Quark Masses from One Anchor in VERSF

Ratio Closure, χ-Halving, Fold-Orientation Access, and the Substrate-Stiffness Step

Keith Taylor VERSF Theoretical Physics Programme — Quark Mass Hierarchy Series

Successor to *From Charge-Sector Maintenance to Quark Masses in VERSF*. That paper audited the five gates between a nonzero quark readout and actual quark masses and identified the absolute current-mass projection as the central zero-anchor vulnerability. The present paper does not attempt to close that gate. It accepts one quark mass as an anchor in a declared scheme and asks the sharper ratio question: given one anchor, are the remaining five quark masses fixed by structural ratios already present in the programme?

Summary for the General Reader

The earlier paper asked whether VERSF could travel from a quark mass *split* to actual quark *masses*. It found that the hardest obstacle was the absolute scale — converting confined quark closure energy into the "current mass" that physicists quote. That conversion needs a projection factor built from the fine-structure constant α, and even after the interface-coupling papers strengthened α's status, the quark-current projection itself stayed a bridge rather than a derivation.

This paper takes a cleaner route. It asks a narrower question:

Suppose one quark mass is supplied as an anchor. Can VERSF calculate the other five?

That reframing removes the absolute scale from the problem entirely. With one anchor in hand, the only surviving question is whether the **ratios** among quark masses are derivable.

The paper chooses the down quark as the anchor and writes every quark mass as a multiple of m_d. The result is a compact calculator:

m_u = (6/13)·m_d
m_s = 20·m_d
m_c = (3080/13)·m_d
m_b = 5·e^(16/3)·m_d
m_t = 5·e^(16/3)·(6/13)·(77/3)^(3/2)·m_d

Relative to m_d = 1 this is approximately

(u, d, s, c, b, t) ≈ (0.4615, 1, 20, 236.9, 1035.6, 62154).

That row sits almost on top of the standing VERSF quark-ratio grid, whose structural estimates were roughly (0.46, 1, 20, 235, 1040, 61500).

But the calculator is more than a ratio table. The moment one real number is supplied for the anchor, every entry becomes an actual predicted mass. As an illustration, suppose the down quark is taken as $m_{d\star} = 4.70$ MeV. Then the five remaining masses are no longer adjustable — they are simply read off:

Quark	Predicted	Measured	Difference
d	4.70 MeV	4.70 MeV	anchor
u	2.17 MeV	≈ 2.16 MeV	+0.5%
s	94.0 MeV	≈ 93.4 MeV	+0.6%
c	1113.5 MeV	≈ 1105 MeV	+0.8%
b	4867.5 MeV	≈ 5076 MeV	−4.1%
t	292.1 GeV	≈ 300.8 GeV	−2.9%

So one supplied quark mass fixes the absolute scale, and the VERSF ratios then determine the other five. The light quarks land within about half a percent — strikingly close. The two heavy down-type comparisons are genuine misses, not just imprecision: the bottom comes out ~4% low and the top ~3% low, and both trace to specific structural ingredients the programme builds (not to the inputs it borrows). The 4.70 MeV anchor is illustrative only; a strict comparison needs every measured mass quoted in the same scheme as the anchor.

There is an uncomfortable pattern worth stating plainly. The calculator rests on four dimensionless ingredients:

$m_u/m_d = 6/13$, $m_s/m_d = 20$, $m_b/m_s = (1/4)\cdot e^{(16/3)}$, $e^{(\Delta\chi_1)} = 77/3$.

The first two are **borrowed** — the light split and the strange baseline — and against current precision data they are essentially exact. The last two are the parts VERSF actually **derives** — and they are the two that miss: the charm-driving increment runs ~2.5σ too high, and the bottom step ~4% too low. The things the programme imports fare best; the things it constructs fare worst. That is diagnostic rather than fatal — it tells you the modelling error lives in the structural machinery, not the imports — but it is the honest headline, and earlier drafts had it backwards by leaning on an older, less precise measurement that happened to agree.

That 2/3 is the load-bearing new step, and it is the weaker of the two derived ratios on the evidence. It is motivated by the ratio of four admissible fold-orientation states to six primitive interface channels, but it is not yet a theorem — and the most precise measurements already lean mildly against it, preferring a value near 0.64 rather than 0.667. Deriving it from closure geometry would turn the heavy up/down split from a fit into a structural output; but a successful derivation would have to land on ~0.64, not 2/3, to match the data. As it stands, 2/3 is a tidy parametrisation that runs slightly high.

So the honest claim is bounded: **given one anchor mass in a fixed scheme, the quark spectrum reduces to four dimensionless structural ratios** — two borrowed and exact, two derived and currently in tension with the data. This is not a zero-anchor derivation. It is a strict narrowing: the problem is no longer six independent masses, but one anchor plus a small auditable set of closure ratios. The remaining risk is not concentrated into a single number — it sits in the two ratios the programme derives, and the data have already begun to adjudicate both.

Abstract

Prior realised-readout and charge-sector-maintenance work reduced the quark mass split to the existence of a signed maintenance functional M_χ, and a subsequent five-gate audit identified the absolute current-mass projection as the central obstacle to a zero-anchor quark mass derivation. This paper changes the target: accept one quark mass as an anchor in a fixed scheme and ask whether the remaining five are determined by structural ratios.

Let the anchor be $m_d^\star$, the down-quark current mass in a declared common scheme and scale. The paper proves a one-anchor ratio formula

$$(m_u, m_d, m_s, m_c, m_b, m_t) = m_d^\star \cdot \left(6/13,\ 1,\ 20,\ 3080/13,\ 5\cdot e^{16/3},\ 5\cdot e^{16/3}\cdot(6/13)\cdot(77/3)^{3/2} \right),$$

numerically $m/m_d^\star \approx (0.4615, 1, 20, 236.9, 1035.6, 62154)$.

The formula combines four dimensionless inputs. **(i)** The light-quark maintenance ratio $m_u/m_d = 6/13$, inherited from the conditional quark mass-ratio grid as the strongest light-sector ratio. **(ii)** The strange baseline $m_s/m_d = 20$, inherited as the strange-sector input held at the resolution limit. **(iii)** The down-sector generation step $m_b/m_s = (1/4)\cdot e^{16/3}$, where $e^{16/3}$ is the inherited localization step and $1/4 = 3/12$ is the colour-channel saturation participation factor. **(iv)** The first up/down susceptibility increment

$$e^{\Delta\chi_1} = 77/3,$$

newly proposed here from fold-orientation access:

$$e^{\Delta\chi_1} = K(K-1)/2 + K\cdot w_\chi,\quad K = 7,\quad w_\chi = 2/3.$$

The value $w_\chi = 2/3$ is the paper's decisive conjectural step. It is motivated by identifying the self-return access weight with the ratio of four admissible fold-orientation states to six primitive interface channels: the fold ontology fixes four admissible orientations, while the interface papers derive six equivalent local channels under a uniform 1/6 allocation.

With $\chi_1 = \ln(6/13)$, $\Delta\chi_1 = \ln(77/3)$, and the half-lazy law $\Delta\chi_2 = \tfrac{1}{2}\cdot\Delta\chi_1$, the heavy up/down ratios become

$m_c/m_s = 154/13 \approx 11.846$, $m_t/m_b = (154/13)\cdot\sqrt{77/3} \approx 60.015$.

The half-step relation is supported empirically by the single-scheme χ-audit ($\rho = 0.503 \pm 0.023$, consistent with ½) and structurally by the half-lazy closure theorem (the odd-sector eigenvalue ½, conditional on explicit hypotheses).

The result is conditional. It is a **one-anchor ratio closure theorem** whose remaining risk is not single but twofold, and is no longer merely "named" — it is measured. The charm hinge $r_\chi = 77/3$ predicts $m_c/m_s = 11.846$, which runs ~2.4–2.7σ *above* the precision lattice average near 11.77; and the derived bottom step $r_b = e^{16/3}/4$ predicts $m_b/m_s = 51.78$, ~2.5σ *below* the concordant lattice pair near 53.8–53.9 (and −18σ from the most precise). Notably, these two — the quantities the programme derives — are exactly the two that miss, while the imported $r_s = 20$ and inherited $r_u = 6/13$ are essentially exact. The algebraic collapse of six masses to one anchor plus four ratios is [Proven]; the structural sourcing of two of those ratios is empirically disfavoured at present and is the live work.

Table of Contents

---

0. Predictive-Content Ledger

This is deliberately a one-anchor paper. It does not derive the absolute quark mass scale. It accepts one quark mass in one scheme as an anchor and computes the remaining five by structural ratios.

The absolute current-mass projection problem from the five-gate paper is therefore *bracketed*, not solved. The α-projection is no longer needed to calculate the other masses once the anchor is supplied; it becomes a separate explanation of the anchor rather than a required input for the ratio calculator.

Quantity	Value	Status
Anchor	m_d⋆	External, declared scheme/scale
Light split	$m_u/m_d = 6/13$	**[Inherited]** — exact to ~1σ
Strange baseline	$m_s/m_d = 20$	**[Inherited]** — import, yet empirically *exact* (≈19.8)
Down-sector step	$m_b/m_s = e^{16/3}/4 = 51.78$	**[Conditional]** — **~4% low** vs lattice 53.94 (QM-5)
First χ-increment	$e^{\Delta\chi_1} = 77/3 \rightarrow m_c/m_s = 11.846$	**[Conditional]** — **~2.5σ high** vs lattice 11.77 (QM-1)
Halving	$\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$	**[Conditional]** — empirically robust under common-scale running (ρ ≈ 0.50)
Ratio closure	six masses fixed by the above	**[Proven]** — pure algebra (Theorem 1)

The predictive content turns sharply on the two derived ratios, and the data have already spoken on both. The charm hinge (77/3) and the bottom step ($e^{16/3}/4$) are the two quantities the programme constructs, and they are the two that miss the precision lattice values — while the imported strange baseline and inherited light split are essentially exact. This grading inversion (§9) is the paper's central diagnostic: it locates the modelling error in the structural machinery, not in the flagged imports.

The strongest result of the paper is **not** that $w_\chi = 2/3$ is proven — it is provisionally disfavoured. It is the following economy, which is independent of the contested magnitudes:

Once $e^{\Delta\chi_1}$ and the half-lazy law are accepted, **all six** quark masses are fixed from one anchor in closed form — and a single number, $e^{\Delta\chi_1}$, drives **both** heavy-generation splits, because the halving law forwards it to the third generation.

The halving law itself (the shape of the χ-profile) is the programme's strongest empirical card; the magnitude fed through it is where the tension lives.

---

1. The One-Anchor Question

The **zero-anchor** problem asks: can VERSF derive (m_u, m_d, m_s, m_c, m_b, m_t) with no mass input? That requires the absolute current-mass projection and remains open.

The **one-anchor** problem asks: given one quark mass, can VERSF derive the other five? This removes the absolute scale and leaves only dimensionless ratios.

Let the anchor be $m_d\star$. The star records that the scheme, scale, and convention are fixed. This is not cosmetic. Quark masses are current masses and depend on the renormalization scheme; the χ-audit makes clear that any meaningful profile test must be carried out in one consistent $\overline{MS}$ convention with proper running and threshold matching.

The one-anchor goal is

$$m_f = m_d\star \cdot R_f,$$

with each R_f dimensionless and structurally sourced.

---

2. Definitions and Scheme Discipline

Index the three generations by their up- and down-type members:

$$(u_1, d_1) = (u, d), (u_2, d_2) = (c, s), (u_3, d_3) = (t, b).$$

Define the same-generation susceptibility as the log up/down mass ratio within a generation:

$$\chi_g = \ln(m_{u_g} / m_{d_g}).$$

Then

$$\chi_1 = \ln(m_u/m_d), \chi_2 = \ln(m_c/m_s), \chi_3 = \ln(m_t/m_b),$$

with increments

$$\Delta\chi_1 = \chi_2 - \chi_1, \Delta\chi_2 = \chi_3 - \chi_2.$$

The half-lazy target is $\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$.

The calculator uses four ratios:

$$r_u = m_u/m_d, r_s = m_s/m_d, r_b = m_b/m_s, r_\chi = e^{(\Delta\chi_1)}.$$

From the definitions,

$$m_c/m_s = e^{(\chi_2)} = e^{(\chi_1 + \Delta\chi_1)} = r_u \cdot r_\chi,$$

and, applying the halving law,

$$m_t/m_b = e^{(\chi_3)} = e^{(\chi_1 + \Delta\chi_1 + \frac{1}{2}\Delta\chi_1)} = r_u \cdot r_\chi^{(3/2)}.$$

Hence the whole spectrum, anchored on the down quark, is

$m_u = m_d \cdot r_u$
$m_s = m_d \cdot r_s$
$m_c = m_d \cdot r_s \cdot r_u \cdot r_\chi$
$m_b = m_d \cdot r_s \cdot r_b$
$m_t = m_d \cdot r_s \cdot r_b \cdot r_u \cdot r_\chi^{(3/2)}$

This is the algebraic skeleton. Note the economy already visible here: r_χ appears in **both** m_c and m_t, the latter through the halving exponent 3/2. One increment magnitude governs two generation splits.

---

3. Prior Inputs

3.1 Conditional quark mass-ratio grid

The standing grid organizes the six quarks against the down quark,

$d = 1$, $u \approx 0.46$, $s \approx 20$, $c \approx 235$, $b \approx 1040$, $t \approx 61500$,

and grades each entry: down as anchor, up as conditional prediction, strange as an import at the resolution limit, charm as a match awaiting the χ law, bottom as a channel-saturation prediction, and top as a by-product. The present paper keeps that grading and sharpens the charm/top entries by supplying the missing χ-increment.

3.2 Fold-selective accessibility and closure orientation

The fold-selective accessibility paper fixes the target: $\chi(g) = \ln(m_{up}(g)/m_{down}(g))$ must arise from **fold-selective access** to a persistent carrier, not from growth of the carrier. The carrier is refinement-persistent and dimension-invariant, so growth alone cannot generate χ. The closure-orientation paper proposes the up/down fold source as forward versus reverse closure, U = ABC and D = CBA, proves the relevant operator class realizable in a toy construction, and leaves the physical identification with the actual up/down bit open.

3.3 Single-scheme χ audit

The single-scheme audit removes the worry that the halving profile is a scheme-mixing artifact. Under consistent $\overline{MS}$ running with threshold matching it finds

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

fully consistent with ½. This paper treats ½ as the structural target, not as a fitted number.

3.4 Half-lazy closure operator

The half-lazy paper supplies the conditional mechanism

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

with the odd sector receiving eigenvalue ½. Under its stated hypotheses — genuine two-fold sector, image-even closure, binary one-step history, equal primitive weights, no branch-cost asymmetry, and log-access intertwining — it proves exact halving. The present paper uses $\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$ as a conditional structural law, not merely a data fit.

3.5 Entropic confinement and the strange baseline

The entropic-confinement paper interprets quark masses as level-2 partial closures — effective quantities arising inside committed hadronic configurations, not intrinsic isolated rest masses — and supplies light- and strange-quark estimates. The grid grades $m_s/m_d \approx 20$ as an import or resolution-limit value rather than a clean independent prediction. This paper inherits $r_s = 20$ and flags it as the weakest input *by provenance*, while noting (§9) that it is empirically the most accurate ratio in the calculator — the provenance grade and the data grade point opposite ways.

3.6 Substrate stiffness and the down-sector step

The substrate-stiffness paper gives the mass architecture

$$m_D = (\gamma_D / D)\cdot S(D)\cdot v, \quad S(D) = S_H \cdot S_L \cdot S_P \cdot S_I,$$

identifying closure-Hessian stiffness, localization compression, persistent distinguishability load, and interface transport complexity as the magnitude-controlling factors. It does not compute exact masses; it turns them into a concrete substrate calculation. The inherited down-sector result is

$$m_b/m_s = (1/4)\cdot e^{(16/3)},$$

the localization factor $e^{(16/3)}$ times a saturated coloured participation fraction 3/12 = 1/4. The grid lists this bottom/strange entry as one of its genuine predictions.

3.7 Fold ontology and interface-channel counting

The new self-return weight draws on two inherited counts. First, the particle/fold ontology gives a fold **four admissible orientation states**, from binary commitment polarity and binary boundary orientation. Second, the interface/capacity papers derive **six equivalent primitive interface channels** with uniform allocation $w_i = 1/6$, inverse participation 1/6, and exclusion of cross-channel covariance from the primitive local observable algebra. The conjecture of §6 is that the closure-orientation self-return weight is the ratio

$$w_\chi = \text{(fold orientation states)}/\text{(interface channels)} = 4/6 = 2/3.$$

---

4. The One-Anchor Quark Mass Theorem

Theorem 1 (One-Anchor Ratio Closure). Let $m_d\star$ be a down-quark anchor in a declared current-mass convention. Assume

$$r_u = 6/13,\ r_s = 20,\ r_b = (1/4)\cdot e^{16/3},\ r_\chi = 77/3,$$

and the half-lazy law $\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$. Then the six quark masses are fixed by

$$m_u = m_d\star \cdot r_u$$
$$m_d = m_d\star$$
$$m_s = m_d\star \cdot r_s$$
$$m_c = m_d\star \cdot r_s \cdot r_u \cdot r_\chi$$
$$m_b = m_d\star \cdot r_s \cdot r_b$$
$$m_t = m_d\star \cdot r_s \cdot r_b \cdot r_u \cdot r_\chi^{3/2}$$

For the stated values,

$$(m_u, m_d, m_s, m_c, m_b, m_t) = m_d\star \cdot \left(6/13, 1, 20, 3080/13, 5\cdot e^{16/3}, 5\cdot e^{16/3}\cdot(6/13)\cdot(77/3)^{3/2} \right),$$

numerically $m/m_d\star \approx (0.4615, 1, 20, 236.923, 1035.636, 62154.0)$.

Proof. The light entries are immediate: $m_u = m_d\star\cdot r_u$ and $m_s = m_d\star\cdot r_s$. The second-generation ratio is

$$m_c/m_s = e^{\chi_2} = e^{\chi_1 + \Delta\chi_1} = r_u \cdot r_\chi,$$

so $m_c = m_s\cdot r_u\cdot r_\chi = m_d\star\cdot r_s\cdot r_u\cdot r_\chi$. The bottom mass is $m_b = m_s\cdot r_b = m_d\star\cdot r_s\cdot r_b$. The third-generation ratio is

$$m_t/m_b = e^{\chi_3} = e^{\chi_1 + \Delta\chi_1 + \Delta\chi_2},$$

and substituting $\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$ gives $m_t/m_b = r_u\cdot r_\chi^{3/2}$. Hence $m_t = m_d\star\cdot r_s\cdot r_b\cdot r_u\cdot r_\chi^{3/2}$. Substituting the four ratios produces the displayed forms. ■

All closed forms above were checked symbolically and numerically (§7); each matches to machine precision.

---

5. The Four Dimensionless Ratios

5.1 The light split — [Inherited]

$r\_u = m\_u/m\_d = 6/13$, equivalently $\chi_1 = \ln(6/13)$.

Inherited from the grid as the strongest light-sector ratio.

5.2 The strange baseline — [Inherited by provenance, exact in fact]

$r\_s = m\_s/m\_d = 20$.

Prior work grades strange as an import or meson/chiral resolution-limit value, not a fully independent structural prediction — so on *provenance* it is the weakest input. On *performance* it is the strongest: the lattice value is $m\_s/m\_d \approx 19.8–20.2$, so $r\_s = 20$ is essentially exact. It multiplies four of the six masses (§8), so this accuracy is load-bearing in the right direction. The lesson, made explicit in §9, is that provenance weakness and empirical weakness are not the same thing, and here they point opposite ways.

5.3 The bottom/strange step — [Conditional] (inherited localization × derived saturation; ~4% low)

$r\_b = m\_b/m\_s = (1/4)\cdot e^{(16/3)}$.

Decompose precisely and hold the decomposition. The factor $e^{(16/3)}$ is the **[Inherited]** localization step and is *not* the suspect. The 1/4 is the **[Conditional]**, derived part: the saturated coloured participation fraction 3/12 — one surviving route in each of three colour channels, divided by the closure-cell suppression unit 12. Because the derived factor is the load-bearing one, r_b as a whole is graded [Conditional]. Numerically $e^{(16/3)} \approx 207.127$, so $r\_b \approx 51.782$ and $m\_b/m\_d = 20\cdot r\_b = 5\cdot e^{(16/3)} \approx 1035.64$. This ratio is anchor-independent and r_s-independent, so it is the cleanest single test in the paper — and it misses: it matches only the lowest of three lattice determinations and is $\geq 2.5\sigma$ below the concordant pair near 53.8–53.9 (§8.1). The discrepancy is carried by the derived factor 1/4, not the inherited localization.

5.4 The first χ-increment — [Conditional, new]

$r\_\chi = e^{(\Delta\chi_1)}$.

The fold-orientation paper localized the gap: the empirical data prefer a self-return weight near 0.64, while the tidy $K^2/2$ proposal corresponds to a diagonal self-return weight ½ and is underived. This paper proposes $w\_\chi = 2/3$, giving

$$r\_\chi = K(K-1)/2 + K\cdot w\_\chi = 21 + 7\cdot(2/3) = 77/3 \approx 25.667 \ (K = 7).$$

This is the one new conjectural ratio. As $m\_c/m\_s = r\_u\cdot r\_\chi = 154/13 = 11.846$, it is directly testable — and currently runs ~2.5σ above the precision lattice value (§6.1).

---

6. The Conjectural Hinge: $w_\chi = 2/3$

The χ-increment magnitude is the central remaining open quantity in the one-anchor calculation. Prior work split the χ-profile into two pieces:

- **Shape** — halving of increments, $\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$. This is the programme's strongest empirical card. Under consistent common-scale running the audit finds $\rho = 0.503 \pm 0.023$, and it is conditionally derived by half-lazy closure. The shape stands on its own legs.
- **Magnitude** — the size of the first increment, r_χ. This is the sole live problem, and the value proposed below is in tension with the data.

The two must not be conflated. The halving regularity is robust and is what the measurements actually endorse; the magnitude is where the modelling risk sits. (One caveat on the shape's strength: $\rho \approx \frac{1}{2}$ emerges only after standard RG running to a common scale. Naive same-reference-scale masses give $\rho \approx 0.31$, so the regularity is a fact about the consistently-run spectrum, not about raw reference-scale numbers — robust, but not analysis-free.)

The closure-orientation paper identified the magnitude with a second-depth closure-access count. The honest orbit count gives $K(K+1)/2 = 28$, while the tidy $K^2/2 = 24.5$ requires a self-return diagonal weight ½. The data prefer a self-return weight close to 0.64.

This paper proposes the structural value $w_\chi = 2/3$, motivated by

w_χ = (admissible fold-orientation states)/(primitive interface channels) = 4/6.

The numerator 4 is inherited from the fold ontology (binary commitment polarity × binary boundary orientation). The denominator 6 is inherited from the interface papers (six equivalent primitive local channels under uniform inverse-participation allocation). The access count is then

r_χ = [off-diagonal oriented access: $K(K-1)/2$] + [self-return access: $K\cdot w_\chi$],

and for $K = 7$, $r_\chi = 21 + 14/3 = 77/3$.

6.1 The data already disfavour 2/3

The proposal $w_\chi = 2/3 \approx 0.667$ is not awaiting a test — the test exists, and it points the other way. The most precise lattice determinations of m_c/m_s are clustered near 11.77, while VERSF predicts 154/13 = 11.846:

Source	m_c/m_s	VERSF 11.846 sits at
FLAG 2024 average ($N_f = 2+1+1$)	11.766 ± 0.030	**+2.7σ**
FNAL/MILC/TUMQCD 18 (precision frontier)	11.783 ± 0.025	**+2.5σ**
HPQCD 09A ($N_f = 2+1$, oldest, least precise)	11.85 ± 0.16	−0.02σ

Inverting the access formula, the data prefer $w_\chi \approx 0.643$–0.647 — essentially the grid's 0.637, not 2/3. The single determination that matches 11.846 is HPQCD 09A, but it is the least precise and a 2+1-flavour result; the precision frontier has since settled near 11.77. Claiming confirmation by leaning on HPQCD 09A would be selecting the weakest data point.

The honest statement is therefore the reverse of a "feature": **$w_\chi = 2/3$ is [Conditional] and in ~2.4–2.7σ tension with the current best m_c/m_s.** It survives only if future determinations migrate up toward 11.85, which would require the FLAG average and FNAL/MILC to be revised upward. The ~0.03 offset between 2/3 and the empirically preferred ~0.64 is not a forward-looking prediction the data will adjudicate — the data have adjudicated it, mildly against the structural value.

6.2 Status of $w_\chi = 2/3$

- **Not fitted continuously** — a rational value built from inherited structural counts 4 and 6.
- **Not yet proven** — the mapping from fold-orientation states to self-return access weight is conjectural (QM-1).
- **In ~2.4–2.7σ tension** with the precision m_c/m_s average; the structural value runs slightly high, and the data prefer ~0.64.
- **Falsifiable, and provisionally disfavoured** — confirmation now requires the lattice average to move up, not merely a future measurement to land where predicted.

A deeper objection precedes the value itself. The weight $w_\chi = 4/6$ is a quotient of counts drawn from two unrelated structures — four fold-orientation states (commitment polarity × boundary orientation) over six interface channels (uniform 1/6 allocation) — and nothing in the programme yet connects the self-return access weight to that quotient beyond numerical proximity. Moreover the decomposition r_χ = [off-diagonal $K(K-1)/2$] + [self-return $K \cdot w_\chi$] is itself a modelling choice that *creates* the slot w_χ occupies; once the slot is granted, 2/3 is one rational among a family. So the genuine open problem (QM-1) is not "derive the value 2/3" but the prior question: why is the access count this particular split, with self-return linear in K? Until that decomposition is forced, $w_\chi = 2/3$ is a tidy parametrisation — and the data say it is a slightly mistuned one.

---

7. The Six-Mass Calculator and Numerical Verification

With $r_u = 6/13$, $r_s = 20$, $r_b = (1/4)\cdot e^{(16/3)}$, $r_\chi = 77/3$:

$m_d = m_{d\star}$
$m_u = (6/13)\cdot m_{d\star}$
$m_s = 20\cdot m_{d\star}$
$m_c = 20\cdot(6/13)\cdot(77/3)\cdot m_{d\star} = 20\cdot(154/13)\cdot m_{d\star} = (3080/13)\cdot m_{d\star}$
$m_b = 20\cdot(1/4)\cdot e^{(16/3)}\cdot m_{d\star} = 5\cdot e^{(16/3)}\cdot m_{d\star}$
$m_t = m_b\cdot(6/13)\cdot(77/3)^{(3/2)} = 5\cdot e^{(16/3)}\cdot(6/13)\cdot(77/3)^{(3/2)}\cdot m_{d\star}$

One-anchor mass-prediction grid

Once the single anchor m_d⋆ is supplied, the calculator returns **predicted masses**, not merely ratios:

$m_f^{pred} = R_f \cdot m_{d\star}$.

Quark	VERSF ratio R_f = m_f/m_d⋆	Predicted mass from anchor m_d⋆	Measured mass	Status
d	1	m_d⋆	m_d⋆	Anchor
u	6/13 = 0.4615	(6/13)·m_d⋆	m_u	Inherited — fits (~1σ)
s	20	20·m_d⋆	m_s	Imported — exact
c	3080/13 = 236.9	(3080/13)·m_d⋆	m_c	Derived — ~2.5σ high (§6.1)
b	5·e^(16/3) = 1035.6	5·e^(16/3)·m_d⋆	m_b	Derived — realised miss (§8.1)
t	5·e^(16/3)·(6/13)·(77/3)^(3/2) = 62154	62154·m_d⋆	m_t	Derived by-product — carries c, b misses

(Closed forms numpy-verified, exact to machine precision.)

Numerical example

For illustration, take the anchor as m_d⋆ = 4.70 MeV. The calculator then predicts:

Quark	Predicted mass	Measured / reference mass	Difference	Status
d	4.70 MeV	4.70 MeV	—	Anchor
u	2.17 MeV	≈ 2.16 MeV	+0.5%	Inherited — fits
s	94.0 MeV	≈ 93.4 MeV	+0.6%	Imported — exact
c	1113.5 MeV	≈ 1105 MeV	+0.8%	Derived — ~2.5σ high (§6.1)
b	4867.5 MeV	≈ 5076 MeV	−4.1%	Derived — realised miss (§8.1)
t	292.1 GeV	≈ 300.8 GeV	−2.9%	Derived by-product

The numerical values depend on the chosen anchor and mass convention. The table above uses m_d⋆ = 4.70 MeV only as an illustrative anchor. A rigorous comparison requires all measured masses to be expressed in the same declared current-mass convention as the anchor. The structural claim is that, once that one anchor is fixed, the other five masses are no longer independently adjustable — they are fixed by the ratio grid.

The one-anchor paper is therefore not just a ratio table. It is a **mass calculator**: one supplied quark mass fixes the absolute scale, and the VERSF ratios determine the remaining five masses.

χ-column outputs:

$\chi_1 = \ln(6/13) \approx -0.77319$
$\Delta\chi_1 = \ln(77/3) \approx 3.24519$
$\chi_2 = \chi_1 + \Delta\chi_1 = \ln(154/13) \approx 2.47200$
$\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1 \approx 1.62260$
$\chi_3 = \chi_2 + \Delta\chi_2 \approx 4.09460$

Hence $m\_c/m\_s = e^{(\chi_2)} = 154/13 \approx 11.846$ and $m\_t/m\_b = e^{(\chi_3)} = (154/13)\cdot\sqrt{(77/3)} \approx 60.015$. The charm ratio runs ~2.5σ above the precision lattice value 11.77 (§6.1), and the bottom step a separate ~4% low (§8) — the two structural outputs are the two that miss, a pattern examined in §9.

---

8. Sensitivity and Error Propagation

Because the calculator is multiplicative, a fractional error δr/r in any ratio propagates as a fractional error in every mass that ratio touches. Reading off the skeleton of §2:

Input	Touches	Fractional leverage
r_u = 6/13	u, c, t	δm/m in those three = δr_u/r_u
r_s = 20	**s, c, b, t**	**four of six masses — highest leverage**
r_b = e^(16/3)/4	b, t	δm/m = δr_b/r_b
r_χ = 77/3	c, t	t carries (3/2)·δr_χ/r_χ — amplified

Two structural facts follow. First, **r_s is the dominant *uncertainty*** in the inherited inputs: it feeds four of six masses. But — and this is the correction the data force — r_s is *not* where the error actually lives. Folded against measurement (§9), the imported r_s = 20 is essentially exact, while the two **derived** ratios, r_χ and r_b, are the two that miss. The leverage table tells us where an r_s error *would* propagate; the data tell us the error is elsewhere.

Second, **r_χ error is amplified into the top quark** by the halving exponent: a fractional shift ε in r_χ moves m_c by ε and m_t by (3/2)ε. The top mass is therefore the most sensitive probe of the w_χ hinge.

8.1 The bottom step is a realised miss, not an open prediction

The cleanest single test in the whole comparison is r_b = m_b/m_s, because it is **anchor-independent and r_s-independent** — it can be checked directly against a scale-invariant lattice ratio with no reference to m_d⋆ or the strange baseline. VERSF gives

$r_b = e^{(16/3)}/4 = 51.78$.

There are three lattice determinations, and they do not form a symmetric spread around the prediction:

Determination	m_b/m_s	VERSF 51.78 sits at
FNAL/MILC 18 (Nf = 2+1+1, directly quoted)	53.94 ± 0.12	**−18σ**
Maezawa–Petreczky (Nf = 2+1, product of $m_b/m_c \cdot m_c/m_s$)	53.78 ± 0.79	**−2.5σ**
ETM (Nf = 2+1+1, directly quoted)	51.4 ± 1.4	+0.27σ

The landscape is two concordant determinations near 53.8–53.9 against a single low value near 51.4. VERSF agrees only with the lowest one. Consistency forbids the move of "rescuing" r_b by leaning on ETM: that is the mirror image of leaning on HPQCD 09A to rescue the charm hinge, which §6.1 rightly refused. The same evidentiary standard applied there applies here — **r_b matches only the lowest of three determinations and is in ≥2.5σ tension with the two that concur.** Against the most precise ratio (FNAL/MILC, 0.2%), the colour-saturation factor 3/12 = 1/4 fails outright. By the paper's own falsification ledger (§11.5) this is a **realised QM-5 failure, not an open prediction**, and it is elevated to co-equal status with the hinge in §10.

The symmetry between the two derived ratios is the uncomfortable diagnostic, and it should be stated plainly: **in both ratios the VERSF value lands on the lattice determination furthest from the precision frontier and misses the concordant cluster** — high in charm (matching only HPQCD 09A / Maezawa, ~2.5σ above the precise 2+1+1 value) and low in bottom (matching only ETM, ~2.5σ below the concordant pair). The two quantities the programme derives both sit on the far side of the most precise measurement, in opposite directions.

This co-location of error with construction is the most useful signal in the paper: it points at the load-bearing machinery — the $K(K-1)/2 + K \cdot w_\chi$ access count and the 3/12 saturation factor — as the site of the modelling error, *not* at the imports (r_s, r_u), which are empirically exact. Provenance risk and empirical risk point in opposite directions here, and the data win.

---

9. Numerical Comparison and Status

The standing grid row and the one-anchor formula sit almost on top of each other:

	u	d	s	c	b	t
Grid (structural estimate)	0.46	1	20	235	1040	61500
One-anchor formula	0.4615	1	20	236.9	1035.6	62154

The agreement with the grid is expected: the formula formalizes the grid rather than replacing it. But the grid is not the test — measured lattice ratios are. Folded against current data, the four inputs sort by accuracy as follows:

Ratio	Provenance	VERSF	Measured	Verdict
r_s = m_s/m_d = 20	**Imported** (flagged "weakest")	20	≈ 19.8–20.2	essentially exact
r_u = m_u/m_d = 6/13	Inherited	0.4615	low edge of m_u/m_d spread	~1σ, fine
r_χ = e^(Δχ₁) = 77/3	**Derived** (new hinge)	11.846 (as m_c/m_s)	11.77	**+2.4–2.7σ high**
r_b = m_b/m_s = e^(16/3)/4	**Derived** (colour saturation)	51.78	53.94 (FNAL/MILC)	**−4% low / −18σ**

This is a grading inversion, and it should be stated without softening: **the two ratios VERSF derives are the two that miss; the two it inherits or imports are the two that fit.** In particular the strange baseline, graded throughout the programme as an import held "at the resolution limit," is the single most accurate entry, while the new structural hinge and the colour-saturation step — the parts the programme actually builds — are where the discrepancy lives.

The miss has a consistent shape across both derived ratios: in each, VERSF agrees only with the lattice determinations furthest from the precision frontier and misses the concordant cluster — ~2.5σ *high* in charm (sitting with HPQCD 09A and Maezawa, away from the precise 2+1+1 value 11.77) and ~2.5σ *low* in bottom (sitting with ETM, away from the concordant pair near 53.8–53.9). The two derived quantities land on the far side of the most precise measurement, in opposite directions (§8.1).

The diagnostic value is high. It relocates the modelling error from where provenance suggested it (the flagged imports r_s, r_u) to where the data actually place it (the derived machinery: the K(K−1)/2 + K·w_χ access count and the 3/12 saturation factor). The next derivation effort should therefore target the structural constructions, not the imports.

---

10. Proved, Inherited, Conjectural, Open

[Proven] here

- Given one anchor m_d⋆, the four ratios (r_u, r_s, r_b, r_χ) and the halving law fix all six quark masses algebraically (Theorem 1).
- For the stated values, the ratio row is (0.4615, 1, 20, 236.9, 1035.6, 62154).
- Charm and top are not independent once r_χ and halving are accepted: m_c/m_s = r_u·r_χ and m_t/m_b = r_u·r_χ^(3/2). A single increment magnitude drives both heavy splits.

[Inherited]

- The quark grid and its status labels (u/d conditional, s/d import, b/s channel-saturation, charm/top χ-dependent).
- The fold-selective accessibility framing: carrier growth ruled out; χ from depth-dependent access to fixed fold projectors.
- The closure-orientation candidate U = ABC, D = CBA.
- The single-scheme χ-audit result $\rho = 0.503 \pm 0.023$.
- The half-lazy operator $W = ½·(I + \Pi_e)$, conditional on its hypotheses.
- The substrate-stiffness architecture $S(D) = S_H·S_L·S_P·S_I$.
- Four fold-orientation states; six primitive interface channels.

[Conditional, and empirically disfavoured] here

- The self-return access weight $w_\chi = 4/6 = 2/3$, giving $e^{(\Delta\chi_1)} = K(K-1)/2 + K·(2/3) = 77/3$. This produces $m_c/m_s = 11.846$, ~2.4–2.7σ above the precision lattice value (§6.1).
- The colour-saturation factor 3/12 = 1/4 in r_b, giving $m_b/m_s = 51.78$ — ~4% below the precision-frontier lattice value (§8.1). This is closer to a realised failure than an open conjecture.
- The identification of four fold-orientation states (numerator) and six interface channels (denominator) as the objects counted by the χ self-return channel — unforced beyond numerical proximity, and the proximity is now mildly negative.

[Open] — named problems *(starred = elevated to decisive)*

- ★ **QM-1 (reframed).** Not "derive the value 2/3," but the prior question: why is the access count split as $K(K-1)/2 + K·w_\chi$ at all, with self-return linear in K? Until that decomposition is forced, w_χ is a parametrisation, and the data prefer ~0.64 over 2/3.
- ★ **QM-5 (elevated).** The down-sector step $m_b/m_s = e^{(16/3)}/4 = 51.78$ fails the most precise lattice ratio (53.94) at ~4%. This is the cleanest, anchor-free test in the comparison and is the sharpest current discrepancy. Re-derive the colour-saturation factor or accept the miss. *(Co-decisive with QM-1.)*
- **QM-2.** Verify that the physical up/down bit is closure orientation ABC/CBA.
- **QM-3.** Verify the half-lazy conditions for the physical χ gate. (The shape $\rho \approx ½$ is empirically robust under common-scale running; this asks for its mechanism, not its existence.)
- **QM-4.** Replace $r_s = 20$ with a structural derivation. Note: r_s is empirically *exact*, so this is a provenance upgrade, not an error repair — lower priority than QM-1/QM-5, contrary to a naive leverage reading.
- **QM-6.** Extend the one-anchor calculator to CKM/PMNS mixing and full Yukawa matrices.

10.1 Scope of "rigorous": the inherited-object boundary

One boundary must be named plainly, because it is where outside scrutiny will concentrate. The two derived ratios that the data disfavour, r_χ and r_b, are built from objects this paper does not audit: the four fold-orientation states and six interface channels (feeding w_χ), the localization step $e^{(16/3)}$, and the half-lazy operator. These are correctly graded [Inherited] — but the §6.2 objection to $w_\chi = 4/6$ propagates upward through them. That objection was that 4/6 is a quotient of counts drawn from two unrelated structures; since the numerator and denominator are themselves outputs of separate upstream derivations, the conjecture's weakness is partly **inherited, not original to this paper**.

This does not lower the present paper's grade — it labels these objects honestly and proves only what follows from them. But it fixes the meaning of "rigorous" here: rigorous *within the programme, conditional on the upstream chain*, not rigorous full stop. For the result to stand to a skeptical outsider, the four inherited objects feeding r_χ and r_b are exactly where that scrutiny will land — and, by the grading inversion (§9), they are also the objects the data now indicate are mistuned. The diagnostic and the boundary point at the same place: the next derivation target is the upstream sourcing of the two derived ratios.

---

11. Falsification Conditions

The calculation fails or downgrades under any of the following.

11.1 Wrong self-return weight — already triggered, mildly. The falsifier is not hypothetical: the precision lattice value $m_c/m_s \approx 11.77$ already favours $w_\chi \approx 0.64$ over the structural 2/3, placing VERSF's 11.846 at ~2.4–2.7σ (§6.1). The discrepancy is magnified 1.5× in the top sector. The hinge survives only if the lattice average migrates upward.

11.2 No fold-orientation identification. If ABC/CBA is not the physical up/down bit, the proposed derivation of r_χ loses its substrate source (QM-2).

11.3 Half-lazy failure. If the physical χ gate violates the half-lazy conditions, then $\Delta\chi_2 = \frac{1}{2}\cdot\Delta\chi_1$ is an empirical pattern, not a theorem, and t/b is no longer forced by c/s.

11.4 Strange baseline failure. If $r_s = 20$ cannot be derived or justified as a resolution-limit input, the calculator still runs algebraically but loses predictive standing for the strange branch and — by §8 — for charm, bottom, and top as well.

11.5 Bottom saturation failure — partly realised. The factor 3/12 = 1/4 gives m_b/m_s = 51.78, which matches only the lowest of three lattice determinations (ETM) and sits ≥2.5σ below the concordant pair near 53.8–53.9 (FNAL/MILC, Maezawa–Petreczky), and −18σ from the most precise (§8.1). The bottom and top outputs carry this miss together. This is the sharpest current discrepancy and is tracked as the elevated QM-5.

11.6 Scheme failure. If the anchor is not supplied in the same scheme and scale as the ratios, the numerical mass outputs are meaningless. This is a ratio calculator in one declared current-mass convention, not a license to multiply mixed-scheme quantities.

---

12. Conclusion

Given one anchor $m_d\star$, the six quark masses are

$m_u = (6/13)\cdot m_d\star$, $m_d = m_d\star$, $m_s = 20\cdot m_d\star$,
$m_c = (3080/13)\cdot m_d\star$, $m_b = 5\cdot e^{16/3}\cdot m_d\star$,
$m_t = 5\cdot e^{16/3}\cdot(6/13)\cdot(77/3)^{3/2}\cdot m_d\star$,

numerically (u, d, s, c, b, t) ≈ (0.4615, 1, 20, 236.9, 1035.6, 62154).

The new content is the candidate closure of the first χ-increment,

$e^{\Delta\chi_1} = 77/3 = K(K-1)/2 + K\cdot(2/3)$, $K = 7$,

with $w_\chi = 2/3$ as the decisive conjectural step — motivated by the ratio of four fold-orientation states to six primitive interface channels, both inherited, but not yet derived as the actual self-return weight of the χ-access operator (QM-1).

The achievement is a real narrowing, stated without inflation:

six quark masses ⟶ one anchor + { 6/13, 20, $(1/4)\cdot e^{16/3}$, 77/3 } + halving.

Two of the four ratios are borrowed (6/13, 20) and against precision data are essentially exact; two are derived ($e^{16/3}/4$, 77/3) and are presently in tension — the bottom step ~4% low, the charm increment ~2.5σ high. One structural economy survives that tension intact: the halving law forwards a single increment magnitude to **both** heavy splits, so $e^{\Delta\chi_1}$ does double duty, and the *shape* of the χ-profile (ρ ≈ ½ under common-scale running) is the programme's strongest empirical card.

The corrected risk diagnosis is the paper's most useful output. A naive leverage reading flags the imports (r_s feeds four masses) as the danger; the data say the opposite — the imports are exact and the two **derived** ratios miss. That co-location of error with construction points the next work squarely at the structural machinery: the access-count decomposition behind w_χ (QM-1, reframed) and the colour-saturation factor behind r_b (QM-5, elevated). The honest position is not "a sharp calculator with one falsifiable offset" but "a proven algebraic collapse whose two derived inputs are currently disfavoured by the best lattice ratios, and whose halving regularity stands."

A final scope statement, because an external referee will reach it immediately (§10.1): "rigorous" here means *rigorous within the programme, conditional on the upstream chain* — not rigorous full stop. The objects feeding the two disfavoured ratios are inherited and unaudited here, so the weakness flagged against $w_\chi = 4/6$ is partly inherited rather than original. The grading inversion and this boundary point at one and the same target: the upstream sourcing of r_χ and r_b is where both the data and a skeptical outsider will press, and it is the honest place to leave the paper.

References

- *From Charge-Sector Maintenance to Quark Masses in VERSF* (five-gate audit; absolute current-mass projection).
- Realised-readout channel paper (signed maintenance functional M_χ).
- Fold-selective accessibility paper (χ from access, not carrier growth).
- Closure-orientation paper (U = ABC, D = CBA; reflection-scope verdict).
- Single-scheme χ-audit (ρ = 0.503 ± 0.023 under consistent $\overline{MS}$ running).
- Half-lazy closure paper ($W = ½·(I + \Pi_e)$; odd-sector eigenvalue ½).
- Entropic-confinement paper (level-2 partial closures; strange baseline).
- Substrate-stiffness paper ($m_D = (\gamma_D/D)·S(D)·v$; $S = S_H·S_L·S_P·S_I$).
- Particle/fold ontology (four admissible fold-orientation states).
- Interface/capacity papers (six primitive channels; uniform 1/6 allocation).
- Conditional quark mass-ratio grid (status-graded reference row).
- Closure-cell derivation of the suppression unit 1/12 (colour-channel saturation 3/12).

Lattice inputs used in the empirical comparison (§6.1, §8.1, §9):

- FLAG Review 2024 (Nf = 2+1+1 average $m_c/m_s = 11.766(30)$; $m_s/m_l = 27.23(10)$).
- Bazavov et al. (FNAL/MILC/TUMQCD), *Up-, down-, strange-, charm-, and bottom-quark masses from four-flavor lattice QCD*, arXiv:1802.04248 — $m_c/m_s = 11.783(25)$, $m_b/m_s = 53.94(12)$; $m_d(2\text{ GeV}) = 4.675(56)$ MeV, $m_s(2\text{ GeV}) = 92.47(69)$ MeV.
- ETM 16B (Bussone et al.), *Mass of the b-quark and B-decay constants from N_f = 2+1+1 twisted-mass lattice QCD*, Phys. Rev. D 93, 114505 (2016), arXiv:1603.04306 — $m_b/m_s = 51.4(1.4)$, directly quoted and continuum-extrapolated (the low determination).
- Maezawa–Petreczky (Nf = 2+1), arXiv:1606.08798 — $m_c/m_s = 11.877(91)$, $m_b/m_c = 4.528(57)$, giving $m_b/m_s = 53.78(79)$ as a product (concordant with FNAL/MILC).
- HPQCD 09A (Nf = 2+1) — $m_c/m_s = 11.85(16)$ (oldest, least precise; the sole charm determination consistent with 154/13).