Assignment, Bath Participation, and Refinement: A Response-Operator Account of the Quark Mass Hierarchy in VERSF

Keith Taylor · AIDA Institute · VERSF Theoretical Physics Programme

---

General Reader Summary

Ordinary matter is built from **quarks** — the constituents of protons and neutrons. Quarks come in six kinds that fall into a tidy pattern: a split into two types ("up-type" and "down-type"), repeated three times across three "generations," each heavier than the last. The first generation makes up everyday matter; the heavier two are unstable and seen only fleetingly.

The masses are the puzzle. Within each generation the up-type and down-type partners are near-equal in weight in the first generation, very unequal in the second, and wildly unequal in the third. There is no clean accepted reason why.

The obvious guess is that a quark's mass is two separate ingredients multiplied together: one ingredient saying "up or down," another saying "which generation." If that were right, the up/down weight gap would be the *same* in every generation. **This paper shows clearly that it is not.** The gap grows enormously across generations, so mass is not two independent ingredients multiplied — the two are tangled. This much does not depend on the technical bookkeeping that bedevils quark-mass measurements, which is what makes it a firm starting point. (Within the framework's own terms it is a sharp constraint; stated for the wider world it is a clean reframing of something long known about the spectrum, not a new discovery about nature.)

What replaces the failed guess is one idea. Treat "up versus down" as a single yes/no switch — one bit. That one bit is fed through an amplifier whose strength changes generation by generation. A single quantity, χ ("kai"), measures how hard the switch is pushed at each depth. The whole story collapses into this one number.

Measured carefully, χ does three things: it **grows** toward heavier generations; it **changes sign** between the first and second generations (in the lightest generation the down-type quark is the heavier one — the famous "why is the up quark lighter than the down?" — and from the second generation on, the up-type wins); and its growth **slows down** higher up. One quantity covers both the first-generation oddity and the heavy-quark ordering, and it leaves the framework's earlier light-quark results intact.

The paper is equally careful about what it does **not** claim. With three generations you cannot tell whether χ eventually hits a ceiling or keeps climbing forever, just more slowly — three points cannot decide, and we say so. We retire an earlier "ladder" of special numbers (13, 42) that

turned out to be an artifact of comparing measurements at mismatched conditions. And we show that a tempting mechanism — locating the growth in a structural "carrier" that survives the framework's refinement process — *cannot* be the source, because that carrier does not itself grow. That is a real result: it rules out a wrong answer and points at the right kind of answer.

The honest bottom line: we have firmly established *that* the up/down switch is processed by a generation-dependent amplifier, measured how it behaves, and ruled out one candidate mechanism. We have not derived *why* it behaves that way. But a vague puzzle has become one sharp, answerable question: what makes the switch change sign between the first and second generations, and then fade?

Table of Contents

Abstract

The Standard Model quark sector exhibits two organizational structures usually treated as independent facts: an up/down species distinction repeated across three generations, and a mass hierarchy that differs sharply between species and grows with generation. This paper organizes these features by three structural properties already present in the VERSF programme — **assignment**, **bath participation**, and **refinement depth** — while making a deliberately narrower and stronger claim about mass than earlier formulations.

The central result is negative in form, positive in content: the quark mass operator is **not separable** across assignment and refinement. Quark mass cannot be written as a product of a species factor and a generation factor. This is scheme-independent: separability would make the log-mass matrix rank 1, and it is rank 2.

The constructive consequence is that a one-bit assignment is processed by a **refinement-dependent response operator**. Through the completion-density anchor $m \propto p_{eff}/K_c$, the log-mass takes the forced linear-in-the-bit form

$$\ln m(A, g) = B(g) + A \cdot \chi(g), \quad A \in \{0, 1\},$$

where A is the up/down fold assignment, B(g) is the bath/refinement baseline, and χ(g) is a **fold-response susceptibility**, identical to the within-generation log-split $\chi(g) = \ln(m_{up}/m_{down})$. Using scheme-clean common-scale ratios, χ(g) is monotone increasing, **changes sign between the first and second generations** ($\chi(1) < 0 < \chi(2)$), and grows **sublinearly** thereafter. The generation-1 mass inversion ($m_d > m_u$) and the heavy-quark hierarchy are the same susceptibility on either side of that sign change.

The light-quark sector — the species ratio 6⁄13, the constituent scale $(20/9)\, m_\pi \approx 310$ MeV, the strange estimate ≈ 97 MeV — survives intact and supplies the programme's points of contact with measurement, of which the sharp, χ-gating one is $\chi(1) = \ln(6/13)$. All three are conditional in the same way: they constrain a derivation only if the underlying confinement-programme outputs (in particular the form $(K - 1)/(2K - 1)$ and the value $K = 7$) were fixed upstream without consulting the measured masses, an independence not yet verified here — so the paper's honest standing is zero confirmed external anchors, with one sharp candidate (6⁄13) and two softer ones in IOU state. An earlier "heavy-quark refinement ladder" with factors 13, 7⁄2, 42 is discarded as a scheme artifact whose rungs cross both structural axes. The open problem is sharpened, and one naive route is retired: the framework's refinement-persistent carrier cannot source χ's growth, since its dimension (the first Betti number) is subdivision-invariant; the depth-dependence must

come instead from a fold-selective, depth-dependent *accessibility* of that sector — a target stated precisely here, modest in what it establishes, and left open.

---

1. Introduction

A central question of the Standard Model programme is whether the observed fermion census is a collection of unrelated empirical facts or can be generated from a smaller set of organizational principles. In the quark sector the standing facts are an up-type / down-type distinction (charges +2⁄3 and −1⁄3), three generations, confinement, and a mass hierarchy that is large and uneven.

Conventionally these are introduced separately. This paper organizes the first three by three VERSF structural properties, then makes a precise, falsifiable claim about the fourth — the mass hierarchy — that is weaker than earlier programme drafts but is defensible under scheme discipline and survives a hostile reading.

The paper is deliberately quark-only. A lepton-sector comparison is a natural discriminator for the mechanism proposed here but is not used, for reasons given in §10.

---

2. The Three Structural Axes

We retain the three-axis classification as an **organizing architecture** for the census, kept cleanly separate from the question of how mass is generated. A quark is characterized by

Quark = (Assignment, Bath Participation, Refinement Depth).

2.1 Assignment

Within the quark bath, **assignment** is the up/down fold distinction — a single bit:

$A \in \{0, 1\}$, $A = 1$ (up-type), $A = 0$ (down-type).

This is narrower than the "charge family" reading used in earlier drafts, and the narrowing is intentional. A genuinely binary assignment cannot, on its own, index several distinct charge families; but a one-bit assignment combined with bath participation does generate the observed doublet structure. The up/down bit plays the role of the weak-isospin label $T_3 = \pm\frac{1}{2}$ within each generation:

Assignment ↔ up/down fold bit (one bit), not charge family.

2.2 Bath Participation

The confinement programme distinguishes complete closures from partial closures embedded in a shared committed environment. Bath-free realizations appear as leptons; bath-participating realizations appear as quarks. Confinement, constituent masses, and colour participation then follow from bath participation rather than from independent assumptions:

Bath Participation ↔ quark–lepton structure.

The present paper works entirely inside the bath-participating (quark) sector.

2.3 Refinement Depth

Generation is interpreted as **realization depth** — how deeply a mode is committed — indexed by $g = 1, 2, 3$. This is the cleanest axis where it can be tested in isolation: the charged-lepton tower $e \to \mu \to \tau$ varies only in depth (same charge, same bath) and is monotone in mass. So at fixed species and fixed bath, deeper refinement means larger mass:

Refinement Depth ↔ generation.

Caveat carried forward. That refinement orders mass *within a fixed species tower* does **not** license treating generation as the sole mass coordinate. The error corrected in this paper is precisely the earlier identification of refinement with a single mass-ordering across species. The three axes are useful for classifying the census; they are **not** separable in the mass operator, as §4 shows.

---

3. The Separability Question and the Null Model

The natural first hypothesis — the **null model** — is that the axes factorize the mass:

Null model: $m(A, g) = S(A) \cdot D(g)$,

a fixed species factor times a generation-dependent refinement factor. In log form this is additive:

$\ln m(A, g) = a(A) + r(g)$, $a(A) = \ln S(A)$, $r(g) = \ln D(g)$.

This has a sharp, scheme-independent test. Define the within-generation **log-split**

$L_g = \ln(m_{up}(g) / m_{down}(g))$.

Under the null model the split is *forced constant*:

$L\_g = a(up) - a(down) = \ln(S\_up / S\_down) = L_0$, independent of g.

The entire empirical question reduces to one line: **does L_g run?** A running split refutes the null model outright; a constant split confirms it. There is no third possibility.

---

4. Non-Separability (Established, Scheme-Independent)

Using scheme-clean common-scale ratios (§7.1 for sources):

g	up-type	down-type	S_g = m_up ⁄ m_down	L_g = ln S_g
1	u	d	0.462	−0.772
2	c	s	11.761	+2.465
3	t	b	60	+4.094

L_g spans about 4.9 in log units. It is not remotely constant. The null model is refuted.

4.1 The obstruction is rank one

The refutation has an exact linear-algebra form. Arrange the masses as a 2 × 3 matrix m(A, g), species as rows, generations as columns. The null model $m = S(A)\cdot D(g)$ makes this matrix an **outer product** of a species vector and a generation vector — hence **rank 1**. Writing the general decomposition $m(A, g) = \exp B(g) \cdot \exp(A\cdot\chi(g))$ (justified in §6), the up-type row is the down-type row multiplied entry-by-entry by $\exp \chi(g)$. That is a single common multiple — rank 1 — **if and only if exp χ(g) is constant in g**, i.e. iff $\chi(g) = \chi_0$. Therefore

separable ⟺ mass matrix rank 1 ⟺ $\chi(1) = \chi(2) = \chi(3)$.

The three sampled values are not equal — (−0.77, +2.46, +4.09) — so the mass matrix is **rank 2** and does not separate as species × generation. The failure of the null model is carried by a single finite-difference fact: the three χ values are not all equal. (We write this informally as $d\chi/dg \neq 0$, but g is a discrete refinement-depth label with three values, not a continuum; there is no derivative here, only the non-constancy of three numbers. The continuum notation is shorthand and nothing rests on χ being a differentiable function of g.)

This is the load-bearing result, and it is robust for a structural reason. In the MS-bar scheme the mass anomalous dimension is **flavour-blind**, so a same-scale mass ratio $m\_i(\mu) / m\_j(\mu)$ is renormalization-group invariant. No choice of common mass scheme or scale can collapse the split to a constant; its variation across generations is RG-invariant. Non-separability is not an artifact of input choices — it is a fact about the spectrum.

The constructive content: assignment and refinement **couple** inside the mass operator. The correct schematic is

$m = F(A, g, \text{bath})$, F not factorizable,

and the irreducible cross-term is exactly the variation of L_g — the running of the species split with depth.

Scope note, to forestall a misreading. The null model refuted here is not a Standard Model expectation — in the Standard Model the Yukawa couplings are free parameters and no factorization of mass into species × generation is asserted. Separability is a live expectation specifically *inside* VERSF's three-axis architecture, where assignment, bath, and refinement are introduced as distinct organizing axes (§2). The result of this section is therefore an internal constraint on VERSF — the three axes, however useful for classifying the census, do not act independently in the mass operator — rather than a discovery about nature. Claimed as such, it is a real and non-trivial constraint; it should not be read as knocking down a hypothesis the wider literature held.

---

5. Completion Density and the Mass Anchor

The VERSF mass anchor associates mass with completion density:

$m \propto C = p_{eff} / K_c$,

where p_{eff} measures distinguishability access and K_c measures commitment burden. In log form,

$\ln m = \ln p_{eff} - \ln K_c + \text{const}$.

This is the interpretive frame for what follows. We flag honestly that the present paper does **not** decompose the response susceptibility $\chi(g)$ (§6–7) into separate p_{eff} and K_c contributions; that is part of the open derivation (§11). The anchor tells us what the response operator must ultimately be built from; it does not by itself supply the numbers.

---

6. The Response Operator

A one-bit input can enter log-mass in only one way. The point is an elementary but load-bearing proposition.

Binary Assignment Decomposition. Let $M(A, g)$ be any observable depending on a one-bit assignment $A \in \{0, 1\}$ and a depth label g. Then there exist **unique** functions $B(g)$ and $\chi(g)$ such that

$M(A, g) = B(g) + A \cdot \chi(g)$.

Proof. Set $B(g) := M(0, g)$ and $\chi(g) := M(1, g) - M(0, g)$. Substituting $A = 0$ and $A = 1$ recovers $M(0, g)$ and $M(1, g)$. Uniqueness is immediate: a function of a two-valued argument is fixed by its two values, so the line through them is unique. ■

The proof is trivial; the consequence is not. $\chi(g)$ is **not an optional modelling choice**: for any observable, the assignment-dependence is captured by exactly one function, and $\chi(g)$ *is* that function. Applied to the log-mass $M = \ln m$ — the representation in which the multiplicative null model $m = S(A)\cdot D(g)$ becomes additive — the decomposition is forced:

$$\ln m(A, g) = B(g) + A \cdot \chi(g), \quad A \in \{0, 1\}. \quad (\star)$$

There is no A^2 term, because $A^2 = A$ for a bit. With $A = 1$ for up-type, $A = 0$ for down-type:

$$\ln m_{down}(g) = B(g), \quad \ln m_{up}(g) = B(g) + \chi(g), \quad \chi(g) = \ln m_{up}(g) - \ln m_{down}(g) = L_g.$$

Here **B(g)** is the common bath/refinement baseline (the down-type tower), **A** is the one-bit fold assignment, and **χ(g)** is the **fold-response susceptibility** — how strongly, and in which direction, one bit of assignment tilts the completion density at depth g.

χ is the entire non-separability. It is the *unique obstruction to separability*: if $\chi(g)$ were constant, the quark mass matrix would factorize into independent assignment and refinement contributions (§4.1); the observed running of χ measures exactly how, and how much, assignment and refinement fail to decouple. Explaining $\chi(g)$ is equivalent to explaining the entire non-separable structure of the quark mass hierarchy.

This is a **conceptual reorganization, not an information compression.** (★) is a lossless reparametrization — $B(g)$ is three numbers and $\chi(g)$ is three numbers, the same six masses in new coordinates, with no reduction in degrees of freedom. Its value is that it isolates the entire assignment-dependence of the spectrum into one object, $\chi(g)$, identified as the unique obstruction to separability. Genuine compression would arrive only with a *derived* $\chi(g)$ carrying fewer than three free parameters; that derivation is the open problem (§11), and until it exists (★) has no predictive content of its own. The physical question becomes:

Why does a one-bit fold assignment produce a refinement-dependent log response inside the confinement bath?

6.1 Gauge Freedom in the Access–Burden Split

The completion-density anchor (§5) invites a decomposition of the susceptibility into an access term and a burden term. Writing each fold's mass through $m \propto p_{eff}/K_c$ gives, identically,

$$\chi(g) = \Delta\ln p_{eff}(g) - \Delta\ln K_c(g), \quad \Delta\ln p_{eff}(g) = \ln p_{up}(g) - \ln p_{down}(g), \quad \Delta\ln K_c(g) = \ln K_{up}(g) - \ln K_{down}(g).$$

This is an **identity, not a theory** — it is $m \propto p\_eff / K\_c$ rewritten, and holds for any χ. Crucially, the masses constrain only the *difference* of the two legs, never the legs themselves. The decomposition carries an exact **gauge freedom**: for any function f(g),

$\Delta \ln p\_eff \to \Delta \ln p\_eff + f(g)$, $\Delta \ln K\_c \to \Delta \ln K\_c + f(g)$,

leaves $\chi = \Delta \ln p\_eff - \Delta \ln K\_c$ unchanged. There is one undetermined function — three real degrees of freedom across three generations — and quark masses fix none of it.

The consequence is a strict limit on interpretation. The split is **not gauge-invariant**, so any statement that the sign change is caused specifically by *access* growth, or specifically by *burden* reduction, is a choice of gauge rather than a measured fact. A narrative such as "the down-fold wins at shallow depth by lower burden, the up-fold overtakes at deeper depth by stronger access" is internally consistent and reproduces χ exactly — but so does "all of it is access," and so does "all of it is burden." No quark mass distinguishes them.

Gauge-fixing requirement. The access–burden language becomes predictive only when one leg is fixed by structure derived **independently of the quark masses**. Absent such a condition, the language is interpretive, not predictive, and no part of the sign change may be attributed to either leg.

This identifies the next proof target exactly — a derivation of the **burden leg** from confinement geometry:

Fold-selective bath burden (open theorem). Derive, from ownership/confinement geometry and without consulting any quark mass, that $K\_c{,}up(g) \neq K\_c{,}down(g)$, yielding $\Delta \ln K\_c(g)$ as a definite function.

If and only if that theorem is established does the identity become a prediction, $\Delta \ln p\_eff(g) = \chi(g) + \Delta \ln K\_c(g)$, with the access leg then independently checkable against fold geometry. Until then the honest statement is that **a fold-selective burden explanation is a candidate gauge-fixing of χ, not an established mechanism**, and the programme does not claim the bath "favours the down-fold by burden." That claim is precisely the open theorem, not an input. (The current bath-burden ledger derives a baseline shared by both folds; a susceptibility cancels any fold-symmetric contribution, so fold-selectivity — $K\_up \neq K\_down$ — is the specific thing that remains to be shown, not a consequence already in hand.)

6.2 The framework's single quantitative anchor

It is worth saying explicitly where, in all of this, the programme makes nontrivial quantitative contact with measurement. Exactly one claim does: $\boldsymbol{\chi(1) = \ln(6/13)}$.

The other results are of a different kind. The non-separability of §4 is *structural* — it refutes a factorization expected within VERSF's own architecture, not a Standard Model claim (§4 scope note), so it constrains the framework rather than measuring nature. The sign change and the deceleration of §7 are convention-free *descriptions* of the spectrum, read from the same masses any derivation would reproduce, so they cannot independently anchor it (§12). The first-generation value is different in kind. The light-quark/confinement programme fixes $m\_u / m\_d = (K - 1)/(2K - 1) = 6/13$ with $K = 7$, and — *provided* that determination did not consult the measured m_u / m_d (the proviso of §12, still to be confirmed) — that number stands entirely outside the χ analysis. It is the single anchor the susceptibility programme rests on. (The programme makes other contacts with light-quark measurement — the constituent scale (20/9) m_π and the strange estimate, §8 — but those do not feed χ, which is the within-generation up/down ratio alone; §13 grades all three together. Here we mean the single contact that *gates the χ derivation*.)

Two clarifications keep the claim at its true strength. First, it is an *anchor*, not a test the χ work itself passes: $\chi(1)$ is *defined* as $\ln(m\_u / m\_d)$, and the 6/13 prediction was established upstream in the confinement programme, which the χ account inherits rather than re-derives. Second, that is exactly what makes it binding on the future. Since 6/13 is fixed elsewhere, it cannot be spent as a fit parameter, so it states the one quantitative requirement any derivation of χ must meet:

Anchor requirement. Any derivation of $\chi(g)$ must reproduce $\chi(1) = \ln(6/13)$ — emerging from the same dynamics that generate the sign change and the deceleration, not imposed as an initial value — conditional on 6/13 having been fixed independently of m_u / m_d.

Everything else the programme currently offers is reorganization of the data into the one variable, $\chi(g)$, that isolates what remains to be explained.

---

7. The Susceptibility $\chi(g)$: Data and Structure

7.1 Scheme-clean inputs

The susceptibility is read off same-scale ratios, the RG-clean quantities:

- **$S_1 = m\_u / m\_d \approx 0.462$** — both masses MS-bar at 2 GeV; a directly determined lattice quantity.
- **$S_2 = m\_c / m\_s \approx 11.761$** — a common-scale lattice ratio, evaluated at a single scale, needing no running argument.
- **$S_3 = m\_t / m\_b \approx 60$** — both masses run to a common scale. An each-at-own-scale reading, $m\_t(m\_t) / m\_b(m\_b) = 163.3 / 4.183 \approx 39$, is **inadmissible** here: it is not a common-scale ratio and is therefore not the RG-invariant object of §4.1. Running both to a common scale (e.g. $m\_t(M\_Z) \approx 172$ GeV, $m\_b(M\_Z) \approx 2.86$ GeV) gives ≈ 60, and because γ_m is flavour-blind this value is the same at every common scale. We use 60. (One honest caveat the light quarks do not carry: the top input ≈ 172.5 GeV is a Monte-Carlo / pole-adjacent quantity with a ~0.5 GeV ambiguity in its MS-bar relation. This is negligible for $\chi(3)$ — it shifts ln 60 at the percent level — but it is the one place the heavy-sector input is less scheme-clean than the rest.)

The same discipline that admits S_2 rejects the each-at-own-scale value 39 for S_3. A ratio of masses defined at *different* scales — top at 163 GeV against bottom at 4.2 GeV — is not RG-invariant; the flavour-blind cancellation of §4.1 holds only at a common scale. Using 39 would commit, for S_3, exactly the inadmissible move §9 charges against the ladder (whose inflated 13 and 42 came from comparing mismatched scales). The fully disciplined values are 11.76 and ≈ 60, not 13 and 42.

7.2 The susceptibility

$\chi(1) = -0.772$, $\chi(2) = +2.465$, $\chi(3) = +4.094$.

Increments:

$\chi(2) - \chi(1) = +3.237$, $\chi(3) - \chi(2) = +1.629$.

Three structural facts, all scheme-proof because all three S_g are common-scale RG-invariant ratios (§7.1), but unequal in how much they constrain a derivation:

1. **χ(g) is monotone increasing**: $\chi(1) < \chi(2) < \chi(3)$. Note the careful attribution. What gen-1 vs gen-2 establishes on its own is *non-constancy* — two unequal values — and that is what carries the rank-2 result of §4.1 and the sign change; it does not need the g=3 point. Full three-point monotonicity is a stronger claim, since $\chi(2) < \chi(3)$ does depend on the g=3 value. So "the mass matrix is rank 2" rests on two points; "χ is monotone increasing" rests on all three.
2. **χ(g) changes sign between the first and second generations**, with $\chi(1) < 0 < \chi(2)$: assignment makes down-type heavier at generation 1 and up-type heavier from generation 2 on. The *existence* of this sign change is convention-free. Its *location* is not: a straight-line interpolation against the integer index puts the zero near g ≈ 1.24, but g is a discrete depth label, and interpolating in ln g, or in any monotone reparametrization of the depth axis, moves the crossing while leaving the sign change between gen 1 and gen 2 untouched. Since the location is fixed by χ(1), χ(2) and a chosen interpolation alone, it carries no information beyond "the sign flips between them." We therefore state only the sign change, not a crossing coordinate.
3. **χ(g) decelerates** in its increments (+3.24 then +1.63; ratio ≈ 0.50). This is the **weakest** of the three claims, and should be flagged as such. With three points there are two first-differences and exactly one second-difference, so "decelerates" is a *single sign bit* — the statement that $\text{increment}_2 < \text{increment}_1$. A generic monotone, sign-changing three-sequence decelerates or accelerates with roughly equal prior, so reproducing deceleration is weak confirmation of any derivation. This is the same thinness §10.1 records from the other side: saturating and unbounded laws fit these three points equally well. With S_3 as the common-scale ratio the deceleration is scheme-proof and its magnitude (increment ratio near 0.5, well below 1) is stable under the few-percent heavy-quark uncertainty — but as a constraint on a future derivation it is worth at most one bit.

7.3 The sign change is the physical content; the sign label is convention

The sign of χ depends on which species is labeled A = 1; it flips under relabeling. What does **not** depend on convention is the **existence of a sign change** between generations 1 and 2 — a depth interval across which assignment ceases to favour one species and begins to favour the other — together with the **magnitude** profile

$|\chi(g)| = (0.772, 2.465, 4.094)$,

which grows and decelerates regardless of labeling. Two distinct convention issues are in play: the *sign* of χ is a labeling choice (this subsection), and the *coordinate* of the crossing is an interpolation choice (§7.2). Neither is physical. The physical content is that the sign changes between gen 1 and gen 2 and that $|\chi|$ grows and decelerates. The paper's claims are stated in those terms so that no convention bites.

7.4 The generation-1 inversion as a boundary value

The familiar puzzle "why is the up quark lighter than the down quark, when every heavier generation has the reverse ordering?" is, in this framing, not an exception but the **value of the same monotone susceptibility on the down-favouring side of its sign change**: generation 1 sits where $\chi < 0$. A single quantity $\chi(g)$ describes both the gen-1 mass inversion and the heavy-sector ordering, as one sequence read on either side of the sign change between gen 1 and gen 2.

The honest weight of this should be stated plainly. Since $\chi(g)$ is *defined* as $\ln(m_{up} / m_{down})$, "the gen-1 inversion is the boundary value of χ" is, at the level of description, close to tautological — it restates that the gen-1 up/down ratio is below one while the others are above. The non-trivial empirical residue is real but modest: three numbers that happen to be monotone-increasing and to cross zero, describable by one sign-changing quantity rather than two unrelated facts. This is a **unification of description, not yet of explanation**; it becomes a unification of explanation only when a single mechanism is shown to *produce* $\chi(g)$ (§11), at which point the gen-1 inversion and the heavy hierarchy would follow from one cause rather than merely sharing one label.

7.5 The top quark

Within (★), the top quark is the up-type tower at maximum depth, where $\chi(3)$ is large and positive. Its largeness is $\chi(3)$ being large, not a special "saturation event." Its prompt decay before hadronization is already accounted for in the Standard Model as a mass effect ($\Gamma \propto m_t^3$ exceeds the hadronization rate); no additional structural label is required, and none is asserted.

---

8. Light-Quark Sector (Established, and the Boundary Value of χ)

The light-quark results are the programme's strongest assets and are unchanged. They also supply the boundary value of the response operator.

8.1 Up–down species ratio

The confinement programme gives

$$m_u / m_d = (K - 1)/(2K - 1) = 6/13 \approx 0.4615 \ (K = 7),$$

against the scheme-clean $S_1 = 0.462$. The match should be stated at its true strength: $6/13 = 0.4615$ sits on the measured central value (lattice determinations cluster around 0.46–0.47, e.g. 0.46(2)(2)), **well within the ≈ 5% measurement uncertainty** on m_u / m_d. The bare number-to-number gap is ≈ 0.1%, but that is not a measured precision — the ratio itself is known only to a few percent, so the meaningful tolerance is ≈ 5%, not 0.1% (see §12). In the language of (★), this is $\chi(1) = \ln(6/13) = -0.773$: the structural prediction lands on the measured boundary value of the susceptibility. This is the single surviving clean K-value, and it lives precisely where the response operator is least processed by refinement.

8.2 Constituent light-quark scale

$$m_{q,const} = (20/9)\, m_\pi \approx 310 \text{ MeV},$$

reproducing the conventional constituent light-quark scale.

8.3 Strange quark

Kaon–pion geometry with active-mode screening gives $m_s \approx 97$ MeV, a credible estimate consistent at the few-percent level with the measured strange-quark mass (the comparison is qualitative; we do not attach a scale to it, since the estimate is not produced at a declared common scale and the paper's scheme discipline applies only to the same-scale ratios of §7.1). The programme therefore gives a coherent account of u, d, s.

---

9. What Is Discarded, and Why

One reading of the heavy sector — the **interleaved heavy-quark refinement ladder** s → c → b → t with factors 13, 7⁄2, 42, taken as successive refinement steps — is inadmissible, for four independent reasons, any one sufficient.

1. **It crosses both axes.** s → c and b → t are *same-generation, different-species* transitions — assignment changes, not refinement. Assigning them refinement factors contradicts §2.
2. **Two of three rungs are species splittings in disguise.** s → c equals the gen-2 up/down ratio; b → t equals the gen-3 up/down ratio. The ladder's two most striking "hits" (13 and 42) are not refinement steps; they are the assignment split χ at generations 2 and 3.
3. **The factors do not survive scheme discipline.** At a single common scale the relevant ratios are 11.76 and ≈ 60, not 13 and 42. The clean integers arise only from comparing

masses at mismatched scales (§7.1), which is inadmissible; under full discipline the gen-3 ratio is *further* from 42 than the inflated value suggested.

4. **No parameter-free law reproduces them.** With three data points and a low-order generating function, neither the integer set (6⁄13, 13, 42) nor the scheme-clean set (0.462, 11.76, ≈ 60) is fit by any single parameter-free K-expression; an exact fit requires one free parameter per generation, a relabeling of the data rather than a law. The obstruction lies in the log-spacing of the data, not in the choice of structural forms.

The constructive replacement is the response operator (★): the same data, organized as a single susceptibility $\chi(g)$ rather than as a cross-species ladder, respects the axes instead of violating them.

---

10. Two Claims We Decline to Make

Honest grading requires stating what the data **cannot** support.

10.1 Saturation is not established

The bounded asymmetry (m_up − m_down) ⁄ (m_up + m_down) runs to ≈ 0.97 at generation 3 and appears to flatten — but that measure lives in [−1, +1], so **any** growing split must flatten near the bound; the deceleration there is a property of the map's range, not of the quarks. On the artifact-free object, the unbounded log-split $\chi(g)$, three data points fit a saturating law (with a ceiling) and an unbounded sublinear law (e.g. $\chi \approx a + b\cdot\ln g$) **equally well**. We therefore claim only that $\chi(g)$ grows and decelerates, and we do **not** claim it saturates or approaches "total up-type dominance." The far-depth behaviour of χ is open and **uncloseable from the quark sector alone**.

10.2 The lepton discriminator is deferred

The mechanism of (★) predicts a sharp test: if χ is intrinsic to the **assignment** operator, it should run in the bath-free (lepton) sector too; if χ is a **bath** (confinement) effect, the lepton split should not run. This would discriminate cleanly between the two candidate theories.

The test is not used here because it is not cleanly available. The neutral partners in the lepton sector are neutrinos, and neutrino **absolute** masses are not determined as charged-lepton and quark masses are. Oscillation experiments measure mass-squared **differences**, leaving the absolute scale and ordering (normal / inverted / quasi-degenerate) unfixed. The lepton split ln(m_ν ⁄ m_ℓ) is therefore not a clean input, and using it would import neutrino-model assumptions this paper avoids. We record the comparison as valuable in principle and defer it.

---

11. The Derivation Target

The open problem is now sharply posed. Derive, from interface dynamics, a fold-response susceptibility $\chi(g)$ with two qualitative features:

1. a **sign change between the first and second generations** — χ negative at gen 1, positive at gen 2, so assignment reverses which species it favours across that interval (no crossing coordinate is required; the location is interpolation-dependent, §7.2);
2. **sublinear (decelerating) growth** in magnitude with depth.

A sign change is more constraining than an arbitrary monotone function, and it *suggests* — without implying — a particular reading: that two contributions to the fold response compete, one dominant at gen 1 and the other from gen 2 on, with the sign change at their crossover. This is a natural interpretation, not a forced one. As §11.1 makes explicit, writing χ as a difference of two functions is by itself vacuous, since any function so decomposes; the two-contribution reading acquires content only under the independence proviso stated there. With that caveat carried, the concrete derivation question is:

What two competing contributions to the fold-assignment response — each motivated independently of the masses — could cross over between the first and second generations, and why would their net effect grow sublinearly with refinement depth?

This is a qualitative, two-feature target (a sign change and a curvature), not a fit to three numbers. A derived $\chi(g)$ that *happened* to change sign between gen 1 and gen 2 and decelerate would be the first object in the heavy-quark sector that was not reverse-engineered from the masses.

11.1 The Fold-Selective Competition Hypothesis

The competition can be given a definite form. Suppose the susceptibility is the difference of two **separately motivated, fold-selective** contributions,

$$\chi(g) = \chi_+(g) - \chi_-(g),$$

where χ_+ arises from structure favouring the up fold and χ_- from structure favouring the down fold, *each derived independently of the masses*. Then the zero-crossing is not fitted but predicted:

$$\chi(g^*) = 0 \Leftrightarrow \chi_+(g^*) = \chi_-(g^*),$$

the depth at which the two fold-selective effects balance. The content of the hypothesis is the word *independently*: writing χ as a difference of two functions is vacuous on its own, since any function so decomposes; the hypothesis has teeth only if χ_+ and χ_- are each fixed by fold-selective structure that never consults m_up ⁄ m_down. The gen-1 mass inversion is then that χ_- wins at the shallowest depth; the heavy hierarchy is that χ_+ wins, and widens its lead, at every greater depth. The access–burden reading of §6.1 is the concrete instance — $\chi_- \leftrightarrow \Delta\ln K\_c$, $\chi_+ \leftrightarrow \Delta\ln p\_eff$ — so the Competition Hypothesis and the fold-selective bath burden open theorem are the same proof obligation seen from two sides.

11.2 Two-channel readout models, and the line they must not cross

A natural class of realizations is a *two-channel readout*, in which a down-favouring χ_- and an up-favouring χ_+ combine, depth by depth, into the observed susceptibility. Any such model produces a sign change automatically — wherever the channels become equal — and that is a genuine structural virtue. But it carries a hazard the standard of §10.1 makes sharp. A combination rule such as a depth-weighted mean,

$$\chi(g) = (1/g)\,[\,\chi_- + (g - 1)\,\chi_+\,],$$

forces monotonicity, deceleration, and saturation for any fixed channels with $\chi_+ > \chi_-$: the weight slides from χ_- toward χ_+ as g grows, so the qualitative shape of $\chi(g)$ is a property of the weighting form, not of the physics.

It is worth pinning this down with the numbers, because it shows the hazard is concrete and not merely a tautology to be warned against in the abstract. Fix the channels to be constant and pin them by the two lightest points: take $\chi_- = -0.772$ (the gen-1 value) and choose χ_+ so the form reproduces the gen-2 value exactly. At g = 2 the rule gives $\chi(2) = \frac{1}{2}(\chi_- + \chi_+)$, so $\chi_+ = 2\cdot(2.465) - (-0.772) = 5.702$. With both channels now fixed, the form *predicts* gen 3:

$$\chi(3) = \tfrac{1}{3}(\chi_- + 2\chi_+) = \tfrac{1}{3}(-0.772 + 11.404) = 3.544,$$

against the measured $\chi(3) = 4.094$ — an underprediction of ~13%. So even the encoding model, with constant channels tuned to the first two points, does not fit the third; the example is not just a stand-in for "any function decomposes," it is *empirically falsified as written*. The lesson cuts the intended way: a model whose form already guarantees monotonicity, deceleration, and saturation still fails the data unless its channels are themselves depth-dependent — and depth-dependent channels are exactly the thing that must be derived rather than fitted.

Such a model does not *explain* the sign change and curvature; it *encodes* them, exactly as the bounded asymmetry of §10.1 manufactures a ceiling from the range of its map. The discriminating content of any two-channel model lies entirely in two places — the channel values χ_+, χ_- and the weighting law — and **both must be fixed by structure derived independently of the quark masses**. A weighting law or channel value inferred from the observed χ profile (for example, a deep-channel volume whose exponent is read off the heavy-quark ratio) is a fit parameter disguised as a structural number, and reproducing χ with it demonstrates nothing. The admissible target is a two-channel model whose channels and combination rule are both derived upstream, with the sign change emerging as a consequence — not a model whose form guarantees a sign change somewhere and whose parameters are then tuned to place it.

11.3 A ruled-out mechanism: carrier persistence is not amplification

A separate strand of the programme bears directly on where χ's depth-dependence can and cannot come from, and it yields a genuine negative result worth recording.

That strand establishes two things, taken here as inputs (each conditional on its own proof, stated where it is derived). First, **scalar observables are refinement-trivial**: for connected admissible substrates the eigenvalue-1 subspace of the refinement transfer operator is one-dimensional ($\ker(W - I) = \mathbb{R}\cdot 1$), so the only refinement-stable scalar mode is the global counting invariant; bulk, gluing, and antichain-supported scalar modes all decay. Second, the first non-trivial refinement-persistent sector is the cohomological transport sector

$$H^1(G(\Lambda)) = C^1(\Lambda) / \mathrm{Im}(d^0),$$

which under midpoint refinement satisfies $\Gamma^* \circ L = 2\cdot\mathrm{Id}$ and so, after the canonical rescaling $\tilde{\Gamma}^* = \frac{1}{2}\Gamma^*$, gives $\tilde{\Gamma}^* \circ L = \mathrm{Id}$ — the sector is **identified across refinement depth**.

It is tempting to locate χ's carrier in H^1: scalars cannot carry an assignment label that survives refinement, so any persistent assignment structure must live in a richer sector, and H^1 is the first such sector. But this cannot be the *source of the growth*, for a reason that is immediate. The dimension of H^1 is the first Betti number — the cycle rank $b_1 = E - V + C$ of the substrate graph — and **subdivision preserves the first Betti number**: midpoint refinement sends $V \to V + E$ and $E \to 2E$, so $b_1 = E - V + C$ is unchanged in one line. The carrier persists but does not grow:

$$\dim H^1(g + 1) = \dim H^1(g).$$

A depth-invariant carrier cannot source a depth-dependent, growing assignment split. Taken literally, the persistence result predicts **no running** of a carrier-size effect — the opposite of what χ requires.

The honest weight of this is modest, and worth grading plainly. It is *valid*: the Betti-invariance is a one-line fact, not a deep theorem. It is *narrow*: it rules out exactly one route — "χ grows because the persistent carrier grows in dimension" — and arguably no careful version of the programme need ever have asserted that route, since a fixed-dimensional sector can obviously carry a depth-dependent functional on it (which is precisely the accessibility move below). And it is *conditional*: the redirect to accessibility inherits the conditionality of the imported $\ker(W - I) = \mathbb{R}\cdot 1$ and persistence results, neither proved here. So this is a genuine internal correction — it retires the carrier-existence reading recorded under §13's discards — but it is a redirect, not a structural elimination of a broad class. Its value is in naming where the depth-dependence must live, not in closing anything.

The growth must therefore come not from the *size* of the persistent sector but from a depth-dependent **accessibility** of it. Formally, $\chi(g) = \ln(6/13) + \ln[\mathcal{A}\_up(g) / \mathcal{A}\_down(g)]$ for some fold-selective accessibility functional $\mathcal{A}\_A(g)$ acting on H^1. We stress what this is and is not: writing χ in this form **renames the target, it does not reduce it** — $\mathcal{A}\_up / \mathcal{A}\_down$ is one-to-one with χ, the same three degrees of freedom in heavier notation. It earns its place only by identifying the *type* of object that must be derived, and by retiring the carrier-size route above.

The natural candidate for a depth-dependent quantity is the **near-marginal transport spectrum**, not the persistent sector itself. For a chain of k diamonds the subleading transfer eigenvalue is

$\lambda_2(k) = \cos^2[\pi/(4k)] \approx 1 - \pi^2/(16k^2)$ (large k),

with mixing time $\tau_{mix} \approx (16k^2/\pi^2) \ln(1/\varepsilon)$. These modes are not persistent, but they remain accessible over long refinement intervals, and — unlike dim H^1 — they *vary with depth*. That is the right type signature for a source of running.

Two cautions discipline this direction, and both are decisive:

- **Fold-blindness.** $\lambda_2(k)$ carries no up/down index; it is a property of the chain, identical for both folds. The entire phenomenon χ captures is the *difference* between folds, and a spectral quantity that behaves identically for both contributes nothing to χ — exactly the cancellation of any fold-symmetric contribution noted in §6.1. Depth-dependence is necessary but not sufficient: the accessibility must be **fold-selective by derivation**, and *why* the two folds project differently onto the near-marginal sector is the part that must be proved, not posited.
- **Inherited saturation.** $\lambda_2(k) \to 1$ monotonically; it saturates. χ does not — §10.1 is emphatic that saturation is not established and must not be assumed. A construction that builds χ's growth from a provably saturating quantity smuggles a ceiling back in through the mechanism after the phenomenology disowned it. Any such construction must either avoid inheriting the asymptote or **explicitly commit to predicting saturation** — a real, falsifiable claim, but one to be made deliberately, not by accident of the carrier's range.

The open theorem is then stated precisely, and narrowly:

Fold-Selective Accessibility (open theorem). Derive, from fold geometry and refinement dynamics — and without consulting any quark mass — a depth-dependent, fold-selective accessibility of the persistent transport sector, i.e. fold projectors $P_{up} \neq P_{down}$ and a refinement-access operator R_g such that $\mathcal{A}_A(g) = \mathrm{Tr}_{H^1}[P_A R_g P_A]$ satisfies $\mathcal{A}_{up} \neq \mathcal{A}_{down}$ with the correct depth-dependence. Success would yield $\chi(g) = \ln(6/13) + \ln[\mathcal{A}_{up}(g)/\mathcal{A}_{down}(g)]$ with the sign change between gen 1 and gen 2 emerging as a consequence.

This is the same proof obligation as the fold-selective bath burden of §6.1, viewed through the refinement spectrum: in both, the undischarged step is fold-selectivity derived independently of the masses. The trace expression above is **not** a derivation — it is a precise statement of the three objects (P_{up}, P_{down}, R_g) that a derivation must supply. The advance recorded here is the negative result (carrier growth is ruled out) and the type identification (depth-dependent fold-selective spectral accessibility), not any closing of the gap.

11.4 A candidate provenance (directional, not a result)

A separate strand derives the down-type baseline $B(g)$ through a bath-burden ledger — a class-owned conserved quantity shared across three colour slots under neutral ownership, read at member scale. That machinery is built for the *baseline*, not the susceptibility, and χ is by construction indifferent to any contribution shared equally by both folds, since $\chi(g) = \ln m_{up} - \ln m_{down}$ cancels it. The bath ledger becomes a candidate source for χ precisely — and only — if the bath burden is **assignment-dependent**: if up- and down-type branches read the bath differently, χ inherits that asymmetry directly, χ(g) = B_up − B_down, and a sign change becomes possible wherever the two readings cross. A reflection-definite one-branch occupancy applying to one fold and not the other is the kind of twofold that could supply the two competing contributions. This is offered as a direction: pursuing it means deriving the up- and down-sector bath corrections *together* and reading χ off their difference, inheriting whatever conditional grade that joint derivation carries. The contested point — whether the bath burden is fold-selective — must be settled by the bath ontology, not by the down-sector ratio it is meant to explain.

---

12. The Boundary-Value Constraint

The derivation target of §11 is qualitative — a sign change and a curvature. There is one further constraint, quantitative and already in hand: the boundary value of χ is fixed independently.

The light-quark/confinement programme derives the first-generation species ratio as m_u ⁄ m_d = (K − 1) ⁄ (2K − 1) = 6⁄13 (K = 7), so that χ(1) = ln(6⁄13) = −0.7732. The structural prediction 6⁄13 = 0.4615 sits essentially on the measured central value (lattice determinations cluster around 0.46–0.47, e.g. 0.46(2)(2)), so χ(1) lands on the measurement, **well within the ≈ 5% measurement uncertainty** on m_u ⁄ m_d. (The bare number-to-number gap is 0.1%, but that is not a measured precision — the ratio is known only to a few percent, so the meaningful tolerance is ≈ 5%, not 0.1%.) The consistency is a genuine asset, and it changes the character of the derivation target: the sign change of §7 does not occur in empty space, but starting from a specific, independently fixed structural anchor at generation 1.

Any successful derivation of χ(g) must satisfy three conditions:

1. **χ(1) = ln(6⁄13)** — the first-generation value reproduces the established light-quark ratio;
2. **χ changes sign between the first and second generations** — down-favouring at gen 1, up-favouring from gen 2 on (no crossing *coordinate* required; location is interpolation-dependent, §7.2);
3. **χ grows sublinearly thereafter** — increments +3.24 then +1.63 (§7).

These are not three independent locks, and they are not even three of equal weight. Only condition 1 is anchored to something fixed *outside* the spectrum being explained — and even that holds only under the K-independence proviso below. Conditions 2 and 3 are descriptions read off the same masses a derivation would be reproducing, so they cannot independently *gate* that derivation; they are consistency checks against the data, not external targets. Condition 3 is the weakest of all: deceleration is a single sign bit (§7.2), reproduced by a generic monotone sign-changing three-sequence with roughly even odds, so it confirms little. The gate is therefore **one external lock (condition 1), one convention-free consistency check (condition 2), and one near-free consistency check (condition 3)**. The "land without being told to" force below applies to condition 1 alone.

The force of condition 1 lies in *how* it must be met; it is the developed form of the anchor requirement of §6.2. The value 6⁄13 is an **upstream result**, derived in the confinement programme as (K − 1)∕(2K − 1) with K = 7. It is a consistency target, not a free boundary condition — *provided* K and the form (K − 1)∕(2K − 1) were fixed without reference to m_u∕m_d. This proviso is load-bearing: if the light-quark ratio were used to tune K, condition 1 is circular and the gate has no external lock at all. [The derivation of K = 7 and the species ratio is given in the confinement / W_7 programme; the claim here requires that those determinations did not consult the measured m_u∕m_d. This cross-reference must be supplied and the independence confirmed.] Granting it, a derivation that satisfies condition 1 by *imposing* χ(1) = ln(6⁄13) as an initial value has not passed the gate — it has spent a parameter the programme already fixed elsewhere. The gate bites only if χ(1) **emerges** as ln(6⁄13) from dynamics that also reproduce the sign change and the deceleration. Because 6⁄13 cannot then be re-used as a fit parameter, condition 1 carries the gate's weight: it prevents a future derivation from merely manufacturing a sign-changing curve, since the curve must also land, without being told to, on a number established before χ was ever defined.

This is what makes χ a better object than the discarded ladder (§9). The ladder rested on 13, 7⁄2, 42 — integers that proved to be scheme artifacts, fragile under the discipline of §7.1. The χ programme rests on χ(1), χ(2), χ(3), RG-clean common-scale ratios. The ladder was a **pattern** noticed in numbers; χ is an **observable** with a fixed boundary value anchored, within measurement uncertainty, to an independently established result. A pattern can dissolve when the inputs are cleaned up, as the ladder did. An observable whose boundary value coincides with a number fixed elsewhere does not — it stands as a target any future derivation must hit.

---

13. Status of the Programme

Quantitative contact with measurement. Before the tiers, the one-line summary that governs them: the paper currently has **zero confirmed external quantitative anchors**. There are three points of contact with measured quantities — all outputs of the same confinement programme, and all therefore carrying the *same* unverified-independence proviso — of which one is sharp and χ-gating and two are softer. The sharp one is χ(1) = ln(6⁄13): a precise ratio that fixes the first-generation value of the susceptibility. The softer two are the constituent light-quark scale (20⁄9) m_π ≈ 310 MeV and the strange estimate m_s ≈ 97 MeV. Every one of the three is an *anchor* (rather than a fit) only if the relevant confinement-programme outputs — in particular the form (K − 1)∕(2K − 1) and the value K = 7 — were fixed upstream without consulting the measured masses, which is not yet verified (§6.2, §12). Everything graded below should be read against that fact.

Established (scheme-independent), but VERSF-internal:

- Non-separability of the quark mass operator: the log-mass matrix is rank 2; separability would require χ(1) = χ(2) = χ(3), and the species split runs with depth. *This refutes a factorization expected within VERSF's own three-axis architecture, not a Standard Model expectation (§4 scope note); for the wider world it is a clean reframing of the long-known fact that L_g runs, not a new empirical result.*
- The three-axis architecture as a classification of the census.

Candidate external contacts, conditional on an unverified independence claim (weaker than "established"):

- *Sharp, χ-gating.* The light-quark species ratio 6⁄13 ≈ 0.462, reproduced as the boundary value $\chi(1)$ within the ≈ 5% measurement uncertainty. This is **established-conditional-on-an-unverified-independence-claim**, not established outright. Because $(K - 1)/(2K - 1)$ is a one-parameter family over integer K, "K = 7 lands on 0.462" is a retrodiction only if both the form *and* K = 7 are forced upstream; absent the confirmed cross-reference (§12), it cannot be distinguished from a one-parameter fit dressed as a prediction. It is a consistency target any χ-derivation must reproduce, not impose (§12) — but only once the proviso is discharged. This is the one contact sharp enough to *gate* a derivation.
- *Softer, additional contacts (same proviso).* The constituent light-quark scale (20⁄9) m_π ≈ 310 MeV reproduces the conventional constituent scale, but as a model-dependent band rather than a precise number. The strange estimate m_s ≈ 97 MeV is consistent with the measured value at the few-percent level — a credible estimate, not a tight hit. Both make genuine contact with measurement, both are outputs of the same confinement programme, and both inherit the identical unverified-independence proviso as 6⁄13; neither is sharp enough to gate a derivation.

Established within the programme, under scheme discipline:

- $\chi(g)$ is monotone increasing (three-point claim, §7.2), changes sign between the first and second generations (convention-free; no crossing *coordinate* claimed, §7.2), and decelerates in its increments (the weakest claim — a single sign bit, §7.2). These are convention-free *descriptions* of the spectrum, not contacts with an independently-fixed external number: they are read from the same masses any derivation would reproduce, so they confirm the reorganization but cannot anchor a derivation (§6.2, §12).

Established as a negative result (valid, modest, conditional):

- The carrier-growth route to χ is retired: the refinement-persistent sector H^1 is dimension-invariant under subdivision (the first Betti number is unchanged), so it cannot source a growing assignment split (§11.3). The depth-dependence must come from fold-selective *accessibility*, not carrier size. This is a one-line fact that rules out one naive route and redirects to accessibility; the redirect inherits the conditionality of the imported $\ker(W - I) = \mathbb{R}\cdot 1$ and persistence results.

Open:

- The functional form of $\chi(g)$ beyond g = 3 (saturating versus unbounded) — uncloseable from quarks alone.
- A derivation of $\chi(g)$ from interface dynamics, including its decomposition into p_{eff} and K_c contributions.
- The access–burden split of χ is not identifiable from masses (gauge freedom, §6.1); it becomes predictive only if one leg is fixed by mass-independent structure.

- The **fold-selective bath burden theorem** (§6.1) and the equivalent **fold-selective accessibility theorem** (§11.3): derive a fold-selective, depth-dependent structure on the persistent transport sector from geometry that never consults a quark mass. These are the same proof obligation seen from the burden side and the spectral side.
- Whether χ is intrinsic to assignment or is a bath effect — the lepton discriminator, deferred pending a clean neutral-lepton mass input.

Discarded:

- The interleaved heavy-quark refinement ladder s → c → b → t.
- The structural-constant reading of the factors as 13, 7⁄2, 42.
- The saturation interpretation of the susceptibility (and of the top quark) as previously stated.
- The carrier-existence reading of the cohomological sector as the source of χ's growth (replaced by the accessibility target, §11.3).

---

14. Conclusion

The quark sector can be organized by three structural axes — assignment, bath participation, and refinement depth — but the mass operator is **not separable** across them, a fact independent of mass scheme. The constructive consequence is that a single bit of fold assignment is processed by a refinement-dependent response operator,

$$\ln m(A, g) = B(g) + A \cdot \chi(g),$$

whose entire content is the susceptibility $\chi(g)$. Under scheme discipline, $\chi(g)$ grows, changes sign between the first and second generations, and decelerates. This single quantity describes both the generation-1 mass inversion and the heavy-quark hierarchy as one sequence read on either side of that sign change, is anchored at $\chi(1) = \ln(6/13)$ — within the measurement uncertainty, and conditional on that ratio having been fixed upstream independently of the masses — by a light-quark result the χ account inherits rather than re-derives, and replaces an interleaved ladder that violated the framework's own axes with an object that respects them.

The picture is not a complete derivation of the quark spectrum. It is a reorganization: the assignment sector's entire mass-dependence is isolated into a single quantity $\chi(g)$, identified as the unique obstruction to separability — so that explaining χ is equivalent to explaining the whole non-separable structure. This is a change of coordinates, not a reduction in degrees of freedom (B and χ are three numbers each); genuine compression would require deriving χ with fewer than three parameters.

Two things sharpen the open problem beyond a blank "derive χ." First, a quantitative gate: any derivation must reproduce $\chi(1) = \ln(6/13)$ without being told to — conditional on that ratio having been fixed independently of the masses, which is the paper's single unconfirmed load-bearing claim (§12). Second, a redirect: the growth of χ cannot come from the size of the framework's refinement-persistent carrier, which does not grow under refinement (the first Betti number is subdivision-invariant) — so it must come from a fold-selective, depth-dependent accessibility of that carrier, with fold-selectivity the precise step still to be derived. The first is one external lock in IOU state; the second rules out one naive route and names where the work must go. Neither adds quantitative weight; both make the grading honest.

In one sentence: the quark hierarchy cannot be written as species × generation, and the entire failure of that factorization is captured by a single observable $\chi(g)$, whose first value is anchored — conditionally — by the independently obtained 6⁄13 light-quark result, whose growth is not a carrier-size effect, and whose remaining behaviour constitutes the open derivation target. The result is not an explanation of the quark masses; it is a precise identification of what still needs explaining.