The Quark Current-Mass Projection Theorem in VERSF

Canonical Chiral Attachment, Typed Schur-Complement Kernel, Two-Gate Interface Readout, Phase–Binary Projection, Additive Colour–Triality Quotient Census, Metric-Compatible Address Anchoring, and Renormalised Current-Mass Matching

Keith Taylor VERSF Theoretical Physics Programme — Fermion Mass Scale Series Gate QCMP-1

General Reader Summary

Every quark has a mass. Physics can measure those masses to good precision, but it has never been able to say why they take the values they do — the numbers are simply read off from experiment and typed into the equations by hand.

The earlier VERSF quark papers made partial progress on this. They showed that if just one quark mass is supplied — the down quark's — the pattern of the other five can be organised around it by fixed structural ratios. That is genuine compression, but it leaves the obvious question standing. A system of ratios still needs a ruler. Where does the first mass itself come from?

This paper proposes an answer, and then spends most of its length testing whether that answer can honestly be called a derivation.

The proposal is that the down quark's mass is not a free number but the output of a readout: the deep structure of the theory is "read" through four filters, and each filter leaves a calculable fingerprint on the result.

1. **A scale.** The electroweak interface — the structure associated with the Higgs field — turns out to be the only object in the theory with the right character to hand a mass of this kind to a down-type quark. It supplies the units: roughly 246 GeV.
2. **Two gates.** The mass record is written through two couplings to that interface, and each contributes one factor of the interface coupling strength — the famous number close to 1/137.

3. **A projection.** Out of everything the interface could express, the readout selects one smooth phase, one handedness, and one branch of the weak pairing. That triple selection contributes a purely geometric factor of 1 in 8π.
4. **A count of nine.** Colour comes in three varieties, and in VERSF the resolution of electric charge into thirds contributes another three, giving nine channels. If those nine channels pour their contributions into a single output with their strengths **added**, the result carries a factor of nine.

Multiply the fingerprints together: nine, times the square of 1/137, divided by 8π, times 246 GeV. The result is 4.695 MeV. The measured down-quark mass, in the convention physicists usually quote, is about 4.7 MeV.

Close. Suspiciously close. And the central discipline of this paper is to refuse to declare victory on the strength of that closeness. Three cautions carry most of the weight.

First, a quark mass is not a single number. Because quarks are permanently confined inside larger particles, their masses can only be defined relative to an accounting convention — a chosen scheme and a chosen energy scale, a little like quoting a salary without saying whether it is before or after tax, and in which currency. The paper has not yet proved which convention its 4.695 MeV belongs to. Until that is settled, the agreement with "about 4.7 MeV" is an encouraging audit, not a match.

Second, counting nine of something does not make it nine times heavier. Nine identical coins are not one ninefold coin. The nine colour–charge channels multiply the mass only if they behave like nine tributaries feeding one river — strengths added — rather than like nine copies of the same river, where the natural bookkeeping averages them back to one. Whether nature adds or averages at this junction is a sharp open question, and the paper states it as an exact either/or: added gives nine, averaged gives one, and nothing in between is allowed without explanation.

Third, numbers may not be multiplied together just because they are all pure numbers. A coupling strength, a geometric density, a fifty–fifty selection, and a channel count are different kinds of quantity. The paper builds a single mathematical machine in which each factor has one designated slot, and it forbids any factor that cannot be shown to occupy its slot. The 4.695 MeV is legitimate only if the whole product comes out of that one machine.

Alongside these sit further checks: the "first down slot" in the theory's bookkeeping must be proved to be the lightest physical down state, not merely labelled as such; and the same honesty is applied to the strange quark, where a promising count of twenty faces exactly the same added-versus-averaged question before it may be called a mass ratio of twenty.

The result is deliberately bounded:

This paper derives a conditional operator architecture for the down current-mass boundary and reduces its numerical value to a finite set of independently testable

projection, coherence, alignment, coupling, and RG premises. It does not yet derive the complete renormalised quark spectrum.

In plain terms: the paper does not claim the down-quark mass has been derived. It claims something more careful and, in the long run, more useful — that the question has been converted into a short list of precise assumptions, each of which can independently pass or fail. If they all pass, the first quark mass follows from structure alone and every other quark mass inherits a ruler. If any one fails, the failure is local, visible, and names exactly what must be rebuilt.

Scope and Status Box

Claim	Status in this paper
A down-type Dirac mass must use the electroweak completion interface	**Conditionally proved** from the inherited chiral representation content
Raw matrix entries are not physical current masses	**Closed**
Canonical current masses are singular values of the renormalised mass operator	**Closed**
Two gate coefficients multiply to α_int^2	**Exact algebra; the per-gate amplitude is the premise (P1), and physical gate identification is open**
The phase–binary kernel equals 1/(8π)	**Conditionally proved** from a zero-mode local-density rule, two unbiased rank-one binary selections, and the canonical phase measure
A mere colour or support degeneracy multiplies a quark mass	**Refuted**
All projection factors are coefficients of one typed Schur-complement kernel	**Structurally proved (Theorem 2A); substrate return open**
The same-side Feshbach blocks are harmless	**Open — External-Block Recanonicalisation Condition stated (P13)**
The compressed gates control the full Schur kernel	**Open — off-line channel closure (P15)**
Theorems 4 and 5 are factors of the same K_int^{-1}	**Open — internal-response factorisation premise (P12)**
The 1/(2π) factor is measure-invariant	**Open — requires the Canonical Phase-Measure Theorem (P14)**
Amplitudes, response weights, and singular coefficients may be freely mixed	**Refuted (Lemma 2B)**

Claim	Status in this paper
The additive internal census gives $\tilde{q}_L D_9 \tilde{q}_R^\dagger = 9 W_{ct}$	**Exact linear algebra given additive quotienting; adjoint category and physical factorisation open**
The census is internal to K_{int}^{-1}; V_d arises only after the gates	**Closed as typing**
A Hilbert-normalised quotient of the same nine channels	**Gives $N_d = 1$; the dichotomy is stated exactly**
The down address equals the smallest down singular line	**Conditionally proved** under address anchoring, reducing-subspace, ordering, and gap premises
$y_d^{(0)} = 9\sqrt{2}\, \alpha_{int}^2 / (8\pi)$	**Conditional boundary theorem**
$m_d^{(0)} = 9\, \alpha_{int}^2\, v_{cl} / (8\pi)$	**Conditional projection-kernel value**
$m_d^{(0)} = m_d^{\overline{MS}}(2\ \text{GeV})$	**Not proved**
The RG/matching map to $m_d^{\mathcal{S}}(\mu)$	**Structurally specified; dynamically open**
A light-quark pole mass	**Firewalled**
A free down quark as an asymptotic observable	**Firewalled**
$m_d/m_e = 9$ for raw measured masses	**Not generally true; RG-transported version derived**
$y_d(\mu\star)/y_e(\mu\star) = 9$ at a common matching point	**Conditional shared-kernel prediction**
The strange support count 20 is automatically a strange mass ratio	**Refuted without a singular-value-twenty realisation**
The audit scale 246.21965 GeV is typed as the observable-normalised $v_F = (\sqrt{2}\, G_F)^{-1/2}$	**Closed as audit typing**
Full six-quark zero-anchor spectrum	**Open**

Abstract

The VERSF quark-mass programme has separated the diagonal quark hierarchy into a dimensionful anchor and a set of dimensionless ratio, accessibility, confinement, and generation structures. The remaining absolute-scale problem is to derive a canonically typed quark current-mass boundary without inserting a measured quark mass.

This paper formulates the Quark Current-Mass Projection Theorem at operator level.

For the raw down-sector closure map $\Gamma_d : \mathscr{H}_D \to \mathscr{H}_{QL}$ with positive kinetic metrics K_Q, K_D, the canonical map is $Y_d^c = K_Q^{-1/2}\, \Gamma_d\, K_D^{-1/2}$, with Hermitian lifts $H_d^L = Y_d^c\, Y_d^{c\dagger}$ and $H_d^R = Y_d^{c\dagger}\, Y_d^c$; after electroweak completion, $M_d^c(\mu\star) = (v_{cl}(\mu\star)/\sqrt{2})\, Y_d^c(\mu\star)$. Physical running current masses are singular values of the renormalised canonical mass operator — not raw entries, address labels, support traces, constituent masses, or hadron masses.

The candidate first down-type projection is built from four factors, and it is typed by a single construction: the effective attachment is the Schur complement

$$E_{d,\mathrm{eff}}^{c} = -\,G_{\mathrm{out}} K_{\mathrm{int}}^{-1} G_{\mathrm{in}}$$

of a block operator in which the gates are off-diagonal amplitudes and the phase, binary, and census factors are spectral values of the internal response K_{int}^{-1}. A type-closure lemma forbids multiplying any dimensionless quantity into the coefficient unless it occupies one of these operator roles. The chiral representation theorem fixes the completion interface as the unique minimal scalar carrier connecting Q_L to d_R. Two source-readable gate maps contribute α_{int}^2 when each gate has interface amplitude α_{int}. A local zero-mode phase-density readout contributes $1/(2\pi)$, while unbiased rank-one chiral and weak-branch selections contribute $1/2$ each, giving

$$\Pi_{\mathrm{cont}} = 1/(8\pi).$$

The colour–triality contribution is formulated not as a degeneracy trace but as an additive quotient census. Nine primitive diagonal attachment channels form the bilateral census operator $D_9 = \Sigma_A |A, L\rangle\langle A, R|$, which by itself has nine unit singular values and produces only degeneracy. If the readout support factors into independent three-dimensional colour and triality-resolution fibres, and the faithful physical quotient carries every primitive channel onto the same canonical down line while preserving additive weight, then

$$\tilde{q}_L D_9 \tilde{q}_R^\dagger = 9 W_{ct}, \quad N_d = 9,$$

on internal census lines — the external down-line map arising only after gate composition, $V_d = V_{\mathrm{out}} W_{\mathrm{int}} V_{\mathrm{in}}$ — whereas a Hilbert-normalised quotient of the same channels gives $N_d = 1$. That additive-versus-normalised dichotomy is stated exactly and promoted to the principal census debt.

Under down address–spectral alignment, reducing-subspace closure, and a non-zero singular gap, the projected canonical Yukawa component is

$$P_d^L Y_d^c P_d^R = y_d^{(0)} V_d, \quad y_d^{(0)} = 9\sqrt{2}\, \alpha_{\mathrm{int}}^2 / (8\pi),$$

where V_d is the partial isometry between the right and left down singular lines. Hence

$$m_d^{(0)} = (v_{\mathrm{cl}}/\sqrt{2})\, y_d^{(0)} = 9\, \alpha_{\mathrm{int}}^2\, v_{\mathrm{cl}} / (8\pi).$$

With the first audit values $\alpha_{\mathrm{int}} = \alpha(0)$ and $v_{\mathrm{cl}} = 246.21965$ GeV, this evaluates to

$$m_d^{(0)} = 4.6952208 \text{ MeV}, \quad y_d^{(0)} = 2.6967973 \times 10^{-5}.$$

This numerical correspondence is graded as a kernel audit, not an equality with a particular quoted $\overline{\mathrm{MS}}$ mass. The renormalised current mass is

$$m_d^{\mathcal{S}}(\mu) = \sigma_{\min}[\, \Phi_{d,\mu\leftarrow\mu_\star}^{\mathcal{S}}[\, M_d^{c}(\mu_\star); \ldots \,] \,],$$

with the coupled flow map $\Phi_d^{\mathcal{S}}$ containing anomalous-dimension evolution, effective-theory changes, threshold matching, canonical re-normalisation, and scheme conversion; a reducing-line transport theorem states exactly when this coupled flow collapses to one scalar factor.

Two corollaries sharpen the programme. First, a support trace does not determine a singular value: the strange ratio $m_s/m_d = 20$ requires the canonical singular coefficients themselves to stand in ratio twenty, with an additive twenty-support quotient ($\tilde{q}_L D_{20} \tilde{q}_R^{\dagger} = 20 W_{ct}^{s}$) one sufficient realisation, not a necessity. Second, a shared colourless/coloured kernel predicts $y_d/y_e = 9$ only at a common matching point; measured low-energy masses satisfy the RG-transported relation

$$m_d^{\mathcal{S}_d}(\mu_d) / m_e^{\mathcal{S}_e}(\mu_e) = 9 \cdot \mathcal{R}_d^{\mathcal{S}_d}(\mu_d, \mu_\star) / \mathcal{R}_e^{\mathcal{S}_e}(\mu_e, \mu_\star),$$

not a scheme-free raw ratio of nine.

The result is a conditional projection theorem and a strengthened audit architecture. It does not yet derive v_{cl}, α_{int}, the physical gate maps, the additive quotient normalisation, the down address frame, or the complete coupled RG map. It converts those debts into finite operator equations and explicit falsifiers — including the external-block recanonicalisation condition on the full Feshbach reduction, the off-line channel closure guarding the compressed Schur kernel, the internal-response factorisation premise connecting the phase–binary and census factors to one K_{int}^{-1}, and the canonical phase-measure requirement behind the $1/(2\pi)$ factor.

Contents

---

1. Inheritance, Notation, and Convention Ledger

1.1 Inherited programme layers

This paper inherits the following structures.

I1 — Fermionic Fock carrier

The matter sector is represented on an antisymmetric Fock space over a positive one-particle spinorial carrier space, conditional on the inherited positivity, antisymmetry, exterior-sector completeness, and adjoint-removal premises. This supplies canonical fermion creation and annihilation operators and a local field candidate.

The present paper does not rederive the CAR algebra. It uses the Fock layer only to type the mass term as the coefficient of a fermion bilinear rather than as an abstract number.

I2 — Chiral matter representation

The one-generation chiral representation skeleton contains

$Q_L : (\mathbf{3}, \mathbf{2}, +1/6)$, $d_R : (\mathbf{3}, \mathbf{1}, -1/3)$,

in the convention

$Q = T_3 + Y$.

The completion interface has the Higgs-doublet representation

$\Phi_{cl} : (\mathbf{1}, \mathbf{2}, +1/2)$.

The gauge-invariant down-type Yukawa route is

$\bar{Q}_L \Phi_{cl} d_R$.

I3 — Canonical normalisation and Hermitian-lift discipline

A raw closure or Yukawa matrix is not a physical prediction. Positive kinetic metrics must be removed, and masses must be read from singular values or Hermitian lifts.

I4 — Address/spectrum separation

An occupancy address is not automatically a mass eigenstate. The address projector and the spectral projector require a bridge and, where necessary, a physically anchored frame.

I5 — Confinement discipline

A colour-triplet quark is an internal field degree of freedom, not an isolated asymptotic state. Current mass, constituent response, and hadron mass are different objects.

I6 — Ordered-commitment χ readout

The up/down contrast, if non-zero, reads a reversal-odd ordered-history residue rather than a symmetric final record. That structure belongs to relative flavour splitting and is not the source of the absolute down anchor.

I7 — Flavour-magnitude factorisation

At gate order, a flavour-dependent Yukawa may be expressed as the projection of a block-diagonal, multiplicatively renormalised operator, with non-factorisable corrections kept as an operator remainder rather than hidden in scalar factors.

I8 — Strange support and confinement candidates

The strange-sector programme supplies a candidate twenty-support readout and a confinement operator selecting a balanced support shell. The present paper uses that work to formulate the trace-to-singular-value bridge, not to assume it.

1.2 Hypercharge convention

This paper uses

$Q = T_3 + Y$.

Some predecessor manuscripts use

$Q = T_3 + Y'/2.$

The conversion is

$Y' = 2Y.$

Thus

$Y(Q_L) = 1/6 \Leftrightarrow Y'(Q_L) = 1/3,$

and

$Y(d_R) = -1/3 \Leftrightarrow Y'(d_R) = -2/3.$

No factor-of-two ambiguity is permitted below.

1.3 Core spaces and maps

Let $\mathscr{H}_{QL}$ be the canonical left-handed down-component generation space inside the quark doublet bundle, and $\mathscr{H}_D$ the right-handed down-type generation space.

The raw down-sector closure map is

$\Gamma_d : \mathscr{H}_D \to \mathscr{H}_{QL}.$

Let the raw kinetic metrics be

$K_Q > 0, K_D > 0.$

The canonical Yukawa map is

$Y_d^c = K_Q^{(-1/2)} \Gamma_d K_D^{(-1/2)}.$

The left and right Hermitian lifts are

$H_d^L = Y_d^c Y_d^{c\dagger}, H_d^R = Y_d^{c\dagger} Y_d^c.$

Their non-zero spectra agree and equal the squared singular values of Y_d^c.

The completion-interface vacuum is

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

so the canonical mass map at a matching point $\mu\star$ is

$$M_d^c(\mu\star) = (v_{cl}(\mu\star)/\sqrt{2}) \cdot Y_d^c(\mu\star).$$

1.4 Meaning of "current mass"

A **quark current mass** in this paper means:

a singular value of the renormalised canonical quark mass operator in a declared renormalisation scheme, at a declared scale, in a declared effective theory.

It is not:

- a free-particle pole mass for a light quark;
- a constituent quark mass;
- a hadron mass;
- an address label;
- a raw Yukawa entry;
- a support-space dimension;
- or a trace, unless a trace-to-singular-value theorem has been supplied.

1.5 Coupling-type firewall

The symbol α_{int} denotes the **per-gate interface readout coefficient** used in the candidate projection architecture.

The theorem does not infer this coefficient from an ordinary QED vertex. In conventional QED a vertex is weighted by e, while

$$\alpha = e^2/(4\pi).$$

Therefore the statement "each gate contributes α_{int}" is a distinct VERSF interface-readout premise. If the microscopic gate contributes e_{int}, $\sqrt{\alpha_{int}}$, $Q \cdot e_{int}$, or another coefficient, the candidate formula changes.

The identification

$$\alpha_{int} = \alpha(0)$$

is used only in the first numerical audit and remains a separate coupling-scale hypothesis.

1.6 Named projection premises

For audit precision, the load-bearing physical premises are named here and referenced throughout. Each carries a falsifier in Section 19.

- **P1 (Gate amplitude).** In compressed form on the selected lines, each interface gate acts as α_int times a partial isometry: P_in G_in P_d^R = α_int V_in and P_d^L G_out P_out = α_int V_out. Whether off-line channels affect the leading line is not asserted here; it is the separate closure premise P15. *(Falsifier F6.)*
- **P2 (Gate count and block structure).** Exactly two gates are necessary and sufficient to close the source-readable scalar mass record, and no direct canonical left–right attachment exists at the same order. *(Falsifiers F5, F26.)*
- **P3 (Charge blindness).** No electric-charge factor is inserted at either gate; charge resolution is carried by the triality structure, not re-weighted per gate. *(Falsifier F7; Debt D2.)*
- **P4 (Local-density readout).** The smooth scalar mass readout samples the local diagonal density of the normalised phase zero mode, not its integrated weight. *(Falsifier F8.)*
- **P5 (Unbiased binaries).** The chiral and weak-branch carriers exist as substrate readout registers distinct from the representation-theoretic doublet and chirality structure, and are unbiased two-state carriers before rank-one selection. *(Falsifier F9.)*
- **P6 (Colour–triality factorisation).** Local colour support and global triality resolution are independent commuting readout indices, giving a $\mathbb{C}^3 \otimes \mathbb{C}^3$ primitive support. *(Falsifier F11.)*
- **P7 (Additive quotient census).** The physical quotient of the nine primitive colour–triality attachment channels preserves additive census weight (c_L = c_R = 1) rather than Hilbert normalisation (c_L = c_R = 1/3). *(Falsifiers F12, F13.)*
- **P8 (Address anchoring).** The raw first down-type address projector Π_𝔇$_1$ is physically anchored, idempotent, and K_D-self-adjoint, so that its canonical image P_𝔇$_1$^c is an orthogonal, frame-covariant, rank-one projector whose relations to the Hermitian lift are frame-invariant. *(Falsifier F14.)*
- **P9 (Anchor ground line).** The anchored down address reduces H_d^R, carries its unique smallest eigenvalue, and all hierarchy operators act trivially on that line at anchor order. *(Falsifiers F15, F16.)*
- **P10 (Reducing-line transport).** The down singular line remains isolated along the coupled renormalisation flow Φ_d^𝑆 and transports multiplicatively on it, up to a mixing remainder bounded by half the running singular gap. *(Falsifiers F17, F18.)*
- **P11 (Type closure).** Every scalar factor in the coefficient is a gate coefficient, a spectral value of the internal response, or a declared compression expectation of that response, per Lemma 2B. *(Falsifier F25.)*
- **P12 (Internal-response factorisation).** The phase, chiral, weak-role, and colour–triality response algebras commute or factorise sufficiently that the compressed resolvent coefficient equals the product of their independently computed coefficients. *(Falsifier F27.)*
- **P13 (Lorentz-typed external-block recanonicalisation).** The same-side blocks of the Feshbach reduction vanish on the physical line, or have leading momentum dependence of Dirac kinetic form A_L(p^2) $\not{p}$ absorbed into independently computed positive metrics

whose canonical removal introduces no additional scalar factor, with every scalar same-chirality contribution B_L vanishing by gauge and chiral selection or higher order. *(Falsifier F28.)*

- **P14 (Canonical phase measure).** The closure phase carrier possesses a substrate-derived measure of total volume 2π, expressed in a physically fixed dimensionless coordinate, and the interface evaluation functional is normalised relative to that measure. *(Falsifier F29.)*
- **P15 (Off-line channel closure).** The internal response does not transport off-line gate components into the down line: $P_d^L G_{out} (1 - P_{out}) K_{int}^{-1} G_{in} P_d^R = 0$ and $P_d^L G_{out} K_{int}^{-1} (1 - P_{in}) G_{in} P_d^R = 0$, exactly or gap-bounded. *(Falsifier F30.)*
- **P16 (Internal spectral gap).** The internal response has a spectral gap at zero on the physical support, $\mathrm{dist}(0, \sigma(K_{int}|_{phys})) \geq \Delta_{int} > 0$, so that $\|K_{int}^{-1}\| \leq 1/\Delta_{int}$ and the compression is stable; since the internal weights are positive response values, the stronger condition $K_{int} > 0$ on the physical support is the natural target, and its derivation (or the derivation of the signature) is owed. *(Falsifier F31.)*

No premise is discharged by the numerical audit. Each requires an independent upstream theorem, and each has a named successor in the bridge-debt ledger of Section 18.

2. The Current-Mass Problem

The earlier one-anchor architecture has the form

$$m_q^{\mathcal{S}}(\mu) = R_q^{\mathcal{S}}(\mu) \cdot m_d^{\mathcal{S}}(\mu),$$

where the dimensionless ratios R_q carry the hierarchy and the down current mass supplies the dimensionful scale.

That architecture separates two questions:

1. Why do quark masses have their relative pattern?
2. Why is the first down-type current-mass coefficient of order a few MeV, rather than zero, the electroweak scale, or a confinement scale?

The present paper addresses the second question.

The candidate scalar kernel is

$$\kappa_0 = \alpha_{int}^2 \, \Pi_{cont}, \quad \Pi_{cont} = 1/(8\pi).$$

The proposed down coefficient is

$\eta_d^{(0)} = N_d \kappa_0, N_d = 9,$

so

$m_d^{(0)} = \eta_d^{(0)} v_{cl}.$

But three additional questions must be answered before this becomes a physical quark mass:

- Does $N_d = 9$ multiply one singular value, or merely count degeneracy?
- Does the first down-type address coincide with the smallest down-sector singular line?
- What renormalisation object does $m_d^{(0)}$ represent?

Those questions are now part of the theorem rather than afterthoughts.

---

3. The Canonical Down-Sector Mass Object

3.1 Raw action and canonical form

Write the down-sector quadratic and Yukawa terms schematically as

$\mathscr{L} \supset \bar{Q}_L i \not{D} K_Q Q_L + \bar{D}_R i \not{D} K_D D_R - \bar{Q}_L \Gamma_d \Phi_{cl} D_R + \text{h.c.},$

with $K_Q, K_D > 0$.

Under the canonical rescaling

$Q_L^c = K_Q^{1/2} Q_L, D_R^c = K_D^{1/2} D_R,$

the Yukawa map becomes

$Y_d^c = K_Q^{-1/2} \Gamma_d K_D^{-1/2}.$

After the completion interface condenses,

$\mathscr{L}_{mass} = - \bar{d}_L^c M_d^c d_R^c + \text{h.c.}, M_d^c = (v_{cl}/\sqrt{2}) Y_d^c.$

3.2 Canonical current-mass theorem

Theorem 1 — Canonical Current-Mass Spectrum

Let M_{d^c} be the canonical down-sector mass map. Then the positive down-sector mass magnitudes at the matching point are the singular values

$m_{di}(\mu\star) = \sigma_i(M_{d^c}(\mu\star)) = (|v_{cl}(\mu\star)| / \sqrt{2}) \sigma_i(Y_{d^c}(\mu\star))$.

Equivalently,

$m_{di}^2(\mu\star) = (|v_{cl}(\mu\star)|^2 / 2) \lambda_i(H_d^R) = (|v_{cl}(\mu\star)|^2 / 2) \lambda_i(H_d^L)$.

Proof

The singular-value decomposition gives

$Y_{d^c} = U_d \Sigma_d W_d^\dagger$,

where Σ_d is non-negative diagonal. Multiplication by $v_{cl}/\sqrt{2}$ scales every singular value by $|v_{cl}|/\sqrt{2}$. The eigenvalues of $Y_{d^c}^\dagger Y_{d^c}$ and $Y_{d^c} Y_{d^c}^\dagger$ are $\sigma_i^2(Y_{d^c})$. ■

Corollary 1.1 — Raw-entry no-go

No coordinate entry $(\Gamma_d)_{ij}$ or $(Y_{d^c})_{ij}$, and no row, column, or diagonal position, is a physical quark mass unless a symmetry-fixed frame and a one-dimensional reducing-subspace theorem make it equal to the corresponding singular value.

3.3 Current mass versus observable particle

The singular value above is a parameter of the renormalised field theory. It does not imply a freely propagating colour-triplet asymptotic state. The confinement firewall remains active:

current mass ≠ constituent response ≠ hadron mass.

---

4. The Chiral Scale-Carrier Theorem

4.1 Representation requirement

In the convention $Q = T_3 + Y$,

$Q_L : (\mathbf{3}, \mathbf{2}, +1/6)$, $d_R : (\mathbf{3}, \mathbf{1}, -1/3)$.

The bilinear $\bar{Q}_L d_R$ transforms as a weak doublet and carries hypercharge

$$-1/6 - 1/3 = -1/2.$$

Therefore a scalar carrier multiplying it must transform as

$$(\mathbf{1}, \mathbf{2}, +1/2),$$

which is the representation of the completion interface Φ_{cl}.

4.2 Theorem

Theorem 2 — Chiral Scale-Carrier Uniqueness

Assume the inherited minimal chiral matter skeleton contains no second scalar carrier with the same Standard-Model quantum numbers as Φ_{cl}. Then every renormalisable, gauge-preserving down-type Dirac mass attachment is proportional to

$$\bar{Q}_L \Phi_{cl} d_R.$$

Consequently, after symmetry breaking, the down-sector mass operator is linear in the completion-interface vacuum scale:

$$M_d^c = (v_{cl}/\sqrt{2}) Y_d^c.$$

Proof

Gauge invariance requires a colour-singlet, weak-doublet scalar with hypercharge +1/2 to contract $\bar{Q}_L d_R$. By the minimality premise, Φ_{cl} is the unique available carrier with those quantum numbers. Renormalisability makes the attachment linear in Φ_{cl}, and evaluation on the vacuum gives the displayed result. ■

What the theorem does not prove

It does not derive:

- the value of v_{cl};
- the dimensionless coefficient Y_d^c;
- or the relation between the VERSF completion scale and a particular renormalised electroweak convention.

It proves only the representation-typed origin of the dimensionful carrier.

4A. The Typed Effective Projection Kernel and Its Schur-Complement Origin

4A.1 The type problem

The candidate coefficient must arise from one operator construction. Gate amplitudes, phase densities, binary response weights, and support multiplicities may not be multiplied merely because they are dimensionless. They may multiply only if they occur as composable coefficients or spectral values inside the same effective chiral attachment map.

4A.2 Reduced attachment operator

Define the reduced canonical attachment operator

$$E_d^c \equiv (1/\sqrt{2})\, Y_d^c,$$

so that

$$M_d^c = v_{cl}\, E_d^c, \quad Y_d^c = \sqrt{2}\, E_d^c.$$

This notation isolates the coefficient appearing directly in the convention $m_d = \eta_d v_{cl}$.

4A.3 Block structure

Let $\mathscr{H}_L^d$ and $\mathscr{H}_R^d$ denote the canonical left and right down-type spaces, and let $\mathscr{H}_{int}$ be the intermediate source-readout space containing the phase, chiral, weak-role, and colour–triality support.

Consider the self-adjoint block operator on $\mathscr{H}_L^d \oplus \mathscr{H}_{int} \oplus \mathscr{H}_R^d$:

$$K_d = \begin{pmatrix} 0 & G_{out} & 0 \\ G_{out}^\dagger & K_{int} & G_{in} \\ 0 & G_{in}^\dagger & 0 \end{pmatrix}$$

where:

- $G_{in} : \mathscr{H}_R^d \to \mathscr{H}_{int}$ is the entry gate;
- $G_{out} : \mathscr{H}_{int} \to \mathscr{H}_L^d$ is the return gate;
- K_{int} is the invertible internal response operator on the physical intermediate support.

Eliminating the intermediate block gives the Schur-complement attachment

$$E_{d,eff}^c = -\, G_{out} K_{int}^{-1} G_{in}.$$

This is the operator-level meaning of the two-gate projection.

Evaluation point. The intermediate-space elimination of a self-adjoint block operator is in general energy-dependent: the effective external operator involves $(E - K_{int})^{-1}$, not K_{int}^{-1}. The reduction above is evaluated at the zero-frequency point of the external chiral problem, $E = 0$, which is the appropriate boundary for a static mass coefficient. This is a declared choice, not a triviality: if the substrate readout fixes a different reference frequency, the compressed coefficient becomes $r_d(E)$ and the projected value shifts. An unfixed E is registered as a hidden continuous direction in the identifiability audit (Section 16.4).

4A.4 Theorem

Theorem 2A — Typed Effective Projection-Kernel Theorem

Assume:

1. there is no direct canonical left–right attachment at the same order;
2. K_{int} is invertible on the selected physical support with a spectral gap at zero, $\operatorname{dist}(0, \sigma(K_{int}|_{phys})) \geq \Delta_{int} > 0$ (P16);
3. the only leading path from $\mathscr{H}_R^d$ to $\mathscr{H}_L^d$ passes through the intermediate source-readout sector;
4. the gate maps have the compressed selected-line forms

$$P_{in} G_{in} P_d^R = \alpha_{int} V_{in}, \quad P_d^L G_{out} P_{out} = \alpha_{int} V_{out} \text{ (P1)},$$

where P_{in} and P_{out} project the selected intermediate lines; 5. the compressed internal response satisfies

$$P_{out} K_{int}^{-1} P_{in} = r_d W_{int}$$

on the physical intermediate line, with W_{int} a partial isometry; 6. off-line channel closure holds (P15):

$$P_d^L G_{out} (1 - P_{out}) K_{int}^{-1} G_{in} P_d^R = 0, \quad P_d^L G_{out} K_{int}^{-1} (1 - P_{in}) G_{in} P_d^R = 0,$$

or both terms are bounded well below the down/strange singular gap in the sense of Section 8.6.

Then

$$P_d^L E_{d,eff}^c P_d^R = -\alpha_{int}^2 r_d V_d,$$

where

$$V_d = V_{out} W_{int} V_{in}$$

is the partial isometry between the physical right and left down singular lines. Under the gap-bounded alternative of assumption 6, the conclusion holds with a controlled remainder,

$$P_d^L E_{d,\mathrm{eff}}^c P_d^R = -\,\alpha_{\mathrm{int}}^2 r_d V_d + \Delta_{\mathrm{offline}},$$

with $\|\Delta_{\mathrm{offline}}\|_2$ bounded relative to the down/strange singular gap in the sense of Section 8.6 — matching the phrasing of falsifier F30.

If the internal-response calculation gives

$$r_d = 9/(8\pi),$$

then

$$\eta_d^{(0)} = 9\,\alpha_{\mathrm{int}}^2 / (8\pi),$$

and consequently

$$y_d^{(0)} = \sqrt{2}\,\eta_d^{(0)} = 9\sqrt{2}\,\alpha_{\mathrm{int}}^2 / (8\pi).$$

Proof

The Schur complement gives

$$E_{d,\mathrm{eff}}^c = -\,G_{\mathrm{out}} K_{\mathrm{int}}^{-1} G_{\mathrm{in}}.$$

Compress to the physical lines and insert the resolutions of identity $1 = P_{\mathrm{out}} + (1 - P_{\mathrm{out}})$ and $1 = P_{\mathrm{in}} + (1 - P_{\mathrm{in}})$ around K_{int}^{-1}:

$$P_d^L E_{d,\mathrm{eff}}^c P_d^R = -\,P_d^L G_{\mathrm{out}} \left[P_{\mathrm{out}} K_{\mathrm{int}}^{-1} P_{\mathrm{in}} + P_{\mathrm{out}} K_{\mathrm{int}}^{-1} (1 - P_{\mathrm{in}}) + (1 - P_{\mathrm{out}}) K_{\mathrm{int}}^{-1} \right] G_{\mathrm{in}} P_d^R.$$

The off-line channel closure conditions (assumption 6; P15) annihilate the second and third terms — without them, $G_{\mathrm{in}} P_d^R$ may carry a component in $(1 - P_{\mathrm{in}})\mathscr{H}_{\mathrm{int}}$, and K_{int}^{-1} can transport that off-line component back into the down line. Under the gap-bounded alternative, those two terms survive as $\Delta_{\mathrm{offline}}$ with the stated norm control. What remains factorises through the selected lines:

$$P_d^L E_{d,\mathrm{eff}}^c P_d^R = -\,(P_d^L G_{\mathrm{out}} P_{\mathrm{out}})(P_{\mathrm{out}} K_{\mathrm{int}}^{-1} P_{\mathrm{in}})(P_{\mathrm{in}} G_{\mathrm{in}} P_d^R) = -\,(\alpha_{\mathrm{int}} V_{\mathrm{out}})(r_d W_{\mathrm{int}})(\alpha_{\mathrm{int}} V_{\mathrm{in}}) = -\,\alpha_{\mathrm{int}}^2 r_d V_d. \quad \blacksquare$$

Sign remark. The overall minus sign is a field-phase convention: rephasing d_R absorbs it, and the physical coefficient is the singular value $|\alpha_{\mathrm{int}}^2 r_d|$.

4A.5 Reading

The theorem solves the type problem.

- α_{int} enters as a gate-map amplitude.

- $1/(2\pi)$, $1/2$, $1/2$, and the support-census coefficient enter as spectral values of the intermediate response.
- Their product is therefore a coefficient of one effective attachment operator.

They are not free-standing probabilities multiplied into an amplitude.

4A.6 The full Feshbach reduction and the external-block condition

The cross-chiral term is not the whole reduction. Eliminating the intermediate space from the stated self-adjoint block operator produces the full effective external operator on $\mathscr{H}_L^d \oplus \mathscr{H}_R^d$:

$$-\begin{pmatrix} G_{out} K_{int}^{-1} G_{out}^\dagger & G_{out} K_{int}^{-1} G_{in} \\ G_{in}^\dagger K_{int}^{-1} G_{out}^\dagger & G_{in}^\dagger K_{int}^{-1} G_{in} \end{pmatrix}$$

The off-diagonal blocks are the cross-chiral attachment used above. The two diagonal blocks cannot simply be omitted: they are same-side response corrections — self-energy and wave-function type terms — and if they alter the left and right kinetic metrics, the cross term must be recanonicalised, and that recanonicalisation could alter the coefficient.

Lorentz-Typed External-Block Recanonicalisation Condition (P13; Debt D18). The Feshbach reduction of K_d generates both cross-chiral and same-side response blocks. The projected coefficient in Theorem 7 is valid only after proving that the same-side blocks (a) vanish on the physical line by chirality, grading, or symmetry, (b) contribute only to already-defined kinetic operators and are fully removed by canonical normalisation, (c) are higher order than the cross attachment, or (d) are explicitly included in revised positive metrics K_Q and K_D whose canonical removal introduces no additional scalar factor.

Absorption into a kinetic metric is not automatic. At zero momentum a same-side block such as $G_{out} K_{int}^{-1} G_{out}^\dagger$ is simply an L → L operator; it becomes a kinetic renormalisation only if the full momentum-dependent self-energy has the Dirac form

$$\Sigma_L(p) = A_L(p^2)\, \not{p} + B_L(p^2).$$

The condition therefore additionally requires that the leading admissible same-side contribution is the $A_L(p^2)\, \not{p}$ part — which renormalises the positive kinetic metrics — and that every scalar same-chirality contribution B_L vanishes by gauge and chiral selection or is higher order. Without this Lorentz typing, "absorbed into the kinetic metric" is too broad to certify.

Until one of (a)–(d) is proved, the Schur-complement coefficient is conditional on the same-side blocks being harmless. This is registered as a named premise and falsifier, not treated as an implicit assumption.

4A.7 Dimensional ledger

The reduced attachment E_d^c is dimensionless, so the block dimensions must combine as

$[G_{out}] \cdot [K_{int}^{-1}] \cdot [G_{in}] = 1.$

The declared typing is: G_{in} and G_{out} are dimensionless amplitudes in substrate units, and K_{int} is a dimensionless response operator on the intermediate support. If a future substrate derivation assigns K_{int} a dimensionful spectrum, the gates must carry the compensating dimension explicitly, and the compensation must be exhibited — K_{int}^{-1} may not hide an unacknowledged scale inside a nominally dimensionless coefficient. This ledger is part of the identifiability audit of Section 16.

4B. Amplitude–Response Type Firewall

The projection architecture contains three distinct mathematical objects:

- **gate amplitudes**, carried by G_{in} and G_{out};
- **positive response weights**, carried by K_{int}^{-1};
- **singular coefficients**, carried by the final left–right attachment map.

These may not be identified casually.

For example, projection of the normalised coherent binary vector

$(|0\rangle + |1\rangle) / \sqrt{2}$

onto one branch has amplitude $1/\sqrt{2}$, whereas compression of the maximally mixed response operator

$\rho_2 = (1/2)\, I_2$

onto one rank-one branch has response eigenvalue 1/2.

Likewise, the normalised phase zero mode has amplitude

$u_0(\varphi) = 1/\sqrt{(2\pi)},$

but local response density

$|u_0(\varphi)|^2 = 1/(2\pi).$

The present construction uses the response-operator values $1/(2\pi)$, $1/2$, $1/2$ — not linear-overlap amplitudes.

Lemma 2B — Type-Closure Lemma

A scalar factor may enter the effective current-mass coefficient only if it is one of:

1. a scalar coefficient of a composable gate map;
2. an eigenvalue or singular value of the internal response operator;
3. an expectation value created by a declared compression of that response operator.

A bare probability, state count, density, or overlap may not be multiplied into the attachment coefficient unless an operator theorem places it in one of these categories (P11).

This lemma is the firewall preventing amplitude, probability, trace, and singular-value quantities from being mixed.

5. The Two-Gate Interface-Readout Theorem

5.1 Gate spaces

Let $\mathcal{H}_\Phi$ be the selected completion-interface line, $\mathcal{H}_J$ the source-readable current line, and $\mathcal{H}_M$ the persistent scalar mass-record line.

Define gate maps

$\mathcal{G}_{in} : \mathcal{H}_\Phi \to \mathcal{H}_J$, $\mathcal{G}_{out} : \mathcal{H}_J \to \mathcal{H}_M$.

Let V_{in} and V_{out} be partial isometries on the relevant one-dimensional channels.

The candidate gate premise (P1) is stated in compressed form on the selected lines:

$P_J \mathcal{G}_{in} P_\Phi = \alpha_{int} V_{in}$, $P_M \mathcal{G}_{out} P_J = \alpha_{int} V_{out}$,

where P_Φ, P_J, and P_M project the selected interface, current, and mass-record lines. Other channels may exist in these spaces; the premise constrains only the leading selected line, and is deliberately not asserted as a global operator equality. Whether those off-line channels leak into the compressed attachment through the internal response is a separate premise (P15, Section 4A.4), not a remark.

5.2 Theorem

Theorem 3 — Compressed Two-Gate Composition

If the selected source-readable mass channel factors through exactly two composable gates with compressed coefficients α_{int} (P1, P2), and if no intermediate projection, charge factor, or wave-function normalisation changes the selected line (P3), then the compressed composition is

$$P_M \mathcal{G}_{out} P_J \mathcal{G}_{in} P_\Phi = \alpha_{int}^2 V_{out} V_{in}.$$

The singular coefficient of the selected line is therefore α_{int}^2. The global equality $\mathcal{G}_{out} \mathcal{G}_{in} = \alpha_{int}^2 V_{out} V_{in}$ is deliberately **not** asserted: it would be stronger than the compressed hypotheses of P1 permit, and off-line behaviour is governed separately by P15.

Proof

Since P_J is idempotent, $(P_M \mathcal{G}_{out} P_J)(P_J \mathcal{G}_{in} P_\Phi) = P_M \mathcal{G}_{out} P_J \mathcal{G}_{in} P_\Phi = (\alpha_{int} V_{out})(\alpha_{int} V_{in})$. Partial isometries preserve the unit singular value on the selected line. ■

5.3 Physical content and burden

The algebra is exact but elementary. The physical theorem still owed must prove:

1. **necessity** — one gate leaves a source-visible but non-scalar record;
2. **sufficiency** — two gates close the scalar mass record;
3. **minimality** — a third gate is not independently required;
4. **gate coefficient** — the coefficient per gate is α_{int}, not e_{int}, $\sqrt{\alpha_{int}}$, $Q \cdot \alpha_{int}$, or another quantity;
5. **charge blindness** — electric charge is not inserted again at each gate if triality already carries the charge-resolution structure;
6. **no hidden normalisation** — the current line does not contribute a compensating Z_J factor.

The numerical match cannot prove any of these.

5.4 Charge-weighted failure

If each gate carries a factor proportional to $Q^2 \alpha_{int}$, then for a down quark $Q = -1/3$,

$$\Pi_{gate} = (\alpha_{int} / 9)^2 = \alpha_{int}^2 / 81,$$

and the proposed scale fails by roughly two orders of magnitude. Charge blindness (P3) is therefore a sharp structural test, not a stylistic choice.

5.5 Relation to the typed kernel

Theorem 3 states the coefficient algebra of the two gates in isolation. Section 4A embeds the same two gates as the off-diagonal blocks of the typed operator K_d, where the composition is mediated by the internal response: the physical composition is not the bare product $\mathcal{G}_{out} \mathcal{G}_{in}$ but the Schur complement $- G_{out} K_{int}^{-1} G_{in}$. The α_{int}^2 law survives unchanged; the internal factors are relocated from multiplied "weights" into the spectrum of K_{int}^{-1}, where Lemma 2B requires them to live.

---

6. The Zero-Mode Phase and Binary-Selection Theorem

The original $1/(8\pi)$ proposal is strengthened here by distinguishing a normalised phase average from a local zero-mode density.

A normalised average of the constant function over S^1 equals one and does **not** generate $1/(2\pi)$. The factor $1/(2\pi)$ arises only if the physical readout samples the local density of the normalised zero mode.

6.1 Access carrier

Let

$\mathcal{A}_d = L^2(S^1_\varphi) \otimes \mathbb{C}^2_{ch} \otimes \mathbb{C}^2_w$.

The normalised phase zero mode is

$u_0(\varphi) = 1/\sqrt{2\pi}$.

Let the pre-readout binary density operators be

$\rho_{ch} = (1/2)\, I_2$, $\rho_w = (1/2)\, I_2$, and let P_LR and P_d be rank-one projectors selecting the physical left–right bridge and the down weak branch.

6.2 Zero-mode evaluation and the canonical phase measure

Let the phase-zero-mode space be the one-dimensional Hilbert space

$\mathscr{H}\_\varphi^{(0)} = \mathrm{span}\{ u_0 \} \subset L^2(S^1, d\varphi)$,

with the zero mode normalised in $L^2(S^1, d\varphi)$, where the substrate-derived angular coordinate has period 2π:

$u_0(\varphi) = 1/\sqrt{(2\pi)}$,

and let

$P_0 = |u_0\rangle\langle u_0|$

be the zero-mode projector.

On the one-dimensional space $\mathscr{H}\_\varphi^{(0)}$, evaluation at the closure-selected phase $\varphi\star$ is a well-defined bounded functional

$\mathscr{E}\_\varphi\star\ u = u(\varphi\star)$.

Then

$\mathscr{E}\_\varphi\star\ P_0\ \mathscr{E}\_\varphi\star\dagger = |u_0(\varphi\star)|^2 = 1/(2\pi)$.

This is the precise operator source of the phase factor: not an average over φ, but the evaluation compression of the normalised zero-mode projector at coincident phase.

Lemma 4A — Evaluation-versus-Integration Dichotomy

For the normalised zero-mode projector,

$\mathrm{Tr}\_\{\mathscr{H}\_\varphi^{(0)}\}\ P_0 = 1$,

whereas

$\mathscr{E}\_\varphi\star\ P_0\ \mathscr{E}\_\varphi\star\dagger = 1/(2\pi)$.

Therefore:

- an integrated zero-mode readout gives 1;
- a local coincident-phase compression gives $1/(2\pi)$.

The Projection-Kernel Theorem (debt D4) must determine which readout the physical scalar mass channel implements.

Measure dependence. The numerical value $1/(2\pi)$ is not invariant under a rescaling of the phase coordinate or measure. With the coordinate measure $d\varphi$ on a circle of length 2π, the normalised constant mode is $u_0 = 1/\sqrt{(2\pi)}$ and its local density is $1/(2\pi)$. With the probability-normalised measure $d\mu = d\varphi/(2\pi)$, the normalised constant mode is instead $u_0 = 1$, and its local density is one. The factor therefore requires an additional owed result:

Canonical Phase-Measure Theorem (owed; P14, Debt D19). The closure phase carrier possesses a substrate-derived measure of total volume 2π, expressed in a physically fixed dimensionless coordinate, and the interface evaluation functional is normalised relative to that measure.

Until that theorem is supplied, the $1/(2\pi)$ factor is conditional on the substrate selecting the coordinate measure $d\varphi$ rather than $d\varphi/(2\pi)$; the choice is not a convention the paper is entitled to make.

Technical note. If the construction is later extended beyond the zero-mode line to all of $L^2(S^1)$, point evaluation is not automatically a bounded functional. It must then be replaced by a Sobolev-space evaluation with sufficient regularity, or by a normalised smeared interface functional. The present one-dimensional zero-mode restriction avoids that functional-analytic ambiguity.

6.3 Theorem

Theorem 4 — Phase–Binary Projection

Assume:

1. the smooth scalar mass readout samples the local diagonal density of the normalised phase zero mode (P4);
2. the phase measure is the substrate-derived coordinate measure of total volume 2π (P14);
3. the chiral and weak branch carriers are unbiased before selection (P5);
4. the physical chiral and down-branch readouts are rank one;
5. the phase, chiral, and weak factors are independent (P12).

Then

$$\Pi\_\mathrm{cont} = |u_0(\varphi)|^2 \cdot \mathrm{Tr}(\rho\_\mathrm{ch}\ P\_\mathrm{LR}) \cdot \mathrm{Tr}(\rho\_\mathrm{w}\ P\_\mathrm{d}) = (1/(2\pi)) \cdot (1/2) \cdot (1/2) = 1/(8\pi).$$

Proof

For a rank-one projector P on a two-dimensional unbiased state $I_2/2$,

$\mathrm{Tr}((1/2)\ I_2\ P) = 1/2$.

Independence makes the three readout weights multiplicative. ■

6.4 What remains open

The theorem is conditional on three physical identifications:

- why the mass coefficient reads the **local zero-mode density** rather than the integrated zero-mode weight;
- why the substrate selects the coordinate measure dφ of total volume 2π rather than a probability-normalised measure (P14);
- why the pre-readout chiral carrier is unbiased;
- why the pre-readout weak-role carrier is unbiased.

If any carrier is already polarised before the projection, the corresponding factor is not 1/2, and the kernel shifts by a computable amount.

There is a sharper danger than polarisation. In the conventional Yukawa construction, selecting the down component of the doublet costs nothing — vacuum alignment does it — and connecting left to right costs nothing — that is the Dirac bilinear itself; the electron mass is $y\,v/\sqrt{2}$ with no residual factor of 1/4. If the "weak-role" and "chiral" binaries were merely the representation-theoretic doublet-component selection and the L–R pairing, both 1/2 factors would double-count structure the canonical mass operator already possesses for free. The premise (P5) is therefore that the binary carriers are **substrate readout registers distinct from the representation-theoretic doublet and chirality structure**. If the discharge theorem (Debt D5) identifies the carriers with the representation structure instead, both factors equal 1 and the kernel fails by a factor of 4 — landing at ≈ 18.8 MeV, exactly as the ablation table of Section 15.2 records. The binaries must be shown to exist as independent structure, not merely to be unbiased. Independence carries a converse obligation: the registers must be auxiliary and nonpropagating — fully compressed or traced out at readout, and absent from the asymptotic one-particle census — or the construction predicts an enlarged matter spectrum (Debt D5A; falsifier F32).

6.5 Non-double-counting with α

Because $\alpha = e^2/(4\pi)$ already contains a conventional 4π, the phase factor $1/(2\pi)$ must arise from an independent closure-phase kernel, not from rewriting the gauge coupling. If the same angular normalisation is counted in both α_{int} and Π_{cont}, the formula double-counts phase-space normalisation.

This is a named audit condition (F10).

7. Additive Quotient Census and the Colour–Triality Coherence Theorem

7.1 The degeneracy no-go

Suppose an operator is the identity on an n-dimensional support:

$O = I_n$.

Then

$\mathrm{Tr}\, O = n$,

but every singular value equals one.

Therefore:

$\mathrm{Tr}\, O = n \nRightarrow \sigma(O) = n$.

A degeneracy count does not multiply a single-particle mass coefficient.

This matters immediately for colour. In an ordinary mass term,

$m_d \Sigma_a \bar{d}^a d^a$ (a = 1, 2, 3),

the three colour states share the same mass coefficient m_d; the coefficient is not multiplied by three merely because the action sums over colour.

Any VERSF factor from colour must therefore arise **before** the canonical one-particle coefficient is formed — and Section 7.5 makes precise what "before" must deliver.

7.2 Why the all-ones operator needs strengthening

An operator such as

$$J_9 = |\mathbf{1}_9\rangle\langle\mathbf{1}_9|, \quad |\mathbf{1}_9\rangle = \Sigma_A |A\rangle,$$

has eigenvalue nine on the normalised symmetric line, but it also contains non-zero matrix elements between every pair of primitive support labels. Its interpretation is therefore not simply "nine primitive contributions."

The cleaner construction begins with nine diagonal primitive attachments and then applies the physical quotient that identifies their output as one canonical down line.

7.3 Primitive bilateral census

Let the right primitive colour–triality support factor as

$$\mathscr{H}_{ct}^R \simeq \mathbb{C}^3_{colour} \otimes \mathbb{C}^3_{triality},$$

and similarly for $\mathscr{H}_{ct}^L$. Let

$$|A, R\rangle, \quad |A, L\rangle, \quad A = 1, \ldots, 9,$$

be orthonormal primitive right and left support states.

The primitive diagonal attachment is

$$D_9 = \Sigma_A |A, L\rangle\langle A, R|.$$

This operator has nine unit singular values. It represents nine orthogonal primitive attachments and, by itself, produces degeneracy rather than a coefficient nine times larger.

Let the internal quotient maps be

$$\tilde{q}_R : \mathscr{H}_{ct}^R \to \mathscr{L}_{ct}^R, \quad \tilde{q}_L : \mathscr{H}_{ct}^L \to \mathscr{L}_{ct}^L,$$

where $\mathscr{L}_{ct}^R$ and $\mathscr{L}_{ct}^L$ are selected internal census lines, spanned by $|\ell_R\rangle$ and $|\ell_L\rangle$. The additive census premise (P7) is

$$\tilde{q}_R |A, R\rangle = |\ell_R\rangle, \quad \tilde{q}_L |A, L\rangle = |\ell_L\rangle$$

for every primitive label A. Then

$$\tilde{q}_L D_9 \tilde{q}_R^\dagger = \Sigma_A |\ell_L\rangle\langle\ell_R| = 9 W_{ct}, \quad W_{ct} = |\ell_L\rangle\langle\ell_R|,$$

with $W_{ct} : \mathscr{L}_{ct}^R \to \mathscr{L}_{ct}^L$ an **internal** partial isometry. This is the exact ninefold enhancement, produced inside the intermediate space.

Internal versus external. The census must not be simultaneously internal and external, and the internal form above is the primary one: it supplies the colour–triality factor of the internal response (Section 9.3), and the external down-line map arises only after the two gates act, $V_d = V_{out} W_{int} V_{in}$ (Theorem 2A). The fully collapsed external representation

$$q_L D_9 q_R^\dagger = 9 V_d, \quad q_{\{L,R\}} = (\text{gate-composed selected-line isometries}) \circ \tilde{q}_{\{L,R\}},$$

remains available as equivalent bookkeeping of the same enhancement, but it is not itself a factor of K_{int}^{-1}.

7.4 Theorem

Theorem 5 — Additive Colour–Triality Quotient Theorem

Assume:

1. the primitive colour–triality support contains nine physically distinct pre-quotient attachment channels (P6);
2. their primitive attachment coefficients are equal;
3. only matching left/right primitive labels attach at leading order;
4. the faithful physical quotient removes the primitive labels while preserving additive current-readout weight (P7);
5. the nine quotient images coincide on the same internal census lines $\mathscr{L}_{ct}^{\{L,R\}}$, which the gates subsequently carry to the canonical down singular lines.

Then the quotient attachment is

$$\tilde{q}_L D_9 \tilde{q}_R^\dagger = 9 W_{ct}.$$

Hence the colour–triality current-readout coefficient is

$$N_d = 9.$$

Proof

Each of the nine primitive diagonal maps becomes the same rank-one internal map after quotienting. Their coefficients add linearly, yielding nine copies of W_{ct}. ■

Grading. Theorem 5 is an exact linear-algebra identity for the stated non-isometric maps; the physical adjoint and quotient category remain open (Section 7.6A).

7.5 The normalisation dichotomy

The factor nine is not forced by the existence of nine supports alone. It depends on what the quotient preserves.

If instead the quotient is Hilbert-normalised,

$$\tilde{q}_R^{norm} |A, R\rangle = (1/3) |\ell_R\rangle, \quad \tilde{q}_L^{norm} |A, L\rangle = (1/3) |\ell_L\rangle,$$

then

$$\tilde{q}_L^{norm} D_9 (\tilde{q}_R^{norm})^\dagger = W_{ct},$$

not $9 W_{ct}$. Therefore there are two sharply distinct possibilities:

additive census quotient $\Longrightarrow N_d = 9$, Hilbert-normalised quotient $\Longrightarrow N_d = 1$.

The decisive physical question is not "are there nine supports?" It is:

Does current-mass readout preserve additive primitive census weight, or normalised Hilbert amplitude, when the colour–triality labels are quotiented?

This is promoted to the principal content of Debt D7.

7.6 What kind of map is the additive quotient?

The additive map $\tilde{q}$, with $\tilde{q}|A\rangle = |\ell\rangle$ for all nine labels, has operator norm

$\|\tilde{q}\| = 3$, since $\tilde{q}\, \tilde{q}^\dagger = 9 |\ell\rangle\langle\ell|$.

It is not an isometry, not a co-isometry, and not an ordinary normalised Hilbert-space quotient map. That does not make it invalid, but it means the word "quotient" is doing substantial physical work, and its category must be declared. The candidate typings are:

- a current summation functional;
- a push-forward of counting measure;
- a coarse-graining map between weighted Hilbert spaces;
- a counit of a finite support algebra;
- an action-level aggregation rather than a state-space identification.

The distinction matters because physical state normalisation naturally favours the $1/\sqrt{n}$ weighting, while an extensive current or action density may legitimately preserve additive weight. The paper's position is therefore stated exactly:

The additive map is not asserted to preserve state norm. It is asserted to preserve an extensive current or action coefficient. Its adjoint and normalisation must therefore be derived from the bilinear pairing or support measure governing the effective action, rather than from the one-particle Hilbert norm.

This pre-empts the objection that $\tilde{q}$ has been made non-unitary in order to manufacture the factor nine: the non-unitarity is a declared property of an extensive readout, owed a derivation (Debt D7), not a hidden choice.

7.6A Fixing the adjoint: two admissible formulations

Theorem 5 uses $\tilde{q}^\dagger$, the composite $\tilde{q}_L D_9 \tilde{q}_R^\dagger$, and its singular value as though $\tilde{q}$ were an ordinary bounded Hilbert-space operator with a specified adjoint. That is mathematically exact only after the inner products or bilinear pairings have been fixed. Two clean formulations are admissible, and the census theorem (Debt D7) must select one.

Option A — Weighted Hilbert spaces. Define counting-weighted source spaces and current-weighted output spaces explicitly, with declared inner products $\langle\cdot,\cdot\rangle_{\text{count}}$ and $\langle\cdot,\cdot\rangle_{\text{current}}$. Then $\tilde{q}^\dagger$ is the adjoint relative to these metrics, and the factor nine is derived inside that declared category.

Option B — Action bilinear. Do not call $\tilde{q}^\dagger$ a Hilbert adjoint at all. Introduce separate push-forward and pull-back maps, $\tilde{q}_*$ and $\tilde{q}^*$, defined by the action pairing, and write

$$\tilde{q}_{L*} D_9 \tilde{q}_R^{*} = 9 W_{ct}.$$

This formulation fits the declared physics most naturally: the quotient preserves an extensive action/current coefficient, not state norm.

Until the category is fixed, Theorem 5 carries the grading stated beneath its proof — exact linear algebra for the stated maps; physical adjoint and category open.

7.7 Equivariance strengthening

Let

$$G = S_3^{\text{colour}} \times S_3^{\text{triality}}$$

act transitively on the nine primitive labels.

If the primitive attachment and quotient maps are G-equivariant, all primitive coefficients are equal. Equivariance therefore removes arbitrary unequal weighting.

It does not decide additive versus normalised quotienting. That final normalisation remains the load-bearing current-census theorem.

7.8 The colour–triality independence debt

The chiral matter representation theorem supplies a threefold colour-support skeleton and a third-unit charge structure. It does not by itself prove an independent 3 × 3 tensor factor.

The required microscopic statement is:

$\mathcal{A}_{\text{colour}}$ and $\mathcal{A}_{\text{triality}}$ are commuting primitive readout algebras whose joint support contains nine independently counted pre-quotient channels, while the final physical quark line is their quotient with a derived normalisation.

This is one of the two most exposed premise pairs in the paper (P6, P7; debts D6, D7), and the additive-versus-normalised dichotomy of Section 7.5 is now its sharpest single question.

8. Metric-Compatible Anchoring and Down Address–Spectral Alignment

The address projector originates before canonical kinetic normalisation, while the Hermitian lift is defined after canonicalisation. Their comparison therefore requires a metric-compatible transport step, which is installed first.

8.1 Metric-compatible address anchoring

Let $\Pi_{\mathfrak{D}_1}$ be the raw first-down address projector on the right-terminal space with kinetic metric $K_D > 0$. Define the metric adjoint

$A^{\ddagger} = K_D^{-1} A^{\dagger} K_D.$

The raw address projector is physically admissible only if

$\Pi_{\mathfrak{D}_1}^2 = \Pi_{\mathfrak{D}_1}, \ \Pi_{\mathfrak{D}_1}^{\ddagger} = \Pi_{\mathfrak{D}_1}.$

Its canonical image is

$$P_{\mathfrak{D}_1}^{c} = K_D^{1/2}\, \Pi_{\mathfrak{D}_1}\, K_D^{-1/2}.$$

Lemma 6A — Canonical Address-Projector Lemma

If $\Pi_{\mathfrak{D}_1}$ is idempotent and K_D-self-adjoint, then $P_{\mathfrak{D}_1}^{c}$ is an ordinary orthogonal projector in canonical coordinates.

Proof

Idempotency follows by similarity transformation. K_D-self-adjointness means $\Pi_{\mathfrak{D}_1}^{\dagger} = K_D\, \Pi_{\mathfrak{D}_1}\, K_D^{-1}$, so

$$(P_{\mathfrak{D}_1}^{c})^{\dagger} = K_D^{-1/2}\, \Pi_{\mathfrak{D}_1}^{\dagger}\, K_D^{1/2} = K_D^{1/2}\, \Pi_{\mathfrak{D}_1}\, K_D^{-1/2} = P_{\mathfrak{D}_1}^{c}. \blacksquare$$

The address–spectrum theorem below compares $P_{\mathfrak{D}_1}^{c}$ — not an uncanonicalised coordinate projector — with the spectral projector P_d^R.

8.2 Frame covariance, not coordinate invariance

A physical address projector is intrinsically defined but frame-covariant. Under a canonical generation-frame rotation U_R,

$$P_{\mathfrak{D}_1}^{c} \mapsto U_R^{\dagger}\, P_{\mathfrak{D}_1}^{c}\, U_R, \quad H_d^R \mapsto U_R^{\dagger}\, H_d^R\, U_R.$$

Therefore the following statements are frame-invariant:

$$[\, P_{\mathfrak{D}_1}^{c} , H_d^R \,] = 0, \quad P_{\mathfrak{D}_1}^{c} = P_d^R, \quad \mathrm{Tr}(\, P_{\mathfrak{D}_1}^{c}\, P_d^R \,).$$

Accordingly, the anchoring premise (P8) demands an intrinsically defined, frame-covariant address projector whose relations to the Hermitian lift are frame-invariant — not a "frame-independent" matrix, which does not exist.

8.3 Address and spectral projectors

Let $P_{\mathfrak{D}_1}^{c}$ be the canonical image of the physically anchored raw projector $\Pi_{\mathfrak{D}_1}$ onto the first down-type right-terminal address in the VERSF occupancy ledger.

It is not enough to define the address as the first coordinate basis vector. The anchoring theorem must supply an intrinsic, frame-covariant Π_𝔇₁ satisfying the admissibility conditions of Section 8.1 (P8).

Let P_d^R be the spectral projector of H_d^R associated with its smallest simple eigenvalue, and let P_d^L be the corresponding projector of H_d^L.

8.4 Exact alignment theorem

Theorem 6 — Down Address–Spectral Alignment

Assume:

1. Π_𝔇₁ is physically anchored, idempotent, and K_D-self-adjoint, so that its canonical image P_𝔇₁^c is an orthogonal projector (Lemma 6A; P8);
2. P_𝔇₁^c has rank one;
3. it reduces the right Hermitian lift,

[P_𝔇₁^c , H_d^R] = 0;

4. its eigenvalue is the unique smallest eigenvalue of H_d^R (P9).

Then

P_𝔇₁^c = P_d^R,

and the identity is frame-invariant in the sense of Section 8.2.

Proof

A rank-one orthogonal projector commuting with a Hermitian operator projects onto an eigenspace. By uniqueness of the smallest eigenvalue, that eigenspace is the range of P_d^R. Frame-invariance follows because both sides transform by the same conjugation. ■

8.5 Left-line transport

On a non-zero simple singular line define the partial isometry

V_d = (1/y_d) P_d^L Y_d^c P_d^R,

so that

$$V_d^\dagger V_d = P_d^R, \quad V_d V_d^\dagger = P_d^L.$$

The projected canonical Yukawa component is then

$$P_d^L Y_d^c P_d^R = y_d V_d.$$

8.6 Approximate alignment and stability

Suppose the unperturbed operator has a down singular value $y_d^{(0)}$ separated from the next singular value by

$$g_d = y_s^{(0)} - y_d^{(0)} > 0.$$

Let a non-factorisable correction satisfy

$$\|\Delta Y_d\|_2 < g_d / 2.$$

Then Weyl's inequality gives

$$| \sigma_i(Y_d^c + \Delta Y_d) - \sigma_i(Y_d^c) | \leq \|\Delta Y_d\|_2,$$

so the perturbation cannot close the gap or reorder the down and strange lines. For the singular vectors, Wedin's sin Θ theorem bounds the rotation of the down singular subspaces by

$$\max(\|\sin \Theta_R\| , \|\sin \Theta_L\|) \leq 2 \|\Delta Y_d\|_2 / g_d,$$

so the address–spectral identification degrades continuously, and controllably, with the declared perturbation, provided the singular gap is declared and non-zero.

The paper does not use a raw "small commutator" as an alignment certificate. The perturbation must be measured relative to the singular gaps.

8.7 Anchor-first is not normalisation-first

Choosing m_d as the unit of a ratio table does not prove that the down address is the undressed physical ground line. The candidate theorem requires a physical statement:

the first down address is the minimal stable colour-triplet chiral mass line, and all hierarchy operators act trivially on that line at anchor order.

This is the **Anchor Ground-Line Premise** (P9). Without it, the projection kernel may acquire an additional factor $\kappa\_d \neq 1$, and Section 16 forbids fitting such a factor after the fact.

---

9. The Interface-Projected Down-Boundary Theorem

9.1 Primitive coefficient

Define

$\kappa_0 = \alpha\_int^2 \, \Pi\_cont = \alpha\_int^2 / (8\pi)$.

Under coherent colour–triality aggregation,

$\eta\_d^{(0)} = 9 \, \kappa_0 = 9 \, \alpha\_int^2 / (8\pi)$.

The equivalent Standard Model Yukawa coefficient is

$y\_d^{(0)} = \sqrt{2} \, \eta\_d^{(0)} = 9\sqrt{2} \, \alpha\_int^2 / (8\pi)$.

9.2 Operator form

In the typed kernel of Section 4A, the compressed internal response on the physical intermediate line is

$P\_out \, K\_int^{-1} \, P\_in = r\_d \, W\_int, \; r\_d = \Pi\_cont \cdot N\_d = (1/(8\pi)) \cdot 9 = 9/(8\pi)$,

where Π_cont is the spectral value assembled in Theorem 4 and N_d is the additive quotient coefficient of Theorem 5, realised through the internal census

$\tilde{q}\_L \, D_9 \, \tilde{q}\_R^\dagger = 9 \, W\_ct$

inside the intermediate space; the external down-line partial isometry arises only after the gates act, $V\_d = V\_out \, W\_int \, V\_in$.

Theorem 2A then gives the projected line operator

$P\_d^\wedge L \; E\_d,eff^\wedge c \; P\_d^\wedge R = - \alpha\_int^2 \, r\_d \, V\_d = - (9 \, \alpha\_int^2 / (8\pi)) \, V\_d$, so, up to the field-phase convention absorbing the sign,

$Y_{d,0}^{\mathrm{proj}} = \sqrt{2}\,|\alpha_{\mathrm{int}}^2 r_d|\,V_d = (\,9\sqrt{2}\,\alpha_{\mathrm{int}}^2 / (8\pi)\,)\,V_d$.

Every factor now occupies a declared operator role: α_{int}^2 is the two-gate amplitude, and $9/(8\pi)$ is a singular coefficient of the compressed internal response.

9.3 The internal-response factorisation premise

Theorem 4 computes $(1/(2\pi)) \cdot (1/2) \cdot (1/2)$, and Theorem 5 computes the additive quotient coefficient nine. What remains to be proved is that these separate operators are factors of the **same** K_{int}^{-1}.

The required statement is an explicit internal-response factorisation,

$K_{\mathrm{int}}^{-1} = R_\varphi \otimes R_{\mathrm{ch}} \otimes R_w \otimes R_{\mathrm{ct}}$,

or an appropriate compressed equivalent, with the selected internal partial isometry factorising in the typed form

$W_{\mathrm{int}} = W_\varphi \otimes W_{\mathrm{ch}} \otimes W_w \otimes W_{\mathrm{ct}}$,

followed by the compressed product identity

$P_{\mathrm{out}} K_{\mathrm{int}}^{-1} P_{\mathrm{in}} = [\,(\,\mathcal{E}_\varphi^\star R_\varphi \mathcal{E}_\varphi^{\star\dagger}\,) \cdot \mathrm{Tr}(\,\rho_{\mathrm{ch}} P_{LR}\,) \cdot \mathrm{Tr}(\,\rho_w P_d\,) \cdot N_d\,]\,W_{\mathrm{int}}$,
$N_d = \sigma(\,\tilde{q}_L D_9 \tilde{q}_R^\dagger\,) = 9$,

where each bracketed factor enters through its declared operator role — the evaluation compression, the two response traces, and the internal census through its singular coefficient directly — so that

$r_d = (1/(2\pi)) \cdot (1/2) \cdot (1/2) \cdot 9 = 9/(8\pi)$

on the selected internal partial isometry.

This is the content of premise P12. It is load-bearing: entanglement or cross-coupling between the phase, chiral, weak-role, and colour–triality sectors could change the compressed coefficient even when every individual marginal factor has the stated value. P12 therefore carries its own falsifier (F27) and successor theorem (Debt D17); it may not be silently absorbed into the word "independence."

The typed chain. The full architecture, with every map in its proper space:

$\tilde{q}_L D_9 \tilde{q}_R^\dagger = 9\,W_{\mathrm{ct}}$ (internal census),

$W_{int} = W_\varphi \otimes W_{ch} \otimes W_w \otimes W_{ct}$ (internal partial isometry),

$P_{out} K_{int}^{-1} P_{in} = (1/(2\pi)) \cdot (1/2) \cdot (1/2) \cdot 9 \cdot W_{int} = r_d W_{int}$, $r_d = 9/(8\pi)$,

$P_d^L E_{d,eff}^c P_d^R = - \alpha_{int}^2 r_d V_{out} W_{int} V_{in} = - (9 \alpha_{int}^2 / (8\pi)) V_d$, $V_d = V_{out} W_{int} V_{in}$.

Nothing on the left of the gates lives in the external space; nothing external is smuggled into K_{int}^{-1}.

9.4 Theorem

Theorem 7 — Interface-Projected Down Boundary

Assume:

1. the canonical chiral scale-carrier theorem (Theorem 2);
2. the typed Schur-complement kernel with no direct left–right attachment, gates as its off-diagonal amplitudes, and type closure of every factor (Theorem 2A, Theorem 3, Lemma 2B; P1–P3, P11, P15, P16);
3. the phase–binary spectral values (Theorem 4; P4, P5, P14) and the additive colour–triality quotient (Theorem 5; P6, P7), combined through the internal-response factorisation (Section 9.3; P12) into the compressed coefficient $r_d = 9/(8\pi)$;
4. the same-side Feshbach blocks vanish on the physical line or are absorbed into recanonicalised kinetic metrics without a residual scalar factor (Section 4A.6; P13);
5. exact down address–spectral alignment through the canonical address projector (Lemma 6A, Theorem 6; P8, P9);
6. reducing-subspace closure of the projected down line;
7. anchor-ground-line triviality of all hierarchy operators at this order (P9);
8. no empirical down-mass datum enters any primitive, projector, quotient weight, normalisation, or gate coefficient (Section 16).

Then the canonical down-sector Yukawa operator contains the one-dimensional component

$P_d^L Y_d^c(\mu\star) P_d^R = y_d^{(0)}(\mu\star) V_d$,

with

$y_d^{(0)}(\mu\star) = 9\sqrt{2}\ \alpha_{int}^2(\mu\star) / (8\pi)$

and the corresponding projected mass boundary

$m_d^{(0)}(\mu\star) = 9\ \alpha_{int}^2(\mu\star)\ v_{cl}(\mu\star) / (8\pi)$.

Proof

The scale-carrier theorem supplies $v_{\mathrm{cl}}/\sqrt{2}$. The typed kernel supplies $-\alpha_{\mathrm{int}}^2 r_d$ on the selected line, with the sign absorbed by field rephasing. The phase–binary spectral values and the additive quotient give $r_d = 9/(8\pi)$. Address–spectral alignment identifies the quotient line with the smallest down singular line. The singular coefficient of the compressed attachment is therefore $|\alpha_{\mathrm{int}}^2 r_d| = 9\,\alpha_{\mathrm{int}}^2/(8\pi)$, and the Yukawa and mass conventions give the displayed results. ■

9.5 Correct epistemic reading

The algebraic proof is short because every physical burden has been moved into the premises, where it is visible and individually falsifiable.

This theorem does **not** yet establish that

$m_d^{(0)} = m_d^{\overline{\mathrm{MS}}}(2\ \mathrm{GeV})$.

It establishes a candidate canonical matching coefficient.

9.6 Bounded non-factorisable remainder

At the next order write

$Y_d^c = y_d^{(0)} V_d + Y_{d,\perp}^c + \Delta Y_d^{nf}$,

where

$P_d^L Y_{d,\perp}^c = 0,\ Y_{d,\perp}^c P_d^R = 0$.

The scalar boundary is stable only if $\|\Delta Y_d^{nf}\|_2$ is smaller than the declared down/strange singular gap, in the sense of Section 8.6. Non-factorisable mixing may not be absorbed into α_{int}, Π_{cont}, or N_d.

10. Renormalised Current-Mass Projection

10.1 Three layers

A complete quark-mass statement has three layers:

projection kernel → renormalised running mass → observable colour-singlet quantities.

The present theorem supplies the first layer.

10.2 Coupled renormalisation-flow map

The full renormalisation evolution of a Yukawa or mass matrix is generally a coupled flow, not a fixed linear operator acting on M_d alone. Write

$M_d^{\mathcal{S}}(\mu) = \Phi_{d,\mu\leftarrow\mu\star}^{\mathcal{S}}[\ M_d^{c}(\mu\star);\ g_i(\mu\star),\ Y_u(\mu\star),\ Y_d(\mu\star),\ v_{cl}(\mu\star),\ \ldots\ ]$,

where $\Phi_d^{\mathcal{S}}$ is the solution map of the coupled renormalisation-group and threshold equations. It includes, as applicable:

- running of gauge and Yukawa couplings;
- running of the completion-interface vacuum parameter;
- wave-function and canonical-normalisation transport;
- threshold decoupling;
- effective-theory changes;
- operator mixing;
- QCD, QED, and electroweak anomalous dimensions;
- matching conventions.

The running current mass remains

$m_d^{\mathcal{S}}(\mu) = \sigma_{min}[\ M_d^{\mathcal{S}}(\mu)\ ]$.

10.3 Prescription and transport theorem

Definition 8 — Renormalised-Mass Prescription

Let $M_d^{c}(\mu\star)$ satisfy Theorem 7, and let the coupled flow map $\Phi_d^{\mathcal{S}}$ be independently fixed without target-mass insertion. The renormalised down-quark current mass is

$m_d^{\mathcal{S}}(\mu) = \sigma_{min}^{+}[\ \Phi_{d,\mu\leftarrow\mu\star}^{\mathcal{S}}[\ M_d^{c}(\mu\star);\ \ldots\ ]\ ]$,

where σ_{min}^{+} denotes the smallest **positive simple** singular value associated with the physically anchored down line. Equivalently, the down-sector map is assumed full rank on the physical generation space; if that rank assumption fails, the down line must be identified by anchoring rather than by minimality, and the failure falls under falsifier F15.

This statement is a prescription — the definition of the running current mass once the flow is supplied — not a theorem. The substantive transport result is Theorem 8A.

Theorem 8A — Reducing-Line RG Transport

Assume the anchored down singular value is positive and simple, and that the down singular line remains isolated throughout the flow (P10), with

$$P_d^L(\mu)\, M_d^{\mathcal{S}}(\mu)\, P_d^R(\mu) = Z_d^{\mathcal{S}}(\mu, \mu\star)\, P_d^L(\mu\star)\, M_d^c(\mu\star)\, P_d^R(\mu\star) + \Delta_d^{mix},$$

with

$$\|\Delta_d^{mix}\| < (1/2)\, g_d(\mu),$$

where $g_d(\mu)$ is the running down/strange singular gap. Then

$$\boldsymbol{m_d^{\mathcal{S}}(\mu) = |\, Z_d^{\mathcal{S}}(\mu, \mu\star)\, |\, m_d^{(0)} + O(\|\Delta_d^{mix}\|),}$$

and the physical down singular line remains continuously connected to the projected boundary line.

Proof

On the isolated reducing line the compressed flow is scalar, so the transported singular coefficient is $|Z_d^{\mathcal{S}}|\, m_d^{(0)}$. Weyl's inequality bounds the displacement caused by Δ_d^{mix} by its norm, and the gap condition prevents level crossing, so $\sigma_{\min}$ continues to track the transported line. ■

Theorem 8A is the precise condition under which the full coupled flow reduces to a scalar multiplicative running factor. The scalar factor $Z_d^{\mathcal{S}}$ is the sector evolution factor written $\mathcal{R}_d^{\mathcal{S}}$ elsewhere in this paper (for example in Theorem 10). Definition 8 supplies the object being transported; Theorem 8A supplies the substance.

10.4 The projection-scale trilemma

The candidate kernel $m_d^{(0)}$ can be physically interpreted in only one of three ways, and the paper refuses to choose among them by numerical preference.

Reading A — Low-energy matching value

The kernel directly equals a running mass in a specific scheme at a projection-selected low scale μ⋆. This requires deriving both the scheme and the scale from the readout itself.

Reading B — Upstream boundary value

The kernel is a boundary mass at an electroweak or completion scale and must be evolved downward. This reading is under direct numerical strain, and the audit is now stated explicitly rather than asserted. Using MS̄ QCD running with the one-loop mass anomalous dimension — evolution exponents 12/23 for n_f = 5 and 12/25 for n_f = 4 — the world-average coupling α_s(M_Z) = 0.1180, and threshold matching at m_b ≈ 4.18 GeV, the evolution factor for a light-quark mass from μ = M_Z down to μ = 2 GeV is ≈ 1.6; including known higher-order corrections it is ≈ 1.7–1.8. A boundary value of 4.695 MeV posited at the electroweak scale would therefore arrive at roughly 7.5–8.5 MeV at 2 GeV, far outside the accepted m_d^MS̄(2 GeV) ≈ 4.7 MeV range. The loop order, thresholds, scheme, and coupling inputs of this audit are declared here and sourced in Appendix E; refining the audit to four-loop running with full decoupling belongs to Debt D11.

Reading C — RG-invariant integration constant

The kernel is an RG-invariant mass parameter m̂_d, from which scheme masses are obtained by standard conversion factors. This is conceptually the cleanest reading, but it must be tested numerically; the near-equality to a 2 GeV MS̄ value would then be a coincidence unless the conversion factor explains it.

The projection-scale theorem (debt D12) must decide among A and C. Reading B is already under numerical strain.

10.5 Scheme-invariance firewall

A formula built from substrate invariants should not know a human-selected MS̄ reference point by accident. Therefore the strongest target is not

m_d$^{(0)}$ = m_d^MS̄(2 GeV),

but either

m_d$^{(0)}$ = m̂_d,

or

m_d$^{(0)}$ = m_d^𝓢⋆(μ⋆)

with (𝓢⋆, μ⋆) derived from the projection kernel.

---

11. Confinement and Observability Firewall

The down quark is a colour-triplet internal field sector. The current mass is a parameter in the quark Lagrangian, not the mass of an asymptotic free particle.

No claim in this paper licenses:

- a light-quark pole mass;
- a free down-quark state;
- direct observation of $m\_d^{(0)}$;
- a hadron mass equal to a sum of projected current masses.

Physical comparison requires gauge-invariant colour-singlet observables or lattice/field-theoretic extraction of the renormalised mass parameter.

The correct chain is

down address → down singular line → renormalised current mass → colour-singlet observable extraction.

---

12. Separation from Ordered χ and Generation Hierarchy

12.1 Absolute scale versus relative split

The current projection kernel supplies an absolute first down-type scale. The ordered-commitment χ structure supplies a relative up/down contrast:

$\chi\_g = \ln(m\_ug / m\_dg)$.

These must not be double-counted.

The first down anchor is constructed without χ:

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

A subsequent up/down relation may take the form

$m\_u1 = m\_d \cdot e\textasciicircum\chi_1$, provided the ordered-history χ readout, its magnitude, and its bridge to the Yukawa operator are independently derived.

12.2 The final-record no-go remains active

A symmetric final-record readout forces

$\chi = 0$.

Therefore the projection kernel itself must not be claimed to generate the up/down asymmetry merely because it contains colour, triality, or weak-branch data. The absolute kernel is common-mode; the relative split requires the reversal-odd ordered-history channel.

12.3 Heavy-generation extension

The present theorem does not justify multiplying the down anchor by every previously proposed quark ratio and calling the full spectrum derived.

For the up-type hinge, the factorised structure has the form

$$Y_u(\mu) = Z_L^{(-1/2)}(\mu)\, B_u\, f_\chi(\hat{\chi})\, Z_R^{(-1/2)}(\mu) + \Delta Y_u^{nf}(\mu).$$

A heavy-quark prediction requires:

- a derived $\hat{\chi}$ spectrum;
- a calibrated first gap;
- an access-to-mass law;
- sector normalisation;
- QCD/electroweak running;
- threshold matching;
- controlled non-factorisable remainders.

The down anchor does not discharge those debts.

13. Projection–Trace Compatibility and the Strange Sector

13.1 Trace-to-singular-value no-go

Let the candidate strange support space have dimension twenty and let

$O_s = I_{20}$.

Then

$\mathrm{Tr}\, O_s = 20$,

but

$\sigma_i(O_s) = 1$ for all i.

Therefore a trace result alone does not prove a strange current mass twenty times larger than the down current mass.

13.2 Additive support quotient

Let $|T, L\rangle$ and $|T, R\rangle$, $T = 1, \ldots, 20$, be orthonormal primitive strange support states, and define the primitive bilateral census

$D_{20} = \Sigma_T |T, L\rangle\langle T, R|$,

with internal quotient maps satisfying

$\tilde{q}_L |T, L\rangle = |\ell_{sL}\rangle$, $\tilde{q}_R |T, R\rangle = |\ell_{sR}\rangle$

for every T, where $|\ell_{sL}\rangle$ and $|\ell_{sR}\rangle$ span the internal strange census lines. Then

$\tilde{q}_L D_{20} \tilde{q}_R^\dagger = 20 W_{ct}^s$, $W_{ct}^s = |\ell_{sL}\rangle\langle \ell_{sR}|$,

an internal partial isometry, with the fully collapsed external form $q_L D_{20} q_R^\dagger = 20 V_s$ available after gate composition exactly as in Section 7.3. An additive support quotient can therefore produce a single singular coefficient twenty. A Hilbert-normalised quotient of the same twenty channels gives unit weight instead.

13.3 Theorem

Theorem 9 — Strange Singular-Value Compatibility

Let y_s^phys and y_d^phys denote the relevant canonical strange and down singular coefficients at the same matching point. Then

m_s / m_d = 20

is compatible with current-mass singular-value discipline **if and only if**

y_s^phys / y_d^phys = 20.

A sufficient operator realisation is the additive twenty-support quotient of Section 13.2 (internal form 20 W_ct^s; collapsed external form 20 V_s).

This is one sufficient construction. It is not mathematically necessary: another independently derived operator could also carry physical singular value twenty. What is necessary is the singular-value ratio; the additive quotient is the proposed structural realisation.

Proof

The current-mass discipline defines m_s and m_d as the corresponding singular values, so the mass-ratio statement is equivalent to the singular-coefficient statement; this gives the equivalence. The displayed quotient computation shows sufficiency of the additive realisation, and the normalised computation shows that the support count alone does not imply it. ■

13.4 Consequence for the strange programme

The strange confinement work has already sharpened:

- why the balanced 3 + 3 shell is selected;
- why occupied/exterior orientation gives twenty rather than ten;
- why internal ordering is quotiented;
- why a canonical invariant line exists inside each support boundary.

The remaining mass-specific bridge is now exact:

Does confinement implement an additive twenty-support quotient — q̃_L D_{20} q̃_R† = 20 W_ct^s — a Hilbert-normalised quotient with unit weight, or some other independently derived operator whose relevant singular value is twenty?

Only an operator with singular value twenty relative to the down unit satisfies Theorem 9. This is a cleaner target than a bare trace normalisation.

---

14. Shared-Kernel Lepton/Quark Relation

14.1 Matching-scale relation

Suppose a colourless charged lepton and the down quark use the same projection kernel

$\kappa_0 = \alpha_{int}^2 / (8\pi)$,

with coherent multiplicities

$N_e = 1$, $N_d = 9$.

At the same matching point,

$y_e^{(0)} = \sqrt{2}\,\kappa_0$, $y_d^{(0)} = 9\sqrt{2}\,\kappa_0$,

so

$y_d^{(0)}(\mu\star) / y_e^{(0)}(\mu\star) = 9$.

This is the true parameter-free shared-kernel statement.

14.2 RG-transport theorem

Theorem 10 — Shared-Kernel Ratio Transport

Let the electron and down coefficients satisfy $y_d^{(0)} = 9\, y_e^{(0)}$ at a common matching point $\mu\star$. Let their independently fixed running and matching factors be $\mathscr{R}_d$ and $\mathscr{R}_e$. Then

$m_d^{\mathcal{S}d}(\mu_d) / m_e^{\mathcal{S}e}(\mu_e) = 9 \cdot \mathscr{R}_d^{\mathcal{S}d}(\mu_d, \mu\star) / \mathscr{R}_e^{\mathcal{S}e}(\mu_e, \mu\star)$.

Proof

Apply each sector's renormalisation map to its matching coefficient and divide. ■

14.3 Consequence

A raw comparison of $m_d^{\overline{MS}}(2\ \mathrm{GeV})$ with the electron pole mass is not an exact test of nine. The two quantities belong to different renormalisation conventions and have different gauge running.

The exact falsifier is:

1. evolve both sectors to the same declared matching description;
2. remove their independently derived RG factors;
3. test whether the boundary-coefficient ratio equals nine.

The earlier raw ratio remains a useful phenomenological clue, not a scheme-free theorem.

---

15. Numerical Audit and Ablation Ledger

15.1 Candidate evaluation

Using

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

the dimensionless coefficient is

$\eta_d^{(0)} = 9\,\alpha_{\mathrm{int}}^2 / (8\pi) = 1.9069237 \times 10^{-5}$.

The Standard Model Yukawa convention gives

$y_d^{(0)} = \sqrt{2}\,\eta_d^{(0)} = 2.6967973 \times 10^{-5}$.

The projected kernel mass is

$m_d^{(0)} = \eta_d^{(0)}\, v_{\mathrm{cl}} = 4.6952208$ MeV.

The colourless kernel is

$m_0^{(0)} = \alpha_{\mathrm{int}}^2\, v_{\mathrm{cl}} / (8\pi) = 0.5216912$ MeV.

These are numerical audits of the candidate kernel, not final renormalised predictions.

15.1A Observable-normalised audit versus running boundary

The numerical value used in the first audit must be typed separately from a running completion vacuum.

Define the observable-normalised electroweak scale

$v_F = (\sqrt{2}\, G_F)^{-1/2}$,

and distinguish it from a scheme-dependent running completion parameter $v_{\mathrm{cl}}^{\mathcal{S}}(\mu)$. The audit input 246.21965 GeV is v_F.

The first numerical correspondence is therefore

$m_{d,\mathrm{audit}}^{(0)} = 9\, \alpha(0)^2\, v_F / (8\pi)$,

not yet

$m_d^{\mathcal{S}}(\mu\star) = 9\, \alpha_{\mathrm{int}}^{\mathcal{S}}(\mu\star)^2\, v_{\mathrm{cl}}^{\mathcal{S}}(\mu\star) / (8\pi)$.

The first formula is an observable-normalised hybrid audit using the Thomson-limit coupling and the Fermi-normalised electroweak scale. The second is the same-scheme running boundary that the projection and RG theorems must eventually derive.

This distinction prevents the numerical audit from quietly mixing an infrared coupling, an observable electroweak scale, and an $\overline{\mathrm{MS}}$ quark mass.

15.1B Colourless kernel audit

The colourless kernel must be placed against the electron, not silently omitted:

$m_0^{(0)} = 0.5216912$ MeV, $m_e^{\mathrm{pole}} = 0.5109989$ MeV,

$m_0^{(0)} / m_e^{\mathrm{pole}} = 1.0209$ — a 2.1% excess,

whereas the down kernel agrees with the conventional $m_d^{\overline{\mathrm{MS}}}(2\ \mathrm{GeV})$ range at the sub-percent level: quoted central values across recent editions lie within ≈ 4.67–4.70 MeV, bracketing the kernel to within ≈ 0.5% on either side.

This asymmetry is direct evidence for the projection-scale trilemma of Section 10.4, because the two audits pull in different directions. If the kernel is a low-scale value (Reading A), the electron comparison is strained at 2.1% while the down comparison is excellent. If it is an upstream boundary (Reading B), the electron comparison must be made against the running mass at the matching point, $m_e^{\overline{\mathrm{MS}}}(M_Z) \approx 0.486$ MeV, giving $0.5217/0.486 \approx 1.07$ — a ≈ 7% excess, worse than the 2.1% pole comparison — while the down-sector Reading-B audit of Section 10.4

moves the down kernel far from its target. No single reading currently makes both audits sub-percent, and the trilemma decision is therefore empirical, not aesthetic.

Both comparisons are graded — as everywhere in this paper — as audits, not derivations. Recording the 2.1% tension is required by the non-insertion discipline of Section 16: an unflattering number is evidence, and suppressing it would be a fit.

15.2 Discrete ablation

Holding α_int and v_cl at the audit values:

Modification	Candidate mass
N = 1	0.5217 MeV
N = 3	1.5651 MeV
N = 9	4.6952 MeV
N = 27	14.0857 MeV
one gate, $\propto \alpha_{int}$	order 10^2 larger
three gates, $\propto \alpha_{int}^3$	order 10^2 smaller
omit phase density $1/(2\pi)$	factor 2π larger
omit one binary 1/2	factor 2 larger
omit both binary factors	factor 4 larger

Each structural ingredient moves the candidate by a discrete, order-separating amount; there is no continuous dial by which a nearby wrong structure can be tuned onto the target. The discrete separation makes the candidate rigid. It does not make it derived.

15.3 Why the match is not statistical proof

A structured menu of small integers, powers of a small coupling, and angular factors can generate accidental near-hits. The numerical value is therefore evidence for prioritising the discharge theorems, not evidence that they are true.

The proof must run from the operator structure to the coefficient without consulting the down mass.

15.4 Coupling-scale sensitivity

Because the result is quadratic in α_{int}, to first order

$$\delta m_d^{(0)} / m_d^{(0)} = 2\, \delta\alpha_{int}/\alpha_{int} + \delta v_{cl}/v_{cl} + \delta N_d/N_d + \delta\Pi_{cont}/\Pi_{cont}.$$

Replacing an infrared coupling by a materially larger electroweak-scale electromagnetic coupling shifts the result at the ten-percent level. The coupling identity and its scale are therefore load-bearing (P1; debt D2).

---

16. Non-Insertion and Identifiability Audit

16.1 Parameter partition

Partition all quantities into:

Prior structural inputs

$$\Theta_{prior} = \{ \Phi_{cl}, v_{cl}, \alpha_{int}, K_Q, K_D, P_{\mathfrak{D}_1}, \mathcal{G}_{in}, \mathcal{G}_{out}, P_{LR}, P_d, \mathscr{H}_{colour}, \mathscr{H}_{triality} \}.$$

QCMP-local constructions

$$\Theta_{QCMP} = \{ u_0, \mathscr{E}_{\varphi\star}, K_{int}, D_9, q_L, q_R, V_d, \Phi_d^{\mathcal{S}} \}.$$

Empirical target class

$$\Theta_{emp} = \{ m_d^{\overline{MS}}(\mu), m_e, m_s, \text{quark mass ratios, hadron masses} \}.$$

The derivation requires

$$\Theta_{emp} \cap (\Theta_{prior} \cup \Theta_{QCMP}) = \emptyset$$

during construction.

16.2 No hidden continuous correction

A residual factor κ_d may not be introduced and fitted to make

$$m_d^{obs} = \kappa_d\, m_d^{(0)}$$

unless κ_d is independently derived and assigned a physical owner.

Otherwise the projection theorem reduces to a one-parameter fit.

16.3 No proxy fitting

The following may not be used to set a QCMP primitive:

- the down current mass;
- the electron mass;
- the strange/down ratio;
- the pion, kaon, nucleon, or constituent-quark scale;
- a lattice-extracted mass;
- a heavy-quark mass;
- a CKM element.

Comparison occurs only after the structure is frozen.

16.4 Sensitivity rank

Let the predicted boundary coefficient be

$$f(\Theta_{prior}, \Theta_{QCMP}).$$

If any unfixed continuous direction $\delta\Theta$ remains such that

$$Df \cdot \delta\Theta \neq 0$$

while all stated structural constraints remain satisfied, the coefficient is not identified.

The current candidate has no explicit continuous local parameter, but upstream ambiguity in α_{int}, the zero-mode normalisation, the quotient weights, or the readout scale functions as hidden freedom until discharged. The Feshbach evaluation energy is one such direction: the reduction of Section 4A.3 is taken at $E = 0$, and until that zero-frequency choice is derived from the static character of the mass boundary, an unfixed E is an undeclared continuous parameter inside $r_d(E)$.

16.5 Quotient-normalisation non-insertion audit

The choice between additive and normalised quotienting must be frozen before any mass comparison.

Let a general quotient carry weight c on every primitive label:

$$\tilde{q}^{(c)} |A\rangle = c |\ell\rangle.$$

Then

$$\tilde{q}_L^{(c_L)} D_9 (\tilde{q}_R^{(c_R)})^\dagger = 9 c_L c_R W_{ct},$$

so the predicted multiplicity is

$$N_d^{eff} = 9 c_L c_R.$$

The desired value $N_d = 9$ corresponds to

$c_L = c_R = 1$ (additive census),

while Hilbert-normalised quotienting corresponds to

$c_L = c_R = 1/3$

and gives $N_d = 1$.

The uniform-weight identity above already presupposes label-independent weights. The fully general hidden direction is per-label: with $\tilde{q}_L |A, L\rangle = c_{L,A} |\ell_L\rangle$ and $\tilde{q}_R |A, R\rangle = c_{R,A} |\ell_R\rangle$,

$$\tilde{q}_L D_9 \tilde{q}_R^\dagger = \left(\Sigma_A c_{L,A} \bar{c}_{R,A} \right) W_{ct},$$

so the unconstrained parameter space is nine complex weight pairs, not one. The equivariance argument of Section 7.7 is what collapses this to the uniform case; the audit therefore cites it as load-bearing rather than assuming uniformity silently.

Any unfixed weights constitute a continuous hidden fit direction. The Coherent Current-Census Theorem must therefore derive the quotient normalisation, not merely the support count. This audit exposes the exact continuous parameters that otherwise hide inside the phrase "coherent aggregation."

---

17. What Is Proved, Conditional, and Open

17.1 Exact mathematical results

The following are exact:

- canonical masses are singular values;
- raw entries are frame-dependent;
- the Schur complement of the typed block operator is − G_out K_int^{-1} G_in, and two scalar gate coefficients multiply within it;
- evaluation compression of the normalised zero-mode projector gives 1/(2π), while its trace gives 1;
- unbiased rank-one selection from a two-state carrier gives 1/2;
- an additive quotient of nine equal channels gives q̃_L D$_9$ q̃_R† = 9 W_ct on the internal census lines (collapsed external form 9 V_d);
- a Hilbert-normalised quotient of the same channels gives W_ct, and a general quotient gives 9 c_L c_R W_ct;
- the canonical image of an idempotent, K_D-self-adjoint raw projector is an orthogonal projector;
- exact address–spectral alignment follows from commuting rank-one reduction plus a simple lowest eigenvalue;
- singular values obey Weyl stability bounds and singular subspaces obey Wedin bounds;
- an additive quotient of twenty equal channels gives q̃_L D$_{20}$ q̃_R† = 20 W_ct^s, while D$_{20}$ itself has twenty unit singular values;
- RG-transported ratios carry the ratio of sector evolution factors.

17.2 Conditional physical theorems

The following are conditional:

- the completion interface is the unique physical mass carrier;
- each source gate contributes α_int (P1);
- exactly two gates are necessary and sufficient (P2);
- the phase readout is the local zero-mode evaluation, not the integrated weight (P4);
- chiral and weak carriers are unbiased at readout (P5);
- colour and triality form independent readout fibres (P6);
- the physical quotient is additive rather than Hilbert-normalised (P7);
- the first down address is physically anchored, idempotent, and K_D-self-adjoint in the raw metric (P8);
- the first down address is the smallest singular line (P9);
- the hierarchy operators act trivially on the anchor line (P9);
- the strange support module is carried by an additive quotient to singular value twenty;
- the projection kernel defines a particular RG boundary object, with the down line isolated along the coupled flow (P10);
- every factor is typed by the substrate readout in a Lemma 2B category (P11);
- the phase, chiral, weak-role, and colour–triality responses factorise inside one K_int^{-1} (P12);
- the same-side Feshbach blocks are absorbed without a residual scalar (P13);
- the substrate selects the coordinate phase measure of total volume 2π (P14);
- the internal response does not leak off-line gate components into the down line (P15);
- the internal response has a spectral gap at zero — or is positive — on the physical support (P16);

- the binary registers are nonpropagating auxiliaries absent from the one-particle census (P5; Debt D5A).

17.3 Open numerical derivations

Still open:

- v_cl from the closure potential;
- α_int and its scale from the coupling theorem;
- the internal response operator K_int and its compressed coefficient r_d;
- the full down-sector canonical Yukawa operator;
- the complete coupled RG and threshold map;
- the low-energy scheme/scale selected by the kernel;
- the full six-quark spectrum;
- the nonperturbative relation to colour-singlet observables.

---

18. Bridge-Debt Ledger

Debt	Required result	Natural successor
D1	Derive v_cl from the closure-interface potential	Closure-Potential Scale Theorem
D2	Derive the physical per-gate coefficient — including whether electric charge enters the gate (P3) — and show whether it is α_int, e_int, or another quantity	Interface Coupling-to-Readout Theorem
D3	Prove exactly two gates are necessary and sufficient	Two-Gate Current-Record Theorem
D4	Derive the local zero-mode density readout	Discrete-to-Continuous Projection-Kernel Theorem
D5	Derive unbiased chiral and weak pre-readout carriers	Binary Selection Theorem
D5A	Prove the chiral and weak binary response registers are auxiliary nonpropagating factors whose compression changes coefficients without enlarging the physical fermion spectrum	Binary Register Non-Doubling Theorem
D6	Prove independent colour and triality-resolution algebras	Colour–Triality Factorisation Theorem
D7	Prove the additive quotient census (c_L = c_R = 1) rather than Hilbert-normalised quotienting, including the categorical typing of the additive map q (Section 7.6)	Coherent Current-Census Theorem
D8	Anchor Π_$\mathfrak{D}_1$ physically as an idempotent, K_D-self-adjoint projector	Down Address-Anchoring Theorem

Debt	Required result	Natural successor
D9	Prove down address–spectral alignment and the singular gap	Down Ground-Line Theorem
D10	Compute Y_d^c from closure response	Full Yukawa Construction
D11	Derive the coupled flow $\Phi_d^{\mathcal{S}}$, thresholds, and the projection scale	VERSF Fermion RG Matching Theorem
D12	Determine whether the kernel outputs a scheme mass or an RG-invariant mass	Projection-Scale Theorem
D13	Prove the strange additive quotient $\tilde{q}_L D_{20} \tilde{q}_R^\dagger = 20 W_{ct}^s$ (or an equivalent singular-value-twenty realisation)	Strange Trace-to-Singular-Value Theorem
D14	Derive the ordered χ magnitude and its Yukawa bridge	Flavour-Magnitude Dynamics
D15	Connect current masses to colour-singlet observables without mass conflation	Confinement Matching Programme
D16	Construct K_{int} from the substrate and verify $P_{out} K_{int}^{-1} P_{in} = r_d W_{int}$ with $r_d = 9/(8\pi)$	Internal-Response Compression Theorem
D17	Prove the internal-response factorisation $K_{int}^{-1} = R_\varphi \otimes R_{ch} \otimes R_w \otimes R_{ct}$ (or a compressed equivalent) and exclude cross-sector entanglement corrections	Internal-Response Factorisation Theorem
D18	Prove the same-side Feshbach blocks have leading Dirac kinetic form $A_L(p^2)\,\not{p}$ absorbed into recanonicalised K_Q, K_D, with scalar same-chirality parts B_L vanishing by gauge and chiral selection	Lorentz-Typed External-Block Recanonicalisation Lemma
D19	Derive the substrate phase measure of total volume 2π in a physically fixed coordinate	Canonical Phase-Measure Theorem
D20	Prove the off-line no-leakage conditions of P15, or bound off-line transport relative to the singular gap	Selected-Line Compression Lemma (feeds D16)
D21	Prove the internal spectral gap Δ_{int}, or the positivity $K_{int} > 0$, on the physical intermediate support	Internal Spectral-Gap Theorem

19. Falsification Conditions

F1 — Fock/field failure

If the spinorial Fock reconstruction does not yield a valid chiral fermion bilinear, the current-mass operator is not physically typed.

F2 — Representation failure

If the inherited matter skeleton does not contain the required Q_L, d_R, and completion-interface representations, the scale-carrier theorem fails.

F3 — Competing scalar carrier

If a second independent scalar with the same quantum numbers contributes at the same order, completion-scale uniqueness fails.

F4 — Kinetic positivity failure

If K_Q or K_D is not positive on the physical sector, canonical normalisation and the ordinary singular-mass interpretation fail.

F5 — Gate-count failure (P2)

If one gate closes the record, or three gates are required, the α_int^2 law fails.

F6 — Gate-coupling failure (P1)

If the per-gate coefficient is not α_int, the candidate formula fails.

F7 — Charge-blindness failure (P3)

If a down charge factor enters each gate, the candidate scale is strongly suppressed and fails.

F8 — Phase-kernel failure (P4)

If the readout uses an integrated zero-mode average rather than local zero-mode density, the 1/(2π) factor fails.

F9 — Binary-carrier failure (P5)

If the chiral or weak pre-readout carriers are not unbiased, one or both 1/2 factors fail. More dangerously: if the carriers do not exist as substrate registers independent of the representation-theoretic doublet-component and chirality structure — which the canonical mass term already provides at no cost — both factors equal 1 and the kernel fails by a factor of 4.

F10 — Double-counting failure

If the 1/(2π) factor is already contained in the definition of α_int, the formula double-counts angular normalisation.

F11 — Colour–triality independence failure (P6)

If triality is not an independent readout fibre from colour, the 3 × 3 support fails.

F12 — Degeneracy/quotient failure (P7)

If the physical attachment remains the diagonal census D_9 — or any direct sum of nine orthogonal channels — without an additive quotient onto one line, the mass coefficient is not multiplied by nine.

F13 — Quotient-weight failure (P7)

If the quotient weights satisfy $c_L c_R \neq 1$ — in particular Hilbert normalisation $c_L = c_R = 1/3$ — then $N_d^{eff} = 9 c_L c_R \neq 9$, and any unfixed weight is a hidden continuous fit direction (Section 16.5).

F14 — Address anchoring failure (P8)

If $\Pi_{\mathfrak{D}_1}$ is only a coordinate projector, or fails idempotency or K_D-self-adjointness, its canonical image is not an orthogonal projector and the down identity is frame-dependent.

F15 — Address–spectral mismatch (P9)

If the first address does not align with the smallest singular line, the projected kernel cannot be called the down mass. Rank failure of the down-sector map on the anchored line — a vanishing down singular value — is the degenerate case of this falsifier.

F16 — Anchor dressing failure (P9)

If hierarchy or localisation operators act nontrivially on the first down line, an additional independently derived factor is required.

F17 — Non-factorisable remainder failure (P10)

If $\|\Delta Y_d^{nf}\|$ is comparable to the down/strange singular gap, the scalar boundary and the address identity are unstable.

F18 — RG interpretation failure (P10)

If no scheme-, scale-, or RG-invariant object corresponds to the kernel, numerical comparison with current masses is undefined.

F19 — High-scale boundary failure

If the kernel is asserted as an electroweak-scale mass but ordinary running produces an incompatible low-energy value, that reading is falsified.

F20 — Light-quark pole-mass failure

Any claim of a clean isolated light-quark pole mass violates the confinement firewall.

F21 — Trace/singular-value failure

If a support trace is treated as a current-mass eigenvalue without a coherent compression, the mass claim is invalid.

F22 — Shared-kernel raw-ratio failure

If m_d/m_e = 9 is compared across unlike renormalisation conventions without RG transport, the purported exact falsifier is invalid.

F23 — Empirical insertion

If any projection factor is selected from m_d, m_e, m_s, or another target observable, the theorem becomes retrodictive.

F24 — Hypercharge-convention failure

If $Q = T_3 + Y$ and $Q = T_3 + Y'/2$ conventions are mixed without conversion, the chiral representation audit is invalid.

F25 — Type-closure failure (P11)

If any scalar factor enters the coefficient without being a gate coefficient, a spectral value of the internal response, or a declared compression expectation, the construction violates Lemma 2B and is invalid.

F26 — Direct-attachment failure (P2)

If a direct canonical left–right attachment exists at the same order, the Schur-complement kernel is not the leading coefficient and the two-gate law fails.

F27 — Internal-response entanglement failure (P12)

If the phase, chiral, weak-role, and colour–triality sectors are entangled or cross-coupled inside K_int, the compressed resolvent coefficient can differ from the product of the marginal factors even when every individual marginal has the stated value, and $r_d \neq 9/(8\pi)$.

F28 — External-block failure (P13)

If the same-side blocks of the Feshbach reduction alter the left or right kinetic metrics such that recanonicalisation introduces a residual scalar factor on the down line — in particular, if a scalar same-chirality contribution $B_L(p^2)$ survives gauge and chiral selection — the projected coefficient of Theorem 7 is modified and the stated value fails.

F29 — Phase-measure failure (P14)

If the substrate selects a probability-normalised phase measure $d\varphi/(2\pi)$ — or any rescaled coordinate without a compensating normalisation — the local zero-mode density is not $1/(2\pi)$ and the phase factor fails.

F30 — Off-line leakage failure (P15)

If the internal response transports off-line gate components into the down line at a level comparable to the down/strange singular gap, the compressed coefficient is not $\alpha_{int}^2 r_d$ and Theorem 2A does not apply as stated.

F31 — Internal near-zero-mode failure (P16)

If an eigenvalue of $K_{int}|_{phys}$ approaches zero, $\|K_{int}^{-1}\|$ diverges and small substrate perturbations change r_d dramatically; the compressed coefficient is then unstable and the prediction fails as a prediction, whatever its central value.

F32 — Register-doubling failure (P5)

If the proposed binary registers add unresolved propagating degrees of freedom, the construction predicts an enlarged matter spectrum and fails.

---

20. QCMP-1 Closure Certificate

QCMP-1 condition	Result
Fermionic field carrier available	Inherited / conditional
Chiral down-type representation available	Inherited / conditional
Hypercharge convention fixed	Closed

QCMP-1 condition	**Result**
Raw-entry mass claims excluded	Closed
Canonical Yukawa map defined	Closed
Hermitian lifts and singular spectrum defined	Closed
Current mass typed by scheme and scale	Closed
Completion interface is the minimal scale carrier	Conditionally proved
Typed Schur-complement kernel (Theorem 2A)	Structurally closed; substrate return open
Full Feshbach reduction stated; Lorentz-typed external-block condition (P13)	Stated; open
Off-line channel closure (P15)	Open
Internal spectral gap or positivity of K_int (P16)	Open
Binary registers nonpropagating (D5A)	Open
Zero-frequency Feshbach evaluation declared	Closed as declaration; derivation open
Dimensional ledger for K_d declared	Closed as typing
Internal-response factorisation (P12)	Open
Type-closure lemma (Lemma 2B; P11)	Closed
Compressed two-gate composition (Theorem 3)	Exact
Exactly two physical gates, no direct attachment (P2)	Open
Per-gate coefficient equals α_int (P1)	Open
Zero-mode evaluation compression gives 1/(2π)	Exact (Lemma 4A) given the canonical phase measure
Evaluation-versus-integration readout derived (P4)	Open
Canonical phase measure derived (P14)	Open
Chiral and weak binary factors give 1/2 each	Exact given unbiased carriers
Unbiased carrier premise derived (P5)	Open
Multiplicity/degeneracy distinction	Closed
Internal additive census q̃_L D_9 q̃_R† = 9 W_ct	Exact linear algebra; adjoint category open
Internal/external census separation (W_ct versus V_d)	Closed as typing
General quotient dependence N_d^eff = 9 c_L c_R	Closed as identity
Independent colour–triality fibre (P6)	Open
Additive versus Hilbert-normalised quotient (P7)	Open — decisive
Categorical typing of the additive map q	Open
Canonical address-projector lemma (Lemma 6A)	Closed
Down address physically anchored (P8)	Open
Exact address–spectral alignment theorem	Closed conditionally
Anchor-ground-line premise (P9)	Open

QCMP-1 condition	Result
Candidate y_d$^{(0)}$ theorem	Conditional
Candidate m_d$^{(0)}$ theorem	Conditional
Numerical value 4.6952208 MeV	Audit only
Equality to MS̄ mass at 2 GeV	Not closed
Renormalised-mass prescription (Definition 8)	Closed as definition
Reducing-line coupled-flow transport (Theorem 8A; P10)	Conditional
High-scale boundary reading numerically audited (Reading B)	Disfavoured at stated loop order
Colourless kernel audited against m_e	2.1% tension recorded; audit only
Binary carriers independent of representation structure (P5)	Open — factor-4 exposure named
Complete coupled RG/matching theorem	Open
Observable-normalised audit typing (v_F versus v_cl^𝑆)	Closed as typing
Confinement/current/hadron distinction	Closed
Ordered χ separated from anchor	Closed
Trace-to-singular-value no-go	Closed
Strange additive quotient target q̃_L D$_{20}$ q̃_R† = 20 W_ct^s	Specified; open physically
Shared-kernel matching-scale ratio y_d/y_e = 9	Conditional
Raw measured mass ratio m_d/m_e = 9	Rejected without RG transport
Numerical non-insertion firewall	Closed as protocol
Full zero-anchor quark spectrum	Open

QCMP-1 grade

QCMP-1 is closed as a canonical current-mass target theorem, a raw-entry and multiplicity no-go, and a conditional projection-boundary architecture. The coefficient 9 α_int^2/(8π) is derived from the architecture only if the substrate returns one typed second-order Schur-complement kernel whose internal response has the required phase, binary, and additive colour–triality quotient coefficients. The remaining decisive debts are the physical gate coefficient, the two-gate block structure, local zero-mode evaluation, additive rather than normalised quotienting, metric-compatible down-address anchoring, and the scheme/RG interpretation of the resulting boundary.

---

21. Conclusion

The quark-mass programme cannot be completed by attaching dimensionless ratios to an observed down-quark anchor. The anchor itself must be derived, and it must be derived as the coefficient of the correct canonical chiral mass operator.

This paper has reformulated that task at the level required for a defensible theorem.

The down current-mass problem is not merely

find a number near 4.7 MeV.

It is the chain

fermionic chiral carrier → canonical Yukawa map → down singular line → interface projection kernel → renormalised current mass → colour-singlet observable extraction.

Within that chain, the candidate kernel is

$$\eta_d^{(0)} = 9\,\alpha_{int}^2 / (8\pi),$$

with

$$m_d^{(0)} = 9\,\alpha_{int}^2\, v_{cl} / (8\pi).$$

The structure of the factorisation is now precise.

- v_{cl} belongs to the unique chiral scale carrier.
- α_{int}^2 belongs to a two-gate interface-readout composition.
- $1/(8\pi)$ belongs to a local zero-mode density and two independent unbiased binary selections.
- 9 belongs not to colour degeneracy but to an additive colour–triality quotient census that carries nine equal primitive channels onto one canonical line with summed weight.

That last distinction is essential. A trace is not a mass. A dimension is not a singular value. A final-state degeneracy does not enhance a one-particle coefficient. The nine channels must be quotiented additively — $c_L = c_R = 1$, not the Hilbert-normalised 1/3 — before the canonical singular line is formed, and that normalisation must be derived, not chosen.

The factors are also now typed. The candidate coefficient is the Schur complement of one block operator: the gates carry α_{int} as amplitudes, and the phase, binary, and census factors are spectral values of the internal response. Nothing is multiplied into the coefficient without an operator role, and the type-closure lemma makes that discipline enforceable.

The same lesson sharpens the strange-sector programme. A twenty-dimensional support space with identity operator has trace twenty but unit singular values. A strange mass ratio of twenty requires the singular coefficients themselves to stand in ratio twenty; an additive support quotient — $\tilde{q}_L D_{20} \tilde{q}_R^\dagger = 20\, W_{ct}^s$ — is one sufficient realisation.

The paper also removes the strongest overstatement of the earlier candidate. The number 4.695 MeV is a candidate projection-kernel value, not yet $m_d^{\overline{MS}}(2\ \mathrm{GeV})$: the complete current mass is $m_d^{\mathcal{S}}(\mu) = \sigma_{\min}[\, \Phi_{d,\mu\leftarrow\mu\star}^{\mathcal{S}}[\, M_d^c(\mu\star); \ldots \,]\,]$, with the coupled flow, thresholds, matching scale, and scheme identity part of the theorem rather than optional corrections. And the shared electron/down kernel predicts nine only at a common matching point, so any measured comparison carries sector-specific RG transport.

The programme has therefore advanced in a meaningful way. The down anchor is no longer an unexplained external ruler, but neither has it been declared derived merely because a rigid formula lands near the right numerical range. It has been reduced to a finite set of operator-level questions:

gate coefficient + two-gate block structure + external-block recanonicalisation + zero-mode evaluation and phase measure + binary neutrality + additive quotient normalisation + internal-response factorisation + address anchoring and alignment + RG transport.

Every one can pass or fail independently. If they pass, the one-anchor grid can become a genuine zero-quark-anchor current-mass construction. If any fails, the failure is local and visible.

The decisive remaining question is no longer whether the numerical formula is attractive. It is whether the independently defined substrate operator returns

$$P_d^L E_{d,\mathrm{eff}}^c P_d^R = \left(9\, \alpha_{\mathrm{int}}^2 / (8\pi) \right) V_d,$$

up to the field-phase convention, as the Schur complement of its physical intermediate readout block. If it does, every factor belongs to one typed operator and the candidate becomes a genuine projection theorem. If it does not, the precise failed block — gate, phase, binary, quotient, alignment, or RG transport — identifies where the construction must be revised.

That is the result of QCMP-1.

Appendix A — Minimal Operator Skeleton

A.0 Typed block operator and Schur kernel

$$K_d = \left(0, G_{\mathrm{out}}, 0 \,;\, G_{\mathrm{out}}^\dagger, K_{\mathrm{int}}, G_{\mathrm{in}} \,;\, 0, G_{\mathrm{in}}^\dagger, 0 \right) \text{ on } \mathscr{H}_L^d \oplus \mathscr{H}_{\mathrm{int}} \oplus \mathscr{H}_R^d,$$

$$E_{d,\mathrm{eff}}^c = -\, G_{\mathrm{out}} K_{\mathrm{int}}^{-1} G_{\mathrm{in}} \text{ (cross block)},$$

with the full external Feshbach reduction

$$-\left(G_{\mathrm{out}} K_{\mathrm{int}}^{-1} G_{\mathrm{out}}^\dagger, G_{\mathrm{out}} K_{\mathrm{int}}^{-1} G_{\mathrm{in}} \,;\, G_{\mathrm{in}}^\dagger K_{\mathrm{int}}^{-1} G_{\mathrm{out}}^\dagger, G_{\mathrm{in}}^\dagger K_{\mathrm{int}}^{-1} G_{\mathrm{in}} \right),$$

whose same-side blocks are subject to the Lorentz-Typed External-Block Recanonicalisation Condition (P13), and

$P_d^L E_{d,\mathrm{eff}}^c P_d^R = -\alpha_{\mathrm{int}}^2 r_d V_d$, $r_d = 9/(8\pi)$,

conditional on the internal-response factorisation $K_{\mathrm{int}}^{-1} = R_\varphi \otimes R_{\mathrm{ch}} \otimes R_w \otimes R_{\mathrm{ct}}$ (P12), off-line channel closure (P15), and the internal spectral gap (P16).

A.1 Canonical down map

$Y_d^c = K_Q^{(-1/2)} \Gamma_d K_D^{(-1/2)}$.

A.2 Hermitian lifts

$H_d^L = Y_d^c Y_d^{c\dagger}$, $H_d^R = Y_d^{c\dagger} Y_d^c$.

A.3 Down singular component

$P_d^L Y_d^c P_d^R = y_d V_d$.

A.4 Compressed two-gate coefficient

$P_M \mathcal{G}_{\mathrm{out}} P_J \mathcal{G}_{\mathrm{in}} P_\Phi = \alpha_{\mathrm{int}}^2 V_{\mathrm{out}} V_{\mathrm{in}}$.

A.5 Continuous projection

$u_0(\varphi) = 1/\sqrt{(2\pi)}$, $\mathcal{E}_{\varphi\star} P_0 \mathcal{E}_{\varphi\star}^\dagger = |u_0(\varphi\star)|^2 = 1/(2\pi)$, $\mathrm{Tr}\, P_0 = 1$.

$\mathrm{Tr}(\, (1/2) I_2 P_{\text{rank-1}} \,) = 1/2$.

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

A.6 Additive quotient census

$D_n = \Sigma_a |a, L\rangle\langle a, R|$, $\tilde{q}\,|a\rangle = c\,|\ell\rangle \Longrightarrow \tilde{q}_L D_n \tilde{q}_R^\dagger = n\, c_L c_R W_{\mathrm{ct}}$.

Additive quotient ($c_L = c_R = 1$), on internal census lines:

$\tilde{q}_L D_9 \tilde{q}_R^\dagger = 9 W_{\mathrm{ct}}$, $\tilde{q}_L D_{20} \tilde{q}_R^\dagger = 20 W_{\mathrm{ct}}^s$.

Collapsed external forms after gate composition: 9 V_d, 20 V_s. Hilbert-normalised quotient ($c = 1/\sqrt{n}$): unit weight.

A.7 Projected Yukawa and mass

$y\_d^{(0)} = \sqrt{2} \cdot (\alpha\_\text{int}^2 / (8\pi)) \cdot 9 = 9\sqrt{2}\ \alpha\_\text{int}^2 / (8\pi)$.

$m\_d^{(0)} = (v\_cl/\sqrt{2})\ y\_d^{(0)} = 9\ \alpha\_\text{int}^2\ v\_cl / (8\pi)$.

A.8 Renormalised current mass

$M\_d\verb|^|\mathcal{S}(\mu) = \Phi\_d{,}\mu{\leftarrow}\mu{\star}\verb|^|\mathcal{S}[\ M\_d\verb|^|c(\mu{\star});\ \ldots\ ]$,

$m\_d\verb|^|\mathcal{S}(\mu) = \sigma\_\min(\ M\_d\verb|^|\mathcal{S}(\mu)\ )$,

and on an isolated reducing line,

$m\_d\verb|^|\mathcal{S}(\mu) = |Z\_d\verb|^|\mathcal{S}(\mu, \mu{\star})|\ m\_d^{(0)} + O(\|\Delta\_d\verb|^|mix\|)$.

Appendix B — Numerical Evaluation

Using

$\alpha\_\text{int} = 1/137.035999084$, $v\_cl = 246.21965$ GeV,

the successive factors are

$\alpha\_\text{int}^2 = 5.3251354 \times 10^{-5}$,

$\Pi\_cont = 1/(8\pi) = 0.0397887358$,

$\eta\_d^{(0)} = 9\ \alpha\_\text{int}^2\ \Pi\_cont = 1.9069237 \times 10^{-5}$,

$y\_d^{(0)} = \sqrt{2}\ \eta\_d^{(0)} = 2.6967973 \times 10^{-5}$,

$m\_d^{(0)} = \eta\_d^{(0)}\ v\_cl = 0.0046952208$ GeV = 4.6952208 MeV.

Colourless kernel:

$m_0^{(0)} = m\_d^{(0)} / 9 = 0.5216912$ MeV.

These values are arithmetic consequences of the candidate kernel and the audit inputs. They do not settle the renormalisation interpretation.

Appendix C — Referee Objections and Replies

Objection 1 — You have simply reverse-engineered $9\alpha^2 v/(8\pi)$

Reply. The numerical hit does not establish the theorem, and the paper never claims that it does. The claim is conditional and survives only if the gate, phase, binary, census, alignment, factorisation, and RG premises — the named premises P1–P16 of Section 1.6 — are independently derived. The numerical audit prioritises those calculations; it does not replace them.

Objection 2 — A QED vertex carries e, not α

Reply. Correct. The per-gate coefficient is not inferred from a standard QED vertex. It is a VERSF interface-readout coefficient denoted α_int. If the microscopic gate carries e_int, the formula changes and the candidate is falsified. This is F6, not a convention.

Objection 3 — The 1/(2π) factor is fake because a normalised phase average equals one

Reply. A normalised average does equal one. The paper does not use that average. It uses the diagonal density of the normalised zero-mode projector, $|u_0(\varphi)|^2 = 1/(2\pi)$, under a local coincident-phase readout. If the physical readout is integrated rather than local, the factor fails (F8).

Objection 4 — Colour degeneracy cannot multiply a quark mass

Reply. Correct, and the paper proves it: degeneracy does not multiply a singular value — the diagonal census D_9 has nine unit singular values. The factor nine requires an additive pre-canonical quotient that carries the nine primitive channels onto one singular line with summed weight. If the colour–triality labels survive orthogonally, the factor fails (F12).

Objection 5 — Triality is already part of colour, so 3 × 3 double-counts

Reply. That is a live objection, and the paper says so. The theorem requires independent commuting local-colour and global triality-resolution readout algebras. The chiral representation paper does not by itself prove that factorisation. D6 is therefore a primary debt.

Objection 6 — Why should the nine contributions add rather than be normalised to one?

Reply. They should not — unless the current-census theorem proves that the physical quotient preserves additive census weight. The general identity is $N_d^{eff} = 9 c_L c_R$: additive quotienting ($c = 1$) gives nine, Hilbert normalisation ($c = 1/3$) gives one, and any unfixed c is a hidden fit parameter (Section 16.5). The paper makes this burden explicit rather than hiding it inside "multiplicity."

Objection 7 — Choosing down as the anchor does not make it the physical ground state

Reply. Correct. Ratio normalisation is bookkeeping. The physical theorem additionally requires address anchoring, reducing-subspace alignment, and identification with the smallest simple singular line. Without those, the kernel is not a down mass.

Objection 8 — Why compare the formula with $m_d^{\overline{MS}}(2\ \mathrm{GeV})$?

Reply. Only as a numerical correspondence audit. The paper expressly rejects equality until the projection-scale and RG theorems determine the output's scheme and scale.

Objection 9 — Does the paper derive a free down-quark mass?

Reply. No. It derives a candidate renormalised Lagrangian coefficient. A light quark is confined and has no clean asymptotic pole-state interpretation.

Objection 10 — Is the electron/down ratio exactly nine?

Reply. The matching-scale boundary coefficients are in the ratio nine under the shared-kernel premise. Low-energy measured masses carry different RG and scheme factors. The exact transported relation is Theorem 10.

Objection 11 — Does a strange support trace of twenty prove $m_s/m_d = 20$?

Reply. No. The identity on twenty supports has unit singular values. What is required is that the canonical singular coefficients stand in ratio twenty; the additive quotient (internal form $\tilde{q}_L D_{20} \tilde{q}_R^{\dagger} = 20 W_{ct}^{s}$) is the proposed sufficient realisation (revised Theorem 9).

Objection 12 — Does this derive all quark masses?

Reply. No. It supplies a candidate first down-type boundary. Up/down splitting, strange coherence, generation amplification, first χ magnitude, heavy thresholds, and full matrix construction remain separate gates.

Objection 13 — You multiply amplitudes, densities, probabilities, and support counts of different mathematical types

Reply. The earlier draft was open to that charge; the present construction is not. Theorem 2A builds the coefficient as the Schur complement of one typed block operator: α_{int} enters twice as a gate amplitude, while $1/(2\pi)$, $1/2$, $1/2$, and the quotient census enter as spectral values of the compressed internal response K_{int}^{-1}. Lemma 2B then forbids any factor that lacks one of the three admissible operator roles. What remains open is not the typing but whether the substrate returns that block structure (D2, D3, D16).

Objection 14 — Your factor 1/4 is structure the Standard Model gets for free

Reply. Correct as a statement about the Standard Model: doublet-component selection costs nothing (vacuum alignment), and the left–right connection costs nothing (it is the Dirac bilinear), so $m_e = y\, v/\sqrt{2}$ carries no residual 1/4. The paper's premise (P5) is therefore explicitly that the two binaries are substrate readout registers distinct from that representation-theoretic structure — registers whose selection is a physical projection event, not a basis choice. If the discharge theorem (D5) identifies the carriers with the representation structure after all, both factors are 1 and the kernel fails by a factor of 4, landing near 18.8 MeV as the ablation table records. The objection is thus not deflected but converted into the sharpest form of falsifier F9: the binaries must be shown to exist, not merely to be unbiased.

Appendix D — Dependency Map and Next Calculations

D.1 Minimal proof chain

Fock/chiral matter ⇒ $\bar{Q}_L \Phi_{\mathrm{cl}} d_R$ ⇒ $M_d^c = (v_{\mathrm{cl}}/\sqrt{2})\, Y_d^c = v_{\mathrm{cl}}\, E_d^c$.

typed block operator + no direct attachment + harmless same-side blocks (P13) + off-line closure (P15) + internal spectral gap (P16) ⇒ compressed kernel $-\,\alpha_{\mathrm{int}}^2\, r_d\, V_d$.

two gates ⇒ α_{int}^2.

canonical phase measure (P14) + local zero-mode evaluation + unbiased chiral binary + unbiased weak binary ⇒ $\Pi_{\text{cont}} = 1/(8\pi)$ inside K_{int}^{-1}.

independent colour–triality support + additive quotient ($c_L = c_R = 1$) ⇒ $\tilde{q}_L D_9 \tilde{q}_R^\dagger = 9 W_{ct}$ ⇒ $N_d = 9$.

Π_{cont} and N_d + internal-response factorisation (P12) ⇒ $r_d = 9/(8\pi)$.

metric-compatible anchoring + smallest-line alignment ⇒ $P_{\mathfrak{D}_1}^c = P_d^R$.

all of the above ⇒ $y_d^{(0)} = 9\sqrt{2}\, \alpha_{\text{int}}^2 / (8\pi)$.

coupled RG/matching on an isolated reducing line ⇒ $m_d^{\mathcal{S}}(\mu) = |Z_d^{\mathcal{S}}|\, m_d^{(0)} + O(\|\Delta_d^{\text{mix}}\|)$.

D.2 Highest-priority computations

1. Construct the block operator K_d from the substrate, verify the absence of a direct left–right attachment, compute the same-side Feshbach blocks to discharge the External-Block Recanonicalisation Condition (P13), and prove the off-line no-leakage conditions (P15).
2. Compute the physical gate coefficient from the interface action and determine whether it is α_{int}, e_{int}, or another coupling.
3. Derive the substrate phase measure and the zero-mode evaluation kernel — local versus integrated — from the emergent-field map (P14, P4).
4. Prove the internal-response factorisation $K_{\text{int}}^{-1} = R_\varphi \otimes R_{ch} \otimes R_w \otimes R_{ct}$ or bound its entanglement corrections (P12).
5. Construct the colour–triality readout algebra and test the 3 × 3 factorisation.
6. Derive the quotient normalisation — additive ($c = 1$) versus Hilbert-normalised ($c = 1/3$) — and the categorical typing of q.
7. Compute the compressed internal response and verify $r_d = 9/(8\pi)$.
8. Construct the physical first-down address projector and verify idempotency and K_D-self-adjointness.
9. Compute the overlap of its canonical image with the smallest right singular projector.
10. Derive the RG-invariant or scheme-specific meaning of the kernel and the reducing-line factor $Z_d^{\mathcal{S}}$, including the four-loop refinement of the Reading-B audit.
11. Replace the strange trace identity by a singular-value-twenty operator and compute it.
12. Only then perform the full six-quark numerical audit.

Appendix E — References and Source Ledger

E.1 Internal dependency ledger (VERSF programme)

The inherited structures I1–I8 and the named imports below must cite exact predecessor papers, sections, and theorem numbers. The series attributions are recorded now; the precise loci are to be pinned against the Master Ledger before circulation. Pinning priority: the loci for I3 (SF-1, canonical normalisation and Hermitian lift) and I6 (ordered-χ readout, FH-1 interface) carry the heaviest load in this paper and should be pinned first. These placeholders are a circulation blocker: the manuscript must not circulate externally until every locus is pinned, because the argument depends on exact inherited results, not broad programme claims.

Inherited structure	Predecessor source (series)	Locus
I1 — Fermionic Fock carrier	VERSF fermionic quantisation and Fock-space construction papers	⟨paper ID, §, Thm — to pin⟩
I2 — Chiral matter representation	Standard Model Census series (SC-1, SC-2); chirality derivation (SM-1)	⟨to pin⟩
I3 — Canonical normalisation and Hermitian-lift discipline	Minimal Hermitian Lift and Frame-Rotation Firewall (SF-1)	⟨to pin⟩
I4 — Address/spectrum separation	SF-1; occupancy-ledger papers	⟨to pin⟩
I5 — Confinement discipline	Confinement and audit papers (CD-1, PHχ-1)	⟨to pin⟩
I6 — Ordered-commitment χ readout	Flavour-magnitude ordered-history papers; charm-hinge interface (FH-1)	⟨to pin⟩
I7 — Flavour-magnitude factorisation	Flavour-magnitude factorisation papers	⟨to pin⟩
I8 — Strange support and confinement candidates	Strange support and confinement papers	⟨to pin⟩
Electroweak completion interface and v_cl	Gauge-census closure (GC-1); electroweak completion interface papers	⟨to pin⟩
Normalisation conventions	Quantitative Normalisation series (QN)	⟨to pin⟩

E.2 External mathematical sources

- Singular-value decomposition, Hermitian lifts, and Weyl perturbation inequalities: R.A. Horn and C.R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press (1991), Chapter 3.
- Singular-subspace perturbation: P.-Å. Wedin, "Perturbation bounds in connection with singular value decomposition," *BIT* 12 (1972) 99–111.

- Eigenvector rotation bounds: C. Davis and W.M. Kahan, "The rotation of eigenvectors by a perturbation. III," *SIAM J. Numer. Anal.* 7 (1970) 1–46.
- Intermediate-space elimination: H. Feshbach, "Unified theory of nuclear reactions," *Ann. Phys.* 5 (1958) 357–390; F. Zhang (ed.), *The Schur Complement and Its Applications*, Springer (2005).

E.3 External physical inputs and conventions

- Light-quark MS̄ mass conventions and m_d^MS̄(2 GeV): Particle Data Group, *Review of Particle Physics* (current edition), quark-masses review.
- Quark-mass anomalous dimension at four loops: J.A.M. Vermaseren, S.A. Larin, and T. van Ritbergen, *Phys. Lett. B* 405 (1997) 327; K.G. Chetyrkin, *Phys. Lett. B* 404 (1997) 161.
- Mass and coupling evolution with decoupling: K.G. Chetyrkin, J.H. Kühn, and M. Steinhauser, "RunDec," *Comput. Phys. Commun.* 133 (2000) 43.
- $\alpha(0)^{-1}$ = 137.035999084: CODATA recommended value (2018 adjustment).
- $G_F = 1.1663787 \times 10^{-5}\ \mathrm{GeV}^{-2}$, hence v_F = (√2 G_F)^(−1/2) = 246.21965 GeV: Particle Data Group.
- α_s(M_Z) = 0.1180: Particle Data Group world average.

Audit rule: every numerical input in E.3 enters only through the observable-normalised audit of Section 15.1A and the Reading-B running audit of Section 10.4. None enters the structural construction, per the non-insertion protocol of Section 16.
