From Charge-Sector Maintenance to Quark Masses in VERSF

A Five-Gate Attempt at the Absolute Scale, Common-Mode Overlap, Generation Recursion, Confinement Projection, and Proportional-Readout Bridge — with a Predictive-Content Audit after the Interface-Coupling and χ-Halving Papers

Keith Taylor *VERSF Theoretical Physics Programme — Standard Model Readout Series*

*Successor to the realised-readout two-part paper. Part I of that work built the access-visible readout channel and reduced the quark question to one inner product; Part II forced that inner product nonzero under the maintenance-conjugacy premise. This paper asks what else must be derived before quark **masses** — not merely the quark **split** — can be computed, attempts the five gates that separate the two, and audits exactly how much of the resulting numerical agreement is prediction and how much is postdiction, now in light of the interface-coupling derivation of α and the single-scheme χ-halving audit.*

Summary for the General Reader

Every fundamental particle has a mass, and in the Standard Model each quark mass is a number physicists measure and write down by hand. The previous paper did not compute these numbers. It reduced the quark problem to one exact test: does the committed record move the underlying closure system in a direction that up-type and down-type quarks read *differently*? In symbols the test was

$B_q = ½ D^T M^{-1} b_m$,

and the paper showed that if the record couples to the same signed maintenance functional that distinguishes up from down, the test is forced nonzero,

$B_q = g_\chi^T H^{-1} g_\chi > 0$.

That proves the split is real under one premise. It still does not give the masses.

A split is not a mass. To get from "up and down differ" to "up weighs this many MeV," four more things are needed: an absolute scale, the part of the mass shared by up and down before the split acts, a rule carrying the first generation into the heavier ones, and a way to turn confined quark energy into the "current mass" physicists quote. A fifth gate ties the readout language to the actual Yukawa couplings. This paper attempts all five.

For the first generation it gets a clean conditional result, $m_u \approx 2.26$ MeV and $m_d \approx 4.90$ MeV, and for the strange quark $m_s \approx 97$ MeV. These are close to the measured values. But the paper is built to be honest about why they are close. The newest interface-coupling papers improve one previous weakness: the number 1/137 is no longer simply borrowed from experiment. Within the VERSF coupling programme, α is structurally derived as an interface-transfer constant — an impedance mismatch between quantum transport and the two-polarization electromagnetic vacuum, with a $K = 7$ / $N_{loop} = 14$ closure derivation and a second-order 1/6 correction. What is still not proven is that this same interface constant is the projection from confined quark closure energy to the current quark mass physicists quote. So the vulnerable step has sharpened: not "why use α?" but "why should the quark current-mass projection be governed by the structurally derived interface coupling?"

Three cautions are stated plainly and carried throughout rather than buried:

- The conversion from confined energy to current mass uses an interface-transfer factor. Simply using the measured electromagnetic fine-structure constant would make the light-quark scale look fitted; the coupling papers strengthen this, deriving α itself structurally as an interface impedance / closure-channel constant. The remaining bridge is narrower but still real — show that the constituent-to-current quark projection is governed by that same interface transfer, possibly with a small scheme/chiral factor η_q.
- The three masses that come out close are not three independent successes. They share one scale and therefore drift together (all ~4–5% high in the same direction). The genuinely strong light-sector result remains the *ratio* $m_u/m_d = 6/13$; the common scale is now "structurally motivated through α_{int}, with a residual projection factor," not fully derived.
- For the heavy quarks, an unconstrained master formula would have too many free slots. A stricter leading-order discipline is imposed: the within-generation up/down split is carried by $M_{\chi,g}$, while the common generation growth is carried by a shared G_g. If the formula is allowed independent $G_{u,g}$ and $G_{d,g}$ without deriving them from PFD stiffness data, it becomes over-saturated and loses predictive content.

In one line: *the first-generation ratio is structurally strong; the common scale is now tied to a structurally derived interface coupling but still requires a quark-current projection theorem; the strange quark follows through inherited meson/chiral relations; and the heavy-quark sector is not a free six-slot fit but a constrained stiffness-plus-halving scaffold whose predictive content depends on computing G_g, $M_{\chi,g}$, and the scheme factors from closure geometry.*

Table of Contents

---

0. The Predictive-Content Ledger

This section is placed first deliberately. Every numerical agreement below is the product of several inherited relations, and the value of the paper cannot be judged without knowing, for each input, whether it is *derived elsewhere*, *postdicted* — a known quantity re-expressed in K = 7 language — or *fitted*. The standard test of predictive content is to count independent inputs against independent outputs: a scheme that consumes as many inputs as it produces outputs has compressed nothing, however close its numbers.

The status of the α-projection in the ledger. Treating the α-projection as a direct empirical insertion is too weak, because the coupling programme gives α a structural origin: one line identifies α as the impedance ratio

$\alpha = Z_0 / (2 R\_K)$,

and the closure/interface sequence supplies the K = 7, N_loop = 14, and second-order 1/6 correction structure that brings the inverse coupling close to the observed low-energy value

(§2.5). The matching papers interpret this as a UV/interface boundary condition propagated by ordinary four-manifold RG flow. So the correct status is no longer "α is simply fitted" but the two-part statement:

- α_int is **structurally derived** as an interface coupling;
- the use of α_int as the **quark current-mass projection** remains a bridge.

First-generation light quarks (outputs: m_u, m_d — two numbers).

Input	Value	Status	Note
m_π	139.57 MeV	Experimental	Sets the hadronic scale used by the entropic-confinement paper
Constituent factor	(20/9)	Inherited (status unadjudicated)	$(20/9)\cdot m\_\pi \approx 310$ MeV $\approx m\_N/3$; this paper inherits it and does not adjudicate whether the companion derivation is independent of the nucleon-third anchor
Interface projection	α_int	Structurally derived as coupling; projection bridge open	α itself is no longer merely borrowed; the open step is whether quark current-mass visibility is governed by the same interface transfer
Residual projection factor	η_q	Open	Needed once the ~4–5% shared high offset is treated as a projection/scheme correction; currently *defined by* that offset, $\eta\_q \approx 0.956$
Maintenance ratio	6/13	Inherited; strongest ratio claim	$m\_u/m\_d = e^{(-2M\_\chi,_1)}$; consistent with the PDG central ratio (the data itself carries ~10–20% asymmetric uncertainty, so no tighter precision is claimable)

The light-quark formula should therefore be written as

$m\_q{,}cur = \alpha\_int\ \eta\_q\ m\_q{,}conf,$

not simply $m\_q{,}cur = \alpha\ m\_q{,}conf$. At leading order **$\eta\_q = 1$** gives the familiar estimates $m\_u \approx 2.26$ MeV and $m\_d \approx 4.90$ MeV — both ~4.6% high. The observed common offset then asks for

$\eta\_q \approx 0.956,$

up to the precise mass inputs and scheme convention. **This must be read honestly: $\eta\_q \approx 0.956$ is at present the offset it is meant to remove (2.16/2.26 = 0.956), i.e. a one-number residual, not a derived correction.** The leading-order prediction is $\eta\_q = 1$; the ~4.6% miss is the residual to be *derived* from partial-closure geometry, not a success already in hand. That residual is now the real scale gap.

Correlation caveat — do not triple-count. The three reproduced light-sector masses (m_u, m_d, m_s) all come out high in the same direction. The independent content is not three masses. It is: one common scale, now tied to $\alpha\_int\ \eta\_q$; one strong ratio, 6/13; one strange-sector meson/chiral extension. The ratio is the strongest piece. The scale is improved (α is no longer borrowed) but not closed (the projection bridge and η_q are open).

Strange quark (output: m_s — one number). Adds the $K = 7$ meson relation $m_K/m_\pi = K/2 = 3.5$ (empirically 3.54, a ~1% **postdiction**) and an active-mode correction. Stated with care to avoid a flipped definition: the active-mode fraction is

$$f_{active} = (2K - 1)/(2K + 1) = 13/15,$$

and the correction *applied* to the leading strange estimate is its **reciprocal**,

$$f_{active}^{-1} = (2K + 1)/(2K - 1) = 15/13 \approx 1.154,$$

which raises $m_s^{(0)} \approx 84.2$ MeV to $m_s \approx 97.2$ MeV. Unless f_{active}^{-1} is derived independently of m_s, it carries a fit signature. Two inputs for one output: no compression at this gate.

Heavy quarks (outputs: m_c, m_s, m_b, m_t). No longer allowed to hide behind a general $G_{a,g}$. At leading order this paper imposes a **shared within-generation amplifier**, $G_{u,g} = G_{d,g} = G_g$, so sectoral splitting inside a generation is carried by $M_{\chi,g}$:

$$m_{a,g\_cur} = m_{q,0\_cur} \cdot G_g \cdot e^{(s_a M_{\chi,g})} \cdot \Pi_{a,g}^{scheme}.$$

Without this restriction, independent $G_{u,g}$ and $G_{d,g}$ factors plus scheme factors can fit the six quark masses without predicting them. Predictive content requires a closure-geometric theorem reducing the G's and M_χ's to PFD stiffness data. The χ-audit and half-lazy theorem (§2.7) further constrain the $M_{\chi,g}$ profile — though the identification that ports the audited ratio onto the within-generation split is itself a bridge, flagged below.

The bottom line.

- α itself: **structurally strengthened** (no longer borrowed).
- α as quark current-mass projection: **still conditional** (the bridge, plus the η_q residual).
- 6/13: **strongest light-sector ratio**.
- m_s: inherited chiral/meson extension, **not an independent third confirmation**.
- Heavy quarks: constrained scaffold; **no predictive content** until G_g, $M_{\chi,g}$, and $\Pi_{a,g}^{scheme}$ are computed from closure data.
- Proportional bridge for the signed split: **reduced to a single common-generator premise** (§6) — the one move that is neither numerical nor vulnerable, and a genuine two-assumptions-to-one reduction.

Calibration note — the K = 7 load. The integer $K = 7$ carries an unusual amount of weight across the construction: 2^7 in α_{int}, $6/13 = (K-1)/(2K-1)$, $K/2 = 3.5$ for the meson ratio, $13/15 = (2K-1)/(2K+1)$ for the active-mode factor. The independence guard applied to α_{int} (§2.5) should be applied *systemically*: a single sensitivity table showing which of these results survive $K = 6$ or $K = 8$ would expose whether they are independent $K = 7$ consequences or a flexible label fitting several numbers at once. This paper does not supply that table, and until it does, the recurrence of $K = 7$ is a structural promise, not a verified structural fact.

---

1. The Five Gates from Quark Contrast to Quark Masses

The previous realised-readout paper made the quark split nonzero under maintenance-conjugacy, but it did not compute masses. The gap is structural, not incremental:

quark split ≠ quark masses.

The split requires only a signed differential factor. The masses additionally require the absolute common-mode scale, the generation amplifier, and the projection to a current-mass scheme. Five distinct gates separate the contrast from the masses:

1. **Absolute overlap scale.** Fix κ_0, or the physical product $\kappa_0 I_{common}$, from closure geometry.
2. **Common-mode quark overlap.** Identify the part shared by up and down before the signed split.
3. **Generation dependence.** Derive (u, d) → (c, s) → (t, b).
4. **Confinement/current projection.** Justify the map from confined partial-closure energy to a current quark mass.
5. **Proportional-readout bridge.** Show the readout log-ratio equals the Yukawa log-ratio.

This paper attempts all five. The order below is reordered for logical flow: the gates that yield genuine results (1, 2, 5) first, the load-bearing vulnerability (4) confronted in the middle now that the α-interface papers have changed its grade, and the generation gate (3) last with its constrained-but-not-yet-computed status.

---

2. Prior Work Used

2.1 BCB fold-overlap Yukawas

The BCB Fold Lagrangian writes the effective Yukawa coupling as

$\kappa_f = \kappa_0 I_f$,

with I_f a dimensionless geometric fold-overlap integral. The quark Yukawa terms are

$\mathscr{L}_Y = -\kappa_u (\bar{Q}_L \tilde{H} u_R + h.c.) - \kappa_d (\bar{Q}_L H d_R + h.c.) + \cdots$, so the quark masses are not free Yukawa labels but closure-overlap readings, once I_f and κ_0 are fixed. **[Inherited]**

2.2 Substrate-stiffness hierarchy

The substrate-stiffness paper gives the mass architecture

$$m_D = \mathcal{D}_{\gamma_D} S(D) v, \quad S(D) = S_H(D) S_L(D) S_P(D) S_I(D),$$

the four factors being closure-Hessian stiffness, localization compression, persistent-distinguishability load, and interface-transport complexity. The product form is equivalent to additive free-energy contributions in log-space. The paper states that quark/lepton separation is driven primarily by additional quark contributions to S_P and S_I — colour-charge distinguishability bookkeeping and multi-channel confinement transport. This matters here because the signed maintenance functional M_χ should live in S_P, S_I, or their coupled sector. **[Inherited]**

2.3 Entropic confinement and effective quark masses

The entropic confinement paper gives the quark-side ontology: quarks are level-2 partial closures, not complete level-4 intrinsic rest-mass defects, and acquire effective mass only inside committed hadronic configurations. It supplies a closure-anchored constituent light-quark scale

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

and proposes a projection to current mass, $m_{q,\mathrm{cur}} = \alpha\, m_{q,\mathrm{const}}$, explicitly marked as model-internal, with first-principles renormalization-group matching left open. It also gives the up/down ratio $m_u/m_d = 6/13$ from $K = 7$ closure-mode coupling asymmetry, and a strange estimate $m_s \approx 97$–100 MeV through the $K = 7$ kaon/pion relation plus an active-mode correction. **[Inherited; α-projection re-graded in §2.5, §5]**

Note on (20/9). The constituent value $(20/9)\cdot 139.57 \approx 310$ MeV coincides with $m_N/3$, the canonical constituent quark mass. It is inherited as $K = 7$ closure structure; this paper does not adjudicate whether the companion derivation of 20/9 is independent of the nucleon-third anchor.

2.4 Realised readout and the proportional bridge

The realised-readout paper marks the proportional-readout hypothesis as load-bearing: the access ratio must be *identified* with the Yukawa log-ratio, and that identification is not automatically supplied by the overlap picture. This paper discharges that bridge for the signed quark split (§6), not in general. **[Inherited; partially discharged below]**

2.5 Interface derivation of α and why it changes Gate 4

The projection $m_{q,cur} = \alpha\, m_{q,const}$ is the natural place to locate the central weak point, and the interface-coupling papers change its grade — a change that must be stated at exactly the strength the companions support, no more.

The fine-structure constant is supported by two compatible structures. First, the **exact impedance interpretation**

$$\alpha = Z_0 / (2 R_K),$$

where Z_0 is the vacuum electromagnetic impedance and R_K the von Klitzing resistance quantum. This reads α as an interface mismatch between one-channel quantum transport and the two-polarization electromagnetic vacuum. *A guard is owed here and stated plainly:* this identity is **exact but is a re-expression, not an independent derivation** — $Z_0 = \mu_0 c$ and $R_K = h/e^2$ together give $Z_0/(2R_K) = \mu_0 c e^2/(2h) = \alpha$ by definition of these constants, so the impedance form re-coordinatises α rather than predicting it. The independent content lives entirely in the second structure.

Second, the closure/interface papers derive the $K = 7$ / $N_{loop} = 14$ formula and its second-order correction,

$$\alpha^{-1} = 2^K [(N+1)/N - 1/(6N^2)],\ K = 7,\ N = 14,$$

which evaluates to 137.034 against the observed 137.036 — agreement to ~15 ppm. The coefficient 1/6 is traced to six-channel local capacity allocation and inverse-participation structure rather than to a fit. The interface-matching paper supplies the QFT interpretation: the interface fixes the UV/boundary normalization of the gauge kinetic term, and ordinary four-manifold RG flow propagates that value to observable scales.

The independence guard, identical in form to the $\kappa = 8/3$ caution. The 15-ppm agreement carries persuasive weight **only if** $K = 7$, $N = 14$, and the 1/6 coefficient are fixed by closure structure *ahead of* the value 137.036, not selected to land on it. The companions assert exactly this ($K = 7$ and $N_{loop} = 14 = 2(3+4)$ are claimed structural, the 1/6 from channel counting); this paper inherits that assertion at the companions' grade and does not re-derive it. The honest status is therefore: **α_{int} is structurally derived assuming its three integers are independently fixed** — a materially stronger position than "α is borrowed," but one whose force is exactly the independence of those integers.

Therefore this paper does not say "α is fitted." It says: α_{int} is structurally derived as an interface coupling; the open bridge is whether quark current-mass visibility uses α_{int}. This sharpens Gate 4 — the open question is no longer the value of α but the **projection theorem**, $\Pi_{cur} = \alpha_{int}\, \eta_q$.

2.6 Fold-selective accessibility, closure orientation, and M_χ

The signed maintenance functional M_χ should not remain an abstract variable. The quark-specific papers give it a natural home.

The fold-selective accessibility paper states the target in terms of fixed fold projectors P_{up}, P_{down} and a generation-dependent access operator R_g. The same-generation susceptibility is

$$\chi(g) = \ln[\, m_{up}(g) / m_{down}(g) \,],$$

and the mechanism cannot be simple carrier growth, because the refinement-persistent carrier does not grow under refinement. The split must come from depth-dependent accessibility or orientation of a fixed carrier.

The closure-orientation paper then proposes a concrete source for the folds: forward closure versus reverse closure,

$$U = ABC, \quad D = CBA,$$

and proves the operator class is buildable in a toy construction. Its strongest grade is an **existence result for the mechanism type**; the physical identification with actual up/down remains conjectural.

The admissibility-maintenance paper supplies the signed energetic interpretation. It argues that no single-signed cost can reproduce both up-type amplification and down-type suppression, so the maintenance factor must be signed:

$$M = M_{loc}\, \Theta_{adm}, \quad \ln M = \ln M_{loc} + \ln \Theta_{adm},$$

where positive $\ln \Theta_{adm}$ is a burden and negative $\ln \Theta_{adm}$ is screening. This paper therefore identifies $M_\chi(q)$ as the up/down-odd component of the admissibility-maintenance coupling,

$$M_\chi(q) = \tfrac{1}{2} [\, \ln \Theta_{adm,d}(q) - \ln \Theta_{adm,u}(q) \,],$$

with sign convention chosen so that $m_u/m_d = e^{(-2M_{\chi,1})}$. This places M_χ in a definite substrate sector: not generation depth alone, not closure class alone, but the charge-sector-odd component of persistent-distinguishability load S_P, interface-transport complexity S_I, or their coupled term. **[Conditional — existence-grade mechanism plus a conjectural physical identification]**

2.7 The single-scheme χ audit and half-lazy closure theorem

The heavy-quark split profile is constrained by the χ-audit and the half-lazy theorem, but the constraint must be stated together with the bridge it depends on, because that bridge is where the data pressure lands.

The **single-scheme χ audit** removes the danger that the apparent halving pattern was a mixed-scheme artifact. Under consistent $\overline{MS}$ QCD running with threshold matching, it finds

$$\rho = \Delta\chi_2 / \Delta\chi_1 = 0.503 \pm 0.023,$$

consistent with ½. The audit does not prove exact equality, but it confirms ½ is the right structural target for **the audited χ**, and that scheme-mixing is no longer the main worry.

The **half-lazy closure paper** gives a conditional structural reduction. Under a genuine two-fold sector, image-even closure, quotient-binary one-step history, equal primitive-class weights, no branch-cost asymmetry, and log-access intertwining, the update operator is

$$W = \tfrac{1}{2} (I + \Pi_e),$$

with odd eigenvalue ½, hence $\Delta\chi_{g+1} = \tfrac{1}{2} \Delta\chi_g$. This is much sharper than a fitted halving rule: it gives the exact list of conditions that force ½.

The bridge that must be named, not assumed. The reduction above governs the *audited* χ. To use it on the heavy-quark up/down profile, one must identify that audited χ with the within-generation maintenance contrast, $\chi_g = \ln[m_{up}(g)/m_{down}(g)] = -2M_{\chi,g}$. **This identification is an undischarged bridge, and the observed quark masses warn against taking it for the naïve within-generation ratio.** Reading χ_g straight off PDG values gives $\chi_1 = \ln(m_u/m_d) \approx -0.78$ (down heavier), but $\chi_2 = \ln(m_c/m_s) \approx +2.61$ and $\chi_3 = \ln(m_t/m_b) \approx +3.66$ at common μ (up heavier): **the up/down ordering inverts after the first generation**, and the increment ratio is $\Delta\chi_2/\Delta\chi_1 \approx 0.33$, not ½. So the audited $\rho \approx 0.503$ is almost certainly *not* the naïve within-generation up/down ratio; it is a scheme-consistent construction in the audit paper whose relation to $-2M_{\chi,g}$ must be supplied explicitly. Until that map is given, the clean recursion below is the structure of the *audited* χ, ported onto $M_{\chi,g}$ **conditionally**.

Granting the bridge, the maintenance contrast inherits the halving, $\Delta M_{\chi,g+1} = \tfrac{1}{2} \Delta M_{\chi,g}$, so the three $M_{\chi,g}$ are not independent: with $M_{\chi,1}$ fixed by 6/13 and one first increment $\Delta M_{\chi,1}$,

$$M_{\chi,g} = M_{\chi,1} + 2\, \Delta M_{\chi,1} \left(1 - 2^{-(g-1)} \right).$$

That is a real reduction of the heavy-sector split freedom from three parameters to two **if and only if** the χ-identification bridge holds; the sign-inversion and the 0.33-versus-½ tension above are exactly what that bridge must reconcile. **[Conditional — on the audited-χ ↔ $-2M_{\chi,g}$ identification, which the masses do not grant for the naïve ratio]**

2.8 Generation and matter-skeleton constraints

The Generation Theorem says generation is class-grain refinement, not a member-grain label, and reduces the three-level count to refinement structure with two marks under stated REF-DIM and REF-BIN conditions. The point here is that generation is not merely a mass tier; it is a refinement-class structure that must be represented at class grain.

The Matter Skeleton Theorem reinforces the status discipline. It gives a skeleton — fractional charges, absence of free fractional members, finite repetition, conditional mass ordering — but explicitly does *not* derive the exact spectrum, CKM/PMNS, or the full Standard Model. That is the right inheritance grade: generation architecture is structurally supported, but numerical quark masses require additional stiffness and projection calculations. The common amplifier G_g

below must ultimately be derived from class-grain refinement and PFD stiffness data, not treated as an arbitrary index. **[Inherited, skeleton-grade]**

2.9 Why ratios are the right targets: renormalization-group invariance

There is a structural reason, deeper than "inputs cancel," that the ratio 6/13 is strong and the absolute scale is weak, and naming it reshapes what the heavy sector should even try to predict. **Quark mass ratios are renormalization-group invariant**: the mass anomalous dimension is flavour-blind, so m_i/m_j is scheme- and scale-independent to all orders. Absolute current masses are not — they are $M\bar{S}$-at-2-GeV conventions. Consequently η_q and $\Pi_{a,g}^{scheme}$ are partly being asked to reproduce a *bookkeeping convention*, not physics, which is why "derive η_q from pure closure geometry" is conceptually strained: geometry has no reason to know about the 2 GeV choice.

The constructive consequence drives the rest of the paper's targeting. In the heavy formula $m_{a,g} = m_{q,0} G_g e^{(s_a M_{\chi,g})} \Pi_{a,g}^{scheme}$, the Π factors should be (nearly) flavour-universal at fixed scale and therefore **cancel in same-scale ratios**. So the hard, partly-conventional task "derive $\Pi_{a,g}^{scheme}$" is the wrong target; the clean and likely-true task is "**show Π^{scheme} cancels in the ratios**." That converts the heavy sector from a scaffold carrying two underived convention factors (η_q, Π^{scheme}) into a set of **scheme-independent ratio predictions plus one experimental scale anchor** — which is exactly the paper's own "ratio strong, scale weak" thesis made structural rather than rhetorical. The targets throughout should accordingly be the scheme-independent ratios, in the *orthogonal basis* that the RGI structure makes available (§10): within-generation ratios $m_{up}(g)/m_{down}(g)$ isolating $M_{\chi,g}$, and cross-generation geometric-mean ratios $\sqrt{(m_{up} m_{down})_{g+1}}/\sqrt{(m_{up} m_{down})_g}$ isolating G_g, with a single measured anchor fixing the overall scale that geometry should not be expected to supply. **[Structural reframing — adopted as the heavy-sector target discipline]**

---

3. Gate 1 — The Absolute Overlap Scale

The BCB formula is

$m_f = (v/\sqrt{2}) \kappa_f = (v/\sqrt{2}) \kappa_0 I_f.$

The universal κ_0 is not separately observable until a normalization convention for the dimensionless I_f is fixed; what is physical is the product $Y_f = \kappa_0 I_f$. For the first-generation light-quark common-mode scale, the entropic-confinement constituent energy $m_{q,const} = (20/9) m_\pi$ projects, under the interface projection of §5, to

$m_{q,proj} = \alpha_{int} \eta_q m_{q,const}, \rightarrow (\eta_q = 1) \alpha_{int} (20/9) m_\pi.$

Numerically, with $m_\pi \approx 139.57$ MeV and $\alpha_{int}^{-1} \approx 137.036$,

$m_{q,const} \approx 310.16$ MeV, $m_{q,proj} \approx 2.26$ MeV (at $\eta_q = 1$).

This is the light-*branch* projected scale, not yet the symmetric common-mode scale. The signed maintenance split (§4) places u and d on opposite sides of a common-mode midpoint. With

$m\_u/m\_d = e^{(-2M\_\chi,_1)} = 6/13 \Rightarrow M\_\chi,_1 = ½ \ln(13/6) \approx 0.3866$,

define the symmetric first-generation common-mode current scale as the geometric mean

$m\_q,0\_cur = \sqrt{(m\_u\ m\_d)} = \alpha\_int\ \eta\_q\ (20/9)\ m\_\pi\ \sqrt{(13/6)} \approx 3.33$ MeV ($\eta\_q = 1$).

Equivalently the absolute common-mode overlap product is

$\kappa_0\ I\_q,0^{(1)} = (\sqrt{2} / v) \cdot \alpha\_int\ \eta\_q\ (20/9)\ m\_\pi\ \sqrt{(13/6)}$,

and with $v \approx 246$ GeV, at $\eta\_q = 1$,

$\kappa_0\ I\_q,0^{(1)} \approx 1.92 \times 10^{-5}$.

This fixes the physical *product* for first-generation light quarks; it does not separately fix κ_0 and I_q,0, which awaits the full BCB overlap calculation.

The branch-assignment choice, flagged. The working formula puts the **bare** interface coupling on the **up** branch — m_u = α_int η_q (20/9) m_π, with m_d larger by 13/6 — so the common-mode scale is √(m_u m_d) = α_int η_q (20/9) m_π · √(13/6), a factor √(13/6) ≈ 1.47 above the bare projection. The symmetric alternative, putting the bare coupling on the *common-mode* scale (m_u = α_int η_q (20/9) m_π / √(13/6)), gives m_u ≈ 1.54, m_d ≈ 3.33 MeV — both ~30% low. **Nothing here derives which branch carries the bare α; it is fixed by the data**, and this ~1.47× assignment is a second piece of unflagged scale work sitting *above* the ~5% η_q residual. The two must not be conflated: the absolute light scale decomposes as a 1.47× branch choice followed by a ~5% η_q correction. The (20/9) ≈ m_N/3 motivation (§2.3) actually points the *symmetric* way — in the constituent-quark model m_N/3 ≈ 310 MeV is the constituent mass of *each* light quark, u and d alike, which is the failing symmetric reading, not the working up-branch one. This is sharper than "underived": it may be in **direct conflict with the inherited definition the formula is built on.** If the entropic-confinement companion (ref 5) intends m_q,const = (20/9) m_π as a *single common* light-quark constituent (the standard ≈ m_N/3 reading the coincidence invokes), then applying α to it yields the *common-mode* scale — the symmetric, ~30%-low result — and the working paper has silently reassigned the bare projection from the mean to the up branch, a √(13/6) ≈ 1.47× move against its own source. So the open item is not merely "data-fixed rather than derived" but "resolve whether (20/9) m_π is the *up* constituent or the *mean* constituent in the source, because the two readings disagree by √(13/6) and only the mean reading is consistent with the m_N/3 motivation." As of this paper it is having it both ways, and that is recorded as open, not resolved.

Gate 1 grade — [Conditional]. Closed at the level of the light-quark common-mode *product*, using the K = 7 constituent scale and the interface projection. **Not** closed as a derivation of κ_0. The entire scale, including this 1.92×10^{-5}, carries the α_int projection bridge and the η_q residual (§5) as multiplicative factors.

4. Gate 2 — Common-Mode Overlap and the Signed Maintenance Split

Let $a \in \{u, d\}$ with signs $s_u = -1$, $s_d = +1$. Define the first-generation signed maintenance amplitude by

$e^{(-2M_{\chi,1})} = 6/13$, so $M_{\chi,1} = \frac{1}{2} \ln(13/6)$.

The first-generation overlap form is $I_{a,1} = I_{q,0}^{(1)} \exp(s_a M_{\chi,1})$, so

$m_{a,1\_cur} = (v/\sqrt{2})\, \kappa_0\, I_{q,0}^{(1)} \exp(s_a M_{\chi,1})$,

with common-mode scale $m_{q,0\_cur} = (v/\sqrt{2})\, \kappa_0\, I_{q,0}^{(1)} = \sqrt{m_u m_d}$. Thus

$m_u = m_{q,0\_cur}\, e^{(-M_{\chi,1})}$, $m_d = m_{q,0\_cur}\, e^{(+M_{\chi,1})}$,

and since $e^{(M_{\chi,1})} = \sqrt{13/6}$,

$m_{u\_cur} = \alpha_{int}\, \eta_q\, (20/9)\, m_\pi$, $m_{d\_cur} = \alpha_{int}\, \eta_q\, (20/9)\, m_\pi\, (13/6)$,

numerically ($\eta_q = 1$)

$m_{u\_cur} \approx 2.26$ MeV, $m_{d\_cur} \approx 4.90$ MeV.

The interpretation of M_χ is fixed, not free. It is not an arbitrary signed parameter. It is the up/down-odd part of the admissibility-maintenance factor (§2.6),

$M_\chi(q) = \frac{1}{2} [\ln \Theta_{adm,d}(q) - \ln \Theta_{adm,u}(q)]$,

justified because a single-signed cost cannot reproduce both up-type amplification and down-type screening. The fold-orientation papers give a candidate substrate source — closure-order orientation, schematically U = ABC versus D = CBA — so that M_χ is the maintenance-cost difference between the two closure orientations, read through fold-selective accessibility.

The signed-maintenance interpretation also unifies two previously separate statements,

$B_q > 0$ and $m_u/m_d = 6/13$,

as consequences of a single nonzero signed charge-sector maintenance coordinate. That unification is real and is independent of whether the *scale* is correct.

Gate 2 grade — [Conditional, strong on ratio]. Closed for the first-generation up/down *ratio* if 6/13 is accepted as the first-generation maintenance amplitude — the paper's best-supported numerical claim (consistent with the PDG central ratio; the ~10–20% data uncertainty forbids a tighter figure) — and structurally sharpened because M_χ now has a proposed substrate home (signed admissibility maintenance in S_P, S_I, or their coupled sector). The *scale* attached to it carries the α_{int} projection bridge.

---

5. Gate 4 — Confinement-to-Current Projection after the α-Interface Papers

Quarks are not isolated rest-mass defects in the way charged leptons are. The entropic-confinement paper treats quarks as level-2 partial closures whose effective masses exist only inside committed hadronic configurations, and gives the light constituent scale $m_{q,const} = (20/9)\, m_\pi$. It then proposes $m_{q,cur} = \alpha\, m_{q,const}$, while openly noting this uses external inputs and is not yet a first-principles scheme matching.

The interface-coupling papers improve the status of the factor α (§2.5): they derive α as an impedance/interface quantity and connect interface coupling to gauge normalization through a boundary-condition / RG-matching picture. The worry "why insert α at all?" is therefore answered to the companions' grade. **What is not answered, and is now the whole of Gate 4's vulnerability, is the projection step: why the *quark current-mass* visibility should be governed by the interface coupling.** The correct updated projection is

$$m_{q,cur} = \alpha_{int}\, \eta_q\, m_{q,conf},$$

with α_{int} the structurally derived interface coupling, $m_{q,conf}$ the confined partial-closure maintenance energy, and η_q a residual current-mass projection factor carrying scheme, chiral, and partial-closure leakage corrections.

The honest accounting of η_q. At leading order $\eta_q = 1$ recovers the old estimates, which are ~4.6% high. The shared offset across m_u, m_d, m_s then *defines*

$$\eta_q \approx 0.956 = 2.16/2.26,$$

up to mass inputs and scheme. This is a sharpened — not eliminated — weakness: the question is not "why α?" but "why should current-mass visibility use α_{int}, and what fixes η_q?" Stating $\eta_q \approx 0.956$ is **not** a second success; it is the same ~5% miss relocated into a named residual that must be *derived*, on pain of being a one-parameter scale fit per the falsification list (§12). One tempting reduction must be ruled out explicitly rather than invoked: the needed $\eta_q < 1$ has **no home in the coupling running.** Current masses are quoted at 2 GeV, where $\alpha_{EM} \approx 1/133.5 > \alpha(0) = 1/137$, and since $m \propto \alpha$ this makes the prediction *larger* ($2.26 \to 2.35$), the wrong direction; and $\alpha(0)$ is the IR endpoint of the QED coupling, so there is no smaller value to appeal to in any case. (QCD mass running is a different object and is not obviously favorable either — $\overline{MS}$ masses rise toward 2 GeV from above.) The reduction $\eta_q \approx 0.956$ therefore requires a genuine partial-closure-leakage mechanism, not a renormalization-group artifact; this is stated so the residual cannot be quietly attributed to running it does not possess.

Current-Mass Interface Projection Hypothesis. If a current quark mass is the primitive local interface-visible component of confined partial-closure maintenance energy, and if interface transfer is governed by the same single-channel-to-two-polarization mismatch that fixes α_{int}, then $m_{q,cur} = \alpha_{int} \eta_q m_{q,conf}$, with leading-order $\eta_q = 1$. A full derivation must compute η_q from partial-closure geometry and map it to a standard current-mass scheme. **[Conditional — load-bearing; η_q currently fit]**

Gate 4 grade — [Strengthened but open]. α is structurally supported (no longer borrowed); its use as the quark current-mass projection remains conditional, and η_q is, as of this paper, the ~5% offset by definition rather than a derived factor. The α-Projection Independence Problem stands: derive $\Pi_{cur} = \alpha_{int} \eta_q$ from closure/interface microphysics *without using* m_u, m_d, m_s, and test whether the result reproduces the observed scale.

---

6. Gate 5 — The Proportional-Readout Bridge for the Signed Split

This gate yields a real conditional result, separated cleanly from the conditional scale material above.

The realised-readout paper marked the bridge $\ln(y_u/y_d) = \ln(a_u/a_d)$ as load-bearing and undischarged. The maintenance framework discharges it **for the signed quark split** by exhibiting a *common generator*. Assume both the overlap odds and the access odds descend from the same maintenance functional M_χ.

Overlaps. With $I_u = I_0 e^{-M_\chi}$ and $I_d = I_0 e^{+M_\chi}$, $\ln(I_u/I_d) = -2 M_\chi$.

Access readout. Take the Gibbs form $R_\chi \propto \exp(-M_\chi \Delta P)$, with ΔP carrying the up/down access contrast. Then $a_u \propto e^{-M_\chi}$, $a_d \propto e^{+M_\chi}$, so $\ln(a_u/a_d) = -2 M_\chi$.

Therefore $\ln(I_u/I_d) = \ln(a_u/a_d)$, and since $y_f = \kappa_0 I_f$ the universal κ_0 cancels, giving

$$\ln(y_u/y_d) = \ln(a_u/a_d).$$

Maintenance Bridge Theorem (conditional). If the access readout and the Yukawa overlap are Gibbs/readout expressions of the *same* signed maintenance functional M_χ, then the proportional-readout hypothesis holds for the up/down log-ratio.

The content of this gate is precise and should not be overstated. The conclusion $\ln(I_u/I_d) = \ln(a_u/a_d)$ is near-tautological once both sides are *posited* as $e^{\mp M_\chi}$; the whole content lives in the **common-generator premise** — that one functional M_χ generates both the overlap odds and the access odds, with the same sign. So the accurate verb is **reduced**, not discharged: the proportional-readout bridge, previously two independent assumptions (that an access contrast exists, and that it equals the Yukawa log-ratio), is reduced to the single premise of a common generator. That is a genuine two-to-one reduction and is worth keeping as the paper's cleanest internal move — but it is a reduction of assumptions, not their elimination.

Gate 5 grade — [Conditional, genuine reduction]. The bridge is *reduced* to a single common-generator premise for the up/down split — a two-assumptions-to-one move, not an elimination. Open for the full matrix and for non-maintenance-generated couplings. It reduces the *bridge* for the split; it does not touch the *scale* (Gate 4) or the *magnitude* of M_χ (Gate 1's 6/13 input).

---

7. Gate 3 — Generation Dependence: Common Growth plus Halved χ-Increments

The heavy-quark sector splits into two different problems: the common generation growth G_g, and the within-generation up/down split $M_{\chi,g}$. The unconstrained form $m_{a,g\_cur} = m_{q,0\_cur}\, G_{a,g}\, e^{(s_a M_{\chi,g})}\, \Pi_{a,g}^{scheme}$ is too loose if $G_{a,g}$ varies freely by both sector and generation. At leading order the stricter form is imposed,

$$m_{a,g\_cur} = m_{q,0\_cur} \cdot G_g \cdot e^{(s_a M_{\chi,g})} \cdot \Pi_{a,g}^{scheme},$$

so G_g carries the common growth of the generation and $M_{\chi,g}$ carries the up/down split.

7.1 The χ-profile is constrained — conditional on one identification

The single-scheme audit gives $\rho = \Delta\chi_2/\Delta\chi_1 = 0.503 \pm 0.023$, consistent with ½, with the main uncertainty from light-quark inputs rather than scheme mixing. The half-lazy closure reduction gives the structural target $W = ½(I + \Pi_e)$, odd eigenvalue ½, hence $\Delta\chi_{g+1} = ½\, \Delta\chi_g$ under its listed conditions.

Porting this onto the maintenance profile requires the identification $\chi_g = -2M_{\chi,g}$ (§2.7), and **that identification is a bridge the masses do not grant for the naïve within-generation ratio.** Read directly off PDG values the within-generation up/down log-ratio inverts sign after the first generation (down heavier at $g = 1$, $M_{\chi,1} = +0.387$; up heavier at $g = 2, 3$, $M_{\chi,2} = -1.305$, $M_{\chi,3} = -1.832$ at common μ — equivalently within-generation ratios {0.46, 13.6, 39}) and gives an increment ratio ≈ 0.328, not ½. The miss is quantitative: feeding the halving recursion the observed $M_{\chi,1}$ and first increment predicts $M_{\chi,3} = -2.15$ against the observed −1.83, ~16% off. So the audited χ — for which ρ ≈ ½ holds — is a scheme-consistent construction distinct from the naïve ratio, the half-lazy ½ and the audit 0.503 forming a self-consistent loop disconnected from the mass object, and its map to $-2M_{\chi,g}$ is owed (and is a live falsification item, §12.1, not a deferred computation). **Granting that bridge**, the recursion $\Delta M_{\chi,g+1} = ½\, \Delta M_{\chi,g}$ follows, giving

$M_{\chi,g} = M_{\chi,1} + 2\,\Delta M_{\chi,1}\,(1 - 2^{-(g-1)})$,

with $M_{\chi,1} = \tfrac{1}{2}\ln(13/6)$ fixed by 6/13 and one first increment $\Delta M_{\chi,1}$ to be derived from closure-access dynamics. This reduces the split profile from three free parameters to two (one initial value, one first increment) — *conditional on the χ-identification, and obliged to reproduce the observed sign inversion*, which the monotone form above does only if $\Delta M_{\chi,1}$ is large and of the sign that flips the ordering. The reduction is real but it is **not** the clean transplant the bare formula suggests; the inversion is the thing the successor derivation must produce, not assume away.

7.2 The common growth G_g

The substrate-stiffness hierarchy gives $m_D = \mathcal{D}_{\gamma_D}\, S(D)\, v$, $S = S_H S_L S_P S_I$, so the common growth factor should be

$$G_g = (\mathcal{D}_g/\mathcal{D}_1) \cdot [\, S_{common}(g) / S_{common}(1) \,],$$

and in log form

$$\ln G_g = \ln(\mathcal{D}_g/\mathcal{D}_1) + \Delta_g \ln S_H + \Delta_g \ln S_L + \Delta_g \ln S_P + \Delta_g \ln S_I.$$

The stiffness paper states quark/lepton separation is structurally forced by additional quark contributions to S_P and S_I, but exact masses still require computing the PFD invariants.

7.3 Leading-order heavy-quark formula

$$m_{a,g,cur} = m_{q,0,cur} \cdot G_g \cdot \exp(s_a M_{\chi,g}) \cdot \Pi_{a,g}^{scheme}, \quad s_u = -1,\ s_d = +1.$$

This has predictive content only if G_g is computed from substrate stiffness, $M_{\chi,g}$ is fixed by the maintenance profile and halving law (under the §2.7 bridge), and $\Pi_{a,g}^{scheme}$ is derived, not fitted. Without those computations it remains a scaffold.

Gate 3 grade — [Strengthened, open]. The χ-profile is empirically clean ($\rho = 0.503 \pm 0.023$) and structurally reduced to half-lazy closure conditions *for the audited χ*; porting it to the within-generation $M_{\chi,g}$ is a named bridge the sign-inverting mass pattern constrains. The common growth is assigned to substrate stiffness. Numerical heavy-quark masses remain open.

---

8. First-Generation Result, Audited

Collecting the ingredients with the interface projection,

$m_{q,const} = (20/9)\, m_\pi$, $m_{q,proj} = \alpha_{int}\, \eta_q\, m_{q,const}$, $M_{\chi,1} = \tfrac{1}{2}\ln(13/6)$,

$m\_u = m\_{q,0}\_cur\ e^{(-M\_\chi,_1)} = \alpha\_int\ \eta\_q\ (20/9)\ m\_\pi$, $m\_d = m\_{q,0}\_cur\ e^{(+M\_\chi,_1)} = \alpha\_int\ \eta\_q\ (20/9)\ m\_\pi\ (13/6)$.

With $m\_\pi \approx 139.57$ MeV, $\alpha\_int^{-1} \approx 137.036$, and $\eta\_q = 1$,

$m\_u \approx 2.26$ MeV, $m\_d \approx 4.90$ MeV.

Audit. Against PDG (≈ 2.16, ≈ 4.70 MeV at 2 GeV), both are ~4.5% high *in the same direction* — a shared scale offset traceable to the $\eta\_q = 1$ projection. (The offsets are 4.6/4.3/4.1% across m_u/m_d/m_s: tight but not literally uniform; the ~0.5% spread sits inside the mass-input uncertainty, so a single η_q is adequate without being exact.) The decomposition is clean (absolute scale × interface projection × signed maintenance split), the **ratio** (6/13, consistent with the PDG central value within its ~10–20% uncertainty) is the strong mechanism-backed claim, and the **scale** carries the α_int projection bridge plus the $\eta\_q \approx 0.956$ residual. The result is conditional on (i) the K = 7 constituent scale, (ii) the interface projection and its η_q residual (§5), and (iii) the 6/13 ratio, and is **not** a first-principles derivation unless all three are derived in one substrate chain without recourse to the masses.

---

9. Strange-Quark Extension, Audited

The entropic-confinement paper uses the chiral square law

$(m\_K/m\_\pi)^2 = (m\_s + m\_q)/(m\_u + m\_d)$, $m\_q = (m\_u + m\_d)/2$,

with the K = 7 meson ratio $m\_K/m\_\pi = K/2 = 3.5$, so $(m\_K/m\_\pi)^2 = 12.25$ and

$m\_s^{(0)} = 12.25\ (m\_u + m\_d) - m\_q$.

With $m\_u \approx 2.26$ and $m\_d \approx 4.90$ MeV, $m\_q \approx 3.58$, so $m\_s^{(0)} \approx 84.2$ MeV. The active-mode fraction is $f\_active = (2K - 1)/(2K + 1) = 13/15$, and the correction *applied* is its reciprocal,

$f\_active^{-1} = (2K + 1)/(2K - 1) = 15/13 \approx 1.154$,

giving $m\_s = f\_active^{-1}\ m\_s^{(0)} \approx 97.2$ MeV. (The fraction is 13/15; the correction applied is its reciprocal 15/13, which raises the estimate — stated this way to forestall a flipped-definition reading.)

Audit. Against PDG (≈ 93.4 MeV at 2 GeV), ~4% high — again sharing the first-generation scale offset, hence not an independent confirmation. Two new inputs enter: $m\_K/m\_\pi = K/2 = 3.5$ is a ~1% **postdiction** of the measured 3.54, and the reciprocal active-mode factor 15/13 **has a fit signature unless derived independently of m_s** — its operational function is to lift 84 → 97 MeV. The "strange enters through two channels (down-type generation-2 *and* mesonic chiral structure)" framing is **aspirational, not realized**: only the chiral channel is computed here, while the generation-2 channel (G_2, $M\_\chi,2$) is open, so the two are not yet cross-checked against each other — the consistency between them is a promissory note, not a performed test. (That $m_s/\hat{m} = 97.2/3.58 = 27.1$ lands on the FLAG value ~27.3 confirms what is happening: the construction is re-deriving the leading-order chiral-perturbation-theory light-quark ratio, i.e. a postdiction, exactly as graded.) The numerical agreement is a postdiction assembled from inherited relations, not a prediction.

Strange grade — [Conditional, postdictive]. Reproduced from inherited mesonic/chiral machinery; full derivation requires deriving $K/2 = 3.5$ and $15/13$ inside the same closure action, independently of m_s.

---

10. Heavy Quarks: The Constrained Master Formula and Its Remaining Computations

The heavy-quark formula is

$$m_{a,g}^{cur} = m_{q,0}^{cur} \cdot G_g \cdot e^{(s_a M_{\chi,g})} \cdot \Pi_{a,g}^{scheme},$$

with G_g the common generation amplifier, $M_{\chi,g}$ the within-generation signed up/down maintenance amplitude, and $\Pi_{a,g}^{scheme}$ the current-mass projection correction.

The χ-audit and half-lazy theorem reduce the freedom in $M_{\chi,g}$ — conditionally. Under the half-lazy conditions and the audited-$\chi \leftrightarrow -2M_{\chi,g}$ identification (§2.7), $\Delta M_{\chi,g+1} = ½\, \Delta M_{\chi,g}$ constrains the split sequence to one initial value plus one first increment rather than one independent parameter per generation. The constraint is real but inherits the §2.7 caveat: the audited χ for which $\rho \approx ½$ holds is not the naïve within-generation ratio (whose increments run ≈ 0.33 and whose sign inverts after $g = 1$), so the reduction holds for the *audited* object and is ported to the maintenance profile across a named bridge, not an identity.

The common growth G_g must come from the substrate-stiffness factors,

$$G_g = (\mathcal{D}_g/\mathcal{D}_1) \cdot [\, S_H(g)\, S_L(g)\, S_P(g)\, S_I(g)\,] / [\, S_H(1)\, S_L(1)\, S_P(1)\, S_I(1)\,],$$

with the PFD stiffness paper giving the architecture but not the numerical evaluation. **The scheme factors $\Pi_{a,g}^{scheme}$** are unavoidable in *absolute* current masses because those are MS̄-at-2-GeV conventions, not intrinsic rest masses — but per §2.9 they are the wrong thing to derive. Since Π^{scheme} is (nearly) flavour-universal at fixed scale, the right target is to **show it cancels in same-scale ratios**, not to compute it.

The RGI structure of §2.9 then hands a *clean orthogonalisation* — not the mixed ratio list a naïve reading would pick (m_c/m_u mixes G_2/G_1 with an M_χ difference and isolates nothing). Two families separate the two unknowns exactly:

- **Within-generation ratios** $m_{up}(g)/m_{down}(g) = e^{(-2M_{\chi,g})}$ isolate **$M_{\chi,g}$ completely** (G_g and Π cancel). The sequence to be predicted is {0.46, 13.6, 39} — equivalently $M_{\chi,g}$ = {+0.39, −1.31, −1.83}, all evaluated at a common μ (the g = 3 value uses $m_t/m_b \approx 39$ at common scale, not the pole-mass ≈ 41, per the convention caveat below; the two differ by 0.46 vs 0.49 in $M_{\chi,3}$ and do not move the halving miss).
- **Cross-generation geometric-mean ratios** $\sqrt{(m_c m_s)}/\sqrt{(m_u m_d)} = G_2/G_1$, and likewise G_3/G_2, isolate **G_g completely** (M_χ and Π cancel).

The orthogonalisation is Branch-A-exact, and degrades under Branch B — which deepens the fork rather than sidestepping it. Both separations above presuppose that G is sector-shared: the within-generation ratio is pure $M_{\chi,g}$ only because G_g and Π cancel between the up and down numerator/denominator, and the geometric-mean equality $\sqrt{(m_c m_s)}/\sqrt{(m_u m_d)} = G_2/G_1$ holds only when the M_χ contributions cancel between numerator and denominator and Π is flavour-universal — both true under **Branch A** (shared G_g). Under **Branch B** (sector-specific up-sector growth, $G_{u,g} \neq G_{d,g}$) the geometric mean $\sqrt{(m_{up} m_{down})}$ is sector-symmetric *by construction* and therefore cannot carry an up-only rise, so the within-generation ratio is no longer pure $M_{\chi,g}$ — it inherits the up/down G asymmetry, and the basis stops being orthogonal. The tool is thus built on the branch the data appear to disfavour. This is not a flaw to paper over; it shows the fork reaches all the way down into *which invariants are even well-defined*. If Branch B is correct, the clean three-point $M_{\chi,g}$ test below is itself only approximate, and the right invariants must be rebuilt around a sector-resolved G — which is one more reason deciding the fork, not computing within a branch, is the heavy-sector task.

This is exactly the G-vs-M_χ separation the paper is trying to achieve, falling straight out of the RGI structure, and it sharpens the §12.1 tension into a scheme-independent statement: the halving law's job is *only* to reproduce the within-generation sequence {+0.39, −1.31, −1.83}, with no G_g or Π in the way, and it fails (sign inversion plus 16% on $M_{\chi,3}$). (A convention caveat travels with the top: m_t is conventionally quoted as $m_t(m_t)$, so any ratio involving it is RG-invariant only when both masses are taken at a genuinely common μ — the within-generation and geometric-mean families above must be evaluated at one scale, not at each quark's own.)

So the heavy-sector result is not "six masses derived" but "the $M_{\chi,g}$ sequence and the G_g sequence reduced to two scheme-independent families, with Π^{scheme} cancelling and one experimental scale anchor." At leading order, with shared G_g, the parameter pressure is reduced — but the within-generation family already shows the shared-G/halving branch in tension with the data (§10 fork, §12.1), so the open heavy-sector question is to *decide the fork*, not to compute within it.

Saturation warning, and the fork it is really half of. If the paper allows $G_{u,g} \neq G_{d,g}$ as independent inputs, the formula becomes over-saturated: it can fit masses without explaining them. So the leading-order rule is $G_{u,g} = G_{d,g} = G_g$ unless a later PFD computation *derives* sector-specific deviations. But this is not an isolated caution — it is one side of a single fork that the heavy data force, and the §12.1 "data tension" is the other side of the *same* fork. The reasoning:

Under the shared-amplifier rule, G_g and Π cancel in the within-generation ratio, so $m_{up}(g)/m_{down}(g) = e^{-2M_{\chi,g}}$ *exactly* — the entire within-generation split is forced onto $M_{\chi,g}$. That is benign at $g = 1$, where the split is gentle (6/13). It is not benign at $g = 2, 3$, where the observed ratios balloon ($m_c/m_s \approx 13.6$, $m_t/m_b \approx 39$): forcing those onto $M_{\chi,g}$ alone drives $M_{\chi,2} \approx -1.31$, $M_{\chi,3} \approx -1.83$, with the sign inverted relative to $g = 1$ and a magnitude the halving law has no reason to produce (it predicts $M_{\chi,3} \approx -2.15$, ~16% off).

Now the physics of *why* those ratios balloon. In the Standard Model the third-generation gap is large mostly because the top Yukawa is O(1) — the **up-sector** Yukawas rise far more steeply across generations than the down-sector — not because the bottom is suppressed. Translated into these variables, a steep up-sector rise is **sector-specific growth**, exactly the $G_{u,g} \neq G_{d,g}$ that the shared-amplifier rule forbids. So the two cautions are one fork seen from two sides:

- *Branch A (shared G_g, the discipline that buys predictive content):* the heavy within-generation splits are a perturbation on a common scale, carried entirely by $M_{\chi,g}$ — but then $M_{\chi,g}$ must grow large, invert sign after $g = 1$, and miss the halving law, doing violence to the clean first-generation picture.
- *Branch B (sector-specific growth, top is special):* the up-sector amplifier rises steeply on its own, $G_{u,g} \neq G_{d,g}$, and $M_{\chi,g}$ is *not* a single mechanism scaled across generations — in which case the shared-G rule is the wrong leading order for $g \geq 2$.

The honest statement is therefore stronger than "a bridge is owed": **the leading-order discipline that is correct for $g = 1$ may be inverted for $g \geq 2$.** The top mass arguably makes sector-specific G the *leading* heavy-sector effect rather than the exception "to be derived if needed." The paper adopts the shared amplifier as the leading-order ansatz because it is the one with predictive content, but flags that the heavy data actively favour Branch B, and that deciding the fork — not computing within a fixed branch — is the real heavy-sector question.

---

11. What Is Proved, Inherited, Conditional, Open

Proved / exact within this paper:

1. The signed maintenance split can be represented as $m_{\{u,d\}} = m_{q,0\_cur}\, e^{\mp M_{\chi,1}}$.
2. The first-generation ratio $m_u/m_d = 6/13$ is equivalent to $M_{\chi,1} = \tfrac{1}{2}\ln(13/6)$.
3. If both access odds and overlap odds are generated by the same M_χ, the proportional-readout bridge is **reduced to a single common-generator premise** for the signed split: $\ln(a_u/a_d) = \ln(I_u/I_d) = -2M_\chi$, the identity being near-tautological once both sides are posited as $e^{\mp M_\chi}$, with the content living in the common-generator premise. **This two-assumptions-to-one reduction is the paper's cleanest internal result.**

Inherited:

1. The BCB fold-overlap form $\kappa_f = \kappa_0 I_f$.

2. The entropic-confinement light-quark scale $m_{q,const} = (20/9)\, m_\pi$ and the partial-closure interpretation.
3. The $K = 7$ light-quark ratio 6/13.
4. The strange-sector meson/chiral extension.
5. The χ-audit result $\rho = 0.503 \pm 0.023$, consistent with ½ (for the audited χ).
6. The half-lazy closure reduction $W = ½(I + \Pi_e)$, conditional on its structural hypotheses.
7. The substrate-stiffness architecture $S(D) = S_H S_L S_P S_I$.
8. The structural derivation of α_{int} as an interface coupling — impedance form (exact but a re-expression), $K = 7 / N_{loop} = 14$ closure counting, and second-order 1/6 correction ($\alpha_{int}^{-1} \approx 137.034$) — at the companions' grade, i.e. assuming K, N, and 1/6 are fixed independently of 137.036.

Conditional (true given a named, unproved premise):

1. The use of α_{int} as the current-mass visibility factor, $m_{q,cur} = \alpha_{int}\, \eta_q\, m_{q,conf}$ — and the claim that η_q is a small derived projection/scheme correction rather than a free fit ($\eta_q \approx 0.956$ is currently the offset by definition).
2. The identification of M_χ with the charge-sector-odd part of admissibility maintenance, $M_\chi = ½(\ln \Theta_{adm,d} - \ln \Theta_{adm,u})$.
3. The physical identification of closure orientation ABC versus CBA with the up/down split (existence-grade mechanism only).
4. The leading-order shared generation amplifier $G_{u,g} = G_{d,g} = G_g$.
5. The audited-$\chi \leftrightarrow -2M_{\chi,g}$ identification that ports the halving law onto the within-generation maintenance profile — constrained by the observed sign inversion and the naïve-ratio increment ≈ 0.33.

Open (named targets for successor papers):

1. Derive η_q from partial-closure geometry and current-mass scheme matching; settle which scale of α_{int} enters the projection.
2. Derive $M_\chi(q)$ from closure geometry, proving up-type constructive coupling and down-type screening *without using the observed quark masses*.
3. Verify the half-lazy closure conditions for the physical quark χ gate, and supply the audited-$\chi \leftrightarrow -2M_{\chi,g}$ map that reproduces the up/down sign inversion.
4. Compute G_g from S_H, S_L, S_P, S_I for quark PFDs.
5. Derive $\Pi_{a,g}^{scheme}$, especially for heavy quarks.
6. Extend the signed-split bridge to the full Yukawa matrices and CKM/PMNS mixing.

---

12. Falsification Conditions

This five-gate attempt fails, or downgrades sharply, if any of the following occurs. These remain predominantly **internal-consistency forks** — failures a successor finds by computation within the programme — with the present numerical agreements being **postdictions** of already-measured quantities.

1. **Half-lazy / data tension — one side of the shared-G fork (promoted; this is data pushing against, not merely a bridge awaiting).** The clean ½ structure and the masses point in different directions. In the orthogonal basis (§10), the within-generation ratios isolate $M_{\chi,g}$ with no G_g or Π in the way, so the halving law's sole job is to reproduce the sequence $M_{\chi,g} = \{+0.387, -1.305, -1.832\}$ (within-generation ratios $\{0.46, 13.6, 39\}$, at common μ); fed the observed $M_{\chi,1}$ and first increment it predicts $M_{\chi,3} = -2.15$ (~16% off), the increment ratio is 0.328 not ½, and M_χ inverts sign after $g = 1$. This is not separable from the §10 saturation warning: under the shared-amplifier discipline the entire ballooning of the heavy within-generation ratios ($m_c/m_s \approx 13.6$, $m_t/m_b \approx 39$) is forced onto $M_{\chi,g}$, which is exactly what drives the sign inversion and the halving miss — whereas the Standard-Model reason those ratios balloon is a steep *up-sector* rise (the top Yukawa is $O(1)$), i.e. sector-specific growth $G_{u,g} \neq G_{d,g}$. So the fork is: either shared G_g (predictive, but $M_{\chi,g}$ must violate the first-generation picture and miss halving) or sector-specific G (top is special, $M_{\chi,g}$ is not one scaled mechanism, and the shared-G leading order is wrong for $g \geq 2$). The half-lazy ½ and the audit 0.503 form a self-consistent loop disconnected from the mass object; deciding this fork — not computing within a fixed branch — is the real heavy-sector task.
2. **Interface projection failure.** If current quark mass is not the interface-visible component of confined partial-closure maintenance energy, then $m_{q,cur} = \alpha_{int}\, \eta_q\, m_{q,conf}$ is not the right projection law.
3. **Large or arbitrary η_q (and the branch-assignment choice).** If η_q cannot be derived and must be tuned separately by flavour, the light-quark scale is fitted, not predicted. *(As of this paper η_q is exactly the ~5% offset.)* Compounding this, the placement of the bare α_{int} on the up branch — a ~1.47× scale choice (§3) — is data-fixed, not derived; if it cannot be derived, the absolute light scale carries a 1.47× *and* a 5% piece of unexplained fitting.
4. **Maintenance-functional failure.** If no closure-geometric M_χ exists, or it does not change sign between up- and down-type partial closures, the signed-maintenance interpretation fails.
5. **Closure-orientation mismatch.** If forward/reverse closure ABC/CBA is not the physical up/down bit, fold orientation remains a mathematical mechanism but not the quark mechanism.
6. **Stiffness-recursion failure.** If G_g cannot be computed from S_H, S_L, S_P, S_I, or if the required values demand free sector-specific $G_{a,g}$, the heavy sector remains non-predictive.
7. **Scheme projection failure.** If Π^{scheme} does *not* cancel in same-scale ratios (§2.9) — i.e. if it is genuinely flavour-dependent rather than a near-universal convention factor — then the RGI-ratio reframing fails and the heavy sector reverts to carrying an underived convention factor per flavour.
8. **Full-matrix failure.** If the signed-split bridge cannot be extended beyond one up/down contrast, the result remains a ratio mechanism, not a full Yukawa theory.
9. **K = 7 over-flexibility.** If the cluster of K = 7 relations — 2^7 in α_{int}, $6/13 = (K-1)/(2K-1)$, $K/2 = 3.5$, $13/15 = (2K-1)/(2K+1)$ — does not survive a K = 6 / K = 8 sensitivity test as *independent* consequences, then K = 7 is a flexible label fitting several numbers at once rather than a single structural input (see the calibration note in §0).

The single move that would convert this from an internally consistent scheme into a prediction is deriving the projection (η_q and the α_{int}-as-projection step) independently of the masses — the quark-sector analogue of the lepton "is $\kappa = 8/3$ independent of m_μ/m_e?" question. This paper does not perform it.

---

13. Conclusion

This paper updates the five-gate route from quark contrast to quark masses.

The strongest light-sector result remains the signed first-generation ratio,

$m_u/m_d = 6/13$, $M_{\chi,1} = ½ \ln(13/6)$,

and the cleanest internal result is the reduction of the proportional bridge for the signed split (§6) to a single common-generator premise: two independent assumptions converted into one. The first-generation masses are

$m_u = \alpha_{int}\, \eta_q\, (20/9)\, m_\pi$, $m_d = \alpha_{int}\, \eta_q\, (20/9)\, m_\pi\, (13/6)$,

giving, at $\eta_q = 1$, $m_u \approx 2.26$ MeV and $m_d \approx 4.90$ MeV, with the strange estimate $m_s \approx 97$ MeV through the inherited $K = 7$ mesonic/chiral extension and the reciprocal active-mode correction.

The status is sharp in two directions at once. The α factor is not merely borrowed — the interface-coupling programme gives it a structural derivation ($\alpha_{int}^{-1} \approx 137.034$, ~15 ppm), conditional on its three integers being fixed independently of 137.036. But the same move relocates, rather than removes, the scale weakness: what remains is the quark-current projection theorem $m_{q,cur} = \alpha_{int}\, \eta_q\, m_{q,conf}$, in which $\eta_q \approx 0.956$ is at present the ~5% offset by definition, not a derived correction.

The heavy-quark formula admits a shared-amplifier leading order,

$m_{a,g\_cur} = m_{q,0\_cur}\, G_g\, e^{(s_a M_{\chi,g})}\, \Pi_{a,g}^{scheme}$,

under which the split amplitudes would obey $\Delta M_{\chi,g+1} = ½\, \Delta M_{\chi,g}$ *if* the physical χ gate satisfies the half-lazy conditions *and* the audited χ maps to $-2M_{\chi,g}$. The RGI orthogonalisation (§10) makes the test scheme-independent — within-generation ratios isolate $M_{\chi,g}$, geometric-mean ratios isolate G_g — and in that clean basis the shared-amplifier/halving branch is in ~50%-level tension with the data: $M_{\chi,g}$ must run {+0.39, −1.31, −1.83}, inverting sign after $g = 1$ and missing the halving prediction by ~16% at $g = 3$. The reason is structural, not accidental — the Standard-Model heavy splits balloon because of a steep *up-sector* rise (the top Yukawa is O(1)), i.e. sector-specific growth $G_{u,g} \neq G_{d,g}$, which the shared-amplifier discipline forbids. So the central heavy-sector finding is a **fork, not a constraint**: either shared G_g (predictive but data-resisted) or sector-specific G (top is special, the shared-G leading order inverted for $g \geq 2$).

Deciding that fork is the heavy-sector task; the common growth G_g must come from substrate stiffness $S = S_H S_L S_P S_I$ either way.

So the quark masses are not derived here. But the problem is far more structured than six arbitrary Yukawa numbers: it is reduced to three named substrate computations —

η_q, $M_{\chi,g}$, G_g —

the current-mass projection, the signed maintenance functional and its (conditional) halving profile, and the common substrate-stiffness generation amplifier. The programme has a conditional light-quark mass scaffold, a structurally strengthened interface projection, a credible signed-maintenance home for the up/down split, and a heavy-sector reduced — via the RGI orthogonalisation — to a clean fork between shared and sector-specific generation growth, with the data favouring the sector-specific branch. If those computations close — and if the α_{int}-as-projection and audited-χ-as-M_χ bridges are discharged without reading either off the masses — the quark masses become outputs. If they do not, the failure point is exact, named, and located at a specific bridge rather than diffused across the whole construction.

References

1. Keith Taylor, *Realised Readout and Charge-Sector Maintenance: a Two-Part Account of Yukawa Structure in VERSF.* (Part I: the access-visible readout channel and the B_q gate; Part II: the maintenance-conjugacy theorem $B_q = g_\chi^T H^{-1} g_\chi > 0$.)
2. Keith Taylor, *The Reflection-Scope Verdict: Do the Carrier, Generation, and Realization Theorems Force S R S = R?*
3. Keith Taylor, *The BCB Fold Lagrangian: Yukawa Couplings as Dimensionless Fold-Overlap Integrals.* (§2.1: $\kappa_f = \kappa_0 I_f$ and the quark fold representations.)
4. Keith Taylor, *The Substrate-Stiffness Hierarchy of Fermion Mass: $S = S_H S_L S_P S_I$.* (§2.2, §7: the generation amplifier and the S_P/S_I maintenance sector.)
5. Keith Taylor, *Entropic Confinement and Effective Quark Mass: Level-2 Partial Closures, the Constituent Scale (20/9)m_π, and the α-Projection.* (§2.3, §3, §5, §9: constituent scale, current-mass projection, 6/13 ratio, strange estimate.)
6. Keith Taylor, *The Yukawa Hierarchy Theorem: Hypercharge Completion, Species–Family Identification, and Mass Ordering from Refinement Depth.* (Charge-blindness $[\hat{Y}, \chi] = 0$.)
7. Keith Taylor, *The Structural Origin of Yukawa Operators in VERSF.* (Closure-norm mediation; the $\kappa_{clos} \approx \ln 14$ gain and its open independence question.)
8. Keith Taylor, *Exact Projected Closure Coordinates and the Dual-Basis Projection Theorem.* (The M^{-1} coordinate correction.)
9. Keith Taylor, *The Record-Current Ontology: Persistent Commitment-Loop Transport as the Substrate Carrier of J^μ.*
10. Keith Taylor, *The Fine-Structure Constant as Electromagnetic Impedance Mismatch: A First-Principles Derivation.* (Source of $\alpha = Z_0/(2R_K)$ and the interface-transfer interpretation; §2.5.)

11. Keith Taylor, *Why the Void Presents a 2D Interface: A First-Principles Derivation.* (Source of the gauge-invariance argument for 2D electromagnetic interface structure; §2.5.)
12. Keith Taylor, *Beyond First-Order Interface Coupling: Second-Order Closure Corrections in the VERSF Framework.* (Source of $\alpha_{int}^{-1} = 2^K[(N+1)/N - 1/(6N^2)]$; §2.5.)
13. Keith Taylor, *Interface Matching and Coupling Normalization in Gauge Theories.* (Source of the interface-as-UV-boundary-condition interpretation; §2.5.)
14. Keith Taylor, *Finite Distinguishability and Local Capacity Competition.* (Source of the local per-channel L^2 selection and IPR structure; §2.5.)
15. Keith Taylor, *Completing the Interface Bridge: Phase Resolution, Symmetry Allocation, and Second-Order Selection in VERSF.* (Source of the six-channel uniform allocation and cross-term exclusion; §2.5.)
16. Keith Taylor, *Uniform Binary Admissibility and the Half-Lazy Closure Operator.* (Source of the conditional half-lazy derivation $W = ½(I + \Pi_e)$; §2.7.)
17. Keith Taylor, *The Single-Scheme χ Audit.* (Source of $\rho = 0.503 \pm 0.023$ under consistent $\overline{MS}$ running; §2.7.)
18. Keith Taylor, *Closure Orientation, Fold Access, and a Candidate Origin of the Same-Generation Quark Mass Ratios.* (Source of the ABC/CBA closure-orientation candidate; §2.6.)
19. Keith Taylor, *Quark Mass Ratios from Fold-Selective Accessibility.* (Source of the fold-access target P_{up}, P_{down}, R_g; §2.6.)
20. Keith Taylor, *The Admissibility-Maintenance Factor.* (Source of the signed maintenance factor Θ_{adm} and the constructive/screening interpretation; §2.6.)
21. Keith Taylor, *The Matter Skeleton Theorem.* (Source of the matter-census skeleton and status discipline; §2.8.)
22. Keith Taylor, *The Generation Theorem.* (Source of generation as class-grain refinement and the two-mark/three-level architecture; §2.8.)
23. Keith Taylor, *The Substrate Stiffness Hierarchy in VERSF.* (Source of $S(D) = S_H S_L S_P S_I$; §2.2, §7.)