

# The Quark Current-Mass Projection Theorem in VERSF

▲ Programme Milestone — Fermion Mass Scale Series

## Deriving the Down-Quark Current Anchor from the Electroweak Completion Scale, Two-Gate Interface Coupling, Phase-Loop Projection, and Colour–Triality Readout

Keith Taylor VERSF Theoretical Physics Programme — Fermion Mass Scale Series Working manuscript — candidate anchor theorem

---

### General Reader Summary

The earlier quark-mass papers reached an impressive but incomplete result. They showed that once the down-quark mass is used as the anchor, much of the quark mass pattern is organised by a short structural grid. The ratios are not random: the strange quark sits at twenty times the down quark; the up/down split is the compact fraction 6/13; the heavier quarks fall into a pattern involving localization growth, channel saturation, and the  $\chi$ -profile.

But there was a serious weakness. The first number — the down-quark anchor itself — was imported from observation. The existing grid was therefore not a first-principles mass derivation. It was a powerful ratio engine attached to one external ruler.

This paper removes that ruler, conditionally.

The central proposal is that the first down-quark current mass is the smallest field-readable mass produced when the electroweak completion interface is projected into a colour-triplet, charge-triality-bearing current channel. In plain language: the Higgs-like completion interface supplies the mass scale; the fine-structure coupling measures how weakly the current can be read by the observable field; the phase-loop geometry spreads the discrete substrate contribution over a smooth continuum field; and colour–triality supplies the down-sector multiplicity.

The candidate formula is

$$\mathbf{m}_d^{(0)} = 9\alpha^2 v / (8\pi)$$

Using the electroweak scale  $v \approx 246.22$  GeV and  $\alpha^{-1} \approx 137.036$ , this gives

$$m_d^{(0)} \approx 4.695 \text{ MeV}$$

That is within 0.6 % of the down-quark current anchor previously used in the one-anchor grid. The significance is not that a number has been matched after the fact. The significance is that the number now has a proposed structural decomposition:

$$9 \times \alpha^2 \times 1/(8\pi) \times v$$

Each factor has an owner. The factor  $v$  is the electroweak completion scale. The factor  $\alpha^2$  is the two-gate source-current projection. The factor  $1/(8\pi)$  is the discrete-to-continuous projection: one phase-loop average and two binary interface halvings. The factor  $9 = 3 \times 3$  is the colour-triality readout multiplicity: three colour supports times three triality-resolved source residues.

Each factor is also **separately falsifiable by a large discrete miss — typically a factor of 2, 3,  $2\pi$ , or 137, depending on which gate fails.** A one-gate coupling law predicts  $\approx 643$  MeV; a three-gate law predicts  $\approx 34$  keV; dropping either factor of 3 in the multiplicity predicts  $\approx 1.57$  MeV. Only the stated combination lands on the observed anchor. The formula is therefore not a flexible fit; it is a rigid structure that either survives its proof gates or fails visibly.

The sharpest consequence involves the electron. Strip the factor of 9 — as one should for a colourless particle — and the same kernel gives  $\approx 0.522$  MeV, within 2 % of the electron mass. If both particles really do read the same kernel, then the Higgs scale and the fine-structure constant **cancel in the ratio**, and the theory makes a completely parameter-free prediction: **the down quark should weigh exactly 9 times the electron.** The measured ratio — comparing the quark's standard quoted mass with the electron's physical mass, a pairing the paper declares and defends — is 9.14 to 9.20 depending on which lattice average is used: one to two standard deviations from the prediction, and testable to ever-higher precision by lattice QCD with no further theoretical input. This is the paper's headline falsifier: it depends on none of the still-owed derivations, and experiment can move on it first.

This draft does not claim the mass programme is finished. It does something more precise: it converts the down-quark anchor from an imported mass into a **conditional projection theorem with named proof gates**. If the four factors are derived independently, the down anchor is no longer free, and the one-anchor quark grid becomes a zero-quark-anchor grid — still conditional on the derivation of  $v$ ,  $\alpha$ , and the RG matching map.

In one sentence:

**The down-quark current mass is proposed to be the minimal colour-triality current readout of the electroweak completion interface, diluted by two-gate coupling and smooth phase projection.**

That is the candidate bridge from ratio physics to absolute mass physics.

---

# Contents

## Abstract

1. Notation, Conventions, and Firewalls 0'. Predictive-Content Ledger 0". Inherited Infrastructure
2. Purpose and Claim Level
3. The Anchor Problem in the One-Anchor Grid
4. The Four Projection Premises
  - 3.1 Premise P1 — Completion-scale linearity:  $m_d \propto v_{cl}$
  - 3.2 Premise P2 — Two-gate coupling:  $\Pi_{gate} = \alpha_{int}^2$  (with Premise P2' — charge-blindness)
  - 3.3 Premise P3 — Continuous projection:  $\Pi_{cont} = 1/(8\pi)$
  - 3.4 Premise P4 — Colour-triality multiplicity:  $N_d = 9$
  - 3.5 Input Provenance and Discharge Ledger (3.5.1–3.5.5)
  - 3.6 Ablation Table
  - 3.7 Premise P5 — Anchor-First Readout
5. The Interface-Projected Down-Anchor Theorem
6. Numerical First Audit
  - 5.1 Accuracy statement
  - 5.2 Immediate significance
  - 5.3 Control signals
7. Inserting the Anchor into the Quark Grid
8. The RG Matching Ledger
  - 7.1 Projection layer
  - 7.2 Running layer and the Reading A / Reading B discrimination
    - 7.2.1 The scheme-invariance problem
  - 7.3 Pole layer
  - 7.4 RG theorem owed
9. Relation to the Strange Current-Mass Operator
10. Relation to the Charged-Lepton Absolute-Scale Problem
  - 9.1 The Parameter-Free Ratio Falsifier (Premise P6)
11. What Has Actually Been Derived
12. Open Debts (Debts 1–8)
13. Falsification Conditions (F0–F11)
14. Milestone Statement
15. Conclusion

Appendix A — Numerical Audit Appendix B — Equivalent Yukawa Form Appendix C — The Electron Clue Appendix D — Claim-Status Ledger Appendix E — Referee Objections and Responses Appendix F — Reproducibility Appendix for the RG Probe (F.1–F.9)

---

## Abstract

The existing VERSF quark-mass programme organises the diagonal quark spectrum through a one-anchor grid. With a declared down-quark current-mass anchor  $m_{d\star}$ , the inherited ratio row takes the schematic form

$$(u, d, s, c, b, t) = m_{d\star} \cdot (6/13, 1, 20, 3080/13, 5e^{\{16/3\}}, 62154),$$

with each cell separately graded as prediction, import, conditional trace, or match awaiting the  $\chi$ -law. This is not yet a first-principles mass derivation, because  $m_{d\star}$  remains an external scale input.

The present paper proposes a candidate **current-mass projection theorem** that replaces  $m_{d\star}$  with an interface-derived anchor. The central claim: the down-quark current anchor is the minimal Standard-Model-facing source-current readout obtained when the electroweak completion interface of scale  $v_{cl}$  is projected through (i) two observable current gates, (ii) the smooth phase-loop dilution of a discrete closure contribution, and (iii) the colour-triality multiplicity of the down-sector carrier. The candidate result is

$$m_{d\wedge proj} = N_d \cdot \Pi_{cont} \cdot \alpha_{int}^2 \cdot v_{cl}, \text{ with } N_d = 9, \Pi_{cont} = 1/(8\pi),$$

so that

$$m_{d\wedge proj} = 9 \alpha_{int}^2 v_{cl} / (8\pi).$$

Under the first readout convention  $\alpha_{int} = \alpha(0)$  and  $v_{cl} = 246.21965$  GeV,

$$m_{d\wedge proj} = 4.6952 \text{ MeV},$$

within +0.54 % of the reference current-quark value  $m_{d\wedge \{M\bar{S}\}}(2 \text{ GeV}) \approx 4.67$  MeV and within 0.1 % of the grid anchor 4.70 MeV.

The result is not presented as an unconditional derivation. Its status is a **conditional projection theorem with four explicit proof gates**:

1. derive  $v_{cl}$  from the closure-interface potential rather than importing the electroweak scale;
2. derive  $\alpha_{int}$ , or the relevant interface coupling, from the VERSF coupling theorem;
3. prove the two-gate  $\alpha^2$  source-current projection;
4. prove the  $1/(8\pi)$  phase-plus-binary projection and the colour-triality multiplicity  $N_d = 9$ .

The paper also installs an **RG firewall**. The projected anchor is a current-mass boundary/readout value, not automatically a pole mass or a high-scale Yukawa. Comparison to observed quark masses requires a renormalisation and matching map

$$m_q^{\{\overline{M\overline{S}}\}}(\mu) = \mathcal{R}_q(\mu \leftarrow \mu_{cl}) \cdot m_q^{\text{proj}},$$

with threshold corrections and scheme conventions specified. That engine is named here but not derived. A preliminary scale audit already discriminates between the two candidate readings of the projection scale  $\mu_{cl}$  and selects the low-current readout (Reading A) as the only reading consistent with the first numerical audit.

Finally, the shared-kernel comparison with the colourless electron ( $N = 1$  versus  $N = 9$ ) yields the paper's **headline falsifier**: since  $v_{cl}$  and  $\alpha_{int}$  cancel in the ratio, the hypothesis predicts  $m_d/m_e = 9$  exactly, parameter-free. Under the declared per-sector readout (P6: pole mass for the colourless electron,  $\overline{M\overline{S}}$  current mass at 2 GeV for the confined quark), the measured ratio is  $9.14 \pm 0.18$  (FLAG  $N_f = 2+1$ ) to  $9.20 \pm 0.10$  (FLAG  $N_f = 2+1+1$ ) —  $1-2\sigma$  from the prediction depending on the lattice average — and the test sharpens with lattice precision alone, independently of every remaining proof gate. Equivalently, the hypothesis plus data predicts the quark readout scale  $\mu_\star \approx 2.1-2.2$  GeV, a number the projection-scale theorem must independently produce. The +2.09 % electron residual and +0.54 % down residual are shown to be a single fact: the measured ratio exceeds 9 by 1.5 %, constraining any future sector-readout correction to  $\kappa_d/\kappa_e = 1.015 \pm 0.020$  ( $N_f = 2+1$ ;  $1.022 \pm 0.011$  for  $2+1+1$ ).

The intended advance is therefore precise: the down-quark anchor is no longer a free empirical input, but is reduced to a candidate substrate-to-field projection formula. The remaining open work is not vague; it is a finite set of projection, coupling, closure-potential, and RG gates.

## 0. Notation, Conventions, and Firewalls

### 0.1 Electroweak completion scale

Let  $v_{cl}$  denote the electroweak completion-interface scale. In Standard Model normalization,

$$\langle \Phi_{cl} \rangle = (0, v_{cl}/\sqrt{2})^T.$$

This paper writes current-mass projection formulae in the convention

$$m_f = \eta_f \cdot v_{cl}.$$

The corresponding Yukawa coefficient in the usual Standard Model convention is

$$y_f = \sqrt{2} \cdot \eta_f,$$

since  $m_f = y_f v_{cl} / \sqrt{2}$ . No factor of  $\sqrt{2}$  is hidden; it is a convention choice, and Appendix B records both forms.

### 0.2 Interface coupling

Let  $\alpha_{\text{int}}$  denote the coupling controlling the field-readable source-current projection of the completion interface into the first-generation charged current channel.

In the first numerical audit this is identified with the low-energy fine-structure constant,

$$\alpha_{\text{int}} \approx \alpha(0).$$

This identification is **not derived here**. The coupling theorem must decide whether the relevant interface coupling is  $\alpha(0)$ , an electroweak-scale  $\alpha$ , a hypercharge/weak mixture, or a VERSF-native coupling whose low-energy limit is  $\alpha$ . The identification is quantitatively load-bearing: because the formula is second order in the coupling, replacing  $\alpha(0)$  by  $\alpha(M_Z) \approx 1/127.95$  would raise the projected anchor by  $\approx 15\%$ , to  $\approx 5.39$  MeV (§7.2). The choice of coupling scale is therefore a falsifiable feature, not a hidden convention.

### 0.3 Projected current mass

The projected current mass is written

$$m_{\text{d}}^{\text{proj}} = N_{\text{d}} \cdot \Pi_{\text{cont}} \cdot \alpha_{\text{int}}^2 \cdot v_{\text{cl}},$$

where:

- $N_{\text{d}}$  — the colour–triality readout multiplicity;
- $\Pi_{\text{cont}}$  — the discrete-to-continuous projection factor;
- $\alpha_{\text{int}}^2$  — the two-gate current-readout penalty;
- $v_{\text{cl}}$  — the electroweak completion scale.

### 0.4 Renormalisation firewall

This paper does not claim that a projected current mass is automatically the same object as a pole mass, a high-scale Yukawa value, or a quoted  $M\bar{S}$  mass at a particular scale.

A complete mass derivation must include the map

$$m_{\text{q}}^{\{M\bar{S}\}}(\mu) = \mathcal{R}_{\text{q}}(\mu \leftarrow \mu_{\text{cl}}) \cdot m_{\text{q}}^{\text{proj}},$$

and, where required,

$$M_{\text{q}}^{\text{pole}} = \mathcal{P}_{\text{q}}(\mu) \cdot m_{\text{q}}^{\{M\bar{S}\}}(\mu).$$

This paper proposes only the projection boundary. It does not derive the running factor  $\mathcal{R}_{\text{q}}$  or the pole-conversion factor  $\mathcal{P}_{\text{q}}$ . It does, however, use standard-QCD running as an external **consistency probe** in §7.2, because that probe already discriminates between the two candidate readings of the projection scale.

### 0.5 Quark-grid firewall

The paper does not rederive the inherited quark ratios

$$m_u/m_d = 6/13, m_s/m_d = 20, m_c/m_d = 3080/13, m_b/m_d = 5e^{16/3},$$

or the inherited top-row entry. It only replaces the external anchor  $m_{d\star}$  with a candidate projection formula.

## 0.6 Higgs-radial firewall

The paper does not derive  $v_{cl}$ , the Higgs radial mass, the scalar self-coupling, or the closure potential. Those remain owed to the closure-interface potential and Higgs-radial theorem.

## 0.7 Coupling firewall

The paper does not derive  $\alpha_{int}$ . It uses  $\alpha(0)$  as the first candidate interface coupling because the projected current is source-charge visible at the readout end. The coupling theorem must eventually derive both the coupling and its scale.

## 0.8 Claim-level firewall

This is not a completed zero-anchor Standard Model mass derivation. The claim is:

**$m_{d\star}$  is replaceable by a structured projection formula, provided the named projection factors are derived.**

# 0'. Predictive-Content Ledger

Object	Value	Status
Down mass anchor $m_{d\star}$	external in prior grid	target of this paper
Completion scale $v_{cl}$	246.21965 GeV in first audit	inherited / not derived here
Interface coupling $\alpha_{int}$	$\alpha(0)$ in first audit	conditional identification
Two-gate coupling factor	$\alpha_{int}^2$	candidate theorem (Premise P2)
Phase-loop projection	$1/(2\pi)$	candidate projection lemma
Chiral bridge halving	1/2	candidate projection lemma

Object	Value	Status
Weak-branch halving	1/2	candidate projection lemma
Total continuous projection $\Pi_{\text{cont}}$	1/(8 $\pi$ )	product of above (Premise P3)
Colour multiplicity	3	inherited from colour triplet
Triality readout multiplicity	3	candidate use of faithful quotient
Down readout multiplicity $N_{\text{d}}$	9	conditional (Premise P4)
Projected anchor	9 $\alpha^2 v / (8\pi)$	central result
Numerical value	4.6952 MeV	first audit; +0.54 % vs $m_{\text{d}}^{\{\text{M}\bar{\text{S}}\}}$ (2 GeV)
Equivalent Yukawa	$y_{\text{d}} = 9\sqrt{2} \cdot \alpha^2 / (8\pi)$ $\approx 2.697 \times 10^{-5}$	first audit
Input provenance	four factors mapped	discharge ledger §3.5; four named theorems owed
Charge-blind gate weight (P2')	premise, explicit	falsifier F2b: charge-weighted gates fail $\times 81$ ; owned by Debt 2
Anchor-first readout (P5)	bookkeeping $\rightarrow$ physics owed	Anchor Ground-State Lemma, Debt 8, F10
Menu scan	144 combos; unique hit	chance $\approx 18$ % within 1 %; §3.6, §5.2
Scheme-invariant output	open	$\hat{m}_{\text{d}}$ check is a named action item; §7.2.1, Debt 6
Projection scale reading	Reading A (low-current readout)	selected by §7.2 audit (+84 % Reading B exclusion, computed); derivation owed
Parameter-free ratio	$m_{\text{d}}/m_{\text{e}} = 9$ exactly	headline falsifier §9.1; under P6: $9.14 \pm 0.18$ (FLAG 2+1) / $9.20 \pm 0.10$ (FLAG 2+1+1), 1–2 $\sigma$ ; predicts $\mu_{\star} \approx 2.1$ –2.2 GeV; immune to $v, \alpha$ debts
Per-sector readout (P6)	pole-e / current-d	premise, §9.1; uniform-M $\bar{\text{S}}$ comparison gives 9.43 (2+1, 2.4 $\sigma$ ) / 9.49 (2+1+1, 4.8 $\sigma$ ), falsifier F11
Full quark table	generated by inherited ratios	conditional synthesis
RG running $\mathcal{R}_{\text{q}}$	—	owed
Pole masses $\mathcal{P}_{\text{q}}$	—	owed
Absolute zero-anchor status	not complete	depends on $v, \alpha, \mathcal{R}$ derivations

## 0". Inherited Infrastructure

This paper inherits five layers.

**First**, the electroweak completion-interface theorem supplies a weak-active, colourless interface of Higgs-doublet form. It is the carrier through which chiral matter records attach to the broken electroweak phase.

**Second**, the broken electroweak phase supplies the scale  $v_{cl}$  as the completion-interface vacuum scale, though its numerical value and potential remain owed.

**Third**, the gauge-closure and faithful-quotient work supplies the colour-triality correlation: colourless fields carry integer electric charge, while colour triplets carry third-shifted charge classes. This is the source of the proposed triality readout multiplicity.

**Fourth**, the one-anchor quark-mass grid supplies the ratio row. This paper does not change that row; it changes the status of the anchor.

**Fifth**, the strange current-mass operator supplies a candidate structural reason for  $m_s/m_d = 20$ . The present paper uses that result only downstream; the derivation of  $m_d$  does not depend on it.

---

## 1. Purpose and Claim Level

The purpose of this paper is to replace the imported down-quark current-mass anchor with a substrate-to-field projection formula.

The previous quark mass-grid has the form

$$m_q = R_q \cdot m_{d\star},$$

where  $R_q$  is the dimensionless ratio attached to each quark. The vulnerability is that  **$m_{d\star}$  is not derived.**

The target replacement is

$$m_{d\star} \rightarrow m_d^{\text{proj}} = 9 \alpha_{int}^2 v_{cl} / (8\pi).$$

The strongest form of the claim is not "the down-quark mass is numerically close." The strongest form is:

**Within the minimal four-factor projection architecture inherited from the VERSF Standard Model stack, the down-quark anchor is fixed to  $m_d^{\text{proj}} = 9\alpha_{int}^2 v_{cl} / (8\pi)$ . The four factors are not adjustable parameters:  $v_{cl}$  is the inherited completion-interface scale;  $\alpha_{int}^2$  is the two-gate source-readout weight;  $1/(8\pi)$  is the normalized phase/chiral/branch**

**projection; and 9 is the product of local colour support and global triality resolution. The remaining work is to prove each discharge theorem, not to choose the factors.**

Claim level: **Conditional / projection-theorem candidate.** It becomes a derivation only when  $v_{cl}$  is derived,  $\alpha_{int}$  is derived, the two-gate coupling law is proven, the continuous projection factor is proven, the colour–triality multiplicity is proven, and the RG map is fixed.

---

## 2. The Anchor Problem in the One-Anchor Grid

The one-anchor grid achieved something important. It separated the quark problem into **absolute scale + dimensionless hierarchy.**

That was the right engineering move: quark masses are scheme-sensitive and scale-sensitive, and trying to derive everything at once would have hidden too much.

Once the ratio grid exists, however, the next question is unavoidable:

### Where does the down anchor come from?

If the answer is "from observation," the grid is a conditional reconstruction, not a first-principles derivation.

The down anchor must therefore be derived from objects already present in the VERSF Standard Model stack. The available ingredients are:

- the completion-interface scale  $v_{cl}$ ;
- the source-charge coupling  $\alpha_{int}$ ;
- the colour triplet;
- the faithful global quotient and charge–triality relation;
- the discrete-to-continuous projection from closure support to smooth field;
- the quark current-mass readout convention.

The present paper proposes that these combine exactly as

$$m_d^{\text{proj}} = (3_{\text{colour}})(3_{\text{triality}}) \cdot (1/2\pi) \cdot (1/2) \cdot (1/2) \cdot \alpha_{int}^2 \cdot v_{cl} = \mathbf{9 \alpha_{int}^2 v_{cl} / (8\pi)}.$$

No further ingredient is available in the inherited stack, and — as §3 shows quantitatively — no subset of these ingredients reproduces the anchor scale. Within the restricted four-factor projection ansatz, the stated combination is the unique structurally available product that lands on the target.

---

### 3. The Four Projection Premises

The candidate formula rests on four premises. Each is named, owned, and separately falsifiable.

$$m_d^{\text{proj}} = N_d \cdot \Pi_{\text{cont}} \cdot \alpha_{\text{int}}^2 \cdot v_{\text{cl}}.$$

#### 3.1 Premise P1 — Completion-scale linearity: $m_d \propto v_{\text{cl}}$

A current mass is a scalar attachment between left and right chiral records. The completion interface is the only object in the inherited electroweak stack that can convert a dimensionless coupling/readout into a fermion mass: it is the unique weak-active, charge-preserving bridge that allows left–right attachment. This is elevated to a named lemma.

**Scale-Carrier Uniqueness Lemma.** In the minimal Standard-Model-facing chiral mass channel, the only available object that can connect a left-handed weak doublet to a right-handed weak singlet while preserving colour, weak charge, and hypercharge is the electroweak completion interface  $\Phi_{\text{cl}}$ . Therefore any first-generation current-mass readout must be proportional to the completion-interface scale:

$$m_d = \eta_d \cdot v_{\text{cl}}.$$

*Proof sketch.* The bridge must be a colour singlet (both chiral records carry matching colour), a weak doublet (to absorb the doublet index of the left record), and carry the hypercharge deficit between the left doublet and right singlet. In the inherited completion stack,  $\Phi_{\text{cl}}$  is the unique object with these quantum numbers. Any alternative scale — the Planck scale, the QCD scale, or an arbitrary substrate mass unit — cannot enter the mass channel directly without introducing a new carrier with the same quantum numbers, or breaking the already-derived electroweak completion architecture. ■ *(conditional on the inherited representation content)*

The consequence for the paper's posture is significant: the linearity  $m_d \propto v_{\text{cl}}$  is **forced by representation closure**, not assumed for convenience.  $v_{\text{cl}}$  enters as an inherited carrier, not a fitted scale. The remaining debt is thereby narrowed: what is owed is the *value* of  $v_{\text{cl}}$  (the Closure-Potential Scale Theorem, §3.5.1), not the justification for the down mass being proportional to it.

**Falsifier (F1).** If the correct completion scale is not  $v_{\text{cl}}$ , or if  $v_{\text{cl}}$  does not enter first-generation current masses linearly, P1 fails and the formula fails with it. Given the lemma, F1 can now fire only by exhibiting a second carrier with the required quantum numbers, or by faulting the inherited representation content itself.

#### 3.2 Premise P2 — Two-gate coupling: $\Pi_{\text{gate}} = \alpha_{\text{int}}^2$

A bare completion interface is not yet a field-readable current mass. It must be projected into an observable source-current channel. The proposed rule:

**A stable current-mass readout requires two source-visible interface gates: an entry gate from completion support into charge-readable current structure, and a return gate from current structure into a persistent mass record.**

**Operator form.** The two gates are defined as maps:

$\mathcal{G}_{in} : \Phi_{cl} \rightarrow J_{source}$  (entry: completion support into charge-readable current structure)  $\mathcal{G}_{out} : J_{source} \rightarrow M_{scalar}$  (return: source-readable current into a persistent scalar mass record)

The scalar mass readout is the composition

$$M_d \sim \mathcal{G}_{out} \circ \mathcal{G}_{in},$$

and if each source-visible gate carries one factor of  $\alpha_{int}$  at the readout level, the composition carries

$$\Pi_{gate} = \alpha_{int}^2.$$

**Minimality and closure.** The two-gate factor is not a perturbative QED vertex count. It is the **minimal source-readout closure count**: one gate makes the completion record source-readable; the second returns the source-readable record to a persistent scalar mass attachment. A record with only one source gate is source-exposed but not scalar-closed — it never re-attaches as a mass. A three-gate record introduces an unnecessary extra source-current traversal and is non-minimal, over-suppressing the first-generation mass. Two is therefore the unique gate count that is both closed and minimal — a claim the Two-Gate Current-Readout Theorem must promote from architecture to proof.

**Premise P2' — Charge-blindness of the gates.** A previously implicit premise is here made explicit, because the formula fails catastrophically without it.  $\alpha$  is defined through unit-charge coupling. If each gate were an electromagnetic readout event weighted by the carrier's charge, the natural per-gate weight for a charge- $Q$  carrier would be  $Q^2\alpha$ , and the down quark ( $Q = -1/3$ ) would carry  $(1/9)^2\alpha^2$  — an 81-fold suppression relative to the electron kernel, fatal to the formula and not rescued by  $N_d = 9$ . The paper's architecture instead asserts:

**each gate counts one source-access event with weight  $\alpha_{int}$ , independent of the carrier's charge  $Q$ ; all charge structure enters the readout only through the triality multiplicity of §3.4.**

This is a coherent division of labour — the triality factor *is* the charge-denominator structure, so charge enters as a global multiplicity rather than a per-gate weight — but it is a premise, not a theorem. The Two-Gate Current-Readout Theorem (Debt 2) must own it explicitly.

**Falsifier (F2b).** If gate weight is shown to carry source charge as  $Q^2$  per gate, the down formula fails by nearly two orders of magnitude, and the entire charge-blind readout architecture fails with it.

**Quantified discrimination.** The gate number is independently fixed by the data: one gate predicts 643.4 MeV, three gates predict 34.3 keV — each excluded by a factor  $\approx 137$  (consolidated in the ablation table, §3.6).

**Falsifier (F2).** If the source-current projection is first order, third order, or governed by a different coupling with materially different value, P2 fails. The theorem must eventually *derive* two gates as necessary and sufficient, not infer them from the match.

### 3.3 Premise P3 — Continuous projection: $\Pi_{\text{cont}} = 1/(8\pi)$

A discrete substrate contribution must be diluted when read as a smooth mass term on the emergent 4-manifold. Rather than presenting this as three intuitive factors, the paper states it as a **normalized projection over a factorized access space**.

**Access space.** The down-current readout has access space

$$\mathcal{A}_{\text{d}} = S^1_{\text{phase}} \otimes \mathbb{C}^2_{\text{chiral}} \otimes \mathbb{C}^2_{\text{weak}},$$

where  $S^1_{\text{phase}}$  is the interface phase loop,  $\mathbb{C}^2_{\text{chiral}}$  carries the two orientations of the left–right bridge, and  $\mathbb{C}^2_{\text{weak}}$  carries the two branches of the weak doublet.

**Normalized projection.** The down-current readout is a uniform phase-density selection over the phase loop and a rank-one selection on each binary factor:

$$\Pi_{\text{cont}} = (1/2\pi) \cdot (\text{rank } P_{\text{chiral}} / \dim \mathbb{C}^2) \cdot (\text{rank } P_{\text{branch}} / \dim \mathbb{C}^2).$$

Since each projector has rank one,

$$\Pi_{\text{cont}} = (1/2\pi) \cdot (1/2) \cdot (1/2) = \mathbf{1/(8\pi)}.$$

The three selections, physically:

1. **Phase-loop normalization,  $1/(2\pi)$ .** The closure contribution carries the uniform phase-loop measure  $d\phi/2\pi$  into the readout as a residual normalization weight. To be precise about the arithmetic: this is a **density reading** — the  $1/(2\pi)$  survives as a multiplicative dilution factor — not an average of an already-normalized quantity, which would leave no residual factor. The Projection-Kernel Theorem must derive the density reading.
2. **Chiral bridge selection,  $P_{\text{chiral}}$  rank one.** A Dirac current mass bridges a left-handed weak doublet to a right-handed weak singlet; the readout selects one orientation of the left–right bridge.
3. **Weak-branch selection,  $P_{\text{branch}}$  rank one.** The down-type channel reads the lower branch through  $\Phi_{\text{cl}}$ ; the up-type channel uses the conjugate interface  $\tilde{\Phi}_{\text{cl}}$ . A single down-current anchor receives one of the two weak-branch supports.

**No-double-count statement.** The chiral and weak-branch halvings act on **different tensor factors** of the access space. The chiral projector selects the left–right bridge orientation on

$\mathbb{C}^2_{\text{chiral}}$ ; the branch projector selects the down component of the weak doublet on  $\mathbb{C}^2_{\text{weak}}$ . Because these projectors act on independent factors, their weights multiply rather than replace one another.

This remains the most exposed factor in the formula and is treated accordingly: the projection-trace form fixes the *structure* of the claim, but the identification of  $\mathcal{A}_d$  as the correct access space, the uniform phase measure, and the rank-one property of each projector all require derivation. The Discrete-to-Continuous Projection-Kernel Theorem (§3.5.3, Debt 1) owns all three.

**Falsifiers (F3, F4).** If the discrete-to-continuous projection carries no  $1/(2\pi)$  phase-loop average, the  $8\pi$  denominator loses its main structural source (a factor- $2\pi \approx 6.3$  miss, predicting  $\approx 29.5$  MeV). If the chiral and branch projections do not supply two independent rank-one selections, the denominator is  $2\pi$  or  $4\pi$  rather than  $8\pi$  (a factor-4 or factor-2 miss, predicting  $\approx 18.8$  or  $\approx 9.39$  MeV). Every failure mode moves the prediction by at least a factor of 2 — well outside the 0.5 % agreement band.

### 3.4 Premise P4 — Colour–trinality multiplicity: $N_d = 9$

The down quark is a colour triplet and a fundamental triality carrier. The faithful global gauge closure ties fractional electric charge to colour triality: a down-type colour carrier is not a single weak attachment but a colour-supported source-current record whose observable electric charge is third-shifted.

The proposed readout multiplicity is

$$N_d = N_{\text{colour}} \cdot N_{\text{trinality}} = 3 \cdot 3 = 9.$$

**The double-counting objection, answered.** The obvious objection is that triality is already part of colour structure, so multiplying  $3 \times 3$  counts the same three twice. The answer is that the two factors count different structures, and the paper states this as a named lemma.

**Colour–Triality Independence Lemma.** The colour factor and the triality factor count different structures. The colour factor 3 counts the **local representation dimension** of the quark colour carrier:  $\dim(\mathbf{3}) = 3$ . The triality factor 3 counts the **global source-grade denominator** imposed by the faithful quotient and charge–trinality congruence,  $Q \equiv -t/3 \pmod{1}$ . The first is local and representation-dimensional; the second is global and source-normalizing. Since they answer different questions — *how many local colour supports?* and *what denominator is required for integer-normalized source readout?* — they are independent and multiply:

$$N_d = 3_{\text{local colour}} \times 3_{\text{global triality}} = 9.$$

Without this lemma, 9 can look like "three colours counted twice." With it, 9 is "local support dimension times global charge-denominator resolution." What remains owed is the composition step itself — proving that the current readout counts the product (the Colour–Triality Current-Readout Lemma, §3.5.4), the independence lemma being its first half.

Note that the  $N_d = 1$  case is not wasted: it is precisely the colourless kernel that reappears as the electron clue in §9 and Appendix C.

**Falsifiers (F5, F6).** If the down readout multiplicity is 3 rather than 9, the predicted scale drops to one third of the target (see §3.6). If colour support does not multiply the current readout at all, the scale drops again by three. This factor must be derived from the faithful quotient and charge-triality theorem; it cannot remain a verbal explanation (Debt 3).

### 3.5 Input Provenance and Discharge Ledger

The projected-anchor formula

$$m_d^{\text{proj}} = N_d \cdot \Pi_{\text{cont}} \cdot \alpha_{\text{int}}^2 \cdot v_{\text{cl}}$$

contains four nontrivial inputs. The formula is not admissible as a first-principles theorem unless each input is either **inherited** from an already-derived VERSF layer or **discharged** by a theorem inside the mass-scale series. This section records the provenance of each factor and names the theorem that discharges it. Together with the falsifiers of §3.1–§3.4, this fixes both directions for every input: where it comes from, and how it dies.

#### 3.5.1 Provenance of $v_{\text{cl}}$ (Premise P1)

The factor  $v_{\text{cl}}$  is inherited from the electroweak completion-interface and broken-phase stack. The completion interface is the unique weak-active, colourless object through which chiral matter records attach to the broken electroweak vacuum; a first-generation current-mass readout must therefore be proportional to the completion-interface scale.

This paper does not derive the numerical value of  $v_{\text{cl}}$ . It inherits  $v_{\text{cl}}$  as the electroweak completion scale and records the closure-potential derivation as an upstream debt. The honest statement of progress is therefore:

**the present formula removes the quark-mass anchor, but it does not yet remove the electroweak-scale anchor.**

**Status:** inherited as scale carrier; numerical derivation owed.

**Discharge theorem (owed):** *Closure-Potential Scale Theorem.* Derive  $v_{\text{cl}}$  from the closure-interface potential, radial stability, and substrate stiffness rather than importing it from electroweak phenomenology.

#### 3.5.2 Provenance of $\alpha_{\text{int}}^2$ (Premise P2)

The coupling  $\alpha_{\text{int}}$  is the source-readable interface coupling. The first audit identifies it with  $\alpha(0)$ , because the projected current mass is read in the low-energy source-current channel — an identification independently supported by the coupling-scale cross-check of §7.2. It remains conditional until the VERSF coupling theorem fixes both the coupling and its scale.

The **power**  $\alpha_{\text{int}}^2$ , however, has a separate structural origin, and its discharge belongs to the present series. A stable current-mass record requires two source-visible gates:

1. an entry gate from completion support into charge-readable current structure;
2. a return gate from current structure into a persistent scalar mass record.

Each gate contributes one factor of  $\alpha_{\text{int}}$ , so the minimal stable current readout is second order:  $\Pi_{\text{gate}} = \alpha_{\text{int}}^2$ . This is not a perturbative QED loop claim; it is a source-readout claim. The mass record is scalar, but the first-generation charged current becomes observable only through source-coupled interface access.

**Status:**  $\alpha_{\text{int}}$  — coupling theorem owed (upstream);  $\alpha_{\text{int}}^2$  — structurally motivated here; full Two-Gate Current-Readout Theorem owed.

**Discharge theorem (owed):** *Two-Gate Current-Readout Theorem*. Prove that a scalar current-mass record requires exactly two source-visible interface gates — not one and not three. (The data already excludes the alternatives by a factor  $\approx 137$  each way; §3.2.)

### 3.5.3 Provenance of $1/(8\pi)$ (Premise P3)

The projection factor  $\Pi_{\text{cont}} = 1/(8\pi)$  is not a numerical decoration. It is the proposed discrete-to-continuous projection from a substrate closure contribution into a smooth scalar current mass on the emergent 4-manifold, and it decomposes into three sub-lemmas:

$$\Pi_{\text{cont}} = \Pi_{\text{phase}} \cdot \Pi_{\text{chiral}} \cdot \Pi_{\text{branch}} = (1/2\pi) \cdot (1/2) \cdot (1/2).$$

1. **Phase-loop normalization,  $\Pi_{\text{phase}} = 1/(2\pi)$ .** A closure contribution with phase access carries the uniform phase-loop measure  $d\phi/2\pi$  as a residual normalization weight (density reading, §3.3) before it appears as a scalar mass readout.
2. **Chiral bridge selection,  $\Pi_{\text{chiral}} = 1/2$ .** A Dirac current mass bridges a left-handed weak doublet to a right-handed weak singlet; only one orientation of the left–right bridge is read in the down-current attachment.
3. **Weak-branch selection,  $\Pi_{\text{branch}} = 1/2$ .** The weak doublet has two branches; the down-type channel reads the lower branch through  $\Phi_{\text{cl}}$ , while the up-type channel uses the conjugate interface  $\tilde{\Phi}_{\text{cl}}$ . A single down-current anchor therefore receives one half of the weak-branch support.

Unlike  $v_{\text{cl}}$  and  $\alpha_{\text{int}}$ , this factor requires no upstream import: it is derivable internally from the substrate-to-field projection architecture. It is therefore the input whose discharge lies most fully within the programme's own control.

**Status:** derivable as an internal projection-kernel theorem; not yet proven.

**Discharge theorem (owed):** *Discrete-to-Continuous Projection-Kernel Theorem*. Prove that scalar current-mass readout from closure support carries exactly one phase-loop average and exactly two binary interface selections.

### 3.5.4 Provenance of $N_d = 9$ (Premise P4)

The down-sector multiplicity  $N_d = 3_{\text{colour}} \times 3_{\text{trinality}} = 9$  rests on two inherited layers plus one owed composition lemma.

**The colour factor 3 is inherited** from the colour census and non-abelian gauge-origin stack: quarks carry threefold colour support, and the  $k = 3$  class fixes the colour group  $SU(3)_C$ . The Non-Abelian Bath-Transport layer derives the colour structure from the  $k = 3$  class and records threefold support as the colour census source.

**The triality factor 3 is inherited** from the faithful global quotient / charge–triality audit. The faithful gauge group requires the congruence

$$2t + 3d + q \equiv 0 \pmod{6},$$

which implies

$$Q \equiv -t/3 \pmod{1}.$$

Colour triplets therefore carry third-shifted electric charge classes, while colour singlets carry integer charge. The triality denominator is not optional bookkeeping; it is the observable face of the  $\mathbb{Z}_6$  quotient.

For a down-type colour-triplet current, the readout must resolve both the 3 colour supports and the 3 triality residues, giving  $N_d = 3 \cdot 3 = 9$ . What is **not** yet proven is the composition rule itself — that the current readout counts the *product* of the two multiplicities rather than either factor alone. The discrimination table of §3.4 shows the product is what the data selects; the lemma must show it is what the structure demands.

**Status:**  $3_{\text{colour}}$  — inherited;  $3_{\text{trinality}}$  — inherited from faithful quotient; product rule — composition lemma owed.

**Discharge theorem (owed):** *Colour–Triality Current-Readout Lemma.* Prove that the first down-type current mass counts the product of colour-support multiplicity and triality denominator, rather than either factor alone.

### 3.5.5 Input-discharge summary

Input	Factor	Provenance	Discharge instrument	Status
Completion scale	$v_{\text{cl}}$	electroweak completion / broken-phase stack	Closure-Potential Scale Theorem	inherited as carrier; numerical value owed
Interface coupling	$\alpha_{\text{int}}$	coupling theorem (upstream)	VERSF Coupling Theorem	owed; $\alpha(0)$ audit-selected (§7.2)

Input	Factor	Provenance	Discharge instrument	Status
Two-gate power	$\alpha_{\text{int}}^2$	this series	Two-Gate Current-Readout Theorem	structurally motivated; theorem owed
Continuous projection	$1/(8\pi)$	this series (internal)	Discrete-to-Continuous Projection-Kernel Theorem	three sub-lemmas named; proof owed
Colour factor	3	$k = 3$ colour census / $SU(3)_C$ closure	inherited	discharged upstream
Triality factor	3	faithful $\mathbb{Z}_6$ quotient / charge-triality audit	inherited	discharged upstream
Composition rule	$3 \times 3 = 9$	this series	Colour-Triality Current-Readout Lemma	owed

The formula  $m_d^{\text{proj}} = 9\alpha_{\text{int}}^2 v_{\text{cl}} / (8\pi)$  is therefore not a free numerical ansatz. It is a compressed statement of four separately owned structures: two factors inherited from proven upstream layers (colour, triality), one factor derivable entirely inside the projection architecture ( $1/(8\pi)$ ), one structural power owned by this series (the two gates), and two upstream scale/coupling debts ( $v_{\text{cl}}$ ,  $\alpha_{\text{int}}$ ) that are shared with the rest of the Standard Model programme rather than specific to the quark sector. The remaining proof burden is finite, named, and explicit.

### 3.6 Ablation Table

The discrimination arguments of §3.2–§3.4 and the scale audits of §7.2 combine into a single ablation table. Each row removes or alters exactly one structural input and records the predicted anchor.

Removed or altered input	Formula effect	Prediction	Failure mode
no colour-triality factor	$N = 1$	0.522 MeV	too low for d (electron kernel, §9)
only one factor of 3	$N = 3$	1.565 MeV	$3\times$ too low
<b>full multiplicity (stated)</b>	<b><math>N = 9</math></b>	<b>4.695 MeV</b>	<b>matches to 0.5 %</b>
extra factor of 3	$N = 27$	14.09 MeV	$3\times$ too high
one-gate coupling	$\propto \alpha$	643.4 MeV	$\approx 137\times$ too high
three-gate coupling	$\propto \alpha^3$	34.3 keV	$\approx 137\times$ too low
no phase-loop factor	denominator 4	29.50 MeV	$2\pi\times$ too high
missing one binary selection	denominator $4\pi$	9.390 MeV	$2\times$ too high
both binary selections missing	denominator $2\pi$	18.78 MeV	$4\times$ too high

Removed or altered input	Formula effect	Prediction	Failure mode
high-scale coupling	$\alpha(M_Z)$	5.386 MeV	15 % too high
interface-scale boundary reading	Reading B + RG	8.59 MeV at 2 GeV ( <i>two-loop</i> , <i>Appendix F.4</i> )	+84 %

The formula is **not continuously tunable**. Each wrong structural choice moves the anchor by a large discrete amount — a factor of 2, 3,  $2\pi$ , or 137 depending on which gate fails — while the smallest available wrong choices (coupling scale, boundary reading) still miss by 15–84 %. The 0.5 % agreement of the stated combination is therefore not reachable by adjusting any single input; it is reachable only by the full stated structure.

**Exhaustive scan.** The one-at-a-time rows above in fact understate the rigidity. Enumerating the paper's full discrete menu —  $N \in \{1, 3, 9, 27\} \times \text{gate count} \in \{1, 2, 3\} \times \text{denominator} \in \{2, 4, 2\pi, 4\pi, 8\pi, 16\pi\} \times \text{coupling} \in \{\alpha(0), \alpha(M_Z)\}$  — gives 144 combinations, of which 91 land in the 0.1–1000 MeV window. Within that menu, the stated combination  $\{9, 2 \text{ gates}, 8\pi, \alpha(0)\}$  is the **unique** combination within 1 % of  $m_d$  — indeed unique within 2.5 % — and **no combination at all** lands within 2.5 % of  $m_u$  or  $m_s$ . Multi-factor compensations fail, not only single-factor ablations. The statistical weight of this scan is assessed honestly in §5.2.

### 3.7 Premise P5 — Anchor-First Readout

Premise P5 — that the projected current mass is read *before* flavour-ratio multiplication — has so far received one line in Theorem 1, while carrying real load. It quietly asserts that the  $\chi$ -profile and localization machinery generating  $m_s/m_d = 20$ ,  $m_c/m_d = 3080/13$ , and the rest of the grid contributes **exactly nothing** to the first-generation down readout: no residual localization factor, no  $O(1)$  overlap coefficient, no structural dressing.

The available argument is this: the ratio grid is written relative to the down quark, so the down cell is the identity cell of the grid — its ratio entry is 1 by construction, and whatever operator content generates the other cells acts as the identity on  $d$ . On this reading, the down quark is the **ground state of the ordering the ratio grid expresses**: the flavour machinery measures departures from the down configuration, and the down configuration itself is the undressed kernel readout.

Stated honestly, that is currently a bookkeeping observation, not a physical theorem. The grid could equally have been written relative to the up quark or the strange quark; what makes  $d$  the *physically* undressed cell — rather than merely the conventionally normalized one — must be shown. The natural candidate argument is that the down quark is the minimal stable colour-triplet configuration of the interface (lowest localization, no channel saturation, no  $\chi$ -excitation), so that all grid operators act trivially on it; but that is an owed lemma, registered as Debt 8.

**Falsifier (F10).** If the down readout carries its own  $O(1)$  structural dressing — a residual localization or overlap coefficient  $\kappa_{P5} \neq 1$  — then the projected kernel value and the grid anchor differ by  $\kappa_{P5}$ , defined by  $m_d^{\text{obs}} = \kappa_{P5} \cdot m_d^{\text{proj}}$ . The measured constraint is  $\kappa_{P5}$

=  $0.995 \pm 0.019$  (FLAG  $N_f = 2+1$ ) or  $1.001 \pm 0.011$  (FLAG  $N_f = 2+1+1$ ) — the prediction sits *high* of the  $N_f = 2+1$  value, so any required dressing is below one. Under the FLAG average the constraint is consistent with exactly 1 at  $0.1\sigma$ , which is the cleanest available statement of the undressed-anchor hypothesis. A dressing theorem producing  $\kappa_{P5}$  outside these bands spoils the agreement.

---

## 4. The Interface-Projected Down-Anchor Theorem

### Theorem 1 — Candidate Interface-Projected Down-Anchor

**Assume:**

- **(P1)** the electroweak completion interface supplies the mass-attachment scale  $v_{cl}$ , entering linearly;
- **(P2)** first-generation down current readout requires exactly two source-current gates, contributing  $\alpha_{int}^2$ ;
- **(P2')** gate weight is charge-blind: each gate counts one source-access event with weight  $\alpha_{int}$  regardless of carrier charge, all charge structure entering only through the multiplicity of P4;
- **(P3)** discrete closure support projects to a smooth scalar current mass with factor  $\Pi_{cont} = 1/(8\pi)$ ;
- **(P4)** the down-sector colour-triality readout multiplicity is  $N_d = 9$ ;
- **(P5)** the projected current mass is read before flavour-ratio multiplication, the down cell being the undressed identity cell of the ratio grid (§3.7).

**Then** the down-quark current anchor is

$$m_d^{\wedge proj} = 9 \alpha_{int}^2 v_{cl} / (8\pi).$$

**Proof.** By P1 the current-mass scale is linear in  $v_{cl}$ . By P2 the source-current readout contributes the multiplicative factor  $\alpha_{int}^2$ . By P3 the discrete-to-continuous projection contributes  $1/(8\pi)$ . By P4 the colour-triality readout contributes the multiplicity 9. By P5 these are the only factors present at the anchor level, and they act independently, so

$$m_d^{\wedge proj} = 9 \cdot \alpha_{int}^2 \cdot (1/8\pi) \cdot v_{cl} = 9 \alpha_{int}^2 v_{cl} / (8\pi). \blacksquare$$

**Reading.** The theorem is algebraically trivial because the real work has been moved into the four projection premises. That is the correct shape for an anchor theorem: it should not hide its load-bearing factors inside a long calculation. It should expose them so that each can be derived or falsified independently. The proof burden is now located entirely in P1–P4, each of which carries a quantified falsifier (§3.1–§3.4), a recorded provenance, and a named discharge theorem (§3.5).

---

## 5. Numerical First Audit

Using the first readout convention

$$\alpha_{\text{int}}^{-1} = 137.035999084, v_{\text{cl}} = 246.21965 \text{ GeV},$$

the projection gives

$$m_{\text{d}}^{\text{proj}} = (9/8\pi) \cdot (1/137.035999084)^2 \cdot 246.21965 \text{ GeV} \quad m_{\text{d}}^{\text{proj}} = \mathbf{0.0046952208 \text{ GeV}} = \mathbf{4.6952208 \text{ MeV}}.$$

The dimensionless coefficient in the convention  $m_{\text{f}} = \eta_{\text{f}} v_{\text{cl}}$  is

$$\eta_{\text{d}} = 9\alpha^2/(8\pi) = 1.9069 \times 10^{-5},$$

and the Standard Model Yukawa convention gives

$$y_{\text{d}} = \sqrt{2} \cdot \eta_{\text{d}} = 9\sqrt{2} \cdot \alpha^2/(8\pi) = 2.6968 \times 10^{-5}.$$

### 5.1 Accuracy statement

The projected value lands within the reference band for the down-quark current mass:

- versus the grid anchor  $m_{\text{d}}^{\star} \approx 4.70 \text{ MeV}$ : agreement to  $\approx 0.1 \%$ ;
- versus the reference current-quark value  $m_{\text{d}}^{\{\text{M}\bar{\text{S}}\}}(2 \text{ GeV}) \approx 4.67 \text{ MeV}$ : deviation  $+0.54 \%$ , comfortably inside the quoted experimental/lattice uncertainty on  $m_{\text{d}}$ .

**Declared comparison convention.** Both comparison values are FLAG 2024 lattice averages of  $m_{\text{d}}^{\{\text{M}\bar{\text{S}}\}}(2 \text{ GeV})$ , differing in sea-quark content: the  $N_{\text{f}} = 2+1$  average gives  $4.67 \pm 0.09 \text{ MeV}$  (agreement  $+0.54 \%$ ), and the  $N_{\text{f}} = 2+1+1$  average, which includes a dynamical charm sea, gives  $4.70 \pm 0.05 \text{ MeV}$  (agreement  $\approx 0.1 \%$ ). The paper declares the  $N_{\text{f}} = 2+1$  value  $4.67 \text{ MeV}$  as its primary comparison throughout, quoting the  $2+1+1$  row alongside where it changes a conclusion. The  $\approx 0.6 \%$  spread between the two averages is a physics difference (sea-quark content), not an evaluation-style difference, and should be regarded as part of the systematic uncertainty on every sub-percent statement in this paper. The deeper convention-ladenness of the  $2 \text{ GeV } \text{M}\bar{\text{S}}$  quote itself is treated in §7.2.1.

### 5.2 Immediate significance

The match is significant because the formula is not a multi-parameter fit. It contains **one scale, one coupling, one phase projection, two binary halvings, and one discrete multiplicity** — and, as §3.6 documents, each discrete choice is separated from its neighbours by at least a factor of 2 and typically a factor of 3 or 137.

The chance-hit probability can now be stated with a number rather than an adjective. The full discrete menu of §3.6 contains 144 combinations, 91 of which land in the 0.1–1000 MeV window — roughly 23 candidates per decade. At that density, the probability that *some* combination lands within 1 % of a single pre-specified target is  $\approx 18\%$ , and within 0.6 % is  $\approx 11\%$ . An order-10 % chance probability is **suggestive, not compelling**. The honest statement is therefore: the numerical match motivates the structure but cannot certify it; the paper's real defence is the discharge-theorem architecture of §3.5 — every factor either inherited from a proven layer or owed to a named theorem — and the match should be read as selecting which theorems to attempt, not as evidence that they will succeed.

What the scan does establish beyond the raw probability is rigidity: within the menu, the stated combination is the unique hit at the 1 % level (and at 2.5 %), and no combination reaches  $m_u$  or  $m_s$  at all. If the discharge theorems land, the formula was never adjustable; if they fail, no neighbouring combination can be retreated to.

The correct status is therefore:

**strong candidate anchor formula, proof-gated** — not a completed first-principles quark-mass derivation.

### 5.3 Control signals

The projection kernel is not isolated to one number. With  $N = 1$ , the same kernel gives  $\alpha^2 v / (8\pi) \approx 0.522$  MeV, near the electron scale (§9, Appendix C). With  $N = 9$ , it lands on the down-quark anchor. With the inherited strange ratio 20 applied to that anchor, it gives  $20 m_d^{\text{proj}} \approx 93.90$  MeV, near the strange current mass at the standard reference scale.

These are not independent proofs — the strange value inherits the down anchor by construction, and the electron relation awaits its own sector readout. But they are useful control signals: the kernel is operating in the correct first-generation charged-fermion scale regime, and it does so at three separate readout points (colourless, coloured, and coloured-with-trace) rather than at the single point it was aimed at.

## 6. Inserting the Anchor into the Quark Grid

Replacing  $m_{d\star}$  by  $m_d^{\text{proj}} = 4.6952208$  MeV gives the raw inherited grid:

	Quark Inherited ratio to d	Raw projected value
u	6/13	2.1670 MeV
d	1	4.6952 MeV
s	20	93.904 MeV
c	3080/13	1.1124 GeV

	Quark Inherited ratio to d	Raw projected value
b	$5e^{\{16/3\}}$	4.8625 GeV
t	inherited top ratio	raw grid value — not a pole mass

The table must be read with the firewalls active. For u, d, s the values sit in the current-mass region where the original anchor grid was designed to operate. For c, b, t the raw numbers must not be compared naïvely to quoted masses without specifying scheme and scale.

The top row especially is a trap. A raw closure-grid top entry is not the observed top pole mass. The earlier top-correction work already showed that direct scale-chaining creates false discrepancies. A real heavy-quark prediction exists only after passage through the RG map of §7.

## 7. The RG Matching Ledger

A first-principles mass derivation needs three layers:

substrate projection  $\rightarrow$  running mass  $\rightarrow$  experimental readout.

This paper supplies only the first layer, but it can already constrain the second.

### 7.1 Projection layer

$$m_q = R_q \cdot m_d^{\text{proj}},$$

where  $R_q$  is the inherited dimensionless quark ratio.

### 7.2 Running layer and the Reading A / Reading B discrimination

The running layer must compute  $m_q^{\{M\bar{S}\}}(\mu) = \mathcal{R}_q(\mu \leftarrow \mu_0) \cdot m_q$ , and the projection scale  $\mu_0$  must be specified. Two readings are possible:

**Reading A — low-current readout.** The projection formula directly outputs the low-current readout value, near the standard light-quark reference scale  $\mu \approx 2$  GeV. Then  $\mathcal{R}_d \approx 1$  by convention for the first audit, and heavy-quark rows still require running.

**Reading B — interface-scale boundary.** The projection formula outputs a boundary condition at  $\mu_{cl} \approx v_{cl}$ . Then all masses, including  $m_d$ , must be run downward using QCD and electroweak RG equations.

Reading B is more orthodox. But a computed scale audit already discriminates between them. Standard QCD running increases light-quark  $M\bar{S}$  masses toward low scales: a two-loop integration with thresholds at  $m_t, m_b, m_c$  and  $\alpha_s(M_Z) = 0.1179$  gives

$$\mathcal{R}_d(2 \text{ GeV} \leftarrow 246 \text{ GeV}) = 1.83, \alpha_s(2 \text{ GeV}) = 0.296$$

(*external-QCD probe, fully reproduced in Appendix F; the RG ledger must eventually own this map internally*). Under Reading B, a boundary value of 4.695 MeV at  $v_{cl}$  therefore runs to 8.59 MeV at 2 GeV — an **+84 % overshoot** against the observed 4.67 MeV. Under Reading A, the same number matches to 0.5 %.

**Provisional conclusion.** The first audit is consistent only with Reading A: the projection outputs a low-current readout, not a  $v_{cl}$ -scale boundary condition. This is recorded as a *selected* reading, not a derived one. The projection kernel must eventually explain *why* its output is anchored at the low-current readout scale — plausibly because the readout is defined by the observable source-current channel, which is a low-energy object — or, alternatively, the coupling identification must shift in a compensating way (see below). Either resolution is a theorem owed, and the tension is stated here so that it cannot be quietly absorbed later.

**Coupling-scale cross-check.** The same audit constrains the coupling identification of §0.2. Replacing  $\alpha(0)$  by  $\alpha(M_Z) \approx 1/127.95$  raises the anchor to  $\approx 5.386$  MeV (+14.7 %). The combination {Reading A,  $\alpha(0)$ } matches; {Reading A,  $\alpha(M_Z)$ } misses by  $\approx 15$  %; {Reading B,  $\alpha(0)$ } misses by +84 %; {Reading B,  $\alpha(M_Z)$ } misses by +111 %. The first audit therefore selects **low-scale readout with low-scale coupling** as the internally consistent convention pair. This coherence is suggestive but must not be overread — see §7.2.1.

### 7.2.1 The scheme-invariance problem

The Reading A comparison contains a deeper exposure that the convention-pair language above does not resolve, and it is stated here so that it cannot be discovered by a hostile reader first.

The target value 4.67 MeV is  $m_d^{\{\overline{MS}\}}(\mu = 2 \text{ GeV})$ : a **scheme-dependent quantity evaluated at a human-chosen reference scale**. The same physical down quark reads

$$m_d^{\{\overline{MS}\}}(1 \text{ GeV}) = 5.85 \text{ MeV (+25 %)}, m_d^{\{\overline{MS}\}}(3 \text{ GeV}) = 4.26 \text{ MeV (-9 %)}$$

(*two-loop, computed; Appendix F.5*). A formula matching the 2 GeV value to 0.5 % is therefore implicitly claiming that the substrate readout knows about the  $\overline{MS}$  scheme at  $\mu = 2 \text{ GeV}$  specifically.

Moreover, the two "low-scale selections" of the convention pair are not symmetric objects.  $\alpha(0)$  has an infrared-fixed, scheme-independent definition — the Thomson limit — and is a genuine physical invariant.  $m_d^{\{\overline{MS}\}}(\mu)$  has no analogue: its running never freezes, and below roughly 1 GeV the perturbative mass ceases to be defined at all. One member of the pair is an IR invariant; the other is a UV-scheme artifact at an arbitrary waypoint. The pair is therefore *coherent as a bookkeeping convention* but not yet *coherent as physics*.

The Reading A selection survives this audit — the +84 % Reading B exclusion is scheme-robust at the level that matters — but the debt it leaves is sharper than "derive  $\mu_{cl}$ ." The projection kernel owes an answer to:

**What is the scheme- and scale-invariant object the kernel outputs, and why does its value numerically coincide with  $m_d^{\overline{\text{MS}}}(2 \text{ GeV})$ ?**

One candidate escape route is that the kernel outputs an RG-invariant mass  $\hat{m}_d$  (the scheme-independent integration constant of the mass RG equation), whose value happens to sit near the 2 GeV  $\overline{\text{MS}}$  quote. Whether it does is a definite, checkable question, and it is registered as a named action item under Debt 6. Until it is answered, the 0.5 % agreement of §5.1 should be read with this proviso attached: the *scale* of the agreement is meaningful (the ablation alternatives miss by factors of 2–137); the final *half-percent* of it is convention-laden.

### 7.3 Pole layer

For heavy quarks, especially top,

$$M_q^{\text{pole}} = \mathcal{P}_q(\mu) \cdot m_q^{\overline{\text{MS}}}(\mu).$$

No top-mass claim is valid until this conversion is performed.

### 7.4 RG theorem owed

This paper therefore creates the following required theorem:

**VERSF Fermion RG Matching Theorem.** Given substrate-projected current masses  $m_q$ , derive or inherit the Standard Model anomalous dimensions, threshold rules, and matching conventions required to compute  $m_q^{\overline{\text{MS}}}(\mu)$  and  $M_q^{\text{pole}}$ , and derive the projection scale  $\mu_{\text{cl}}$  of the readout from the projection kernel itself.

Without this theorem, the anchor formula is promising but not peer-review complete.

---

## 8. Relation to the Strange Current-Mass Operator

The strange current-mass operator derives  $m_s/m_d = 20$  from the twenty oriented balanced support atoms of a six-channel interface. If that result closes, the present paper immediately gives

$$m_s^{\text{proj}} = 20 \cdot m_d^{\text{proj}} = 180 \alpha^2 v / (8\pi) \approx \mathbf{93.90 \text{ MeV}},$$

consistent with the reference  $m_s^{\overline{\text{MS}}}(2 \text{ GeV}) \approx 93.4 \text{ MeV}$  to  $\approx 0.5 \%$  — the same deviation band as the anchor itself, as expected since the ratio is exact by construction.

The two papers play different roles:

- the strange paper derives a **ratio trace**;

- this paper derives a **current-mass anchor**;
- together they give an **absolute strange current mass**.

The hierarchy is:

current anchor + support trace = absolute current mass.

## 9. Relation to the Charged-Lepton Absolute-Scale Problem

The charged-lepton ratio law has its own absolute-scale problem: it describes  $m_e : m_\mu : m_\tau$  but still needs an electron anchor or a derived electron projection.

The present paper suggests a general first-generation pattern:

$$m_{\text{first}} = N_{\text{sector}} \cdot \Pi_{\text{cont}} \cdot \alpha_{\text{int}}^p \cdot v_{\text{cl}},$$

with, for down quarks,  $N_{\text{sector}} = 9$  and  $p = 2$ .

Charged leptons are colourless and carry no colour triality, so a naive analogue has  $N_{\text{sector}} = 1$ . The colourless kernel — already computed as the  $N = 1$  row of the §3.4 discrimination table — is

$$m_0 = \alpha^2 v / (8\pi) \approx 0.5217 \text{ MeV},$$

which lies within **+2.1 %** of the electron mass  $m_e = 0.51100 \text{ MeV}$ .

This observation is handled with the statistical care the exhaustive enumeration of §3.6 demands. That enumeration finds that at the 2 % level, hits in this discrete menu are cheap: a structurally meaningless combination,  $27 \cdot \alpha(M_Z)^3 \cdot v / (2\pi) = 0.5051 \text{ MeV}$ , lands  $-1.2 \%$  from the electron mass — *closer* than the kernel's  $+2.1 \%$ . The electron clue is therefore graded honestly: **consistent, but within the chance-hit band until the lepton readout theorem exists**. It is not "too suggestive to ignore"; it is suggestive enough to record and test.

What survives the downgrade is the structural point, and it is worth stating precisely. The junk combination shares no factor with the down formula — different multiplicity, different gate count, different coupling scale — and therefore predicts nothing. The kernel hypothesis, by contrast, uses the *identical* structure for both particles with only the sector multiplicity changed, and that identity is what generates the parameter-free ratio prediction of §9.1. The clue's evidential weight is low; its *testability* is what distinguishes it from numerology, and §9.1 is the test. The residual  $+2.1 \%$  is four times the down-sector residual and must eventually be explained (lepton readout corrections, coupling-scale refinement, or failure of the shared-kernel picture). The present paper records it as a consistency signal:

$$m_e^{\text{candidate}} \sim 1 \cdot \alpha^2 v / (8\pi), \quad m_d^{\text{candidate}} \sim 9 \cdot \alpha^2 v / (8\pi).$$

If the lepton absolute-scale theorem derives the first relation and the present theorem derives the second, the first-generation charged-fermion scale is unified by a **common projection kernel with sector-dependent multiplicity**. That would be a major programme convergence — and a shared point of failure: if the electron relation collapses under proper lepton readout, the down-anchor interpretation weakens with it (F8).

## 9.1 The Parameter-Free Ratio Falsifier

The shared-kernel hypothesis contains a prediction sharper than either absolute value, and it deserves to be stated as the paper's headline falsifier.

If the down quark and the electron read the same projection kernel with  $N = 9$  and  $N = 1$  respectively, then  $v_{cl}$  and  $\alpha_{int}$  cancel in the ratio, and the hypothesis predicts, with no free inputs of any kind:

**$m_d / m_e = 9$ , exactly.**

**Premise P6 — Per-sector IR-observable readout.** Before the comparison, the readout convention must be stated, because the natural comparison is *not* a same-scheme comparison and pretending otherwise would be a hidden inconsistency. The declared rule: **the kernel outputs the IR-observable mass of each sector**. For the colourless electron, that is the pole mass — IR-complete, scheme-invariant, and consonant with the  $\alpha(0)$  Thomson-limit selection of §7.2. For a confined quark, no pole exists, and the readout defaults to the conventional current mass,  $m_d^{\{M\bar{S}\}}(2 \text{ GeV})$ . The rule is asymmetric but physically motivated: each sector is read at the point where it is an invariant, and the  $\{\alpha(0), \text{pole-e}, \text{current-d}\}$  triple becomes coherent rather than accidental. P6 is a premise, owned by the RG/readout theorem (Debt 6).

**The exposure P6 exists to answer, stated before a hostile reader finds it:** a genuinely uniform-convention comparison — both masses in  $M\bar{S}$  at 2 GeV, using  $m_e^{\{M\bar{S}\}}(2 \text{ GeV}) \approx 0.495 \text{ MeV}$  (pole→ $M\bar{S}$  conversion  $-0.23 \%$ , QED running  $\gamma_m = 3\alpha/2\pi$  over  $\approx 8.3$  e-folds) — gives a uniform ratio of **9.43 ( $N_f = 2+1$ ),  $2.4\sigma$  from 9**, or **9.49 ( $N_f = 2+1+1$ ),  $4.8\sigma$  from 9**. The  $1-2\sigma$  headline below is therefore a property of the specific per-sector pair that P6 declares, not of the ratio in every convention. If the correct lepton readout is the  $M\bar{S}$  mass at the common scale, the naive shared kernel fails today (falsifier F11).

Measured, under P6 ( $m_d$  in  $M\bar{S}$  at 2 GeV;  $m_e = 0.51100 \text{ MeV}$  pole):

$m_d/m_e = 9.14 \pm 0.18$  (FLAG 2024  $N_f = 2+1$ :  $4.67 \pm 0.09$ ) —  **$0.8\sigma$  from 9**;  $m_d/m_e = 9.20 \pm 0.10$  (FLAG 2024  $N_f = 2+1+1$ :  $4.70 \pm 0.05$ ) —  **$2.0\sigma$  from 9**.

The uncertainty is dominated entirely by  $m_d$ , and the honest headline is: **the prediction sits  $1-2\sigma$  low, depending on the lattice average**.

Three properties make this the sharpest statement in the paper.

**First, it is immune to every scale debt.** Debts 4 and 5 — the derivations of  $v_{cl}$  and  $\alpha_{int}$  — do not touch it: both quantities cancel identically. What remains is the readout-scale question, and it is best stated affirmatively rather than defensively. Because  $m_e$  does not run under QCD while  $m_d$  does (§7.2.1:  $m_d = 5.85 / 4.67 / 4.26$  MeV at 1 / 2 / 3 GeV, so the P6 ratio reads  $11.4 / 9.14 / 8.33$ ), the shared-kernel hypothesis plus current data **predicts the quark readout scale:**

**$\mu_\star = 2.13$  GeV, band 1.97–2.30 (FLAG  $N_f = 2+1$ ) or  $2.18 \pm 0.10$  GeV (FLAG  $N_f = 2+1+1$ )** — the scale at which  $m_d^{\overline{MS}}(\mu_\star) = 9 m_e$  (exact two-loop integration, Appendix F.6).

This converts the convention question into a target: the projection-scale theorem (Debt 6) must independently derive a value in that window. If it produces  $\mu_\star \approx 2.1\text{--}2.2$  GeV, the shared kernel is strikingly confirmed; if it produces a materially different scale, the shared kernel is falsified regardless of how the lattice numbers move. (A crossing scale trivially exists for any nearby prediction — the content is not its existence but whether the theory independently lands in this window.)

**Second, it unifies the two residuals.** The +2.09 % electron miss (Appendix C) and the +0.54 % down miss (§5.1) are not two independent facts. They are one fact:

$1.0209 / 1.0054 = 1.0154$  — the measured P6 ratio is 9.14, i.e. **+1.5 % above 9**, and that single number *is* both residuals.

Any eventual explanation of the electron residual — a lepton readout correction, a coupling-scale refinement — must move the down residual coherently, because only their ratio is observable in this test. Writing the eventual sector corrections as  $m_e = \kappa_e \cdot m_0$  and  $m_d = 9\kappa_d \cdot m_0$ , the measurement requires  $\kappa_d / \kappa_e = 1.015 \pm 0.020$  ( $N_f = 2+1$ ;  $1.022 \pm 0.011$  under  $2+1+1$ ). That is a joint target for the charged-lepton absolute-scale theorem, tighter than either residual alone.

**Third, it is testable now, at lattice precision.** No new theory input is required to sharpen the test: every improvement in the lattice determination of  $m_d$  directly tightens both the ratio and the  $\mu_\star$  window. The FLAG-average tension is already  $2.0\sigma$  — informative, not yet decisive. A future determination pinning the P6 ratio at, say,  $9.14 \pm 0.03$  would falsify the naive shared kernel and *quantitatively fix* the required sector correction  $\kappa_d / \kappa_e$ ; a determination consistent with 9.00 would be striking confirmation.

This is the correct headline falsifier because it inverts the paper's usual epistemic direction: everywhere else, conclusions wait on the discharge theorems; here, experiment can move first — and the theory has been handed a specific number,  $\mu_\star$ , that it must produce.

---

## 10. What Has Actually Been Derived

Conditionally on the stated projection premises P1–P6 (including the explicit P2'), this paper derives:

1. a candidate replacement for the imported down-quark anchor;
2. the formula  $m_d^{\text{proj}} = 9\alpha_{\text{int}}^2 v_{\text{cl}}/(8\pi)$ ;
3. the Scale-Carrier Uniqueness Lemma:  $m_d \propto v_{\text{cl}}$  forced by representation closure (§3.1);
4. the operator form of the two-gate readout,  $M_d \sim \mathcal{G}_{\text{out}} \circ \mathcal{G}_{\text{in}}$ , with the minimality/closure exclusion of one and three gates (§3.2);
5. the normalized projection-trace form of  $1/(8\pi)$  over the factorized access space  $\mathcal{A}_d = S^1 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$ , with the no-double-count statement (§3.3);
6. the Colour–Triality Independence Lemma: local representation dimension  $\times$  global source-grade denominator (§3.4);
7. the numerical first audit  $m_d^{\text{proj}} = 4.6952 \text{ MeV}$  (+0.54 % vs the 2 GeV reference value);
8. the dimensionless coefficient  $\eta_d = 9\alpha^2/(8\pi) \approx 1.9069 \times 10^{-5}$  and the Yukawa form  $y_d = 9\sqrt{2} \cdot \alpha^2/(8\pi) \approx 2.6968 \times 10^{-5}$ ;
9. the conversion of the inherited quark ratio grid from one external anchor to one projected anchor;
10. the consolidated ablation table showing the formula is not continuously tunable (§3.6);
11. the Reading A / Reading B scale discrimination and the coherent {low-scale readout, low-scale coupling} convention pair (§7.2);
12. the explicit RG firewall required before heavy-quark pole comparisons;
13. the parallel colourless-kernel clue for the charged-lepton absolute-scale theorem, and the strange control signal (§5.3);
14. the parameter-free ratio prediction  $m_d/m_e = 9$  exactly, immune to the  $v$  and  $\alpha$  debts, measured under P6 at  $9.14 \pm 0.18$  (FLAG 2+1) /  $9.20 \pm 0.10$  (FLAG 2+1+1) —  $1-2\sigma$  (§9.1) — with the two sector residuals unified into the single constraint  $\kappa_d/\kappa_e = 1.015 \pm 0.020$ , and recast as the readout-scale prediction  $\mu_\star = 2.13$  (band 1.97–2.30) /  $2.18 \pm 0.10 \text{ GeV}$  (Appendix F).

This paper does **not** derive:

- $v_{\text{cl}}$ ;  $\alpha_{\text{int}}$ ; the coupling values or their running;
- the phase-loop projection from first principles;
- the two binary halvings from a microscopic operator;
- the colour–triality multiplicity 9 from a complete current-readout theorem;
- the  $u/d$ ,  $s/d$ ,  $c/d$ ,  $b/d$ , or  $t/d$  ratios;
- the RG matching map or the projection scale  $\mu_{\text{cl}}$ ;
- the top pole mass; the full Yukawa matrices; CKM or PMNS;
- confinement or  $\Lambda_{\text{QCD}}$ .

---

## 11. Open Debts

**Debt 1 — Discrete-to-Continuous Projection-Kernel Theorem (§3.5.3).** The factor  $1/(8\pi)$  must be derived from the substrate-to-field projection: exactly one phase-loop average and exactly two binary interface selections. The decomposition  $1/(8\pi) = (1/2\pi) \cdot (1/2) \cdot (1/2)$  is

structurally plausible but not yet a theorem. This is the debt most fully internal to the programme, requiring no upstream import.

**Debt 2 — Two-Gate Current-Readout Theorem (§3.5.2).** The power  $\alpha^2$  must be derived from the current-readout architecture: two gates necessary and sufficient, not merely selected by the data (which already excludes one and three gates by factors  $\approx 137$ ). The theorem must additionally own the charge-blindness premise P2' explicitly: it must show that gate weight counts source-access events independently of carrier charge, with all charge structure entering through the triality multiplicity — or the formula fails by nearly two orders of magnitude (F2b).

**Debt 3 — Colour–Triality Current-Readout Lemma (§3.5.4).** Both factors of 3 are inherited (colour census; faithful  $\mathbb{Z}_6$  quotient), but the composition rule — that the readout counts their *product* — must be proven. It cannot remain "three colours times three triality slots" as a verbal explanation.

**Debt 4 — Closure-Potential Scale Theorem (§3.5.1).** The scale  $v_{cl}$  must be derived from the closure-interface potential, radial stability, and substrate stiffness. Until then, the quark anchor has been reduced to the electroweak scale, not to absolute substrate units.

**Debt 5 — Coupling theorem.** The relevant  $\alpha_{int}$  must be derived and its scale fixed. The first audit selects  $\alpha(0)$  (§7.2), but a completed theory must say why the infrared coupling is the correct readout coupling.

**Debt 6 — Scheme and RG matching.** The paper must be joined to a full RG matching ledger, including a derivation of the projection scale  $\mu_{cl}$ . The debt is sharper after §7.2.1: the kernel owes not merely "why a low scale" but a statement of the **scheme- and scale-invariant object it outputs**, and an explanation of why that object's value coincides numerically with  $m_d^{\overline{MS}}(2 \text{ GeV})$  — a scheme-dependent quantity at a human-chosen waypoint. **Named action item:** compute the RG-invariant mass  $\hat{m}_d$  and check whether the kernel output identifies with it; this is a definite, checkable question. The theorem additionally owns premise P6 (per-sector IR-observable readout, §9.1) and inherits a quantitative target from the ratio falsifier: the derived quark readout scale must land in the window  $\mu^\star = 2.13 \text{ GeV}$ , band 1.97–2.30 (FLAG  $N_f = 2+1$ ) /  $2.18 \pm 0.10 \text{ GeV}$  (FLAG  $N_f = 2+1+1$ ) (Appendix F.6). Until this debt is discharged, the Reading A selection remains an audit result, and the final half-percent of the §5.1 agreement remains convention-laden.

**Debt 7 — Electron anchor comparison.** The relation  $m_e \sim \alpha^2 v / (8\pi)$  (+2.1 % residual) must be developed in the charged-lepton absolute-scale theorem, including an account of the residual. The target is now quantified by the ratio falsifier (§9.1): the lepton-readout correction must satisfy  $\kappa_d / \kappa_e = 1.015 \pm 0.020$  ( $N_f = 2+1$ ;  $1.022 \pm 0.011$  for  $2+1+1$ ) — it cannot explain the electron residual without moving the down residual coherently. If the shared-kernel relation fails, the down-anchor interpretation requires revision.

**Debt 8 — Anchor Ground-State Lemma (§3.7).** P5 must be promoted from bookkeeping to physics: prove that the down quark is the minimal stable colour-triplet configuration of the interface — lowest localization, no channel saturation, no  $\chi$ -excitation — so that all flavour-grid

operators act as the identity on it and the down cell carries no  $O(1)$  structural dressing ( $\kappa_{P5} = 1$ ). Without this lemma, the identity of the projected kernel value and the grid anchor is a normalization convention, not a derived fact.

---

## 12. Falsification Conditions

**F0 — Headline falsifier: the parameter-free ratio (§9.1).** The shared-kernel hypothesis predicts  $m_d/m_e = 9$  exactly, with  $v_{cl}$  and  $\alpha_{int}$  cancelling, under the per-sector IR-observable readout P6. Measured:  $9.14 \pm 0.18$  (FLAG  $N_f = 2+1$ ) to  $9.20 \pm 0.10$  (FLAG  $N_f = 2+1+1$ ) — **1–2 $\sigma$  from the prediction depending on the lattice average.** Equivalently, the hypothesis plus data predicts the quark readout scale  $\mu_\star = 2.13$  (band 1.97–2.30) GeV ( $N_f = 2+1$ ) /  $2.18 \pm 0.10$  GeV ( $2+1+1$ ), which the projection-scale theorem must independently produce. This falsifier is immune to Debts 4 and 5 and to the coupling-scale identification, and it sharpens automatically with every improvement in the lattice determination of  $m_d$ . Unlike F1–F10, this condition waits on no discharge theorem beyond the readout premise P6 itself (whose failure mode is F11).

**F1 — Wrong scale.** If the correct completion scale is not  $v_{cl}$ , or  $v_{cl}$  does not enter linearly, the formula fails.

**F2 — Wrong coupling power.** If the source-current projection is one-gate (predicting 643 MeV) or three-gate (predicting 34 keV), or governed by a materially different coupling, the formula fails or must be replaced.

**F2b — Charge-weighted gates.** If gate weight is shown to carry source charge as  $Q^2$  per gate, the down formula acquires an 81-fold suppression relative to the electron kernel and fails by nearly two orders of magnitude; the charge-blind readout architecture (P2') fails with it.

**F3 — No phase-loop factor.** If the discrete-to-continuous projection carries no  $1/(2\pi)$  phase-loop average, the prediction shifts by a factor  $\approx 6.3$  and the  $8\pi$  denominator loses its structural source.

**F4 — No binary halvings.** If the chiral and branch projections do not supply two independent halvings, the denominator is  $2\pi$  or  $4\pi$ , a factor-4 or factor-2 miss.

**F5 — Triality multiplicity absent.** If the down readout multiplicity is 3 rather than 9, the predicted scale is 1.565 MeV — one third of the target.

**F6 — Colour multiplicity absent.** If colour support does not multiply the readout, the scale again drops by three.

**F7 — RG mismatch.** If the RG ledger overturns the Reading A selection — i.e. if the projection output is proven to be a  $v_{cl}$ -scale boundary value — the first numerical audit fails at the +84 % level (two-loop, computed) and the formula must be reinterpreted or abandoned.

**F8 — Electron inconsistency.** If the colourless first-generation kernel fails badly for the electron once proper lepton readout is derived, the shared projection-kernel interpretation weakens, and the "minimal readout" reading of the down anchor with it.

**F9 — Quark-grid ratio failure.** If the inherited quark ratios fail under scheme-consistent treatment, the anchor formula may survive but the full mass-grid programme fails.

**F10 — Anchor dressing.** If the down readout carries a residual  $O(1)$  structural dressing  $\kappa_{P5} \neq 1$  from the flavour machinery (§3.7), the kernel value and the grid anchor decouple. The measured constraint is  $\kappa_{P5} = 0.995 \pm 0.019$  (FLAG  $N_f = 2+1$ ) or  $1.001 \pm 0.011$  (FLAG  $N_f = 2+1+1$ , consistent with exactly 1 at  $0.1\sigma$ ); the dressing theorem must independently reproduce a value in these bands for the present agreement to survive.

**F11 — Lepton readout convention.** If the correct lepton-sector readout is shown to be the  $M\bar{S}$  mass at the common quark scale rather than the pole mass (i.e. if P6 fails), the uniform-convention ratio is 9.43 ( $N_f = 2+1$ ,  $2.4\sigma$ ) or 9.49 ( $N_f = 2+1+1$ ,  $4.8\sigma$ ), and the shared-kernel prediction  $m_d/m_e = 9$  fails on current data (§9.1).

---

## 13. Milestone Statement

The previous quark mass-grid milestone was

$$m_q = R_q \cdot m_d \star.$$

This paper proposes the upgrade

$$m_q = R_q \cdot 9 \alpha_{int}^2 v_{cl} / (8\pi).$$

The down anchor is no longer a free empirical input. It is replaced by a structured projection:

completion scale + two-gate current coupling + phase/binary field dilution + colour-triality readout.

The programme therefore moves from

one-anchor quark ratio calculator

to

**candidate zero-quark-anchor current-mass calculator.**

The qualifier matters. The quark anchor is not derived absolutely, because  $v_{cl}$ ,  $\alpha_{int}$ , the projection scale, and the RG matching remain owed. But the arbitrary down-quark mass input

has been eliminated in favour of four named structural gates, each carrying a quantified falsifier separated from the target by a factor of 2 or more.

That is the intended milestone. In Master Ledger terms, this paper is a **gate conversion**: it converts one gate (the down anchor) from *open / empirically imported* to *structured / conditional*, and in doing so it opens a finite set of narrower, named gates — the discharge theorems, premise proofs, and readout lemmas of the debt register (§11), currently eight, plus the  $\mu^\star$  and  $\hat{m}_d$  targets of §7.2.1 and §9.1. One broad debt is retired; a finite, named, individually falsifiable set is created. That trade — replacing an unstructured import with explicit proof obligations — is the same trade the programme has made at every previous milestone.

---

## 14. Conclusion

The one-anchor quark mass programme was always meant to be temporary. It isolated the dimensionless structure of the quark hierarchy without pretending the absolute scale had been derived. That move worked: it produced a graded ratio grid, exposed the strange/down trace as the safest first target, and identified the top-scale trap as a scheme problem rather than a new mystery.

But a first-principles mass derivation cannot keep the anchor forever.

This paper proposes the first serious replacement:

$$m_d^{\text{proj}} = 9 \alpha^2 v / (8\pi) \approx 4.695 \text{ MeV.}$$

The formula has the right size — within 0.5 % of the reference current mass. More importantly, it has the right architecture. It uses the electroweak completion scale already required by the mass-attachment programme. It uses the source-current coupling in second order, with one-gate and three-gate alternatives excluded by two orders of magnitude each. It uses a discrete-to-continuous phase and binary projection whose absence would shift the prediction by factors of 2 to  $2\pi$ . It uses the colour–triality multiplicity forced by the fact that down quarks are colour-supported third-charge carriers, with the multiplicity alternatives 3 and 27 excluded by a factor of 3 each. And its colourless kernel independently lands within 2 % of the electron mass, suggesting a common first-generation projection origin with sector-dependent multiplicity.

The sharpest single statement follows from that last point: because  $v$  and  $\alpha$  cancel in the ratio, the shared-kernel hypothesis — under the declared per-sector readout P6 — predicts  $m_d/m_e = 9$  exactly. The measured value sits  $1-2\sigma$  high depending on the lattice average, and the prediction can be recast in its strongest form: the shared kernel plus current data **predicts the quark readout scale,  $\mu^\star \approx 2.1-2.2 \text{ GeV}$**  — a specific number the projection-scale theorem must independently produce. The test tightens with lattice precision alone; on this one prediction, experiment and the  $\mu^\star$  derivation race each other, and everything else in the paper waits on neither.

If the projection theorem is proven, the prior quark grid becomes

$$(u, d, s, c, b, t) = [9\alpha^2v/(8\pi)] \cdot (6/13, 1, 20, 3080/13, 5e^{16/3}, 62154),$$

with the heavy entries passed through an RG and pole-mass ledger before comparison to experiment.

The paper is therefore not the end of the mass programme. It is the correct next bridge. It converts the largest visible defect of the quark grid — the imported down anchor — into a finite theorem with finite, quantified ways to succeed or fail.

The proof sequence is now explicit:

derive  $v_{cl} \rightarrow$  derive  $\alpha_{int} \rightarrow$  derive  $1/(8\pi) \rightarrow$  derive  $N_d = 9 \rightarrow$  derive  $\mu_{cl}$  and run RG.

That is hard, but it is no longer vague.

In one sentence:

**The down-quark current anchor is proposed to be the minimal colour–trinality current projection of the electroweak completion interface, and its value is  $9\alpha^2v/(8\pi) \approx 4.695$  MeV.**

That is the strongest current version of the zero-anchor quark-mass bridge.

---

## Appendix A — Numerical Audit

Inputs:

$$\alpha^{-1} = 137.035999084, v = 246.21965 \text{ GeV}.$$

Then:

$$\alpha = 7.29735257 \times 10^{-3} \quad \alpha^2 = 5.32513545 \times 10^{-5} \quad \alpha^2v/(8\pi) = 5.216912 \times 10^{-4} \text{ GeV} = 0.5216912 \text{ MeV (colourless kernel)} \\ m_d^{\text{proj}} = 9 \times 0.5216912 \text{ MeV} = \mathbf{4.6952208 \text{ MeV}}.$$

Grid outputs:

$$m_u = (6/13) m_d = 2.167025 \text{ MeV} \quad m_s = 20 m_d = 93.904416 \text{ MeV} \quad m_c = (3080/13) m_d = 1.112406 \text{ GeV} \\ m_b = 5e^{16/3} m_d = 4.862541 \text{ GeV}$$

The top entry is not printed as a physical prediction: its interpretation is the most scheme-sensitive and belongs exclusively inside the RG/pole ledger.

Discrimination values (§3):

one-gate:  $9\alpha v/(8\pi) = 643.41 \text{ MeV}$  three-gate:  $9\alpha^3 v/(8\pi) = 34.26 \text{ keV}$   $N = 1$ :  $0.52169 \text{ MeV}$ ;  $N = 3$ :  $1.56507 \text{ MeV}$ ;  $N = 27$ :  $14.0857 \text{ MeV}$   $\alpha(M_Z)$  variant ( $\alpha^{-1} = 127.952$ ):  $5.3856 \text{ MeV}$ .

## Appendix B — Equivalent Yukawa Form

In the Standard Model convention  $m_d = y_d v/\sqrt{2}$ , the projected anchor gives

$$y_d v/\sqrt{2} = 9\alpha^2 v/(8\pi) \Rightarrow y_d = 9\sqrt{2} \cdot \alpha^2/(8\pi) \approx 2.6968 \times 10^{-5}.$$

This is the form to use when integrating the result into a full Yukawa-matrix paper.

## Appendix C — The Electron Clue

The projection kernel without colour-triality multiplicity is

$$m_0 = \alpha^2 v/(8\pi) = 0.5217 \text{ MeV},$$

versus  $m_e = 0.51100 \text{ MeV}$ : a residual of +2.09 %.

The paper does not claim this as a proof; the charged-lepton readout has its own sector rules, and the lepton residual is four times the down-sector residual, which itself demands explanation. But the structure is suggestive:

$$m_e \sim 1 \cdot \alpha^2 v/(8\pi), m_d \sim 9 \cdot \alpha^2 v/(8\pi).$$

If the lepton absolute-scale theorem derives the first relation with its residual accounted for, and the present theorem's premises are proven for the second, the first-generation charged-fermion scale is unified by a common projection kernel with sector-dependent multiplicity. That would be a major programme convergence.

The sharpest form of the clue is the ratio (§9.1):  $v$  and  $\alpha$  cancel, leaving the parameter-free prediction  $m_d/m_e = 9$  against the measured  $9.14\text{--}9.20$  ( $1\text{--}2\sigma$ , under the per-sector readout P6). In that form the two residuals of this appendix and §5.1 collapse into one number —  $1.0209/1.0054 = 1.0154$  — so any account of the electron residual is simultaneously an account of the down residual. The eventual lepton-readout correction is thereby constrained to  $\kappa_d/\kappa_e = 1.015 \pm 0.020$  ( $N_f = 2+1$ ;  $1.022 \pm 0.011$  for  $2+1+1$ ) — a single joint constraint replacing what previously read as two independent residuals, and one that any proposed correction must satisfy in both sectors at once.

## Appendix D — Claim-Status Ledger

Claim	Status	Comment
$m_d$ ★ must be removed	established programme need	SM-23 gate

Claim	Status	Comment
$m_d \propto v_{cl}$ (P1)	conditional / strong	Scale-Carrier Uniqueness Lemma: forced by representation closure, conditional on inherited representation content
two-gate current projection (P2)	conjectural theorem	operator form $M_d \sim \mathcal{G}_{out} \circ \mathcal{G}_{in}$ ; Two-Gate Current-Readout Theorem owed; alternatives excluded $\times 137$ each way
phase-loop projection $1/(2\pi)$	conjectural sub-lemma	uniform phase-density selection over $S^1_{phase}$
two binary halvings	conjectural sub-lemmas	rank-one selections on independent tensor factors $\mathbb{C}^2_{chiral}, \mathbb{C}^2_{weak}$
$\Pi_{cont} = 1/(8\pi)$ (P3)	conditional	normalized projection trace over $\mathcal{A}_d$ ; Projection-Kernel Theorem owed; internal to programme
colour factor 3	inherited	$k = 3$ colour census / $SU(3)_C$ closure; local representation dimension
trality factor 3	inherited	faithful $\mathbb{Z}_6$ quotient: $2t+3d+q \equiv 0 \pmod{6} \implies Q \equiv -t/3 \pmod{1}$ ; global source-grade denominator
colour-trality independence	stated lemma	different structures counted: local vs global; answers double-counting objection
$N_d = 9$ (P4)	conditional	Colour-Triality Current-Readout Lemma owed; alternatives excluded $\times 3$ each way
charge-blind gates (P2')	premise, newly explicit	if gate weight $\propto Q^2$ , formula fails $\times 81$ (F2b); owned by Debt 2
anchor-first readout (P5)	bookkeeping observation	down cell is grid identity by construction; physical ground-state status owed (Debt 8, F10)
scheme-invariant kernel output	open question	§7.2.1: $M\bar{S}(2 \text{ GeV})$ is convention-laden; $\hat{m}_d$ identification is the named action item (Debt 6)
chance-hit probability	quantified	$\approx 18\%$ within $1\%$ over the 144-combo menu; match motivates, does not certify (§5.2)
menu uniqueness	scan result	stated combination unique within $2.5\%$ of $m_d$ ; no combo within $2.5\%$ of $m_u, m_s$ (§3.6)
$m_d = 9\alpha^2 v / (8\pi)$	central conditional theorem	numerical audit strong
$m_d = 4.6952 \text{ MeV}$	first audit	$+0.54\%$ vs FLAG $N_f = 2+1$ ( $4.67 \pm 0.09$ ); $\approx 0.1\%$ vs FLAG $N_f = 2+1+1$ ( $4.70 \pm 0.05$ ); convention declared §5.1
projection scale $\mu_{cl}$ (Reading A)	audit-selected	Reading B excluded at $+84\%$ (two-loop, computed); derivation owed
coupling scale $\alpha(0)$	audit-selected	$\alpha(M_Z)$ variant misses by $\approx 15\%$ ; derivation owed
full quark grid absolute values	downstream	inherits ratios

Claim	Status	Comment
RG matching	owed	essential for heavy quarks
top pole mass	not claimed	requires RG/pole conversion
electron relation	clue, downgraded	+2.1 % residual; within chance-hit band until lepton readout theorem (§9); ratio test §9.1 is the discriminator
$m_d/m_e = 9$ exactly	headline falsifier	parameter-free; $v, \alpha$ cancel; under P6: $9.14 \pm 0.18$ (FLAG 2+1, $0.8\sigma$ ) / $9.20 \pm 0.10$ (FLAG 2+1+1, $2.0\sigma$ ); §9.1
per-sector readout (P6)	premise, newly explicit	pole-e / current-d; uniform- $M\bar{S}$ ratio $9.43$ (2+1, $2.4\sigma$ ) / $9.49$ (2+1+1, $4.8\sigma$ ) (F11); owned by Debt 6
readout-scale prediction	derived target	$\mu_\star = 2.13$ , band $1.97\text{--}2.30$ (FLAG 2+1) / $2.18 \pm 0.10$ GeV (FLAG 2+1+1); Appendix F.6; projection-scale theorem must land in window
residual unification	derived observation	electron and down residuals are one fact: ratio +1.5 % above 9; constrains $\kappa_d/\kappa_e = 1.015 \pm 0.020$ (FLAG 2+1) / $1.022 \pm 0.011$ (FLAG 2+1+1)
strange control signal	consistency check	$20 m_d^{\text{proj}} \approx 93.90$ MeV; inherits anchor by construction
final zero-anchor derivation	not yet	requires $v, \alpha, \Pi, N_d, \mu_{\text{cl}}, \mathcal{R}$ proofs

## Appendix E — Referee Objections and Responses

**Objection 1 — "You fitted  $9\alpha^2v/(8\pi)$  to a known number."** Response: each factor has a recorded provenance and a named discharge theorem (§3.5); the alternatives to each factor miss the target by large discrete amounts (§3.6); and the paper's claim level is explicitly conditional on the discharge theorems, not on the numerical match. A fit has adjustable parameters; this formula has none. What it has instead is unproven premises, which is a different and honestly declared weakness.

**Objection 2 — "v is still imported, so nothing has been derived."** Response: correct in part, and stated in the paper's own words (§3.5.1): the quark-mass anchor has been reduced to the electroweak scale, not to absolute substrate units. What has changed is the *number of independent scale inputs*: previously the programme owed both  $v$  and  $m_{d\star}$ ; now it owes only  $v$ , with  $m_{d\star}$  reduced to structure. The Scale-Carrier Uniqueness Lemma (§3.1) additionally shows that the proportionality to  $v$  is forced by representation closure, so the residual debt is the value of  $v$ , not the form of the mass law. Moreover, the paper's headline falsifier (§9.1) is entirely immune to this objection: in the ratio  $m_d/m_e = 9$ , both  $v$  and  $\alpha$  cancel identically, so the sharpest prediction in the paper survives even if the  $v$  and  $\alpha$  derivations never arrive. The Closure-Potential Scale Theorem remains owed for the absolute values.

**Objection 3 — " $\alpha(0)$  was chosen because it works."** Response: partially conceded and quantified. The audit of §7.2 shows that the {low-scale readout, low-scale coupling} pair is the

unique internally consistent convention among the four candidates, with the alternatives missing by 15–111 %. That is a selection, not a derivation, and the paper says so — and §7.2.1 further concedes that the pair is coherent as bookkeeping but not yet as physics, since  $\alpha(0)$  is an IR invariant while  $m_d^{\overline{\text{MS}}}(2 \text{ GeV})$  is not. The coupling theorem must derive the readout coupling; until then the identification is flagged as conditional in §0.2, §3.5.2, and Appendix D.

**Objection 4 — "Colour and triality are double-counted:  $3 \times 3$  counts the same three twice."** Response: the Colour–Triality Independence Lemma (§3.4) answers this directly. The colour factor counts the local representation dimension of the carrier,  $\dim(\mathbf{3}) = 3$ . The triality factor counts the global source-grade denominator imposed by the faithful  $\mathbb{Z}_6$  quotient,  $Q \equiv -t/3 \pmod{1}$ . One is local and representation-dimensional; the other is global and source-normalizing. They answer different questions and act at different structural levels; hence they multiply. What remains owed is the composition lemma proving the readout counts the product — but the double-counting objection specifically is answered by the independence of the two structures.

**Objection 5 — "The RG running is not completed, so the numerical agreement is meaningless."** Response: the paper claims a projected current anchor only — not pole masses, not high-scale Yukawa outputs (§0.4, §7). The RG dependence is not ignored but used as a discriminator: the running audit of §7.2 excludes the interface-scale boundary reading at the +84 % level (two-loop, computed) and selects the low-current readout, which is itself a nontrivial structural finding. The full RG Matching Theorem is owed and is listed as Debt 6; heavy-quark and top comparisons are explicitly firewalled until it exists.

## Appendix F — Reproducibility Appendix for the RG Probe

### F.1 Purpose and status

This appendix records the external-QCD calculation used in §7.2, §7.2.1, and §9.1. Its purpose is limited: to make the RG probe reproducible and to show how the quoted values

$$\mathcal{R}_d(2 \text{ GeV} \leftarrow 246 \text{ GeV}) = 1.83, \alpha_s(2 \text{ GeV}) = 0.296, m_d(1 \text{ GeV}) = 5.85 \text{ MeV}, m_d(3 \text{ GeV}) = 4.26 \text{ MeV}, \mu_\star = 2.13 \text{ GeV} (N_f = 2+1) / 2.18 \text{ GeV} (N_f = 2+1+1)$$

were obtained.

This appendix is **not** the VERSF RG Matching Theorem. It uses standard two-loop QCD running as an external diagnostic. Finite threshold corrections, higher-loop running, electroweak corrections, nonperturbative low-scale issues, and scheme-invariant mass reconstruction remain part of Debt 6.

### F.2 Inputs

$$\alpha_s(M_Z) = 0.1179, M_Z = 91.1876 \text{ GeV}, m_t = 172.76 \text{ GeV} \text{ (pole, used as threshold)}, m_b = 4.18 \text{ GeV}, m_c = 1.27 \text{ GeV}, v_{cl} = 246.21965 \text{ GeV}.$$

For the light-quark comparison convention,  $m_d^{\overline{\text{MS}}}(2 \text{ GeV}) = 4.67 \pm 0.09 \text{ MeV}$  is the FLAG 2024  $N_f = 2+1$  average (the paper's declared primary convention, §5.1), and  $4.70 \pm 0.05 \text{ MeV}$  is the FLAG 2024  $N_f = 2+1+1$  average. The electron pole mass used in the ratio test is  $m_e = 0.51099895 \text{ MeV}$ .

### F.3 Running conventions

Define  $a_s = \alpha_s/\pi$  and  $L = \ln \mu^2$ . The two-loop QCD beta function is

$$da_s/dL = -\beta_0 a_s^2 - \beta_1 a_s^3, \beta_0 = (11 - 2n_f/3)/4, \beta_1 = (102 - 38n_f/3)/16,$$

and the two-loop quark-mass anomalous-dimension equation is

$$d \ln m / dL = -\gamma_0 a_s - \gamma_1 a_s^2, \gamma_0 = 1, \gamma_1 = 101/24 - 5n_f/36.$$

Step-function flavour thresholds:  $n_f = 6$  above  $m_t$ ; 5 between  $m_t$  and  $m_b$ ; 4 between  $m_b$  and  $m_c$ ; 3 below  $m_c$ . At this diagnostic level,  $\alpha_s$  and the light-quark mass are taken continuous across thresholds; finite matching corrections are not included.

### F.4 Reproducing the Reading B exclusion

Running  $\alpha_s$  from  $M_Z$  up to  $v_{cl}$  (crossing  $m_t$ ) gives  $\alpha_s(v_{cl}) = 0.103$ . Running the light-quark mass factor from  $v_{cl}$  down to 2 GeV (crossing  $m_t$  and  $m_b$ ) gives

$$\mathcal{R}_d(2 \text{ GeV} \leftarrow v_{cl}) = 1.83, \alpha_s(2 \text{ GeV}) = 0.296.$$

If the projected mass  $m_d^{\text{proj}} = 4.695 \text{ MeV}$  were a boundary value at  $v_{cl}$ , standard running would give  $m_d^{\overline{\text{MS}}}(2 \text{ GeV}) = 1.83 \times 4.695 = 8.59 \text{ MeV}$ , an overshoot of

$$(8.59 - 4.67)/4.67 = +84 \%$$

against the declared reference value. This is the basis for the §7.2 statement that the first numerical audit selects Reading A over Reading B.

### F.5 Reproducing the low-scale drift

Starting from  $m_d^{\overline{\text{MS}}}(2 \text{ GeV}) = 4.67 \text{ MeV}$ , the same two-loop running gives

$$m_d^{\overline{\text{MS}}}(1 \text{ GeV}) = 5.85 \text{ MeV}, m_d^{\overline{\text{MS}}}(3 \text{ GeV}) = 4.26 \text{ MeV}.$$

This is the basis for the scheme-and-scale warning of §7.2.1: the value matched by the projection formula is a running  $\overline{\text{MS}}$  mass at a conventionally chosen scale, not a fixed physical pole mass.

### F.6 Reproducing the readout-scale target $\mu_\star$

The shared-kernel ratio test predicts  $m_d(\mu_\star) = 9 m_e = 4.59899$  MeV. Solving by bisection along the two-loop trajectory:

$\mu_\star = 2.127$  GeV, band [1.968, 2.303] over  $4.67 \pm 0.09$  (FLAG  $N_f = 2+1$ );  $\mu_\star = 2.184$  GeV, band [2.090, 2.283] over  $4.70 \pm 0.05$  (FLAG  $N_f = 2+1+1$ ).

Allowing for the input uncertainties and small threshold/convention shifts, this is the source of the paper's rounded target  $\mu_\star \approx 2.1\text{--}2.2$  GeV. The interpretation is important and is stated in §9.1: the existence of *some* crossing scale is not a prediction — a running mass crosses nearby values somewhere. The prediction is that the future VERSF projection-scale theorem must independently derive a readout scale in this narrow region.

## F.7 Uniform- $\overline{M\overline{S}}$ lepton comparison

The paper's headline ratio uses the electron pole mass under the per-sector readout premise P6. A uniform  $\overline{M\overline{S}}$  comparison gives a different result. At one-loop QED level,

$$m_e^{\{\overline{M\overline{S}}\}}(2 \text{ GeV}) \approx m_e^{\{\text{pole}\}} \cdot (1 - \delta_{\{\text{pole} \rightarrow \overline{M\overline{S}}\}}) \cdot \exp[-(3\alpha/2\pi) \ln(2 \text{ GeV} / m_e)],$$

with  $\delta_{\{\text{pole} \rightarrow \overline{M\overline{S}}\}} \approx \alpha/\pi \approx 0.0023$ , giving  $m_e^{\{\overline{M\overline{S}}\}}(2 \text{ GeV}) \approx 0.495$  MeV. Then

$$m_d^{\{\overline{M\overline{S}}\}}(2 \text{ GeV}) / m_e^{\{\overline{M\overline{S}}\}}(2 \text{ GeV}) \approx 4.67/0.495 \approx \mathbf{9.43} \text{ } (N_f = 2+1), \text{ or } 4.70/0.495 \approx \mathbf{9.49} \text{ } (N_f = 2+1+1).$$

This is why the paper states that the parameter-free ratio test is not convention-free: under P6 the comparison is pole-electron versus current-quark; under a uniform  $\overline{M\overline{S}}$  convention at 2 GeV, the ratio moves to 9.43 ( $N_f = 2+1$ ,  $2.4\sigma$  from the prediction) or 9.49 ( $N_f = 2+1+1$ ,  $4.8\sigma$ ). **P6 is not cosmetic.** It is a real readout premise that the future RG/readout theorem must justify (F11).

## F.8 Reference implementation

The following minimal Python code reproduces the values in this appendix at the diagnostic level used in the paper. Its outputs have been verified against an independent two-loop implementation; agreement is exact to the quoted precision.

```
import math

def coeffs(nf):
    beta0 = (11.0 - 2.0 * nf / 3.0) / 4.0
    beta1 = (102.0 - 38.0 * nf / 3.0) / 16.0
    gamma0 = 1.0
    gamma1 = 101.0 / 24.0 - 5.0 * nf / 36.0
    return beta0, beta1, gamma0, gamma1

def deriv(y, nf):
    a, log_m = y
    beta0, beta1, gamma0, gamma1 = coeffs(nf)
    da_dL = -beta0 * a*a - beta1 * a*a*a
```

```

dlogm_dL = -(gamma0 * a + gamma1 * a*a)
return da_dL, dlogm_dL

def rk4_segment(mu0, mu1, alpha0, mass_factor0, nf, steps=10000):
    L0 = math.log(mu0 * mu0)
    L1 = math.log(mu1 * mu1)
    h = (L1 - L0) / steps
    a = alpha0 / math.pi
    log_m = math.log(mass_factor0)
    for _ in range(steps):
        k1a, k1m = deriv((a, log_m), nf)
        k2a, k2m = deriv((a + 0.5*h*k1a, log_m + 0.5*h*k1m), nf)
        k3a, k3m = deriv((a + 0.5*h*k2a, log_m + 0.5*h*k2m), nf)
        k4a, k4m = deriv((a + h*k3a, log_m + h*k3m), nf)
        a += h * (k1a + 2*k2a + 2*k3a + k4a) / 6.0
        log_m += h * (k1m + 2*k2m + 2*k3m + k4m) / 6.0
    return math.pi * a, math.exp(log_m)

# Inputs
MZ, alpha_s_MZ = 91.1876, 0.1179
mt, mb, mc = 172.76, 4.18, 1.27
vcl = 246.21965

# Run alpha_s from MZ up to vcl
alpha_mt, _ = rk4_segment(MZ, mt, alpha_s_MZ, 1.0, nf=5)
alpha_vcl, _ = rk4_segment(mt, vcl, alpha_mt, 1.0, nf=6)

# Run mass factor from vcl down to 2 GeV
alpha_mt_down, R = rk4_segment(vcl, mt, alpha_vcl, 1.0, nf=6)
alpha_mb, R = rk4_segment(mt, mb, alpha_mt_down, R, nf=5)
alpha_2, R = rk4_segment(mb, 2.0, alpha_mb, R, nf=4)

print("alpha_s(vcl) =", alpha_vcl)           # 0.103
print("alpha_s(2GeV) =", alpha_2)           # 0.296
print("R_d(2 <- vcl) =", R)                 # 1.83
print("Reading B m_d =", R * 4.6952208)     # 8.59 MeV

# Low-scale drift from 2 GeV
m2 = 4.67
alpha_mc, R_2_to_mc = rk4_segment(2.0, mc, alpha_2, 1.0, nf=4)
alpha_1, R_2_to_1 = rk4_segment(mc, 1.0, alpha_mc, R_2_to_mc, nf=3)
alpha_3, R_2_to_3 = rk4_segment(2.0, 3.0, alpha_2, 1.0, nf=4)
print("m_d(1 GeV) =", m2 * R_2_to_1)        # 5.85 MeV
print("m_d(3 GeV) =", m2 * R_2_to_3)        # 4.26 MeV

# Solve for mu_star where m_d(mu_star) = 9 m_e
me = 0.51099895
target = 9.0 * me

def mass_at_mu_from_2(mu, m2):
    # bracket may extend below 2 GeV; cross mc if needed
    if mu >= mc:
        _, Rf = rk4_segment(2.0, mu, alpha_2, 1.0, nf=4)
    else:
        a_mc, R1 = rk4_segment(2.0, mc, alpha_2, 1.0, nf=4)
        _, R2 = rk4_segment(mc, mu, a_mc, R1, nf=3)
        Rf = R2

```

```

return m2 * Rf

def solve(m2, lo=1.5, hi=4.0):
    for _ in range(80):
        mid = 0.5 * (lo + hi)
        if mass_at_mu_from_2(mid, m2) > target:
            lo = mid
        else:
            hi = mid
    return 0.5 * (lo + hi)

print("mu_star Nf=2+1   =", solve(4.67))    # 2.13 GeV
print("mu_star Nf=2+1+1 =", solve(4.70))    # 2.18 GeV

```

Expected output, up to small rounding differences:

$\alpha_s(v_{cl}) \approx 0.103$ ;  $\alpha_s(2 \text{ GeV}) \approx 0.296$ ;  $\mathcal{R}_d \approx 1.83$ ; Reading B  $m_d \approx 8.59 \text{ MeV}$ ;  $m_d(1 \text{ GeV}) \approx 5.85 \text{ MeV}$ ;  $m_d(3 \text{ GeV}) \approx 4.26 \text{ MeV}$ ;  $\mu_\star \approx 2.13 \text{ GeV}$  ( $N_f = 2+1$ ),  $\approx 2.18 \text{ GeV}$  ( $N_f = 2+1+1$ ).

## F.9 Limitations of this appendix

This appendix deliberately omits: three-loop and four-loop QCD running; finite threshold matching corrections (including the choice of pole versus  $\overline{\text{MS}}$  top-threshold placement, which shifts  $\mathcal{R}_d$  at the sub-percent level); electroweak corrections; QED running for quark masses; nonperturbative uncertainty near 1 GeV; a derivation of the projection scale from VERSE; and a scheme-invariant reconstruction of the kernel output.

These omissions are acceptable for the present diagnostic purpose because the RG probe is used only to distinguish large alternatives — Reading A versus Reading B, a difference of order 84 % — and to locate  $\mu_\star$  within a  $\pm 0.1\text{--}0.2 \text{ GeV}$  window whose width is dominated by the  $m_d$  input uncertainty, not by the running approximations. They are not acceptable for a final mass theorem. The full RG Matching Theorem must replace this appendix with a scheme-complete and threshold-complete treatment.