). If completion density is a positive monotone function of cumulative refinement efficiency, then the generation completion densities are strictly ordered:

$$C_1 < C_2 < C_3.$$

Grade: Conditional, on ECB. This is a *second* bridge, distinct from ANCH-COMP, and it is named so the dependency is explicit rather than buried: ECB links the efficiency ordinal (Proven, Theorem 5.7) to the completion-density ordinal; ANCH-COMP then links the completion-density ordinal to mass (Corollary 5.9). ECB is not established here.

It is worth stating precisely what ECB asks, because the named form is more defensible than the old bare assumption it replaces. Efficiency $E(c) = L_{(\leq c)}/K$ is a register-occupancy fraction — the share of the $K = 7$ closure register filled by generation depth c — while completion density $C = p\_v/K\_c$ is a per-tick flip rate. These are different quantities, and ECB is the claim that the second increases with the first. The question a reviewer should press is therefore: *why should completion density increase with the fraction of the closure register already saturated?* That is a sharper and more tractable question than the one the paper faced before §5d ("why should completion density be ordered at all?"), because it asks for monotonicity against a quantity that already measures progressive saturation, rather than asking for an ordering out of nothing. ECB is not discharged by being named; but naming it isolates the single remaining structural assumption on this link and states it in a form where the burden is to argue a co-monotonicity, not to posit an order.

Corollary 5.9 — From Completion-Density Ordinal to Mass Ordinal

If ANCH-COMP holds, so that the mass-side amplification is a positive monotone function of completion density, then

$m_1 < m_2 < m_3$.

Grade: Conditional, on ANCH-COMP (and, per §6a, on the anchoring relation $m \propto C$ and a bounded ε).

The layered chain

This separates the hierarchy argument into three distinct steps:

dyadic refinement ⟹ efficiency ordinal ⟹ completion-density ordinal ⟹ mass ordinal.

The first step is arithmetic once the inherited (1, 2, 4) loads and $K = 7$ capacity are granted (Theorem 5.7). The second step is the Efficiency–Completion Bridge (ECB, Corollary 5.8, Conditional). The third is ANCH-COMP plus the anchoring relation (Corollary 5.9, Conditional). So the dependency graph is transparent: Proven (Y7) → Conditional (ECB) → Conditional (ANCH-COMP) → conclusion. This is stronger than treating $C_1 < C_2 < C_3$ as a bare input, because the ordinal now has a structural origin in the refinement/capacity architecture rather than appearing as an assumption — and the one remaining structural assumption on the upstream side, ECB, is named rather than hidden. Two honesty constraints remain explicit. First, the chain is Proven only at its first link; the two bridges are Conditional, and the result should not be quoted as a derived mass ordering. Second, the relation of this efficiency-derived ordinal to the *per-sector* ordering required by §6a is not yet settled: Theorem 5.7 supplies a single ordinal over refinement depth, while §6a (via Result 7.1) requires an ordered triple *per charged sector* carried by sector-dependent $p_{v,a}(c)$. Whether the single efficiency ordinal induces the three per-sector orderings, or whether the sectoral sensitivity perturbs them, is part of the upstream problem the eigenmode programme must close. What §5d establishes is that the *depth* ordinal is no longer bare; what it does not establish is that the depth ordinal alone fixes every sector's completion-density triple.

---

6. The Species–Mass Correspondence

Let the durable species be ordered by completion density:

$\sigma_1 < \sigma_2 < \sigma_3$,

where σ_1 has the lowest completion density and σ_3 the highest. This is a statement about the spectral quantity C.

The anchor–completion bridge converts that spectral order into a mass-amplification order. Because $C \mapsto A$ is monotone,

$C(\sigma_i) < C(\sigma_j)$ implies $A(\sigma_i) < A(\sigma_j)$.

The species–mass correspondence then states:

If σ_i has lower completion density than σ_j, then the effective Yukawa readout of σ_i is lower than that of σ_j, all else equal.

Therefore, for a fixed fermion sector a, and subject to sectoral corrections and mixing effects:

$y_{(a,1)} < y_{(a,2)} < y_{(a,3)}$.

This gives the qualitative pattern:

electron < muon < tau, up < charm < top, down < strange < bottom.

These three ladders are all charged sectors. The neutral-lepton sector is held back deliberately. The roster $L = (1, 2)_{-1/2}$ contains the left-handed neutrino, so a neutral-lepton seat is present at each refinement depth, but the mass formula $m_f = y_f v / \sqrt{2}$ may not be the operative one there: depending on whether neutrino mass is Dirac, Majorana, or seesaw, the relevant operator and the meaning of "Yukawa readout" differ. The within-sector ordering claim below is therefore asserted for the three charged sectors only; the neutrino sector is excluded pending its mass mechanism (see §13), not silently assumed to follow the same ladder.

The paper does not claim that each inequality is numerically exact without corrections. It claims that the mass ordinal is structurally supplied by the species ordinal — completion density (spectral) ordering them, ANCH-COMP carrying that order to anchoring strength (mass).

Theorem 6.1 — Species–Mass Correspondence

Under the anchor–completion bridge, completion-density order supplies the generation mass order.

Grade: Conditional on ANCH-COMP.

Proof. The spectral arc orders durable species by completion density C. The Replication and Hypercharge Completion results identify each durable species with a complete Standard Model family. The anchor–completion bridge maps C monotonically to the mass-amplification factor A. The ordering of durable species by completion density therefore induces the ordering of generations by mass.

6a. Mass Ordering Within a Sector: Downstream Proven, Upstream Open

Theorem 6.1 gives the correspondence at the level of the ordinal. Sharpening it into a strict inequality on the masses themselves requires separating two steps that earlier drafts ran together under a single "Conditional on ANCH-COMP." They have different grades, and conflating them mislabels where the risk lives.

The two steps are:

- **Downstream:** *given* the completion-density ordering $C_1 < C_2 < C_3$, derive the mass ordering $m_1 < m_2 < m_3$.
- **Upstream:** establish $C_1 < C_2 < C_3$ in the first place.

The downstream step is Proven (modulo the anchoring relation and ε). The upstream step is Open. The risk is entirely upstream, and it is not a risk that the *bridge* fails — it is that the *ordering of the species by completion density* is uncomputed.

Theorem 6.2 (downstream) — Mass ordering from completion-density ordering. Fix a charged fermion sector a. Assume the anchoring relation $m \propto C$ of §5c (so within the sector $m_{(a,c)} = Y\_a \cdot A(c) \cdot v/\sqrt{2}$ with $A(c) \propto C\_a(c)$, positive constants) and assume the off-diagonal correction $\varepsilon_{(a,c)}$ of §9 is below the gaps it would close. Then

$$C\_a(1) < C\_a(2) < C\_a(3) \Longrightarrow m_{(a,1)} < m_{(a,2)} < m_{(a,3)}.$$

Grade: Proven, modulo (i) the anchoring relation $m \propto C$, itself Conditional on the action postulate $P\hbar$ and the Δt calibration (§5c), and (ii) ε below the gaps. No separate monotonicity hypothesis is needed.

Proof. By the anchoring relation, within sector a the mass is a positive constant times the completion density: $m_{(a,c)} = \kappa\_a \cdot C\_a(c)$ with $\kappa\_a = (\eta\hbar/c^2\Delta t) \cdot$ (sectoral/EWSB factors) > 0 and independent of c. Multiplication of a strictly ordered positive sequence $C\_a(1) < C\_a(2) < C\_a(3)$ by a positive c-independent constant is strictly order-preserving, so $m_{(a,1)} < m_{(a,2)} < m_{(a,3)}$ — provided the dropped ε does not exceed the gaps. Monotonicity is *automatic* here: it is not assumed of the bridge, it is the elementary fact that scaling by a positive constant preserves order. This is exactly the §5c finding that F ≈ identity-up-to-constants doing useful work — the near-identity is precisely what makes this step Proven rather than Conditional-on-monotonicity.

The upstream problem (Open, but partially derived). What the theorem does *not* establish is its own hypothesis: $C\_a(1) < C\_a(2) < C\_a(3)$. That ordering was, in earlier drafts, bare Input D. It is now partially derived: §5d (Theorem 5.7) supplies a refinement-efficiency ordinal $E_1 < E_2 < E_3$ from the inherited dyadic loads in the $K = 7$ register, and the Efficiency–Completion Bridge (ECB, Corollary 5.8) carries that to the completion-density ordinal. So the *depth* ordinal has a structural origin. What remains Open is twofold. First, ECB is itself Conditional. Second, §5d gives a single ordinal over refinement depth, whereas the result below needs an ordered triple *per charged sector*: whether the VERSF interface dynamics force

$p_{v,1}/K_{c,1} < p_{v,2}/K_{c,2} < p_{v,3}/K_{c,3}$

in *each* sector is the output of the eigenmode programme (anchoring paper §16, future work). Because the sectoral sensitivity forced by Result 7.1 most naturally lives in a sector-dependent $p_{v,a}(c)$, the completion density is in general per-sector, $C_a(c) = p_{v,a}(c)/K_c(c)$, with K_c shared across sectors. So the upstream target is sharp: show that the §5d depth ordinal induces, for each charged sector, an ordered triple $(C_a(1), C_a(2), C_a(3))$ — three ordered triples, not one, and the link from the single depth ordinal to the three is not established here.

Scope of the downstream result.

1. **Within charged sector only.** The theorem orders masses *inside* each charged sector. It does not order masses *across* sectors: the tau is heavier than the up quark, but A(c) says nothing about that — cross-sector scale is carried by Y_a, not by $\hat{Y}$. It is silent on the neutral-lepton sector (excluded in §6). Read globally the statement is false; read per-charged-sector it follows.
2. **ε below the gaps.** The ordering drops the §9 correction $\varepsilon(a,c)$. A sufficiently large ε could perturb it; the result holds only where ε is below the gaps it would close. That bound is *assumed*, not established.
3. **Anchoring relation assumed.** The step $m \propto C$ is Conditional on $P\hbar$ and $\Delta t = t_P$ (§5c). The downstream theorem inherits that conditionality and nothing more — in particular it does *not* additionally assume monotonicity, which is free.

The honest one-line summary: *the mass ordinal follows from the completion-density ordinal (Proven, modulo Pħ and ε); whether the three species are actually ordered by completion density is Open, pending the eigenmode programme.*

---

7. The Spacing Constraint on the Third Generation

The most conspicuous feature of the Yukawa spectrum is not merely that masses increase. It is that the third generation is disproportionately heavy, especially in the up-type quark sector. This is not a remark inviting speculation; it is a constraint, and it can be stated as one.

Consider the simplest candidate amplification: a one-parameter exponential

$A(c) = \exp(\lambda \cdot K(c)),$

with K(c) linear in the generation index and λ a single sector-independent constant. This predicts a *geometric* mass spectrum — equal ratios between successive generations within a sector, and the same ratios across sectors. The observed spectrum is neither. The approximate within-sector ratios are:

- charged leptons: $\mu/e \approx 207$, then $\tau/\mu \approx 17$;
- up-type quarks: $c/u \approx 580$, then $t/c \approx 140$;
- down-type quarks: $s/d \approx 20$, then $b/s \approx 45$.

Two facts read directly off these numbers. First, within each sector the successive ratios are non-constant (207 then 17; 580 then 140; 20 then 45). Second, across sectors the pattern differs — the leptons and up-type quarks show the larger jump at the first step, the down-type quarks at the second. Either fact alone refutes a constant-ratio spectrum.

Result 7.1 — Sectoral sensitivity is forced (a determination, not a disjunction). Within the exponential schema $A_a(c) = \exp(\lambda_a \cdot K(c))$, and with the §9 correction ε below the gaps, the observed spectrum forces the sectoral sensitivity λ_a to be sector-dependent. Non-uniform completion-depth spacing $K(c)$ is neither sufficient on its own nor forced. The two candidate carriers of the irregularity are:

(i) non-uniform completion-depth spacing $K(c)$ across refinement classes; (ii) sector-dependent sensitivity λ_a;

and the result is that (ii) is forced while (i)-alone is excluded.

Grade: Proven (follows from the measured ratios plus the architectural fact that c is one shared index), as a negative-plus-determination result, within the exponential schema and the ε-below-gaps regime.

Proof. The key architectural fact is that $K(c)$ — the completion-depth spacing — is a property of the refinement class c, and c is a single index shared across all sectors (the generations are one index; this is the whole construction of the paper). So $K(c)$ is sector-independent *by construction*: the same three values $K(1)$, $K(2)$, $K(3)$ enter every sector.

Now suppose, for contradiction, that the irregularity were carried by (i) alone — non-uniform $K(c)$ with $\lambda_a = \lambda$ uniform across sectors. Then within sector a the successive mass ratio is

$$m_{(a,c+1)}/m_{(a,c)} = \exp(\lambda \cdot [K(c+1) - K(c)]),$$

which depends on c only through the sector-independent quantity $\lambda \cdot \Delta K(c)$. Every sector would therefore exhibit the *same* sequence of jumps — the same ratio shape — because both λ and $\Delta K(c)$ are sector-independent. But the data show different shapes: leptons and up-type quarks peak at the first step ($\mu/e \approx 207 > \tau/\mu \approx 17$; $c/u \approx 580 > t/c \approx 140$), while down-type quarks peak at the second ($s/d \approx 20 < b/s \approx 45$). A single shared shape cannot be simultaneously front-peaked and back-peaked. Contradiction. Hence (i)-alone is impossible: non-uniform K with uniform λ cannot produce cross-sector divergence in the ratio shape.

Removing the contradiction requires the c-dependence of the ratio to differ by sector. With K(c) shared, the only remaining sector-dependent handle in the schema is λ_a. Therefore λ_a must be sector-dependent: (ii) is forced. Non-uniform K(c) may additionally be present, but it is neither sufficient alone (just shown) nor required (a sector-dependent λ_a with uniform K already breaks both the within-sector constancy and the cross-sector agreement). This collapses the earlier disjunction to a determination.

Scope — where the determination holds, and its one escape. The determination is conditional on two things, and both must be stated or it is false. First, the **exponential schema**: the proof tests $A_a(c) = \exp(\lambda_a K(c))$; it forces sector-dependent λ_a *within that schema*, which is the schema under examination. Second, and the genuine escape hatch, **ε below the gaps**: the §9 off-diagonal correction $\varepsilon_{(a,c)}$ is sector-and-generation dependent, so if ε is doing real work it can manufacture different cross-sector ratio shapes even with uniform λ and shared K. In that case the irregularity could live in ε instead, and the conclusion weakens to "sector-dependent λ_a *or* a non-negligible ε." So the honest determination is:

Within the exponential schema and with ε below the gaps, sector-dependent λ_a is forced; non-uniform K is neither sufficient alone nor required. Outside the ε-below-gaps regime, the fork is "sector-dependent λ_a or non-negligible ε."

What this hands the next paper. This is a strictly stronger result than a disjunction over (i) and (ii), and it sharpens the handoff to the §9 sectoral-anchor paper into a positive instruction: the law replacing the uniform exponential must carry its generation-to-generation irregularity in a *sector-dependent sensitivity* (the natural home of which is a sector-dependent $p_{v,a}(c)$, per Input D and §6a), *or* in the §9 correction ε — and those are the only two doors. The completion-depth spacing K(c), being shared across sectors, cannot by itself be the source of the cross-sector divergence.

The third generation is heavy because it sits at the highest completion density among the stable replicated families — the largest p_v/K_c, hence the largest effective Yukawa readout (see §5c for the completion-density convention and the caution against reading "deep" as large anchoring depth K_c). The top quark is then not an arbitrary large entry in the Yukawa table but the up-type readout of the highest stable refinement class — and Result 7.1 fixes that this readout's *sectoral* peculiarity (why the up-type jump structure differs from the down-type) cannot come from the shared depth spacing, but must come from the sectoral sensitivity or the §9 correction.

7a. Why the Hierarchy Terminates at Three

The reader who has followed the ladder asks immediately: if higher completion density means heavier mass, why does the ladder stop at three rungs rather than four? The honest answer connects three results without claiming to derive the count here.

The generation count is **not proven in this paper.** It is Input D, inherited from the species-count theorem of the spectral arc, which fixes the number of *stable* durable refinement classes at three.

This paper takes that count as given; what it adds is the reason the ladder it sits on terminates physically as well as structurally. Two distinct supports converge, and they should be kept distinct rather than blended into a single "therefore three."

The internal support (structural). The inherited species-count theorem admits three stable refinement classes. A fourth would have to be a fourth *stable* depth in the same catalogue, and the count says there is none. This is the structural terminus, and it is imported, not derived here.

The external support (empirical). Even granting the question of whether a fourth stable depth could exist, §11 shows that any fourth species satisfying the species–family identification would be a chiral, sequential fourth generation. It would be expected heavier than the third *under continued monotone anchoring* — that is, if the monotonicity that orders the known three rungs extends to a fourth. This is an extrapolation, not a consequence of Theorem 6.2: 6.2 is established over the inherited three classes, and asserting its hypothesis (positive monotone λ_a) persists to a fourth rung whose anchor behaviour is not in evidence is an additional assumption. With that caveat, a chiral sequential fourth family is tightly constrained by Higgs production and electroweak precision data. So the completion-density mass ladder, far from hiding a fourth family, would on the natural extrapolation place it heavy and visible, exactly where experiment most strongly excludes it.

The two supports are different in kind — one is a structural fact about the inherited catalogue, the other an empirical bound on what a further rung would imply — and the paper claims only their conjunction, not a first-principles derivation of three. Read together, replication (why families repeat), hierarchy (why they ascend in mass), and termination (why the ascent stops at three) form one connected picture: the same completion-density coordinate that orders the masses also runs out of stable rungs at the third, and the experiment closes off the fourth from the other side.

---

8. Why Charges Do Not Change When Masses Do

This is the conceptual payoff.

Ordinary intuition might expect that if the third generation is structurally deeper, it should also differ in charge. But that would destroy family replication. VERSF avoids this because the charge map does not see the coordinate the mass map uses.

Charge is a horizontal identity within the gauge bundle. Mass is a vertical readout that additionally reads completion depth.

The same winding/seat slot repeated at different refinement depths gives:

- same electric charge;
- same color representation;
- same weak representation;

- different completion-density anchoring;
- different effective Yukawa coupling.

Thus:

electron, muon, tau = same charged-lepton slot, three refinement depths.

up, charm, top = same up-type quark slot, three refinement depths.

down, strange, bottom = same down-type quark slot, three refinement depths.

The hierarchy does not break replication because the charge map is blind to the refinement coordinate it acts on.

This is the deep reason the paper matters. It explains why the Standard Model's repeated families can be identical under the forces while radically different in mass.

---

9. Sectoral Structure: Up, Down, Charged Lepton

The preceding sections explain the generation ordering. They do not yet explain why the up-type, down-type, and charged-lepton sectors sit at different scales. That requires sectoral anchors.

The proposed decomposition is:

$y_{(a,c)} = Y\_a \cdot A(c) \cdot \varepsilon_{(a,c)}$,

where:

- Y_a is the sectoral anchor, read from (w, s);
- A(c) is the generation/refinement amplification, the image of C(c) under ANCH-COMP;
- $\varepsilon_{(a,c)}$ is a correction term from mixing, channel leakage, or incomplete diagonalization.

The role of this paper is primarily to establish A(c), the generation hierarchy factor. The sectoral anchors Y_a remain partly inherited and partly open.

This division matters. Without it, the paper would be forced to explain too much at once. The correct first target is not the full numerical Yukawa matrix. It is the existence of a repeated mass hierarchy across sectors.

The programme can then proceed:

1. this paper: generation hierarchy;
2. next paper: sectoral anchors;
3. later paper: mixing matrices and off-diagonal terms.

---

10. Yukawa Matrices and Diagonal Approximation

In the Standard Model, Yukawa couplings are matrices. After electroweak symmetry breaking and diagonalization, they produce physical masses and mixing.

A complete VERSF treatment must therefore eventually explain not only the diagonal hierarchy but also the off-diagonal structure that gives CKM and PMNS mixing.

This paper works in the diagonal approximation. It treats each sector as though the dominant mass readout is:

$Y_a \cdot A(c)$,

with off-diagonal terms temporarily neglected. This is justified at this stage because the question being addressed is the existence and ordering of the hierarchy, not the exact mixing angles.

In matrix form, the leading approximation is:

$Y_a \approx Y_a^0 \cdot \mathrm{diag}(A_1, A_2, A_3)$, with $A_1 < A_2 < A_3$.

The next correction would allow small off-diagonal terms:

$Y_a \approx Y_a^0 \cdot [D + \Delta]$,

where D is the diagonal completion-density matrix and Δ encodes cross-refinement leakage or overlap between completion channels.

CKM and PMNS mixing would then arise not from arbitrary matrix entries, but from the mismatch between the diagonalization bases of different sectoral completion matrices. That is not derived here. It is the next natural problem.

---

11. Fourth Generation Constraint Revisited

The Replication Theorem showed that a fourth durable species, if it exists and satisfies the same species–family identification, would be a fourth sequential generation.

This paper sharpens the interpretation. If the mass hierarchy follows completion density and the monotone anchoring of the known three classes continues to a fourth, a fourth species would not merely be a charge-identical fourth family; it would be expected heavier, sitting beyond the third refinement depth in completion-density order. The "expected heavier" is an extrapolation of the §6a monotonicity past the classes over which it is established, and is flagged as such.

So World B — the existence of a fourth durable species — would imply:

- a fourth complete anomaly-free family;
- the same gauge charges as the known generations;
- higher completion-density anchoring (by extrapolation of monotone anchoring);
- under that extrapolation, heavier fermions.

That is exactly the kind of chiral sequential fourth generation already strongly constrained by experiment. A chiral fourth family enhances Higgs production through gluon fusion by the heavy extra quark doublet and shifts the electroweak precision observables; both are tightly bounded by the data. The constraint bites here precisely because the World-B species is chiral and sequential, not vector-like — a vector-like family would evade these particular bounds, but it is not what replication-plus-completion predicts.

The Yukawa hierarchy therefore does not rescue the fourth species. It sharpens the problem. A fourth species would be expected to be heavy, chiral, and sequential, not hidden by charge alteration. The internal fourth-mode question remains open mathematically but externally pressured physically.

---

12. Non-Circularity

The argument is not circular because the three key ingredients are independent.

First, hypercharge completion is a within-family check. It asks whether the inherited roster and charges are the Standard Model family (they are, by assignment) and whether they pass the anomaly conditions (they do). The identification rests on the roster and charges; the cancellation is the consistency confirmation, not the certifier.

Second, replication is an across-family charge-identity result. It asks whether different refinement classes carry the same charges.

Third, mass hierarchy is an across-family completion-density result. It asks whether different refinement classes have different anchoring strengths.

None of these implies the others:

- a chiral set could be anomaly-free without being the Standard Model family (cancellation is necessary, not sufficient) and without being replicated;
- bundles could be replicated without having different masses;
- masses could differ without the bundles being Standard Model families.

The conjunction is therefore substantive:

Standard Model roster and charges (anomaly-consistent) + charge-blind replication + completion-density mass readout = Standard Model generations with Yukawa hierarchy.

---

13. What This Paper Does Not Do

This paper does not derive the Standard Model gauge group. That is inherited.

It does not derive the hypercharge lattice from first principles. It uses the inherited winding/seat charge assignment and checks that the completed bundle cancels anomalies.

It does not derive exact numerical Yukawa values.

It does not derive CKM or PMNS mixing.

It does not derive neutrino masses.

It does not close the fourth-species operator question.

It does not prove the anchor–completion bridge from below. It uses ANCH-COMP as the conditional, monotone link between completion density and mass readout.

What it does is narrower and important:

1. completes the species–family identification through hypercharge and anomaly cancellation;
2. preserves charge identity across replicated species;
3. identifies refinement depth, read through completion density, as the source of generation mass ordering;
4. interprets Yukawa couplings as emergent readouts rather than fundamental arbitrary constants.

13a. What Has Been Eliminated as Fundamental

It is worth totalling the ledger, because the cumulative achievement is larger than any single paper's contribution and because honesty requires marking *where* each elimination was made. Some of these inputs cease to be fundamental in the programme but are inherited by this paper; one is eliminated here; the rest remain free.

Ceased to be fundamental in the programme (inherited here, not eliminated by this paper):

- **Generation count.** No longer a free integer chosen to fit data. Fixed by the species-count theorem of the spectral arc as the number of stable refinement classes. Inherited as Input D.
- **Family replication.** No longer an unexplained repetition. Derived as the charge-blindness of the refinement coordinate by the Replication Theorem. Inherited as Input A.

Eliminated here:

- **Generation mass ordering.** No longer an accident of which Yukawa happens to be larger — but the elimination is precise about which half it achieves. *Given* the completion-density ordering, the mass ordering follows: $m \propto C$ is order-preserving, so this downstream step is Proven (modulo the anchoring relation, Conditional on $P\hbar$, and a bounded ε; Theorem 6.2). The completion-density ordering itself is no longer bare either: §5d derives the depth ordinal from refinement efficiency ($E_1<E_2<E_3$, third generation saturating the $K=7$ register), leaving only two conditional bridges and the per-sector question between that and the mass ordinal. So the contribution is: mass ordering ceases to be a free choice, replaced by a layered argument whose first link (the efficiency ordinal) is Proven and whose remaining links (ECB, ANCH-COMP, and the per-sector triples) are Conditional or Open.

Remaining free (open, deferred to later papers):

- **Exact Yukawa magnitudes** — the spacing of the eigenvalues $A(c)$, not just their order (Y4, Open).
- **The CKM matrix** — quark mixing from off-diagonal completion-channel mismatch (Y5, Open).
- **The PMNS matrix** — lepton mixing, likewise (Y5, Open).
- **Neutrino masses** — depending on the mass mechanism, not addressed here.

The shape of the achievement is therefore: the *qualitative* structure of the family sector — how many, why repeated, in what order — has moved from fundamental input to derived consequence; the *quantitative* structure — the magnitudes and the mixing angles — remains free and is the explicit target of the papers downstream. The ledger makes the boundary between the two exact, so neither the claim nor its limits can be overread.

The parameter count, made auditable. A skeptical physicist will want the reduction as a number, not a narrative. Here it is, scoped precisely. The Standard Model's charged-fermion Yukawa sector has, on the diagonal, **nine independent magnitudes**: three generations × three charged sectors, $\{e, \mu, \tau\}$, $\{u, c, t\}$, $\{d, s, b\}$. To these add the **generation count** (a free integer, set to 3 by fiat) and the **replication fact** (a free postulate — that the families repeat at all). So the SM starting point in this sector is *9 magnitudes + 1 count + 1 postulate*.

VERSF as of this paper changes each of the three:

- the **count** moves from free to fixed (Input D, the species-count theorem — inherited, not eliminated here);

- **replication** moves from postulate to theorem (a corollary of the charge-side factorization F1 / Replication Theorem);
- the **9 magnitudes** factor as $\{A(c) : c = 1,2,3\} \times \{Y\_a : a = 1,2,3\}$ = **6 structured quantities** (plus the ε corrections of §9).

So the headline reduction is: *from 9 magnitudes + 1 count + 1 postulate, down to 6 structured quantities (+ ε), with the count fixed and replication derived.*

Two honesty constraints on that headline, both gradeable:

1. **The 6 are not yet free parameters either.** The three A(c) are the image of an inherited spectral quantity (Input D through the anchoring relation, §5c) — not independently tunable. The three Y_a are the explicit target of the §9 sectoral-anchor paper — not delivered here. So "6 structured quantities" means *6 quantities with structure already constraining them*, not 6 knobs; the eventual free count, once §9 and the eigenmode programme land, is intended to be smaller still.
2. **The factorization 9 → 3×3 is itself a claim, not an identity.** Asserting that every diagonal magnitude is exactly $Y\_a \cdot A(c)$ with no irreducible per-entry freedom is the content of Y3, and it carries Y3's grade — **Conjectural/Conditional**, contingent on the explicit mass-map construction. If a residual per-entry term survives beyond $Y\_a \cdot A(c)$ and ε, the 6 undercounts. So "6" is the count *if the factorization holds*; it is not yet established that it does.
3. **Scope: diagonal only.** The 9 counts the diagonal magnitudes. The full charged-Yukawa sector also carries the CKM mixing (three angles + one phase), which this paper parks in Y5 (Open). The reduction claim is therefore explicitly *9 diagonal magnitudes → 6 structured quantities*; it does not touch the off-diagonal freedom, which is neither counted nor reduced here.

With those three constraints attached, the count is exact and auditable: **9 diagonal magnitudes + 1 count + 1 postulate → 6 structured quantities (Y3-conditional factorization) + ε, count fixed, replication derived, mixing untouched.**

---

14. Conclusion

The Standard Model repeats one family of matter three times and assigns each repeated field a different Yukawa coupling. It thereby records two facts: charges repeat, masses differ. It does not explain why these facts coexist.

VERSF explains the coexistence by separating the coordinates. Charge is read from winding and seat structure, and refinement class is invisible to it. That is why families replicate. Mass is read from winding and seat structure together with completion density, and refinement class is visible to it. That is why replicated families form a hierarchy.

The paper began by completing the species–family identification. A durable species bundle contains the five Standard Model Weyl multiplets with the hypercharges Q = +1⁄6, u^c = −2⁄3, d^c = +1⁄3, L = −1⁄2, e^c = +1. These cancel the linear, cubic, mixed non-abelian, and Witten anomaly conditions. Under the inherited winding/seat assignment, one durable species is therefore a complete anomaly-free Standard Model family.

With that completion in place, the Yukawa question becomes well-posed. Since all species are gauge-identical, their mass differences cannot come from charge. They come from refinement depth. Completion-density order supplies the generation ordinal, and the anchor–completion bridge converts that spectral ordinal into effective mass-side Yukawa strength.

Stated at its sharpest, the separation is a charge-side factorization theorem: the charge map is independent of refinement class, so gauge identity and refinement depth are separable coordinates and charge cannot read depth (Theorem 4.1). This much is structural. In operator form, the amplification operator $\hat{Y} = F(\mathcal{C})$ — built as a positive monotone function of the already-existing completion-density operator $\mathcal{C}$ — commutes with every charge operator, $[\hat{Y}, \chi] = 0$ (Theorem 5.1). That commutator is not a restatement: it forbids generation-dependent charge, so charge universality across generations (the equality of electron and muon charge, measured to the limits of precision) is in this framework a theorem rather than a coincidence, and any non-commuting amplification is already falsified. The mass ordering then splits cleanly by grade. *Given* the completion-density ordering, the masses order: m ∝ C is order-preserving up to a positive constant, so the downstream step is Proven, modulo the anchoring relation (Conditional on P$\hbar$) and a bounded ε (Theorem 6.2). Whether the three species are *actually* ordered by completion density is Open — the eigenmode programme must compute, per charged sector, an ordered triple (C_a(1), C_a(2), C_a(3)), and it has not been run. The measured ratios contribute one firm positive determination: because the completion-depth spacing is shared across sectors, the cross-sector divergence in the jump pattern forces the sectoral sensitivity to be sector-dependent (Result 7.1) — neither a uniform exponential nor a shared depth spacing can produce the observed sector-by-sector differences. And the completion-density ordering that the downstream step consumes is no longer bare: §5d derives a refinement-efficiency ordinal $E_1 < E_2 < E_3$ from the inherited dyadic loads (1, 2, 4) in the K = 7 closure register, identifying the third generation as the saturation generation ($E_3 = 1$, the family that completes the register rather than adding a new charge). The full hierarchy argument is then layered — dyadic refinement ⟹ efficiency ordinal ⟹ completion-density ordinal ⟹ mass ordinal — with a transparent dependency graph: Proven at the first link (Y7), Conditional at the two named bridges (ECB, then ANCH-COMP). That is a stronger and more defensible spine than hanging the ordering directly on ANCH-COMP alone, and naming ECB rather than leaving it implicit puts the one remaining upstream assumption openly on the table. None of this touches the Higgs mechanism, which continues to convert these couplings into masses unaltered; what VERSF supplies is the structural origin of the couplings the Higgs sector leaves free.

The result is not yet a table of numerical masses, and it is not a closed functional form for the hierarchy: the observed mass ratios are irregular and sector-specific, which already excludes a one-parameter exponential in generation number. What the paper delivers is the structural explanation of why the Yukawa hierarchy exists at all. The Yukawa couplings are not fundamental labels attached by hand, but emergent effective readouts of the completion density

p_v/K_c at which each replicated family sits in the realized structure (higher completion density, heavier family — see §5c).

The first generation is light because it sits at the lowest completion density among the stable replicated families. The second is heavier because it sits at higher completion density. The third is heaviest because it sits at the highest completion density of any stable family, before the fourth-mode question collides with experiment. Identical charges and unequal masses are therefore not two unrelated facts. They are the two faces of one factorized structure: winding and seat fix what a particle is under the forces; completion density p_v/K_c fixes how strongly that particle is realized as mass.