Closure Orientation, Fold Access, and a Candidate Origin of the Same-Generation Quark Mass Ratios

Keith Taylor · VERSF Theoretical Physics Programme

General Reader Summary

Everything you can touch is built out of tiny particles called quarks. There are six kinds, and they sort neatly into two families of traits: every quark is either an "up-type" or a "down-type," and each comes in one of three "generations" — think of the generations as light, heavy, and heavier versions of the same idea. Here are the six, with their approximate weights (in MeV, a physicist's mass unit) and, in the right-hand column, how heavy the up-type partner is *relative to* its down-type partner in the same generation:

Generation	up-type	down-type	up ÷ down	which is heavier
1 (lightest)	up ≈ 2	down ≈ 5	**0.46**	down-type heavier
2	charm ≈ 1,300	strange ≈ 95	**11.8**	up-type heavier
3 (heaviest)	top ≈ 173,000	bottom ≈ 4,200	**60**	up-type much heavier

You'd expect the two types to keep the same weight relationship in every generation. They don't, and the right-hand ratio is the whole story. Look at the **first row**: in the lightest generation the down-type is the heavier one (ratio below 1), whereas in both heavier generations the up-type wins (ratio above 1), by a margin that keeps growing. So somewhere between generation 1 and generation 2, the heavier partner *switches sides*. Why? (One technical aside for the careful reader: physicists must compare the two partners using a specific, consistent definition of "weight" called the MS-bar scheme — and this is forced, not a convenience, because the light quarks have no weight at all in the main alternative scheme, so MS-bar is the only one in which all six can be compared on the same footing. In that scheme the top-to-bottom figure is near 60 rather than the ~40 a quick division of the rough weights above would give. The bold ratios are the cleanly-defined ones, and the precise results later in the paper hold specifically because MS-bar is the only consistent choice.)

It helps to put these three ratios on a scale where "equal weight" sits at zero: below zero means down-type heavier, above zero means up-type heavier. On that scale the three numbers are about **−0.77, +2.5, +4.1** — one quantity, which the paper calls **χ** ("chi"), that starts negative, crosses zero between the first and second generations, then keeps climbing but in *shrinking* steps (the jump from generation 1 to 2 is about twice the jump from 2 to 3). If you could explain where χ comes from, you'd explain the whole pattern at once — the sign flip, the growth, and the slowdown. So the real question is: what produces χ?

One tempting answer is that some underlying structure simply *grows* as you move to heavier generations, and the growing structure drives the weight gap. Earlier work in this programme proved that can't be it — the relevant structure stays the same size no matter how far you go. (More than that: *nothing* of its mathematical kind can grow in this setting, so a "bigger structure" explanation is off the table entirely.)

The alternative is subtler, and is what this paper pursues. Picture a fixed object that doesn't change. What changes from generation to generation is how strongly each type — up versus down — can *reach into* that object. In the lightest generation the down-type reaches in more effectively, so it wins. Deeper in, the up-type reaches in harder and overtakes. The switch isn't something growing; it's a change in *direction* — access tilting from one side to the other.

That still leaves the deepest question: where does the up-versus-down distinction itself come from? The proposal is striking in its simplicity: it's the difference between running a certain process **forwards** and running it **backwards**. Forward gives you one type; reverse gives you the other.

For that idea to mean anything, you first have to show such a forwards-versus-backwards split can actually exist as a clean mathematical object — that "reverse" is a genuine mirror image, not just "the same thing in a different order." **This paper proves it can.** Using three simple, fully written-out operations, running them forward versus backward really does land you in two distinct places, and the split falls right out of the operations themselves with nothing added by hand. That's a solid result: the machinery the idea needs is real and buildable. What the paper is careful *not* to claim is that the actual physical process is this exact example — giving the toy pieces physical names is a hopeful guess, not a proof, and the paper labels it as such.

Then it tries for real numbers, and here the honesty matters most. The growth pattern splits into two parts: the *shape* of the steps (how each generation's gap relates to the next) and the *size* of each step. The shape works beautifully — a specific simple rule predicts how the third generation relates to the second with no adjustable knob, matching the measured world to about two parts in a thousand (0.2%); once a single overall size is fixed from generation 2, generation 3 then lands within about 1%. The size is the unfinished part. A neat-looking formula for it (written $K^2/2 = 24.5$) gets within about 4%, but on close inspection it quietly rests on a choice — a "weight" given to a certain self-referencing channel — picked to make the answer come out right. The honest count of that channel gives a different number, and the measurements actually point to a weight of about **0.64**, which no clean rule yet produces. The number the formula "reproduces" for generation 1, meanwhile, turns out to be a value already fixed by earlier work, rewritten — not a fresh, independent success.

So the bottom line is unusually precise about its own limits. The *kind* of mechanism is now proven to be buildable; the *shape* of the pattern is confirmed to a fraction of a percent; and the entire remaining gap has been squeezed down to a single number near 0.64, which the underlying theory must now either produce on its own (in which case the pattern would be fully explained) or fail to produce (in which case the theory makes a clear, testable prediction that's off by a few percent, and we'd know it). Either way, the unknown has gone from "the whole quark mass pattern" to "one number we can go compute." That narrowing is the point.

Table of Contents

Epistemic conventions

Claims are graded inline: **[Proven]** (established here by a non-trivial explicit construction or computation), **[Definitional]** (true by definition, scope, or elementary identity — carries no evidential weight), **[Conditional]** (holds given a stated upstream result or gate), **[Conjectural]** (a proposal not yet derived), **[Open]** (a named proof obligation). The [Definitional] tier is kept distinct from [Proven] deliberately: statements like "a varying weight ratio is a licensed reorientation" or "$(ABC)^\dagger = C^\dagger B^\dagger A^\dagger$" are true but tautological, and folding them under [Proven] would let that grade stop carrying information exactly where it matters most. The grade discipline is strict: a claim inherits the weakest grade among its premises, and no numerical agreement upgrades a claim's grade without an independent derivation of the structure that produced it.

Abstract

The companion programme established that the same-generation quark mass ratios are carried by a single observable, the fold-response susceptibility $\chi(g) = \ln(m\_up(g)/m\_down(g))$, and that χ cannot be sourced by growth of the refinement-persistent carrier H^1, whose dimension is refinement-invariant. The carrier-based route that survives is **fold-selective accessibility**: fixed fold projectors P_up, P_down on H^1, accessed depth-dependently by an operator R_g, with the up/down folds themselves to be produced by **closure-order orientation** (forward closure U = A B C versus reverse closure D = C B A). That target was stated under two gates — OP0 (are the closure operators of a suitable operator class?) and OP0.5 (is closure orientation the up/down bit at all?).

This paper does three things, each at its true grade.

First **[Proven]**, it discharges the *existence* half of OP0: it exhibits an explicit self-adjoint triple (A, B, C) whose forward closure U is non-normal, whose forward and reverse closures occupy distinct subspaces, and whose orthogonal fold projectors $P_+ = UU^\dagger$, $P_- = U^\dagger U$ are produced by the closure operator alone, with no added structure. This proves the operator class the companion paper required is non-empty and realizable.

Second **[Conjectural]**, it proposes that the VERSF admissibility–transport–completion operators realize this class, and exhibits the resulting accessibility observable. This is a *labeling* of the toy structure with VERSF roles, not a demonstration that the VERSF operators are of the proven type; that demonstration remains the substance of OP0.

Third, it subjects a concrete quantitative ansatz — a one-parameter refinement law for $\chi(g)$ — to the programme's own **parameter ledger**, and separates a profile from a magnitude that fare oppositely. The first-generation value it "reproduces" is the upstream ratio $(K-1)/(2K-1)$ relabeled in fold-census language, inheriting that result's conditionality rather than constituting an independent second derivation. The growth **profile** is parameter-free *in the increment ratio* (the gen-1→2 log-step is twice the gen-2→3 step, 0.6%; cumulative 3/2, 0.2%) while carrying *one fitted scale* in the magnitude; the single genuine prediction is gen-3 from gen-2, confirmed to 1%. This increment ratio is the document's most robust quantitative asset — it contains no K, no scale, and no anchor, so it survives even if $K = 7$ is calibrated. Critically, it lives only in the MS-bar scheme: light-quark pole masses do not exist, so MS-bar is the *only* scheme in which all three $\chi(g)$ are defined, which forces $\chi(3) \approx \ln 60$ by consistency (the pole-scheme ratio ≈ 41 would break the result at 8%). "Harmonic" is one of two equally-good profile families (the other, geometric-½, is better-motivated, deriving the deceleration from the fold binary); the shared content is the single increment ratio, not the form. The **magnitude** factor proposed as $K^2/2 = 24.5$ is the weak link: the honest reversal-orbit count of second-depth closure access is $K(K+1)/2 = 28$, and recovering $K^2/2$ requires a self-return diagonal weight of ½, whereas the masses prefer ≈ 0.64. So $K^2/2$ is suggestive at ~4% but underived.

The net result is a proven existence theorem for the fold mechanism's operator type, a clearly-graded candidate identification, an empirically confirmed (but not yet derived) growth profile, and the isolation of the entire remaining quantitative gap into one computable number — the closure self-return weight, which the data have already pinned near 0.64.

Part I — An Existence Proof for the Fold Mechanism

1. Introduction and inherited target

The same-generation ratios and their logarithms are

$m_u/m_d \approx 0.462$, $m_c/m_s \approx 11.76$, $m_t/m_b \approx 60$,

$\chi(1) = \ln(0.462) = -0.773$, $\chi(2) = \ln(11.76) = +2.465$, $\chi(3) = \ln(60) = +4.094$.

The increments are $\chi(2) - \chi(1) = +3.238$ and $\chi(3) - \chi(2) = +1.629$, a ratio of 1.988. Three qualitative features organize the target: a sign change between the first and second generations ($\chi(1) < 0 < \chi(2)$), monotone growth, and deceleration of the increments.

The scheme is forced, not chosen — and this is load-bearing for the headline result. Same-scale mass *ratios* are RG-invariant (γ_m is flavour-blind), so scale is a non-issue; but ratios are *not* scheme-invariant at the precision the profile result (§8) relies on. The values above are the **MS-bar** ratios, and MS-bar is not a preference but the only admissible choice: light-quark pole masses do not exist (u, d, and barely s never reach a perturbative pole below Λ_{QCD}), so MS-bar is the only scheme in which all three $\chi(g)$ — and in particular $\chi(1)$ and $\chi(2)$ — are uniformly defined at all. $\chi(g)$ is therefore *defined* as the MS-bar log-ratio, and $\chi(3) \approx \ln 60$ follows by consistency with that definition, not by selection to fit. The contrast makes the stakes explicit: in a pole-mass scheme (where the light quarks are not even defined) the top/bottom ratio is ≈ 41, giving $\chi(3) = \ln 41 = 3.689$ and a cumulative increment ratio of $(3.689 + 0.773)/3.238 = 1.378$ — an 8% miss of the 3/2 profile result of §8. So the flagship 0.2% agreement lives entirely in MS-bar; that is not a vulnerability to be hidden but a forced consistency, since MS-bar is the only scheme in which the three-point sequence exists. (One honest wrinkle at the top: m_t is itself only quasi-defined — extracted from a Monte-Carlo mass — so "uniformly defined" is cleanest for u, d, s; the residual top-scheme ambiguity shifts ln 60 at the percent level and does not reach the 8% pole/MS-bar gap.)

The companion programme reduced the explanatory problem to one object and one mechanism. The reduction (taken here as upstream input):

- **[Conditional, upstream]** $\chi(g)$ is the unique obstruction to separability of the quark mass operator; explaining $\chi(g)$ is equivalent to explaining the within-generation hierarchy.

- **[Conditional, upstream]** Carrier *growth* cannot source χ: for refinement realized as subdivision, no topological invariant grows, and dim H^1 is in particular invariant. The depth-dependence must enter through fold-selective *accessibility* of a fixed carrier, not its size.
- **[Open, upstream]** The fold projectors P_up, P_down are the one object the accessibility target leaves underived. Closure-order orientation was proposed as their source, under two gates: **OP0** (operator realization — are the closure operators a suitable class with a meaningful adjoint?) and **OP0.5** (assignment identification — is closure orientation the up/down bit specifically?).

The present paper attacks the *existence* half of OP0 directly: not "are the VERSF operators of the right type?" (still open) but "does the right type exist at all, with the demanded properties?" A negative answer would have killed the mechanism outright. The answer is positive, and the proof is elementary and explicit.

---

2. A realizable closure triple

Work on a two-dimensional Hilbert space with orthonormal basis $\{e_1, e_2\}$. Define three operators:

A = (1 0 ; 0 0), B = (0 1 ; 1 0), C = (0 0 ; 0 1).

All three are self-adjoint by inspection: $A = A^\dagger$, $B = B^\dagger$, $C = C^\dagger$. They are, respectively, a projector onto e_1, the basis swap, and a projector onto e_2 — a minimal stand-in for a selection / transport / selection sequence. **[Proven]** the triple is self-adjoint, which is the favourable input class H0 named in the companion paper.

---

3. Forward closure, reverse closure, non-normality, and the folds

Forward and reverse closure. Direct multiplication gives

U = A B C = (0 1 ; 0 0), D = C B A = (0 0 ; 1 0).

Because A, B, C are self-adjoint and reversal of an operator product conjugates it, $D = (A\,B\,C)^\dagger = U^\dagger$. **[Proven]** in this model reverse closure *is* the adjoint of forward closure — the Tier-1 specialization the companion paper could only assume is here a computed fact, because H0 holds by construction.

Distinct fold subspaces. The ranges are

$\mathrm{ran}(U) = \mathrm{span}(e_1)$, $\mathrm{ran}(D) = \mathrm{span}(e_2)$, so $\mathrm{ran}(U) \neq \mathrm{ran}(D)$.

[Proven] forward and reverse closure occupy genuinely distinct subspaces — the operative non-triviality condition (★) of the companion paper, here satisfied outright, and strictly stronger than the near-vacuous $A B C \neq C B A$.

Non-normality. Computing the two products,

$U U^\dagger = (1\ 0 ; 0\ 0)$, $U^\dagger U = (0\ 0 ; 0\ 1)$, so $U U^\dagger \neq U^\dagger U$.

[Proven] U is non-normal, and this is exactly what permits distinct folds: for a normal operator $\mathrm{ran}(U) = \mathrm{ran}(U^\dagger)$, which would collapse the two folds. The companion paper's claim "the mechanism lives in the failure of normality" is here instantiated with a witness.

Fold projectors from the closure operator alone. Define

$P_+ = U U^\dagger = (1\ 0 ; 0\ 0)$, $P_- = U^\dagger U = (0\ 0 ; 0\ 1)$.

Each is Hermitian, idempotent, and positive semi-definite — orthogonal projectors onto the forward and reverse fold subspaces. **[Proven]** the fold split $P_+ \neq P_-$ is produced by U with no structure added beyond the closure operator. (For this U the singular-value decomposition is trivial — U has a single non-zero singular value 1 — so the orthogonal range projectors P_+, P_- and the closure-canonical idempotents coincide; the orthogonal/oblique distinction the companion paper flagged does not bite in this minimal model. It will require attention in any model where U is non-normal *and* has non-orthogonal singular structure.)

The fold carrier is

$\Delta P = P_+ - P_- = (1\ 0 ; 0\ -1)$,

Hermitian with spectrum $\{+1, -1\}$: positive on the forward-only subspace, negative on the reverse-only subspace. **[Proven]** ΔP is the real, traceable Hermitian avatar of the anti-Hermitian residue $\Omega = U - D = (0\ 1 ; -1\ 0)$, $\Omega^\dagger = -\Omega$.

---

4. What the existence proof establishes — and what it does not

[Proven] (existence). The operator class the companion paper required is non-empty: there exists a self-adjoint triple whose forward closure is non-normal, whose forward and reverse closures reach distinct subspaces, and whose orthogonal fold projectors are generated by the closure operator alone. Every Tier-1 statement the companion paper conditioned on H0 — $D = U^\dagger$, anti-Hermitian residue, Hermitian fold projectors, real sign-definite traces — holds in this model by direct computation. The *type* is realizable; the mechanism is not vacuous.

[Open] (identification — this is OP0 proper). Nothing above shows that the VERSF closure operators are this triple, or any triple of this type. The model is an existence witness, not an identification. OP0 in its full form still demands: define the VERSF operations A, B, C; prove they act on H^1; prove they live in a unital ∗-algebra with a meaningful adjoint; determine whether they are self-adjoint. The witness lowers the risk that OP0 is unsatisfiable; it does not solve it.

This distinction is the hinge of the paper and is maintained throughout: **the operator type is proven to exist; its VERSF realization is conjectural.**

Part II — Candidate VERSF Identification

5. The admissibility–transport–completion reading

[Conjectural]. The minimal model invites the VERSF identification

A = Π_adm (admissibility selection), B = T (transport exchange), C = Π_comp (completion selection),

so that forward closure U = Π_adm T Π_comp reads "select admissible, transport, select complete," and reverse closure U† = Π_comp T Π_adm reverses the order. The proposal has the right shape — two selection idempotents bracketing a transport exchange is exactly the structure that produced non-normality in the model, since selection projectors onto different subspaces do not commute with an exchange between them.

But the grade is **[Conjectural]**, and three obligations separate it from a derivation. (i) That Π_adm, Π_comp, T act on H^1 and not some larger space — the companion paper's "same H^1" bridge. (ii) That they are self-adjoint (or that whatever weaker involution VERSF supplies still yields D = U† and Hermitian folds; absent that, the model's clean Tier-1 results do not transfer). (iii) That their ranges are *distinct* in the precise sense ran(U) ≠ ran(U†) — the physical analogue of (★) — rather than merely non-commuting. None of these is established by writing the labels; each is part of OP0.

A second, independent gate sits above all of Part II: **OP0.5**, whether closure orientation is the up/down bit specifically rather than one of the other binary distinctions the framework already carries (quark/lepton bath participation; weak isospin, with colour three-valued). The natural discharge would identify closure orientation with the weak-isospin label $T_3 = \pm\frac{1}{2}$ used elsewhere in the programme. **[Open]** that identification is not made here.

---

6. Refinement access and the fold-imbalance observable

[Conditional on H0]. Introduce a refinement-access operator diagonal in the fold basis,

$R\_g = N_+(g) P_+ + N_-(g) P_-$,

with $N_+(g), N_-(g) \geq 0$ the depth-dependent access weights, their depth-dependence carried by the transient (non-persistent) near-marginal spectrum of the refinement dynamics, not by the dimension of H^1 (which is invariant). The fold accessibilities are

$\mathcal{A}_+(g) = \mathrm{Tr}[R\_g\ P_+] = N_+(g)$, $\mathcal{A}_-(g) = \mathrm{Tr}[R\_g\ P_-] = N_-(g)$,

and the response observable, in the un-normalized log-ratio form the mass anchor requires, is

$\chi(g) = \ln \mathcal{A}_+(g) - \ln \mathcal{A}_-(g) = \ln(N_+(g) / N_-(g))$.

The fold imbalance is $\delta(g) = \mathrm{Tr}[R\_g\ \Delta P] = N_+(g) - N_-(g)$, and **[Proven, given δ real]** the sign law sign $\chi(g)$ = sign $\delta(g)$ holds exactly: which fold is heavier at depth g is the sign of one real number.

A structural note on reorientation (correcting a possible misreading). Because R_g is diagonal in the fold basis, the sign flip is produced by the ratio $N_+(g)/N_-(g)$ crossing 1 — i.e. by shifting access weight from the reverse-fold eigenspace of ΔP toward the forward-fold eigenspace. This *is* the companion paper's "reorientation across the eigenspaces of ΔP"; it is not in tension with the no-go that "uniform scaling cannot flip the sign." The no-go forbids $R\_g = c\_g R_0$ with *fixed shape* (constant N_+/N_-); the present R_g has a varying ratio and so changes shape, which is precisely the licensed mechanism. **[Definitional]** a diagonal access operator with varying weight ratio is a legitimate realization of reorientation — true by the scope of the no-go, not a result. What remains entirely open is whether the *specific* weights $N_+(g)$, $N_-(g)$ can be derived without consulting the masses — which is the whole content of Part III.

Part III — The First Quantitative Law, Held to the Ledger

7. The first-generation value is the upstream ratio, relabeled

[Conditional, inheriting upstream]. The companion confinement programme fixes the first-generation species ratio as $m\_u/m\_d = (K - 1)/(2K - 1)$, and for K = 7 this is 6/13, so $\chi(1) = \ln(6/13) = -0.773$. The fold-census reading proposed here is

$N_+(1) = K - 1 = 6$ (forward fold accesses the six boundary channels of the K = 7 closure), $N_-(1) = 2K - 1 = 13$ (reverse fold accesses the full loop census $N\_loop = 2K = 14$, less one stabilizing mode), giving $N_+(1)/N_-(1) = 6/13$ and $\chi(1) = \ln(6/13)$.

It must be stated plainly what this is. Since $6 = K - 1$ and $13 = 2K - 1$, the fold-census ratio is *algebraically identical* to the upstream $(K - 1)/(2K - 1)$ — it is the same formula written in fold-access language, **not an independent second derivation of 6/13**. Two consequences follow.

- **[Conditional]** $\chi(1)$ here inherits the upstream result's exact conditionality. It is a genuine external anchor *iff* $K = 7$ and the form $(K - 1)/(2K - 1)$ were fixed upstream without consulting the measured m_u/m_d — the same status question (parameter-free check versus calibration) that governs $\kappa = 8/3$ in the lepton arc, still unresolved. The fold relabeling does not improve this status; it transports it.
- **[Open]** The one piece of new content is the *interpretation* — that $K - 1$ is forward-fold access and $2K - 1$ is reverse-fold access at $g = 1$. The unforced step is the "−1 stabilizing mode" subtraction that turns N_loop = 14 into 13: no closure-geometric reason is given for why the reverse fold sees the loop census minus exactly one mode while the forward fold sees boundary channels. Until that subtraction is derived, the decomposition is a chosen reading consistent with 6/13, not a consequence of it.

The honest summary: §7 reproduces $\chi(1)$ because it *is* $\chi(1)$ rewritten, and claims of "no fitted parameter" must be read in that light — there is no new parameter because there is no new derivation, only a relabeling that inherits whatever grade the upstream 6/13 carries.

---

8. The refinement profile: one prediction, family-free, and K-independent

The information ceiling, stated up front. The data are three points. The anchor $\chi(1)$ consumes one (§7). A single magnitude G is fixed on $\chi(2)$. So **exactly one number, $\chi(3)$, is a genuine prediction** — and no analysis of three within-generation ratios can do better than one parameter, one prediction. Every framing below is bounded by that ceiling; it is stated here once, plainly, rather than conceded in pieces.

The family-free content is a single statement. Write the law as $\chi(g) = \chi(1) + W(g)\cdot\ln G$, with W a depth-profile ($W(1) = 0$) and G a magnitude. The one prediction can be isolated free of *both* G and the choice of profile family: it is the **ratio of log-increments**,

$[\chi(3) - \chi(1)] / [\chi(2) - \chi(1)]$,

which depends only on $W(3)/W(2)$ and cancels G and the anchor entirely. The data give $(4.867)/(3.238) = 1.503$, or equivalently the consecutive form $(\chi(2)-\chi(1))/(\chi(3)-\chi(2)) = 3.238/1.629 = 1.988$. So the entire empirical nugget is one number: **the gen-1→2 increment is twice the gen-2→3 increment** (predicted 2, observed 1.988, 0.6%; cumulative 3/2 vs 1.503, 0.2%). That is the result. Everything else is which family one wraps around it.

"Harmonic" is one label among many; do not let it carry weight. Infinitely many one-parameter decelerating families reproduce a 2:1 first-to-second increment and agree on all three points; they differ only at $g \geq 4$, which is unmeasurable. So the 0.2% tests the *value* 3/2 at one point, not the *form* "harmonic." Two natural families both pin the second increment to half the first:

- **Harmonic** ($W(g) = H_{g-1}$): increments 1, ½, ⅓, … — gives next ($g = 4$) increment $\frac{1}{3} \cdot \ln G$. Its weakness is motivation: "why 1/k weighting" requires importing a separate hypothesis (a bath-share story in which k participants sharing a bath give 1/k — worth pursuing, but not native here).
- **Geometric-½** ($W(g) = 1 + \frac{1}{2} + \ldots + (\frac{1}{2})^{g-2}$, i.e. decay ratio $r = \frac{1}{2}$): increments 1, ½, ¼, … — gives next increment $\frac{1}{4} \cdot \ln G$. Its strength is that the halving is *native*: it is the fold count. If each refinement depth bifurcates the accessible fold-channels along the forward/reverse binary the whole paper is built on, the new access added at depth k is half the previous, and $r = \frac{1}{2}$ falls out of the two-fold structure with **nothing imported**.

The two are observationally identical on three generations and diverge only at $g = 4$ (⅓ vs ¼ of ln G), beyond measurement. So harmonic should be demoted from its title role to one candidate of two, and geometric-½ is, if anything, the better-motivated reading — it derives the deceleration ratio from the same orientation binary that produces the folds, rather than from an unmotivated 1/k. Grade: **[Conjectural]**, with the observational degeneracy stated honestly; neither family is established, and the shared, family-free content is the single 2:1 increment ratio.

The increment ratio is the paper's most foundationally robust asset — more robust than the anchor. The number 3/2 (equivalently 2) contains no K, no G, and no anchor: it is pure profile. It therefore survives *even if K = 7 turns out to be calibrated to light-quark data* (the open gate on the upstream 6/13, §7). That is strictly more robust than the §7 anchor itself, which is K-dependent and conditional, and more robust than anything in §9–§10A, all of which sit downstream of K. So the one solid quantitative result in the document is also the one least exposed to the programme's deepest open gate — worth stating as such.

What is genuinely established, then: a single prediction ($\chi(3)$ from $\chi(2)$ with one fitted scale, confirmed to 1%), whose family-free content is one number (the 2:1 increment ratio, 0.6%), K-free and anchor-free. The magnitude G is the separate, weaker ingredient, and §9 dissects it.

---

9. The magnitude factor: combinatorial proposal and the diagonal-weight problem

A combinatorial origin is proposed for $G = K^2/2$. The reading: second-depth refinement compares closure sites against closure sites, sampling ordered pairs from the $K = 7$ closure positions, giving K^2 raw access; quotienting by forward/reverse fold orientation (which identifies a pair with its reverse) then halves this to $K^2/2 = 49/2 = 24.5$. The intent is to derive G from the same $K = 7$ census that fixes 6/13, with no new parameter.

[Open] the quotient is not forced — the halving hides a chosen diagonal-weight rule. Quotient an ordered-pair set $S \times S$ ($|S| = K$) by the reversal involution $\sigma(i,j) = (j,i)$ honestly, via orbit counting. Off-diagonal pairs fall into 2-element orbits $\{(i,j),(j,i)\}$, contributing $K(K-1)/2$. Diagonal pairs (i,i) are *fixed* by σ — each is its own orbit — so a genuine orbit count gives them weight 1, contributing K. The true reversal-orbit count is therefore

$$|(S \times S)/\sigma| = K(K-1)/2 + K = K(K+1)/2,$$

which for $K = 7$ is **28**, not 24.5. To recover exactly $K^2/2$ one must instead assign the diagonal pairs weight **½**:

$$K(K-1)/2 + (\tfrac{1}{2})\cdot K = (K^2 - K + K)/2 = K^2/2.$$

So the factor $K^2/2$ is *equivalent to* the half-weight diagonal postulate; the "÷2 for two folds" has been relocated into a weight on the self-return channels, not removed. The decomposition $49/2 = 21 + 7/2$ (off-diagonal bridges plus half-weighted self-returns) is correct arithmetic, but its load-bearing hypothesis — diagonal weight ½ — is precisely the undischarged step. The three principled diagonal conventions give three different magnitudes, and the data adjudicate sharply.

[Proven] the data reject the combinatorial value and prefer no clean convention. With the profile granted, the gen-2 ratio ($H_1 = 1$) is m_c/m_s = $(6/13)\cdot G$, so the measured 11.76 requires $G = 11.76 \times 13/6 =$ **25.48**. Comparing:

Diagonal weight	G	gen-2 = (6/13)G	vs 11.76	gen-3 = (6/13)G^1.5	vs 60
0 (drop self-returns) → K(K−1)/2	21	9.69	−18%	44.4	−26%
½ (the K²/2 rule) → K²/2	24.5	11.31	−3.8%	55.98	−6.7%
1 (true orbit count) → K(K+1)/2	28	12.92	+9.9%	68.4	+14%
empirical best (gen-2 exact)	25.48	11.76	0	59.4	−1.0%

The proposed $K^2/2 = 24.5$ is the nearest of the three integer conventions, but it still undershoots the required 25.48 by ~4%, and the value the data actually want corresponds to a diagonal weight of $(25.48 - 21)/7 =$ **0.64**, not ½, not 0, not 1. No clean reversal-quotient rule produces 0.64. So $K^2/2$ is **[Conjectural], suggestive at the ~4% level, and not even the best member of its own family** — the half-weight postulate is both undischarged and slightly disfavoured by the masses.

---

10. The corrected ledger, and what it teaches

The combinatorial analysis changes the ledger verdict in both directions, and the honest reading separates the profile from the magnitude.

The profile is the asset; the magnitude is the liability.

Ingredient	Free parameters	Against	Verdict
Profile (harmonic or geometric-½)	0 (the increment ratio is G-independent)	the increment-ratio datum	**passes** as a parameter-free check (1.503 vs 1.500, 0.2%) — *conditional on deriving the profile form (§8)*
Magnitude via $K^2/2$ (derived)	0 claimed, but the ½ weight is a chosen convention	gen-2, gen-3	**fails** — misses gen-2 by 3.8% and gen-3 by 6.7%, beyond measurement error; the data want weight 0.64, not ½
Magnitude G (fitted)	1	—	**passes** — G = 25.48 is *fixed on* gen-2 (not a prediction of it); gen-3 is then *predicted* to 1%

Two conclusions follow, and they reverse the earlier all-negative reading.

- **[Proven] the gen-3 postdiction is real and was hidden by $K^2/2$.** Fixing the single magnitude G from gen-2 alone forces gen-3 to 59.4 against 60 — a 1.0% agreement, inside the few-percent error on m_t/m_b. Equivalently, gen-2 and gen-3 independently imply almost the same G (25.48 and 25.65), which is *why* the profile works. The 3.8%/6.7% misses reported for $K^2/2$ are entirely the cost of forcing G = 24.5 when the data want 25.5; a fitted G removes them. The honest accounting carries no implied leverage: **one parameter is fixed on gen-2, and one prediction (gen-3) is confirmed to 1%** — not "1 parameter against 2 targets," since gen-2 *defines* G rather than testing it. This is the one-prediction ceiling of §8, reached, and conditional on the profile.
- **[Open] the magnitude is the load-bearing unknown, and the data have already constrained it.** The combinatorial route was meant to drive N_free to 0 by deriving G, but it lands on 24.5 (weight ½) while the masses want 25.48 (weight 0.64). So the live obligation is sharp and quantitative: derive the diagonal self-return weight from the completion-density / self-return machinery and check whether it yields ≈ 0.64 (vindicating a derived G and reaching N_free = 0) or yields ½ (in which case the mechanism *predicts* ratios ~4–7% low, a clean falsifiable claim to own rather than hide). The data have ruled out diagonal weights 0 and 1 decisively; the open question is a single number near 0.64.

On the boxed lemma. A "Two-Step Fold-Quotient Lemma" stating $K(K-1)/2 + K/2 = K^2/2$ is a true identity but not a derivation of the factor: its hypothesis (diagonal weight ½) *is* the ÷2 it purports to explain, so it is an identity wearing a derivation's clothes — the same failure mode this arc has flagged at every level. The defensible lemma is the honest one: **the reversal quotient of second-depth closure access is $K(K+1)/2 = 28$; recovering $K^2/2$ requires assigning self-return channels weight ½, and the masses prefer weight ≈ 0.64; the diagonal weight is therefore an undischarged parameter that closure geometry must supply.** Stated that way it is [Open], correctly graded, and points at exactly the number to compute.

On a proposed weight w_self = (K+2)/2K. A natural next move is to "derive" the data-preferred weight by writing it as a K-expression: with self-return support K + 2 (K closure-site support plus two fold-orientation endpoints P_+, P_-) over loop census 2K, w_self = (K + 2)/2K, giving for K = 7 w_self = 9/14 = 0.6429, G = K(K−1)/2 + (K+2)/2 = 25.5.

The magnitude is then well-targeted: G = 25.5 sits essentially on the gen-2-required 25.48, predicting m_c/m_s = 11.77 (vs 11.76) and m_t/m_b = (6/13)·25.5^{3/2} = 59.4 (vs 60, a 1.0% miss — the residual is the harmonic profile's own ~1%, since gen-3 alone prefers G = 25.65; no choice of G removes it). This is recorded, but at grade **[Conjectural], reverse-targeted, not derived**, for three independent reasons, any one sufficient.

- **It writes the fitted number in K-notation.** The weight 0.64 was extracted from the masses in §9 (solving 21 + 7w = 25.48); the form (K+2)/2K is then offered because it equals 0.6429 at K = 7. By the grade discipline of this programme, numerical agreement cannot purchase a structural claim without an independent derivation of the structure that produced it, *with the target value quarantined*. Here 9/14 was derived *toward* 0.64, not while blind to it — the same move as the ½-rule one level up, with a better-aimed free choice.
- **The "+2" double-spends the orientation quotient.** The off-diagonal count is K(K−1)/2 rather than K(K−1) precisely because the two fold orientations (forward/reverse, P_+/P_-) were already identified in the quotient. The diagonal (i,i) is the σ-fixed set under that same involution. Crediting it +2 *for those two orientations* spends the orientation degrees of freedom a second time. Under a consistent orbit count a σ-fixed point carries weight 1, returning G = K(K+1)/2 = 28 — the honest value already stated. There is no single quotient in which the orientation pair both halves the off-diagonals and augments the diagonal.
- **It stacks unverified census readings rather than discharging one.** The denominator 2K = N_loop and the g = 1 "−1 stabilizing mode" are already undischarged (§7); "self-return support = K + 2" adds a third chosen reading. The numerator also mixes units — off-diagonal channels are counted one per unoriented bridge, while the diagonal is assigned a site-census number K plus 2, then normalized differently. The combination lands at 0.64 at K = 7, but at, say, K = 2 it gives w_self = 1, a self-return weighted as heavily as a full bridge — geometrically backwards for a channel with no source/target separation, which signals the form is fitted at K = 7 rather than structurally motivated across K.

The honest one-line record: **the data fix w_self ≈ 0.64; the form (K+2)/2K reproduces that at K = 7 and gives the well-targeted G = 25.5, but it is reverse-targeted, re-spends the orientation quotient (conflicting with the K(K+1)/2 = 28 count this section endorses), and adds two further census posits — so the diagonal weight remains the open quantity, not a derived one.** The genuine discharge is the one §10's first [Open] item names and this shortcut skips: derive the self-return weight from the completion-density / self-return dynamics directly, with 0.64 quarantined, and *then* check whether it lands near 9/14. Producing (K+2)/2K by noticing it equals the fitted value is exactly what the programme declines to count.

The right type of object, and why no closed form can settle it. The diagonal channel (i,i) is the only one whose source and target coincide, which makes it the trajectory most exposed to completion saturation — so the natural *type* for w_self is not a combinatorial count at all but a **completion-survival probability**, w_self = 1 − σ, with σ a saturation loss. This explains, structurally, why the data sit strictly between the orbit-count weight 1 and the half-weight ½: a survival probability for the most-saturated channel lands naturally in (½, 1), and the measured ≈ 0.64 is a partial suppression, neither full nor half. **[Open]** this is recorded as the correct *direction*, not a result.

But the direction comes with a sharp negative result that retires an entire line of attack. The magnitude target is low-precision: fixing G from gen-2 gives 25.48, from gen-3 gives 25.65, so w_self is pinned only to roughly 0.64–0.66, a band of order a few percent (the harmonic profile's own ~1% floor, §8). The neighbourhood of 0.64 is *dense with simple closed forms* — 9/14 = 0.643, 16/25 = 0.640, 2/π = 0.637, e^{−0.446} = 0.640 all fall inside the band — so **hitting 0.64 with a closed form carries essentially no evidential weight**, and any "derivation" whose only check is that it equals 0.64 is indistinguishable from the others by the data. A construction such as $\sigma = (3/5)^2 = 9/25$ (giving w_self = 16/25 = 0.640 and G = 25.48 exactly) is the same reverse-targeting once more: the factors 3 and 5 are named after the value is known, and "5 = closure degrees minus the two fold endpoints" re-spends, for a third time, the orientation pair already spent in the off-diagonal quotient, while "3 = generation depth" smuggles the generation count into a depth-*independent* magnitude where the harmonic profile H_{g−1}, not w_self, is supposed to carry it.

The correctly-posed test. Because a single number matching a single number is never a test, and matching a closed form to a low-precision target is dead on arrival, a derivation of w_self can count only if it meets one of two bars, stated in advance:

- **Over-determination.** The completion/saturation mechanism predicts w_self *and* at least one further independently measurable quantity, with the mechanism's parameters fixed on the other observable; w_self then follows with no remaining freedom and lands in the data band. One mechanism predicting two-or-more numbers is a test; one number tuned to one number is not.
- **Precision beyond the target's width.** Failing over-determination, a blind derivation's win condition is landing *inside* the 0.64–0.66 band, not on a specific decimal — and claiming an exact value like 16/25 over-states a precision the data cannot reward.

So the open problem is now correctly isolated and, importantly, *closed-form coincidence is ruled out as a route to it*: w_self must come from completion-survival dynamics, derived with the masses quarantined and validated by over-determination — not by band-membership (which §10A retires as nearly automatic) and not by any further search for the elegant fraction that happens to equal 0.64.

- **[Proven]** The diagonal access ansatz $R\_g = N_+(g)P_+ + N_-(g)P_-$ realizes reorientation correctly (§6); the mechanism and the profile are sound. The entire remaining burden is one magnitude — equivalently, one diagonal weight near 0.64, of completion-survival type — not a functional form.

The companion paper's first milestone stands as the right next target: derive that the access weights *force exactly one* sign crossing (gen-1 inverted, all higher generations not) from the geometry of P_+, P_-, R_g, before any magnitude is matched. The harmonic law achieves a single crossing trivially (any monotone weight ratio rising through 1 crosses once), so the single-crossing property here is not yet evidence of mechanism — it is a property of monotone weights, not of derived geometry.

10A. The obstruction-triad route, and a quarantined derivation program

The completion-survival framing (§10) gives the *type* of object but not a way to compute it. The most promising concrete route reduces the weight to rank data plus a geometric angle, and — unlike the closed-form attempts — its ingredients can in principle be derived with the masses sealed off. It is recorded here as a program, with each leg at its true grade and each candidate number held in the band rather than boxed.

The reduction. Let $\mathscr{H}_{\text{self}}$ be the diagonal self-return sector of the second-depth closure-pair space, and write the survival weight as

$w_{\text{self}} = (a + n - o)/K$, $a = \mathrm{Tr}(A)$, $n = \mathrm{Tr}(N)$, $o = \mathrm{Tr}(A\,N)$,

with A the admissible self-return projector and N the null/stabilising projector on $\mathscr{H}_{\text{self}}$. For a rank-one overlap, $o = \cos^2\theta$, where θ is the principal angle between the admissible self-return sector and the null/stabilising direction. The magnitude problem becomes a *geometric* one: **determine θ.** This reduction is the genuine advance of the route — it converts an arbitrary parameter into a principal angle that closure geometry might actually supply. Grade: **[Open]**, but a well-posed open.

The triad. A diagonal self-return (i,i) survives only by passing three logically distinct gates — it must not collapse to the null/stabilising mode (E_0), must not exceed completion capacity (E_+, over-completion), and must not fall below admissibility (E_-, under-admissibility). These three failure modes are *conceptually* motivated by the requirement that a self-return produce a distinguishable, admissible, non-null return; they are not introduced numerically. Grade: **[Conjectural, motivated]** — the first ingredient in this whole search that is not a number in disguise. If each gate is a rank-one exclusion, $a = K - 3 = 4$ and (with $n = 1$) the weight reduces to $w_{\text{self}} = (5 - \cos^2\theta)/7$.

But the cardinality is target-shaped, and this must be faced first. The reachable range of $w_{\text{self}} = (a + 1 - \cos^2\theta)/7$ as $\cos^2\theta$ runs over [0, 1] is the interval [a/7, (a+1)/7]. For $a = K - 3 = 4$ (three gates) this is **[0.571, 0.714]**, which *contains* the empirical 0.64. For four gates ($a = 3$) it is [0.429, 0.571] — 0.64 unreachable; for five ($a = 5$), [0.714, 0.857] — also unreachable. So three gates is the *unique* small cardinality whose reachable band straddles the target. This is the strong form of "Worry 3," and it is worse than the generation-coincidence version: the triad's *size* is shaped by reachability of 0.64 independently of any generation count. Until the cardinality is derived from structure that never sees the masses (the $\{\text{hub}, K_c, p_{\text{eff}}\} = 2 + 1$ decomposition, made the first deliverable below), "three gates" is itself a target-shaped choice, and the section must own that — it is exactly the move §10 retires, here applied to the section proposing the quarantine. Apply the quarantine rule to the quarantine.

Three worries keep this honest, and each is converted below into a concrete, mass-quarantined test rather than left as an objection.

The standing result is the reduction, not a number. $w\_self = (5 - \cos^2\theta)/7$ with θ the principal angle between the admissible self-return sector A and the null direction N. Every candidate value below is an *illustrative special case* of this formula under some assumption about θ; the objective is to derive θ, and no decimal is boxed as the answer.

Worry 1 — the symmetric special case ($\cos^2\theta = 1/3$) is the wrong one, and points at the right calculation. Assuming the null direction is the symmetric diagonal of the three obstruction axes gives $\cos^2\theta = 1/3$ and $w\_self = (5 - 1/3)/7 = 2/3 = 0.6667$ (G = 25.67, gen-3 ≈ 60.0, gen-2 ≈ 11.85, +0.7%). But this imports a *symmetry the mechanism elsewhere requires to be broken* — the folds are asymmetric, and that asymmetry is what produces χ at all — so 2/3 is recorded as the symmetric special case and *another member of the dense band* (alongside 9/14, 16/25, $2/\pi$), not a result. The tell: 2/3 lands at the gen-3 end of the band while 0.64 lands at the gen-2 end, so "which fraction" is just "which generation you fit."

The fix (over-determination via an already-committed asymmetry). The system carries an asymmetry that is fixed upstream and not free: the g = 1 anchor commits R_1 to weight the folds as 6 : 13. *If* the principal angle were set by that same asymmetry, $\cos^2\theta$ would be forced — e.g. $6/19 = 0.3158$ (giving $w\_self = 89/133 = 0.6692$, G = 25.68) — with no new parameter, the anchor asymmetry doing two jobs at once (the sign-change *and* the magnitude). That is the right *kind* of move. But "$\cos^2\theta$ equals the fold-access ratio" is itself an undischarged geometric claim: θ is the overlap between A and N on $\mathscr{H}\_self$, while 6:13 is a fold-*access* weighting in R_1, and there is no a priori reason the two coincide. Writing $\cos^2\theta = 6/19$ *because* it lands in the band would be the same reverse-targeting one level up, merely better-disguised by the ratio's upstream pedigree. So the over-determination logic is sound and worth stating — *if* the angle is fixed by a committed asymmetry it is not a fit — but *which* asymmetry, and whether the A–N overlap genuinely equals it, must be computed, not assumed. This collapses into the model calculation below.

Worry 2 — rank-one ($a = K - 3$) is asserted, not derived; and the angle must be computed, not assumed. Both are settled by one concrete, mass-quarantined exercise.

The model calculation (forces the rank and the angle together). Extend the explicit OP0 toy operators (§§2–3) from the 2-site model to the full K = 7 closure graph, and build A (admissible self-return projector), N (null/stabilising projector), and the three obstruction operators E_0, E_+, E_- as actual matrices on $\mathscr{H}\_self$. Two independent readouts follow from the *same* model, neither touching the masses:

- **Ranks.** Diagonalize each obstruction operator and read how many directions each excludes. If each is rank-one, $a = K - 3 = 4$ is derived; if one is rank-two, the formula corrects to $a = K - 4$ and the route is falsified before any fit.

- **The angle, and the over-determination test.** Compute the A–N principal angle directly (the SVD of P_A P_N), giving $\cos^2\theta$ as a *number out of the model*. Independently, read the g = 1 fold-access ratio off R_1. If these two independently-computed quantities coincide — i.e. the A–N overlap equals the anchor's 6 : 13 asymmetry — then $\cos^2\theta$ is over-determined by the anchor (the 6/19-type value is derived, not asserted); if they differ, Test 1 is falsified and the angle is shown *not* to be the anchor ratio, for free. Either outcome is informative, and neither consults m_c/m_s or m_t/m_b. **(Tests 1 and 2, now a single calculation.)**

Cross-sector transfer — the kill-shot. The triad, if geometric, must appear with the *same* rank-one structure in the bath-free lepton sector, predicting the charged-lepton analogue weight from the identical (a + n − o)/K form with the lepton's K, *before* looking at lepton masses. If rank-one is real it transfers; if it was reverse-fit to K = 7 quarks it will not. This is the strongest test available and uses exactly the lepton discriminator the companion programme already named (and deferred). **(Test 3.)**

Worry 3 — the "3" of the triad may be the generation count smuggled in. Three obstructions, three generations, and three keeps appearing where a fit needs it. Break the coincidence by deriving the triad's cardinality from structure that has nothing to do with generations: the completion-density anchor m ∝ p_eff/K_c is a *two-leg* structure (access and burden), and the K = 7 census carries *one* hub/stabilising channel. The proposed identification is

E_- (under-admissibility) ↔ p_eff floor, E_+ (over-completion) ↔ K_c ceiling, E_0 (null) ↔ the hub/stabilising mode,

so that "three" = 2 (the access/burden dichotomy) + 1 (the hub), each fixed for reasons independent of the generation count. If this mapping holds, the triad stops being a posit and becomes a consequence of the completion anchor already in the programme. **(Test 4.)**

The program, reordered so cardinality gates the rest. The standing result is the reduction w_self = $(5 - \cos^2\theta)/7$ *conditional on a = 4*; but since the reachability of 0.64 is what makes a = 4 look right (above), the cardinality must be derived **before** the rank/angle calculation is meaningful — if a ≠ 4, the band shifts and the rest is moot. Three deliverables, gated in order, each independently falsifying:

1. **Cardinality first.** Derive the triad's size from the {null, over-completion, under-admissibility} = {hub, K_c, p_eff} = 2 + 1 decomposition (the access/burden dichotomy of the completion anchor, plus the one hub channel), with no reference to the generation count or the masses. If this fails, a = K − 3 is dead and deliverables 2–3 are moot. This is promoted to first because it gates everything downstream.
2. **One model calculation** (the K = 7 closure graph): given a = 4, read the obstruction ranks (confirming rank-one or falsifying it) *and* compute the A–N principal angle, testing whether it equals the anchor's 6 : 13 fold-access asymmetry (over-determination if so; falsified if not). Masses quarantined.

3. **The cross-sector lepton prediction**, run blind — the kill-shot, and the one deliverable whose success would actually establish the route, since it predicts a number the construction was not built on.

Band-membership is retired as a win condition. Two facts collapse it. First, the reduction spans [0.571, 0.714] over the whole angle range, so landing *somewhere* in the empirical neighbourhood is nearly automatic — one free angle suffices, and both 6/13 → 0.648 and 6/19 → 0.669 already land, so "the anchor asymmetry" does not even pick a unique value. Second, the empirical band is wider than the internal 0.64–0.66 profile residual once the measurement error on m_c/m_s (dominated by m_s, several percent) is propagated: the required G is more like 25.5 ± 1.5 and w_self ≈ 0.64 ± 0.1, filling most of [0.571, 0.714]. So band-membership tests almost nothing. The **only** admissible bar is over-determination — the A–N angle forced to equal the anchor asymmetry (deliverable 2), or the blind lepton transfer (deliverable 3). A number that merely lands in band counts for nothing.

Honest grade of §10A: a better-posed program of equal current standing. At present a = 4, n = 1, and rank-one are *assumed*, the cardinality is target-shaped, and θ is *not computed* — so the section contains **zero numbers actually downstream of mass-free computation**; it contains a plan to produce them. Its claimed superiority over the closed-form attempts is therefore entirely *promissory*: the reduction-to-an-angle is a genuine structural advance, but the route's empirical standing today equals the closed-form attempts', not exceeds it, and becomes superior only if the gated deliverables — especially the blind lepton transfer — succeed. Grade: **[Conjectural], better-posed but not yet better-standing**; the candidate values (2/3 symmetric, ≈ 0.669 anchor-angle) remain members of the dense band, held, not boxed. The quarantine rule the section invokes against others applies to it: until a number falls out of a mass-free computation, §10A has reduced the problem, not solved any part of it.

10A.1 The cyclic limit: a blind, falsifiable near-miss

The model calculation (deliverable 3b) can be run in its simplest, cyclic form immediately, and it returns the route's first genuinely blind output.

[Proven, within the cyclic model] A pure cyclic K = 7 closure diagonalises in the Fourier basis. In that basis the hub/null mode E_0 is orthogonal to the admissible complement of $E_0 \oplus E_+ \oplus E_-$, so the overlap vanishes:

A N = 0, hence w_self = (4 + 1 − 0)/7 = 5/7 = 0.7143.

This is computed with the masses quarantined — it falls out of the cyclic operator structure, not from any fit. And it **misses the band**: G = 21 + 7·(5/7) = 26, giving m_c/m_s = (6/13)·26 = 12.0 (vs 11.76, +2.0%) and m_t/m_b = (6/13)·26^{3/2} = 61.2 (vs 60, +2.0%), outside the ~1% gen-2 error and at the edge of the gen-3 error.

The miss is the result, and it is the strongest magnitude-branch fact in the document — because for the first time the route *committed to a number before looking and could have been right*. Every earlier candidate (½, 9/14, 16/25, 2/3) was aimed at 0.64; the cyclic model, computed blind, lands at 0.714 and is wrong by a definite, signed 2%. A construction that can produce a falsifiable miss has empirical content; one that can only produce 0.64 does not. So 5/7 is recorded as a near-miss, not buried.

The completion-density correction (the fork). The cyclic operator is, however, not the VERSF completion operator. Completion is density-weighted: a trajectory is evaluated by the completion demand it places on realised structure, so the relevant operator is ρ_comp = ρ_cyc + ρ_dens, of which only ρ_cyc is diagonal in the Fourier basis. The self-return load is L_self = P_self ρ_comp P_self, and the question is whether the density correction couples the admissible sector to the hub:

A N = 0 (cyclic) vs A N ≠ 0 (after density weighting).

The structural claim is conditional and carries no number: **diagonal self-returns concentrate completion demand at a single site, while the hub is the global balancing mode of the K = 7 closure; a density-weighted operator must compare local concentration against global stabilisation, which couples the local diagonal sector to the global hub mode unless completion density is uniform.** Uniform density returns ρ_comp = ρ_cyc and the failed 5/7; non-uniform density forces A N ≠ 0 and w_self = (5 − Tr(A N))/7. The claim is the implication (non-uniform density ⟹ non-diagonal completion operator ⟹ A N ≠ 0), not any value of Tr(A N).

Two things make this a tighter test than fitting a correction to 0.64, and one guards against the obvious failure mode.

- **The discrepancy is signed and pre-set.** Reaching the band requires Tr(A N) ≈ 5 − 7·(0.64 to 0.66) ≈ 0.4–0.6 — a *positive* overlap that *lowers* w_self from 5/7. The correction is not free: it must reduce the cyclic weight by ~8–10%, in a definite direction. And that direction is the structurally-predicted one: a self-return concentrating demand at one site should survive *less* than the cyclic count suggests, i.e. completion density *suppresses* self-return survival. So the sign of the required correction and the sign of the structural argument agree *before any magnitude is computed* — a (weak) piece of over-determination in its own right.
- **The cyclic limit is a zero-parameter prediction already on the table and already wrong.** So the completion correction is explaining a *known, blind, signed discrepancy* (5/7 is 2% high), not producing 0.64 from scratch. That is a far stronger position than a free function aimed at the target.
- **The line that must not be crossed.** Tr(A N) ≈ 0.4–0.6 is *read off the discrepancy*, so a ρ_dens shaped to produce that value proves nothing — it is 16/25 again, one layer down. The valid procedure is fixed: **derive ρ_dens from the completion-density programme first, compute P_A ρ_dens N and hence Tr(A N), and only then compare w_self to the band.** An independently-derived ρ_dens that lands the weight in band explains the discrepancy; one that overshoots, undershoots, or has the wrong sign falsifies the route. The number is earned only if ρ_dens never sees the masses.

So the magnitude branch now has a genuine fork with a derived starting point: **$\rho_{dens} = 0 \Longrightarrow w_{self} = 5/7$ (blind, ~2% high); $\rho_{dens} \neq 0 \Longrightarrow w_{self} = (5 - \mathrm{Tr}(A\,N))/7$, a completion-density prediction to be derived and then checked.** The cyclic near-miss converts "derive 0.64" into "derive the completion correction that lowers a known 5/7 by ~9% in the predicted direction" — sharper, signed, and falsifiable.

Part IV — Assessment

11. Status of the programme

[Proven] here, unconditionally

- An explicit self-adjoint triple (A, B, C) exists whose forward closure $U = A\,B\,C$ is non-normal, with $\mathrm{ran}(U) \neq \mathrm{ran}(U^\dagger)$, and whose orthogonal fold projectors $P_+ = UU^\dagger$, $P_- = U^\dagger U$ are generated by U alone. The operator type the companion paper required is realizable.
- In this model the companion paper's Tier-1 results hold by computation: $D = U^\dagger$, Ω anti-Hermitian, Hermitian fold projectors, Hermitian carrier ΔP with spectrum $\{\pm 1\}$.
- The diagonal access ansatz realizes reorientation; the sign law $\mathrm{sign}\,\chi(g) = \mathrm{sign}\,\delta(g)$ holds. (That a varying weight ratio is a licensed, non-no-go-violating flip is [Definitional] — true by the scope of the no-go, not a proof.)
- The refinement profile yields one prediction, parameter-free *in the increment ratio*: the gen-1→2 log-step is twice the gen-2→3 step ($[\chi(3)-\chi(1)]/[\chi(2)-\chi(1)] = 3/2$, observed 1.503, 0.2%; consecutive form 2 vs 1.988, 0.6%), independent of magnitude, anchor, and K. With one scale G fixed on gen-2, gen-3 is *predicted* to 1% (59.4 vs 60) — one parameter fixed, one prediction confirmed (not "1 against 2 targets"; gen-2 defines G). The increment ratio is the document's most K-robust asset, surviving even if K = 7 is calibrated, and it holds only in MS-bar (the pole scheme, where light quarks are undefined, would break it at 8%).
- The true reversal-orbit count of second-depth closure access is $K(K+1)/2 = 28$; the proposed magnitude $K^2/2 = 24.5$ requires assigning self-return channels weight ½, and the masses prefer weight ≈ 0.64 (required G = 25.48). The combinatorial $K^2/2$ is therefore suggestive at ~4% but not derived.
- The cyclic K = 7 closure model gives, with masses quarantined, $w_{self} = 5/7 = 0.714$ ($A\,N = 0$ in the Fourier basis) — the obstruction-triad route's first blind, falsifiable output. It misses the band (G = 26, both ratios +2%), and the miss is the point: the route can commit to a wrong number, which is what gives it empirical content (§10A.1). The VERSF completion correction must lower 5/7 by ~9%, in the structurally-predicted direction; whether it does is open.

[Conditional]

- χ(1) = ln(6/13) as a fold-census ratio is the upstream (K − 1)/(2K − 1) relabeled ($N_+(1)$ = K − 1, $N_-(1)$ = 2K − 1), inheriting that result's unverified independence from the measured m_u/m_d. It is not an independent second derivation.

[Conjectural]

- The identification A = Π_adm, B = T, C = Π_comp — a labeling of the proven toy structure with VERSF roles.
- The magnitude factor: that the self-return diagonal weight is ≈ 0.64 (hence G ≈ 25.5), and that this emerges from closure self-return / completion-density structure rather than being fitted. The combinatorial $K^2/2$ (weight ½) is the nearest clean candidate and misses by ~4%; the K-expression w_self = (K+2)/2K = 9/14 hits 0.643 and gives the well-targeted G = 25.5, but is reverse-targeted to the fitted 0.64, re-spends the orientation quotient, and adds two census posits (§10), so it is logged, not accepted.
- That the refinement profile is harmonic *or* geometric-½ (observationally identical on three generations) — the shared content, a 2:1 increment ratio, is empirically supported to 0.6%, but neither family's origin is derived; geometric-½ is the better-motivated reading (halving = the fold binary), so "harmonic" is demoted to one candidate of two.
- The obstruction-triad route to the self-return weight (§10A): the standing result is the *reduction* w_self = (5 − $\cos^2\theta$)/7 (conditional on a = 4), with θ the principal angle between admissible and null structure. The cardinality a = 4 is *target-shaped* — only three gates give a reachable band [0.571, 0.714] containing 0.64 — so deriving the cardinality (the {hub, K_c, p_eff} = 2+1 decomposition) gates everything and is the first deliverable. Candidate values (2/3 symmetric; ≈ 0.669 anchor-angle) are held in the band, not boxed. Band-membership is retired as a win condition (the band is nearly fully reachable, and m_s error widens it to ≈ 0.64 ± 0.1); only over-determination — the A–N angle forced to the anchor asymmetry, or the blind lepton transfer — can count. Honest grade: **better-posed but of equal current standing** to the closed-form attempts; it contains zero numbers yet downstream of mass-free computation, only a plan to produce them.

[Open]

- **OP0** (full): show the VERSF closure operators act on H^1, live in a unital ∗-algebra with a meaningful adjoint, and are self-adjoint. The existence witness does not discharge this.
- **OP0.5**: show closure orientation is the up/down bit (candidate route: identify with $T_3 = \pm\frac{1}{2}$).
- The "−1 stabilizing mode" subtraction in $N_-(1)$ = 2K − 1.
- The self-return diagonal weight: derive it from closure / completion-density structure *with the value 0.64 quarantined*. Its right *type* is a completion-survival probability w_self = 1 − σ (the diagonal channel, source = target, is the one most exposed to saturation, which is why the data sit in (½, 1)). A negative result retires the formula-hunt: the target is low-precision (gen-2 vs gen-3 fix G at 25.48 vs 25.65, so w_self ≈ 0.64–0.66), and many simple closed forms fall in that band — 9/14, 16/25, 2/π, e^{−0.446} all ≈ 0.64 — so matching a closed form carries no evidential weight. The valid bar is over-determination (the mechanism also predicts another measured quantity) or landing inside

the band from a blind derivation, not hitting a decimal. The data have ruled out weights 0 and 1.
- A derivation of *why* the refinement profile is harmonic, from the completion-density / transient-spectrum programme. The profile's empirical content (increment ratio 3/2 to 0.2%) is established; its origin is not.
- The forced single sign-crossing as a consequence of derived geometry, not of monotone weights.
- Whether χ saturates or grows unbounded at deeper refinement (uncloseable from three generations alone).
- Why exactly three generations (generation count), not assumed settled.

---

12. Open proof obligations (gates first)

1. **OP0 — operator realization (the gate).** The existence proof (Part I) shows the type is non-empty; OP0 asks whether the *VERSF* operators are of that type. Until solved, Part II's identification stands at conjecture and every Tier-1 transfer is suspended.
2. **OP0.5 — assignment identification.** Is closure orientation the up/down bit? Above all of Part II.
3. **Magnitude / diagonal-weight derivation.** Obtain the self-return weight from closure / completion-density structure without consulting m_c/m_s. Closed-form matching to 0.64 is retired (§10); the standing result is the reduction w_self = (5 − $\cos^2\theta$)/7, conditional on the cardinality a = 4. The cyclic limit is already computed (§10A.1): w_self = 5/7 = 0.714, blind and ~2% high — the route's first falsifiable output. Deliverables, **gated in order**: (3a) **derive the cardinality first** — the {null, over-completion, under-admissibility} = {hub, K_c, p_eff} = 2+1 decomposition, with no reference to the generation count or masses; if it fails, a = K − 3 is dead and the rest is moot; (3b) **derive ρ_dens from the completion-density programme** and compute Tr(A N), checking whether the completion correction lowers 5/7 in the predicted (suppressing) direction — earned only if ρ_dens never sees the masses; the same K = 7 model reads the obstruction ranks and the A–N principal angle, testing over-determination by the 6 : 13 anchor; (3c) run the identical construction in the bath-free lepton sector and predict that weight blind (the cyclic limit already predicts a lepton 5/7-analogue) — the cross-sector kill-shot, and the only deliverable that would establish the route. **Band-membership is not a win condition** (the band is nearly fully reachable, and m_s error widens it to ≈ 0.64 ± 0.1); only over-determination counts.
4. **Profile derivation.** Derive *why* the depth-profile takes the form it does — harmonic (1/k weighting, needing an imported bath-share story) or geometric-½ (decay ratio ½, native to the fold binary), the two being observationally identical on three points. The profile is empirically supported, not rejected — the family-free 2:1 increment ratio matches to 0.6% — so the obligation is to derive a confirmed ratio from structure, with geometric-½ the better-motivated target.
5. **Forced single crossing.** Derive that the fold geometry forces exactly one sign change, independent of magnitude.

6. **Upstream anchor status.** Resolve whether 6/13 (hence K = 7) is derived independently of m_u/m_d. Sets whether $\chi(1)$ is an external lock or a calibration.

---

13. Conclusion

The mechanism the companion programme proposed for the up/down folds — closure-order orientation producing non-normal closure with distinct fold subspaces — is **[Proven]** realizable: an explicit self-adjoint triple instantiates every structural property demanded, with the fold projectors falling out of the closure operator alone. This discharges the existence half of OP0 and removes the worry that the mechanism is empty. It does not identify the VERSF operators with the witness; that identification is **[Conjectural]** and is the substance of OP0 proper.

The attempt to convert the picture into a quantitative ratio law splits into a profile and a magnitude, which fare oppositely. The first-generation value is the upstream $(K - 1)/(2K - 1)$ rewritten in fold-census language, carrying that result's conditionality rather than adding to it. The **profile** is parameter-free *in the increment ratio* — the gen-1→2 log-step is twice the gen-2→3 step (0.6%; cumulative 3/2, 0.2%), independent of magnitude, anchor, and K — while carrying one fitted scale in the magnitude; the single genuine prediction is gen-3 from gen-2 (one parameter fixed, one prediction confirmed to 1%), and it lives only in the MS-bar scheme, the only one in which the light quarks are defined. The form is one of two equally-good families (harmonic, or the better-motivated geometric-½ whose halving is the fold binary); only the shared increment ratio is tested. The **magnitude** is the weak link: the combinatorial $K^2/2 = 24.5$ is equivalent to a self-return diagonal weight of ½, whereas the honest reversal-orbit count is $K(K+1)/2 = 28$ and the masses prefer ≈ 0.64. So $K^2/2$ is suggestive at ~4% but not derived.

The honest standing, in one sentence: the *operator type* behind the fold mechanism is proven to exist and the reorientation reading is sound; the *growth profile* is empirically confirmed to sub-percent precision but not yet derived; and the *one remaining free quantity* is the closure self-return weight, which the data pin near 0.64–0.66 and which closed-form matching cannot settle (the band is dense with simple fractions). The live route to it is the obstruction triad (§10A), whose standing result is the reduction $w\_self = (5 - \cos^2\theta)/7$ — the magnitude recast as a principal angle between admissible and null structure, with no decimal boxed as the answer. It is, for the first time, a correctly-posed and falsifiable derivation program rather than a fit: three mass-quarantined deliverables — one model calculation on the K = 7 closure graph that forces the obstruction ranks *and* tests whether the A–N angle is over-determined by the 6 : 13 anchor; the completion-anchor identification of the triad; and a cross-sector lepton prediction run blind — any one of which can kill it. The route from closure orientation to fold structure is explicit and proven realizable; the route to the numbers is down to a single weight whose derivation is now posed sharply enough to fail.