Quark Mass Ratios from Fold-Selective Accessibility

The χ Response Function and Closure-Order Orientation

Keith Taylor · VERSF Theoretical Physics Programme

General Reader Summary

Matter is built from a small number of fundamental particles. Among them are *quarks*, which come in six kinds. Nature lays them out in a tidy grid: three *generations* — think of them as three weight classes, from light to heavy — and within each generation a pair, one *up-type* quark and one *down-type* quark.

The puzzle lives in the pairs. In the lightest generation the down-type quark is the heavier of the two. But in both heavier generations this reverses, and the up-type quark becomes far heavier than its partner. The relationship inside a pair is not fixed — it flips as you climb the weight classes. Physicists have measured this for decades, but no one can say *why* it happens.

The whole pattern can be captured by a single number, written χ (the Greek letter "chi"). At each generation, χ measures how lopsided the up/down pair is. A real theory should *predict* χ from something deeper, rather than merely record the measured masses. This paper asks where χ could come from.

One natural guess is that some underlying structure simply *grows* as you move to heavier generations, carrying the masses up with it. Earlier work in this programme showed that guess cannot work: the relevant structure stays the same size — it does not grow. The idea pursued instead is subtler. The structure stays fixed, but each generation *reaches into* it differently. In the lightest generation the "down" side has the upper hand; deeper in, the "up" side reaches in more strongly and overtakes it.

Where does the "up versus down" distinction itself come from? The paper proposes that it is the difference between running a certain process *forwards* and running it *backwards* — two readings of the same operation, like a film played in either direction. Whether this can be made rigorous rests on two honest preconditions, which the paper keeps strictly separate and does not wave away. The first is to pin down exactly what kind of mathematical objects these forwards and backwards processes are; until that is settled, "backwards" is only *a* candidate opposite, not yet a clean mirror image. This first gate is called OP0. The second is to show that it is *this* particular forwards/backwards distinction that lines up with up versus down, and not some other two-way distinction the theory already contains; that is a separate assumption, named here rather than smuggled in. Throughout, the paper marks which conclusions hold before these gates are passed and which only after.

The sharpest result concerns the flip itself. If each generation simply reached into the structure *harder* than the last, the up/down balance could never reverse — it would only grow more lopsided in one direction. To produce the reversal nature actually shows, the reaching-in must change *direction*, swinging toward the "up" side as generations get heavier. "Reach in harder" is not enough; "reach in differently" is required. That is a more demanding and more testable statement than the programme began with — and any genuine derivation must produce it without peeking at the measured masses to get the answer it wants.

Table of Contents

Abstract

The VERSF analysis of the quark hierarchy established that quark masses are not separable into an independent species factor and an independent generation factor. The entire obstruction to separability is carried by a single observable, the fold-response susceptibility χ(g), the within-generation log-split between up-type and down-type quarks. That result eliminated one mechanism — the refinement-persistent transport carrier cannot source the growth of χ, because its dimension is invariant under refinement — and left a precise target: derive χ(g) from **fold-selective accessibility** of the persistent carrier, that is, from fold projectors P_up, P_down and a refinement-access operator R_g, without fitting them to the masses.

This paper does two things. First (Part I) it states the target rigorously: the defining accessibility relation in un-normalised form, with the first-generation value as a genuine *check* rather than an inserted constant; the argument that within-generation **ratios**, not absolute masses, are the renormalisation-stable objects; and the underdetermination that the access leg must resolve. Second (Part II) it proposes a source for the one object the target leaves open — the fold projectors — in **closure-order orientation**: forward closure U = A B C versus reverse closure D = C B A. This proposal is made under a single opening proof obligation, OP0, the Operator Realisation Problem, which fixes the type of every later expression. The construction is presented in two tiers: what holds purely multiplicatively (Tier 0) and what unlocks only if OP0 establishes a meaningful adjoint (Tier 1).

The paper does not claim a completed derivation. It states the precise objects to be derived, the bar for explanatory success as an explicit parameter ledger, and the structural result that — *under a favourable resolution of OP0* — the observed sign inversion between the first two generations cannot come from growth of the access operator but only from its **reorientation** across the closure subspaces. (This reorientation result is Tier 1: its sign content rests on the adjoint OP0 must supply; the linearity half survives unconditionally.) Meeting the bar would convert the

quark hierarchy from an observed pattern into a structural consequence of fold-selective refinement dynamics. Naming OP0 as the first gate is the discipline that keeps the rest honest.

Part I — The Target: What Must Be Derived

1. Introduction

The six quarks arrange into three generations and two species branches: up, charm, top; and down, strange, bottom. The masses do not increase uniformly — the up/down split changes character across generations, inverted in the first and strongly up-heavy thereafter.

The preceding paper established that this hierarchy is **not separable**: quark mass cannot be written as $m(A, g) = S(A) \cdot D(g)$, with A the up/down assignment and g the generation, for then the up/down ratio would be constant across generations, which it is not. The replacement object is the fold-response susceptibility

$\chi(g) = \ln[m\_up(g) / m\_down(g)]$,

in terms of which the log-mass reads

$\ln m(A, g) = B(g) + A \cdot \chi(g)$, (1)

with B(g) the common bath/refinement baseline and $A \in \{0, 1\}$ the one-bit up/down fold assignment.

Equation (1) is, by itself, an identity rather than a result: at fixed g there are only two log-masses, and any pair can be written this way. The physical content lives downstream, in two claims (1) does not contain — that the bit A denotes the *same* up/down assignment at every generation (so χ is one coherent response function), and that this χ is generated by a specific, independently fixable mechanism. Part I states the target that mechanism must hit; Part II proposes the mechanism.

The preceding paper also showed χ cannot be sourced by growth of the refinement-persistent carrier: it persists under refinement but its dimension is invariant. This paper therefore pursues:

χ(g) is generated by **fold-selective accessibility** of a persistent transport carrier — and the up/down folds are generated by **closure-order orientation**.

---

2. Carrier Growth Is Ruled Out — Generically, Within a Stated Scope

The carrier-growth route runs: assignment needs a structure surviving refinement; the first non-trivial refinement-persistent sector is the cohomological transport sector H^1; so χ might grow with H^1. The last step fails. Under midpoint refinement the cycle rank $b_1 = E - V + C$ is invariant, since $V \to V + E$ and $E \to 2E$ give $b_1' = 2E - (V + E) + C = E - V + C = b_1$, so $\dim H^1(g + 1) = \dim H^1(g)$. A dimension-invariant carrier cannot supply a growing susceptibility.

The no-go is stronger than it looks, within a stated scope. Midpoint refinement is a subdivision — a topology-preserving operation — so the invariance of b_1 is an instance of homotopy invariance: **for refinement realised as subdivision, no purely topological invariant can grow.** The scope qualifier is essential and is carried to the conclusion: should any VERSF refinement move do more than subdivide — adding gluings, relations, or commitment structure that changes homotopy type — the genericity would overreach and b_1-invariance would require separate verification for that move.

Within the subdivision scope the claim has two consequences. The elimination of carrier growth is robust, not a coincidence about one sector. And — against convenience — *persistence under refinement cannot be what privileges H^1*, since every topological carrier persists. H^1 is singled out because it is the **first non-trivial assignment-bearing sector**, not because it survives. Survival is necessary but generic; assignment-bearing capacity is what selects it.

So the carrier remains necessary as the stable arena in which fold-assignment information lives, but the growth of χ must come from *how the folds access that arena*, not the arena's size:

carrier persistence supplies the arena; fold-selective accessibility supplies the mass response.

This raises a question answered fully only in Part II: if H^1 is dimension-invariant, what makes the access operator R_g vary with depth? The depth-dependence does not live in the persistent (topological) sector; it enters through the **transient, non-persistent spectrum** of the refinement dynamics — the near-marginal modes (the $\lambda_2(k)$ modes of the companion paper) whose eigenvalues drift with depth while the homology rank stays fixed. R_g is the depth-dependent modulation of access to a depth-invariant arena.

---

3. The Accessibility Hypothesis

Let P_{up} and P_{down} be fold projectors on H^1 for the up- and down-type branches, and R_g the refinement-access operator. Define the fold accessibility by

$$\mathcal{A}_A(g) = \mathrm{Tr}_{H^1}[P_A R_g P_A], \quad A \in \{up, down\}. \quad (2)$$

For an orthogonal projector P_A ($P_A^2 = P_A$) and the cyclic trace, (2) equals

$$\mathcal{A}_A(g) = \mathrm{Tr}_{H^1}[P_A R_g], \quad (2')$$

an identity that removes an illusory modelling freedom: rewriting the trace does not reduce the freedom; only constraining the ingredients does. The susceptibility is **generated** by the accessibility ratio. The defining relation is the un-normalised form

$$\chi(g) = \chi(1) + \ln[\mathcal{A}\_up(g) / \mathcal{A}\_down(g)] - \ln[\mathcal{A}\_up(1) / \mathcal{A}\_down(1)], \quad (3)$$

which carries two distinct obligations that must not be collapsed:

- **The running obligation.** The g-dependent difference must reproduce the *change* of χ relative to the first generation.
- **The anchor obligation.** The *raw, un-normalised* first-generation ratio must itself reproduce the light-quark value,

  $$\ln[\mathcal{A}\_up(1) / \mathcal{A}\_down(1)] = \chi(1) = \ln(6/13), \quad (3a)$$

  as a genuine output of the g = 1 fold geometry, not a constant inserted by hand.

These are separate tests. A construction satisfying (3) but failing (3a) has explained the *shape* of the hierarchy while leaving its first-generation *value* unexplained.

For brevity define the normalised running ratio $\boldsymbol{\mathcal{R}}(g) = [\mathcal{A}\_up(g)/\mathcal{A}\_down(g)] / [\mathcal{A}\_up(1)/\mathcal{A}\_down(1)]$, with $\boldsymbol{\mathcal{R}}(1) = 1$, so (3) reads $\chi(g) = \ln(6/13) + \ln \boldsymbol{\mathcal{R}}(g)$. This convenience **silently discharges the anchor obligation**: dividing out the g = 1 ratio makes $\chi(1) = \ln(6/13)$ true by construction and removes (3a) from view. $\boldsymbol{\mathcal{R}}$ is therefore used below only as shorthand for the *running*; whenever the anchor is at stake, the operative statements are (3) together with (3a).

The derivation target is now precise: derive both the running $\boldsymbol{\mathcal{R}}(g)$ **and** the raw first-generation ratio (3a) from fold geometry and refinement dynamics, **without fitting either to the quark masses.** The observed χ requires $\chi(1) = \ln(6/13)$, exactly one sign change between the first and second generations, $\chi(3) > \chi(2)$, and sublinear growth ($\chi(3) - \chi(2) < \chi(2) - \chi(1)$).

---

4. Why Fold-Selectivity Is Essential

A refinement-access operator R_g alone is not enough. If R_g acts identically on the two folds it cancels from the ratio $\mathcal{A}\_up/\mathcal{A}\_down$: any fold-blind effect contributes equally to both branches and cannot generate χ. The discipline:

depth-dependence is necessary, but fold-selectivity is decisive.

A mechanism that varies with g but carries no up/down distinction cannot explain the split; one that distinguishes the folds but does not vary with g cannot explain the running. A non-trivial result requires $[P\_up, R\_g] \neq [P\_down, R\_g]$ as operators on H^1, not merely $R\_g \neq R\_{g'}$. Near-marginal refinement modes are suggestive — they have the right depth-sensitivity — but unless

P_{up} and P_{down} project differently onto them they remain fold-blind and cancel out of χ. This is why the folds must be genuinely distinct objects, which is exactly what Part II proposes to supply.

---

5. Targets Must Be Ratios: Flavour-Universality and Scheme Dependence

A first objection is bookkeeping: quark masses are scheme- and scale-dependent, so *which* m_t/m_b is targeted? The answer is also the construction's strongest defence. The targets are within-generation **ratios**, not absolute masses, and ratios are far better-defined under QCD running. The mass anomalous dimension γ_m is **flavour-universal** — it depends on the gauge coupling, not the quark's identity — so to leading order every running mass at a common scale carries the same multiplicative renormalisation, which cancels in any same-scale ratio:

$m_{up}(\mu) / m_{down}(\mu)$ is RG-invariant to good approximation.

This is what makes $\chi(g)$ a well-posed target at all: the absolute baseline $B(g)$ in (1) is scheme- and scale-laden, while the ratios captured by χ are comparatively robust. Two caveats belong in any honest statement. First, scheme dependence of the anchor: the identification $6/13 \approx m_u/m_d$ holds in the light-quark MS-bar ratio (in a pole-mass scheme the light-quark numbers are not even well-defined), so the construction must target that same ratio. Second, breaking of flavour universality: electroweak and quark-threshold effects break exact universality of γ_m at sub-leading order — small for the ratios of interest but not zero, so a complete treatment specifies the scale and scheme. Targeting χ rather than absolute masses is not a retreat from ambition; it is the move that makes the target survive renormalisation.

---

6. Access and Burden: Underdetermination, Not Gauge

The VERSF mass anchor writes mass through completion density, $m \propto p_{eff} / K_c$, so the up/down log-split is

$$\chi(g) = \Delta\ln p_{eff}(g) - \Delta\ln K_c(g). \quad (4)$$

Equation (4) is exact but not yet a theory: the masses fix only the *difference* of the access leg $\Delta\ln p_{eff}$ and the burden leg $\Delta\ln K_c$, never the two legs separately. It is tempting to call the freedom $\Delta\ln p_{eff} \to \Delta\ln p_{eff} + f(g)$, $\Delta\ln K_c \to \Delta\ln K_c + f(g)$ a "gauge freedom," but the phrase overstates the structure. There is no symmetry; this is one equation in two unknowns, and a single observable cannot determine two functions. The accessibility proposal is best understood not as fixing a gauge but as **supplying the missing equation** — deriving the access leg independently through the fold-selective action of R_g, $\Delta\ln p_{eff}(g) \equiv \ln[\mathcal{A}_{up}(g)/\mathcal{A}_{down}(g)]$. If that succeeds, the burden leg becomes a checkable output, $\Delta\ln K_c(g) = \Delta\ln p_{eff}(g) - \chi(g)$, and (4) acquires predictive content. Until P_up, P_down, R_g are fixed without consulting the masses, this is a programme that *adds* an equation, not one that exploits a symmetry.

Part II — The Mechanism: Closure-Order Orientation as the Source of the Folds

7. The Fold Projectors Are the Open Object

Part I leaves exactly one object underived: the fold projectors P_up, P_down. It requires P_up ≠ P_down "from assignment geometry, not from quark masses," but does not say what produces them. Closure orientation — forward closure versus reverse closure — is a binary geometric distinction of the right kind. Its role is narrow and precise: **it is a candidate source of P_up and P_down; everything downstream is Part I unchanged** — provided the opening gate (OP0) is passed. Crucially, the susceptibility must remain a difference of logarithms, as (4) demands. A bare trace carries no logarithm and cannot reconcile with the anchor's log structure; so closure orientation must feed the *projectors inside the log-ratio* (3), not replace χ with a linear trace.

Two posits are bundled in "candidate source of the right kind," and only one of them is OP0. OP0 asks whether A, B, C are operators of a suitable class. A second, logically independent posit is that closure-order orientation is the binary distinction that maps to **up/down at all** — rather than one binary among several the framework already carries (quark/lepton bath participation; weak isospin; with colour three-valued). The word "natural" does not discharge this, and given the discipline the rest of the paper applies to free riders, it should be named:

OP0.5 — Assignment Identification. Establish that the up/down fold bit *is* closure-order orientation specifically, and not merely one available reversible/irreversible distinction. The natural route to discharging OP0.5 would be to identify closure orientation with the weak-isospin label $T_3 = \pm\frac{1}{2}$ used elsewhere in the programme; that identification, if it can be made, closes OP0.5 — but it is not made here, and until it is, the closure-order/up-down correspondence is an open posit additional to OP0.

Everything in Part II is therefore conditional on *two* gates, OP0 and OP0.5, not one. The tier marks below track OP0 (the operator class); OP0.5 (the identification) sits above the whole of Part II and is recorded in §18.

---

8. The First Gate — OP0, the Operator Realisation Problem

Nothing below denotes a physical observable until A, B, C are given a definite realisation.

OP0 — Operator Realisation Problem. Define the VERSF closure operations A, B, C: specify the space on which they act, prove whether that space is H^1, and determine whether they belong to a unital $*$-algebra with a *meaningful adjoint* †. If so, determine whether A, B, C are individually self-adjoint.

OP0 gates the type of every later expression. An adjoint requires an inner-product structure, and that same structure is what makes a projector *orthogonal* (hence Hermitian) rather than merely idempotent. So OP0 controls three things at once: whether reverse closure is the adjoint of forward closure, whether the fold projectors are Hermitian, and whether the accessibility traces are real and sign-definite. The section is therefore written in two tiers, kept strictly apart:

- **Tier 0 (unconditional).** Uses only the associative product of A, B, C; holds regardless of OP0.
- **Tier 1 (conditional on OP0).** Requires a meaningful adjoint, orthogonality, positivity, or reality.

The favourable outcome has a name for convenience:

H0 (favourable resolution of OP0). A, B, C are self-adjoint operators on H^1 in a unital $*$-algebra with a meaningful adjoint.

H0 is a hoped-for outcome of OP0, not a premise. Where a result needs it, this is flagged "(Tier 1: assumes H0)". The bridge to Part I — closure operators acting on the *same* H^1, so the projectors they generate are literally P_up, P_down — is itself part of what OP0 must establish.

---

9. Forward Closure, Reverse Closure, and the Residue

(Tier 0.) Define forward and reverse closure by ordering alone, and their difference:

$U = A\,B\,C$ (forward), $D = C\,B\,A$ (reverse), $\Omega = U - D$.

This is all that is licensed before OP0. In particular we do **not** write $D = U^\dagger$, because the adjoint is exactly what OP0 has not yet supplied. The honest reading:

Reversal of closure order defines a **candidate opposite fold** — not necessarily an **adjoint fold**.

D is the opposite-ordered product; $\Omega = U - D$ is a difference whose Hermitian/skew type is undetermined at Tier 0. This is why χ cannot be a bare trace of Ω: $\mathrm{Tr}(R_g\,\Omega)$ is real only if Ω and R_g have compatible types, which is unknown until OP0. The construction must route the asymmetry through an object whose type it controls, and defer the type claim to OP0.

(Tier 1: assumes H0.) With self-adjoint A, B, C, reverse closure *is* the adjoint of forward closure, $D = C\,B\,A = (A\,B\,C)^\dagger = U^\dagger$, and the residue is anti-Hermitian, $\Omega^\dagger = -\Omega$, with a clean Hermitian/skew decomposition. Only here is it legitimate to argue that $\mathrm{Tr}(R_g\,\Omega)$ is imaginary for self-adjoint R_g. If OP0 returns a different class (non-self-adjoint generators, a weaker involution, or no meaningful adjoint), this specialization is withdrawn and §9 stops at $\Omega = U - D$.

---

10. Fold Projectors from Closure Orientation (the bridge)

The projector construction survives OP0-agnostically in its *idempotent* form, with orthogonality deferred to OP0.

(Tier 0.) Forward and reverse closure each have a range; form idempotents projecting onto them,

E_{up} onto ran(U) = ran(A B C), E_{down} onto ran(D) = ran(C B A),

well-defined using only the multiplicative structure. These are the **candidate opposite folds**.

(Tier 1: assumes H0.) With a meaningful adjoint there are *two* projector conventions onto the closure ranges, and they coincide iff U is normal. (i) The **orthogonal range projectors** P_{up} onto ran(U) and P_{down} onto ran($U^\dagger$) are Hermitian and positive semi-definite **regardless of whether U is normal** — orthogonality is fixed by the inner product, not by U's spectral behaviour. (ii) The **closure-canonical idempotents** (range ran(U) along ker(U), etc.) are *oblique* — non-Hermitian — precisely when U is non-normal. Both are consistent with the non-triviality condition below; which one is the physical fold is a sub-question of OP0. §§11–12 are written for convention (i), where everything is manifestly real; §11 then shows convention (ii) changes none of the conclusions, at the cost of a symmetrisation.

The real non-triviality condition (Tier 0). "A B C ≠ C B A" is nearly content-free — non-commutativity of three generic operators is automatic. The operative requirement is that the folds occupy **different subspaces**,

$E_{up} \neq E_{down} \Leftrightarrow$ ran(A B C) ≠ ran(C B A), (★)

stated without any adjoint, and strictly stronger than A B C ≠ C B A.

(Tier 1: assumes H0.) Under an adjoint, (★) acquires a structural reading. A normal operator has ran(U) = ran($U^\dagger$), so for the **orthogonal** range projectors of convention (i), $P_{up} \neq P_{down} \Rightarrow [U, U^\dagger] \neq 0$. **Non-normality of the closure operator is a necessary condition for distinct folds** — "forward and reverse closure reach the same subspace" is exactly normality, and the mechanism lives in its failure. Note this is fully consistent with the clean Hermitian picture: distinct subspaces give distinct *orthogonal* (hence Hermitian) projectors, so non-normality does not by itself force a non-Hermitian carrier — it only does so if convention (ii) is adopted.

The carrier. Define the fold difference $\Delta E = E_{up} - E_{down}$ (Tier 0). Under H0 with the orthogonal projectors of convention (i) this is the Hermitian $\Delta P = P_{up} - P_{down}$, spectrum in

$[-1, 1]$: positive on the forward-only subspace, negative on the reverse-only subspace, zero on shared subspaces. ΔP is the Hermitian avatar of the anti-Hermitian residue Ω — same asymmetry, real and traceable. (The convention-(ii) oblique case, where ΔE is non-Hermitian, is handled in §11 and shown not to change the conclusions.)

---

11. The Accessibility Functional in Log-Ratio Form

The architecture of §3's functional survives at Tier 0 as a definition; its interpretation as real, non-negative observables is Tier 1.

(Architecture, Tier 0.) With the fold idempotents and R_g (depth-dependence in the transient $\lambda_2(k)$ spectrum),

$\mathcal{A}_{up}(g) = \mathrm{Tr}_{H^1}[R_g E_{up}]$, $\mathcal{A}_{down}(g) = \mathrm{Tr}_{H^1}[R_g E_{down}]$,

and χ keeps the log-ratio form (3), $\chi(g) = \ln \mathcal{A}_{up}(g) - \ln \mathcal{A}_{down}(g)$ — a difference of logarithms, as the mass anchor (4) demands, holding as a definition regardless of OP0.

(Tier 1: assumes H0.) Only with a meaningful adjoint and $R_g \succeq 0$ are $\mathcal{A}_{up}$, $\mathcal{A}_{down}$ guaranteed **real and non-negative**, so (3) denotes a real susceptibility. At this tier the sign of a log-ratio equals the sign of the difference, giving the exact **sign law**

$\mathrm{sign}\, \chi(g) = \mathrm{sign}(\mathcal{A}_{up}(g) - \mathcal{A}_{down}(g)) = \mathrm{sign}\, \mathrm{Tr}[R_g \Delta P] \equiv \mathrm{sign}\, \delta(g)$, (5)

with $\delta(g) = \mathrm{Tr}[R_g \Delta P]$ the real **fold imbalance**. The whole question of which fold is heavier at depth g reduces to the sign of one real number.

The functional form $\mathrm{Tr}[R_g P_A]$ is the forced one under linearity + positivity + basis-independence: requiring $\mathcal{A}_A$ to be linear in R_g, non-negative for non-negative R_g (the projector sandwich (2) gives $P_A R_g P_A \succeq 0$), and basis-independent (the trace) singles it out essentially uniquely, so it carries no free parameter. But "positivity" and "basis-independence" are Tier-1 notions, so even the forcing argument is downstream of OP0; pre-OP0 the functional form is a discrete modelling choice to be logged.

Robustness to the projector convention (Tier 1). The sign law (5) and the result of §12 are written for convention (i), where ΔP is Hermitian. They survive convention (ii), where the physical fold carrier $\Delta E = E_{up} - E_{down}$ is the non-Hermitian difference of oblique idempotents, with only a symmetrisation. For self-adjoint R_g, cyclicity gives $\mathrm{Tr}[R_g \Delta E]^* = \mathrm{Tr}[(R_g \Delta E)^\dagger] = \mathrm{Tr}[\Delta E^\dagger R_g] = \mathrm{Tr}[R_g \Delta E^\dagger]$, so the real part is

$\mathrm{Re}\, \mathrm{Tr}[R_g \Delta E] = \mathrm{Tr}[R_g \cdot \tfrac{1}{2}(\Delta E + \Delta E^\dagger)] = \mathrm{Tr}[R_g \Delta E_{sym}]$, $\Delta E_{sym} = \tfrac{1}{2}(\Delta E + \Delta E^\dagger)$ Hermitian.

Defining the observable accessibilities as real quantities ($\mathcal{A}$_A = Re Tr[R_g E_A], equivalently using ΔE_sym in the imbalance), the imbalance δ(g) = $\mathcal{A}$_up − $\mathcal{A}$_down = Tr[R_g ΔE_sym] is real, so the sign law (5) holds verbatim with ΔE_sym in place of ΔP, and the §12 growth-cannot-flip conclusion (δ(g) = c_g·const has fixed sign) is unchanged. What convention (ii) does *not* hand over for free is **positivity** of the individual $\mathcal{A}$_A: for oblique E_A, $\mathcal{A}$_A ≥ 0 is an extra condition rather than an automatic consequence of R_g ⪰ 0. So the headline result is robust to which projector OP0 selects; only the positivity leg distinguishes the two conventions, and that distinction folds into OP0.

Anchor as a genuine check. There is no additive ln(6/13) in (3). The anchor is a constraint to *satisfy*: χ(1) = ln(6/13) ⟺ Tr[R_1 E_up] / Tr[R_1 E_down] = 6/13 (3a). Because (3) carries no separate constant, the first-generation residue does real work — reproducing χ(1) = ln(6/13) < 0 requires R_1 to weight the down-fold over the up-fold in the ratio 6 : 13. An additive form χ = ln(6/13) + ω(g) would instead force ω(1) = 0, making the residue invisible exactly where the first-generation split is measured; the un-normalised form keeps it visible.

---

12. The Sign Change Is Reorientation, Not Growth

(Tier 1: assumes H0, so δ(g) is real.) From (5), a single sign change between the first two generations is exactly δ(1) = Tr[R_1 ΔP] < 0 and δ(2) = Tr[R_2 ΔP] > 0, with no further crossing for g ≥ 2. The key impossibility result:

Proposition (growth cannot flip). If the access operator grows only in scale, R_g = c_g R_0 with c_g > 0 and fixed shape R_0, then δ(g) = c_g Tr[R_0 ΔP] has the sign of Tr[R_0 ΔP] for every g, so δ never changes sign and χ never flips.

The argument is linearity of the trace (Tier 0); the conclusion about sign needs δ real (Tier 1, via either the Hermitian ΔP of convention (i) or the symmetrised ΔE_sym of convention (ii) — §11). The consequence is structural: a positive, uniformly growing access operator — the natural reading of "accessibility increases with refinement" — actively **prevents** the sign change. To flip, R_g must **reorient** across the eigenspaces of the (Hermitian) carrier, shifting weight from the negative subspace (reverse-accessible) toward the positive subspace (forward-accessible) between g = 1 and g = 2. The flip lives in the *direction* of R_g in the carrier's eigenbasis, not its magnitude. This is sharper and harder than "make accessibility increase," and it is the content of OP4.

---

13. The Corrected Reduction

A weaker statement — assuming only that the running is increasing and sublinear — does *not* force a sign change: take running increments 0, 0.1, 0.15 (monotone, sublinear) and no crossing occurs; and once the missing quantitative premise (the running must pass |ln(6/13)| ≈ 0.77 by g =

2) is added, the "conclusion" merely restates it. The honest statement separates bookkeeping from what must be derived, and tiers it.

Reduction (honest form). Adopt §§3, 9–11 and define $\delta(g) = \mathrm{Tr}[R_g \Delta P]$.

(a) **Anchor (Tier 0 form).** $\chi(1) = \ln(6/13)$ iff (3a) — a constraint on R_1, not an output of the form. (b) **Sign law (Tier 1).** sign $\chi(g)$ = sign $\delta(g)$, exactly. (c) **Single flip ⇔ single zero-crossing of δ (Tier 1).** χ changes sign exactly once between $g = 1$ and $g = 2$ iff $\delta(1) < 0 < \delta(2)$ and $\delta(g) > 0$ for $g \geq 2$. (d) **Necessity of reorientation (Tier 0 algebra, Tier 1 sign).** If $R_g = c_g R_0$, δ has constant sign and (c) fails; hence any construction satisfying (c) must have R_g reorient across the eigenspaces of ΔP.

Parts (a), (b), (d) carry genuine deductive content — (a) makes the anchor a test, (b) is an exact identity, (d) is a real impossibility result — and (b)–(d) are downstream of OP0. Part (c) is the **target**, not a theorem: precisely what must be *derived* from R_g and the fold geometry. The magnitude and deceleration requirements ($\chi(2) > 0$, $\chi(3) > \chi(2)$, $\chi(3) - \chi(2) < \chi(2) - \chi(1)$) are likewise targets for the derived R_g.

Part III — Assessment

14. What Would Count as Success, and the Parameter Ledger

A completed derivation must establish: fold projectors $P_{up} \neq P_{down}$ from closure geometry, not masses; a depth-dependent R_g motivated by refinement dynamics; fold asymmetry $[P_{up}, R_g] \neq [P_{down}, R_g]$; the qualitative χ behaviour as an output — the raw $g = 1$ ratio reproducing the anchor via (3a) (not a normalised $\chi(1)$ true by construction), $\chi(1) < 0 < \chi(2)$ with exactly one sign change, $\chi(3) > \chi(2)$, and decelerating growth; and no smuggled saturation (if R_g is built from a saturating spectral quantity, either keep that saturation out of χ or *predict* that χ saturates at deeper refinement and own it).

The first milestone is qualitative, not quantitative. Magnitudes are cheap to fit; structure is not. The single most valuable early output — before any number is matched — is to show that closure-fold accessibility forces **exactly one** crossing of χ (first generation inverted, every higher generation not), via the reorientation result §12 / reduction (c)–(d). A construction in which the inversion at $g = 1$ and its single reversal at $g = 2$ are *forced* by the geometry of P_{up}, P_{down}, R_g demonstrates real content even before the curve is calibrated.

The "no fitting" discipline becomes binding only when free parameters are counted against the facts they must reproduce. Counted in the un-normalised form (3)+(3a), the construction must produce **three** χ values, one of which (the $g = 1$ anchor) is a hard check.

Quantity	Count	Notes
χ values the mechanism must reproduce	$\chi(1)$, $\chi(2)$, $\chi(3)$ → **3**	All outputs via (3) and (3a); the g = 1 value is not divided out.
Externally fixed in value	$\chi(1) = \ln(6/13)$ **(conditional — §15)**	If 6/13 is derived upstream of m_u/m_d, the g = 1 ratio is a genuine check that *raises* the bar, not one that removes a target.
Ingredients to be fixed upstream	{A, B, C} → P_up, P_down; and R_g	Conditional on OP0 that these are well-typed operators on H^1.
Free parameters after upstream fixing	**N_free**	Must be reported honestly.
Accessibility functional Tr[R_g P_A]	**0** if forced (Tier 1), else **+1 discrete**	"Forced" needs the linearity + positivity + basis-independence argument, which presumes OP0.

The rule is sharp: **N_free ≥ 3 ⟹ guaranteed fit, zero explanatory content** — three measured ratios re-encoded as unknown operators. A genuine derivation needs N_free strictly less than the number of χ values reproduced; the ideal is N_free = 0, with the single flip and the anchor both emerging. The ledger is meaningful only once OP0 certifies that {A, B, C, R_g} are genuine operators; until then N_free is uncounted because the objects are unfixed.

---

15. The First-Generation Anchor and the Upstream Question

The first-generation ratio is special because it is not merely read off the χ sequence. In the VERSF confinement programme the light-quark species ratio is $m_u/m_d = (K - 1)/(2K - 1)$, and for K = 7, $m_u/m_d = 6/13$, so $\chi(1) = \ln(6/13)$. This anchor is the boundary condition and cannot double as a fit parameter. The correct standard is not "choose $\mathcal{A}$_up(1), $\mathcal{A}$_down(1) so that $\chi(1) = \ln(6/13)$"; it is "derive P_up, P_down, R_1 so that their trace ratio *lands on* the 6/13 value already fixed upstream" — i.e. pass (3a).

One question, which the programme should resolve rather than leave conditional:

Is K = 7 — and hence (K−1)/(2K−1) = 6/13 — derived upstream independently of the measured light-quark ratio, or was K chosen to land near it?

Per §14, neither answer removes $\chi(1)$ from the values to reproduce; they change what reproducing it *means*. If 6/13 is independently derived, passing (3a) tests the construction against a value it had no hand in setting — a hard external lock, raising the bar. If K = 7 was selected to match the measured ratio, the construction must still reproduce $\chi(1)$ via (3a), but doing so no longer tests it against anything independent, and all three χ values must come from N_free. This is the same status question that governs $\kappa = 8/3$ in the lepton-sector arc (parameter-free check vs calibration), and should be answered with the same explicitness. The paper's logic is exactly as strong as the answer.

---

16. Consequence for Quark Mass Ratios

If $\chi(g)$ is derived, the within-generation ratios follow immediately:

$m\_up(g) / m\_down(g) = \exp[\chi(g)]$,

so $m_u/m_d = \exp[\chi(1)]$, $m_c/m_s = \exp[\chi(2)]$, $m_t/m_b = \exp[\chi(3)]$. Deriving χ is therefore equivalent to deriving the non-separable within-generation mass-ratio hierarchy. The absolute masses would still require the baseline $B(g)$; the within-generation ratios require only $\chi(g)$, and §5 explains why the ratios — not the absolute masses — are the renormalisation-stable objects to target first.

---

17. What This Paper Does Not Claim

This paper does not claim to have derived all six quark masses, that the absolute masses are fixed, or that $\chi(g)$ has been computed from first principles. It does not claim that near-marginal transport modes automatically explain the hierarchy, that closure orientation has been *proven* to source the folds, that closure orientation has been *identified* with the up/down bit (OP0.5, open), or that the closure operators have been realised as adjoint-bearing operators on H^1 (OP0, open).

The claim is narrower: given the elimination of carrier growth (a generic topological fact within the subdivision scope), the only viable carrier-based route to χ is fold-selective, depth-dependent accessibility; closure-order orientation is a candidate source of the folds whose identification with up/down is itself a named posit (OP0.5); and *if* it is realised in a suitable operator class (OP0) *and* it is the right binary (OP0.5), the observed inversion is forced to be a reorientation rather than growth. The paper states the route, the renormalisation sense in which its targets are well-posed, the two gates (OP0, OP0.5) that must be passed first, and the ledger by which it would or would not become a derivation.

---

18. Open Problems (OP0 first)

OP0 — Operator Realisation (the gate). Define A, B, C; specify their space; prove whether they act on H^1; determine whether they live in a unital $*$-algebra with a meaningful adjoint, and whether they are self-adjoint. Until OP0 is solved, §§9–10 stand only at Tier 0 ($D = C\,B\,A$ is the opposite-ordered product, $\Omega = U - D$, folds candidate-opposite not adjoint), and everything marked Tier 1 in §§9–13 is suspended. A sub-question of OP0: which projector convention —

orthogonal (i) or oblique closure-canonical (ii) — is the physical fold; the headline result survives either (§11), but the positivity leg distinguishes them.

OP0.5 — Assignment Identification. Establish that the up/down fold bit *is* closure-order orientation specifically, and not merely one available binary distinction among those the framework carries. The natural route is to identify closure orientation with the weak-isospin label $T_3 = \pm½$ used elsewhere in the programme; that identification, if it can be made, discharges OP0.5. It is not made here, so the closure-order/up-down correspondence is an open posit sitting above the whole of Part II, additional to OP0.

OP1 — real non-triviality. Show ran(A B C) ≠ ran(C B A) (★) in the closure algebra — not the near-trivial A B C ≠ C B A. If OP0 supplies an adjoint, sharpen to the necessary non-normality $[U, U^\dagger] \neq 0$.

OP2 — access operator. Derive R_g from refinement dynamics, depth-dependence in the transient $\lambda_2(k)$ spectrum, not the persistent sector.

OP3 — magnitude and running. Show the derived δ(g) yields $\chi(2) > 0$ and ln-sublinear growth thereafter.

OP4 — forced single flip (sharpened). Show R_g reorients across the eigenspaces of ΔP so δ(g) crosses zero **exactly once**, between g = 1 and g = 2 — the flip derived as reorientation, not granted as growth.

OP5 — generation count. "Refinement gives three levels" accommodates three generations; it does not explain why exactly three. Same open status as OP0/OP1 and the upstream anchor; not to be counted as settled.

OP6 — no smuggled saturation. If R_g derives from a saturating spectral quantity, either keep the saturation out of χ or make χ-saturation an explicit prediction.

---

19. Conclusion

The preceding quark-hierarchy paper identified χ(g) as the unique obstruction to separability and ruled out carrier growth. This paper sets the target precisely and proposes the missing ingredient. The carrier H^1 persists under refinement but does not grow — generically, within the subdivision scope — so the growing up/down split must come from how each fold accesses the carrier. The up/down folds themselves are proposed to come from closure-order orientation, forward versus reverse closure, under the discipline that nothing is claimed about reality, adjoints, or Hermiticity until the Operator Realisation Problem OP0 is solved.

The derivation target, stated with its gate:

Solve OP0 to realise A, B, C as operators on H^1; thereby derive P_up, P_down (forward/reverse closure ranges) and R_g (transient $\lambda_2(k)$ spectrum) without consulting quark masses, such that $\mathcal{A}$_A(g) = Tr_{H^1}[R_g P_A] satisfies (3) with the anchor obligation (3a) — the raw g = 1 ratio reproducing ln(6/13) as a check, not an inserted constant — with **exactly one** sign change between the first and second generations and sublinear growth thereafter, and report N_free (and the functional-form status) so the result can be told apart from a fit.

In one sentence: the quark mass ratios will be derived when the χ response function is shown to arise from a refinement-access operator that *reorients* across the forward/reverse closure subspaces of a refinement-persistent carrier — expressed in log-ratio form, anchored by a genuine 6/13 check, with a published parameter count and a forced single inversion — rather than from the carrier's size; and the first thing the programme owes a proof of is OP0.

Dictionary (tier-marked)

Structure	Object	Tier	Meaning
Forward closure	U = A B C	0	Forward-ordered closure
Reverse closure	D = C B A	0	Opposite-ordered closure (candidate opposite fold)
Reverse = adjoint	D = U†	1 (H0)	Only if OP0 gives self-adjoint generators
Residue	Ω = U − D	0	Asymmetry; Hermitian/skew type **unknown** until OP0
Fold idempotents	E_up, E_down onto ran(U), ran(D)	0	Candidate opposite folds
Orthogonal fold projectors	P_up, P_down	1 (H0)	Hermitian folds; reality of traces
Carrier of asymmetry	ΔE (0) / ΔP = P_up − P_down (1)	0 / 1	Real, traceable asymmetry once OP0 holds
Refinement depth	g = 1, 2, 3	0	Generation
Access operator	R_g (transient $\lambda_2(k)$ spectrum)	1 for R_g ≥ 0	Refinement readout
Fold imbalance	δ(g) = Tr[R_g ΔP]	1	sign χ(g) = sign δ(g)
Response	χ(g) = ln($\mathcal{A}$_up/$\mathcal{A}$_down)	form 0 / real 1	Observable mass-ratio split

---

Status

Established upstream

- Quark mass is non-separable across assignment and refinement; $\chi(g)$ captures the whole within-generation mass-ratio hierarchy.
- The g = 1 ratio is checked against the 6/13 light-quark result via (3a) — a genuine check, not a normalised input (its independence is conditional on §15).
- Carrier growth is ruled out, generically within the subdivision scope: no topological invariant grows under refinement realised as subdivision.

Established here, unconditionally (Tier 0)

- Forward/reverse closure defined by ordering alone: U = A B C, D = C B A, Ω = U − D. Reverse is a candidate opposite fold, not necessarily an adjoint fold.
- The fold idempotents and the operative non-triviality condition ran(U) ≠ ran(D) (★), stated without any adjoint.
- χ in log-ratio form (3) as a definition, consistent with the mass anchor; the g = 1 value a genuine check (3a).
- The impossibility argument "scalar growth has fixed sign" (linearity of trace).

Conditional on OP0 / H0 (Tier 1)

- $D = U^\dagger$; Ω anti-Hermitian; Hermitian fold projectors and carrier (convention (i)), or the symmetrised carrier ΔE_sym (convention (ii)).
- Reality and non-negativity of $\mathcal{A}$_up, $\mathcal{A}$_down; the exact sign law sign $\chi(g)$ = sign $\delta(g)$; the **headline reorientation result** (growth cannot flip the sign — its sign content is Tier 1, the linearity half Tier 0); the forcing of the trace functional. The reorientation result is robust to the orthogonal/oblique projector convention (§11); only the positivity leg distinguishes them.

Still open

- OP0 first, then OP0.5 (assignment identification: is up/down = closure orientation, e.g. via $T_3 = \pm\frac{1}{2}$?), then OP1–OP6 (§18). OP0 is the gate for the operator class; OP0.5 sits above all of Part II.
- Derive R_g from refinement dynamics; exhibit fold-selectivity without using quark masses; obtain magnitude and sublinear running; force the single flip as reorientation; resolve whether 6/13 is derived or calibrated; settle whether the accessibility functional is forced or a logged choice; settle which projector convention is physical (positivity leg); confirm refinement is subdivision-only or verify b_1-invariance for any non-subdivision move; explain the generation count; determine whether χ saturates at deeper refinement.

The corrected conjecture (stated with its gates)

If OP0 establishes the VERSF closure operators as genuine adjoint-bearing operators on H^1, *and* OP0.5 identifies closure orientation with the up/down bit, *then* the quark mass hierarchy is produced not by a growing carrier and not by a bare residue, but by a refinement-access operator that **reorients** across the forward/reverse closure subspaces — turning the down-fold's first-generation advantage into the up-fold's advantage thereafter — expressed in log-ratio accessibility form and anchored by a genuine 6/13 check at g = 1. Naming OP0 and OP0.5 as the gates does not weaken this; they are the first things the programme owes a proof of.