The Admissibility-Maintenance Factor

Localization, Maintenance, and What Mass Measures

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The charged-lepton hierarchy law factorizes mass into two pieces. One piece, *localization*, comes from how tightly a closure compresses across refinement depth, and it supplies a common geometric ladder shared by all matter. Once the localization constant κ is imported — derived in the lepton programme independently of the muon, not calibrated to it — the electron-to-muon ratio follows from this ladder essentially exactly, and the 207-vs-206.8 match is then a genuine parameter-free check rather than a definition. This step sets the scale against which everything else is measured (it serves below as the maintenance-neutral anchor). The second piece, *completion*, measures everything left over. For the electron and muon the completion factor is essentially one; for the tau it is a suppression near 1/12, recovered by two independent derivation chains.

This paper has two halves. The first poses a problem; the second proposes what solves it.

The problem. It is tempting to carry the factorization straight across to quarks. The first half argues the step cannot be taken naively, and explains precisely why. The obstacle is not that quarks lack a completion factor — it is that the factor refuses to behave uniformly. Read against the same localization ladder, completion *amplifies* the up-type hierarchy and *suppresses* the down-type hierarchy, while separately suppressing the tau. One object pulls in opposite directions across sectors, and no single one-signed cost can do that. A quieter point carries equal weight: the localization ladder was calibrated on leptons, so reusing it for quarks rests on an assumption the lepton data cannot by itself justify. The first half closes by isolating completion as an independent, well-posed object and eliminating an entire class of mechanisms as the source of the sign reversal.

The resolution. The second half proposes what completion *is*. It is not an added cost. It is the signed energetic price a closure pays to remain *admissible* — to keep being a coherent, allowed structure. A complete closure, such as a charged lepton, maintains itself: the price is internal, and the tau's suppression is read as a saturated closure spending most of its structure on holding its own admissible state. A partial closure, such as a quark, cannot remain admissible alone; it pays through coupling to the shared confinement environment — the bath. That coupling carries a sign. Where the bath *destabilizes* a mode, mass is amplified (a burden); where it *screens* or *stabilizes*, energy is returned and mass is suppressed. This single signed quantity does what an

unsigned cost cannot: it explains why up-type quarks sit above their localization expectation while down-type quarks sit below it — same generation, same closure class, opposite sign, because the two charge sectors couple to the same bath with opposite phase.

The reframing does not solve the hierarchy. It converts a placeholder into a physically interpretable object with a sharp, falsifiable demand: derive from closure geometry *why up-type partial closures couple constructively while down-type partial closures couple screening-wise*, and show the same maintenance mechanism reduces to internal self-maintenance in the lepton limit. That last requirement is the paper's central gamble: if the tau's price and the quarks' price are *not* two faces of one mechanism, the second half reduces to a renaming.

The deepest consequence, if the gamble holds, is that the paper is no longer really about quarks. It is about what mass measures. The earlier picture, "mass = localization cost," gives way to "mass = localization × admissibility maintenance" — localization fixing where a structure sits, maintenance measuring how hard reality must work to keep it admissible. Quarks were simply the sector where the maintenance term could not be ignored; the electron paid it all along, only at a vanishing rate.

Contents

---

1. Introduction

The charged-lepton programme established a hierarchy law of factorized form

$M = M_{loc} \cdot \Theta$.

The localization component M_{loc} follows from localization compression composed with sector rescaling. The completion component Θ collects all departures from pure localization behaviour. For the electron and muon, $\Theta \approx 1$; for the tau, $\Theta \approx 1/12$.

This paper does two things, kept deliberately separate.

Part I — The Completion Problem establishes, at grade [Proven] or [Conditional], that completion is an independent object with definite structure. A cancellation result shows only the *variation* of completion is observable, so the existence of a completion contribution explains nothing by itself. An extraction against the localization ladder shows completion varies with opposite sign across charge sectors. That extraction is shown to rest on an undischarged assumption (sector-universal localization), and the extracted numbers are reinterpreted accordingly. Finally, an elimination proposition removes a whole class of mechanisms from contention. Part I commits to no interpretation of what completion is.

Part II — The Admissibility-Maintenance Hypothesis proposes, at grade [Conjectural], that completion is the signed energetic price of admissibility: paid internally by complete closures, through the shared confinement bath by partial ones. The whole of Part II rests on one governing claim, stated at its head as the **Central Conjecture** — that the tau suppression and the quark maintenance couplings are two limits of a single functional. Granting that, Part II shows the signed structure is forced by the data, advances a self/bath unification that resolves the soft lepton/quark dichotomy, identifies the sign of the coupling with charge sector — the first falsifiable prediction the mechanism must reproduce — extracts the residual maintenance couplings from all observed ratios and finds they show a pattern — one amplifying entry against a set of screening ones, with a maintenance-neutral anchor — rather than a scatter (§11), and closes (§13) with the deepest consequence: a reframing of what mass measures, from "localization cost" to localization composed with signed admissibility maintenance, valid across every matter sector.

The firewall between the parts is intentional. Part I stands whether or not Part II is correct: if the maintenance hypothesis is later falsified, the problem statement, the sign-reversal extraction, and the elimination are untouched. Part II borrows none of its conjectural weight from Part I's proven results; it imports their conclusions as constraints it must meet.

PART I — THE COMPLETION PROBLEM

2. Imports and Inheritance Discipline

This paper derives no new localization input and introduces no new measured ratio. It inherits the following and is capped at their grade. No result is graded above **[Conditional]** except purely logical steps, which are **[Proven]** but inherit the conditional status of any datum they consume.

Import L1 — Localization compression.

$L_g = L_0 \cdot e^{(-\kappa g)}$, with $\kappa = 8/3$. **[Conditional — inherited from the charged-lepton programme]** *Fallback:* if κ is mis-stated, every localization ratio rescales by the corrected exponent; the structure of the argument is unaffected but the numerical Θ-ratios shift.

Import L2 — Sector rescaling.

$M_g = \tilde{E}_g / L_g^2$, hence $M_{loc}(g) \propto e^{(2\kappa g)} = e^{(16g/3)}$. **[Conditional — inherited]** Adjacent-generation localization ratio: $M_{loc}(g+1) / M_{loc}(g) = e^{(16/3)} \approx 207$.

Import L3 — Tau completion.

$\Theta_\tau \approx 1/12$. **[Conditional — inherited, two convergent derivation chains]** *Fallback:* if one chain is later found defective, the tau datum weakens but the up/down sign reversal in §4 stands independently, since it does not use Θ_τ.

Conjectural import C1 — Sector-universal compression rate.

The compression rate κ = 8/3 calibrated on the lepton sector applies unchanged to the up-type and down-type quark sectors. **[Conjectural — Gate G-κ open, §5]** *Fallback:* if κ is sector-dependent (κ_{up}, κ_{down}, κ_{lep} distinct), then part of what §4 attributes to completion variation is in fact localization variation, and the extracted Θ-ratios must be reinterpreted as the product of a localization-rate mismatch and a residual completion factor. This fallback is load-bearing.

Ontology/confinement imports.

Complete closures are admissible in isolation; partial closures are not, acquiring physical meaning only within a confinement-complete composite. **[Conditional — inherited from the confinement and ontology programmes]** These are used only in Part II; they are listed here for completeness.

3. The Cancellation Result: Only Relative Completion Is Observable

Assume the factorization M = M_loc · Θ. For any two structures indexed 1 and 2,

$M_2 / M_1 = (M\_loc,_2 / M\_loc,_1) \cdot (\Theta_2 / \Theta_1)$.

The mass ratio depends on completion only through the ratio Θ_2/Θ_1. A completion factor constant within a comparison cancels identically and leaves no trace in any mass ratio.

Result 3.1 [Proven]. A sector-uniform completion factor is unobservable. The hierarchy problem depends not on the existence of completion but on its variation.

This is algebra and carries no inheritance cost. Its consequence is structural and decisive: the mere observation that quarks carry an additional completion contribution (because they are partial closures requiring confinement completion, where leptons are complete closures) explains *nothing* about why quark ratios differ from lepton ratios. A constant additional factor would cancel. The entire explanatory burden falls on the generation- and flavour-dependence of Θ.

Corollary 3.2 [Proven]. The quark–lepton closure distinction, taken alone, is not explanatory for the hierarchy. It is a precondition for completion existing, not an account of how completion varies.

This corrects the natural but mistaken expectation that establishing "quarks are partial closures" would by itself yield a different hierarchy law. It would not.

---

4. Completion Carries Structure: The Sign-Reversal Extraction

Read the observed mass ratios against the localization ladder of Import L2. By Result 3.1 the residual after removing the localization factor *is* the completion ratio, conditional on the localization input.

Throughout, m_q denotes the running (MS-bar) current-quark mass. This choice is forced and is what makes the targets well-posed:

Well-posedness note [Conditional]. Same-charge quark mass ratios are approximately renormalization-group invariant. The QCD mass anomalous dimension γ_m is flavour-universal, so under QCD running *all* mass ratios are invariant, cross-charge included — QCD therefore cannot be what privileges same-charge ratios. The residual scale-sensitivity is electromagnetic: γ carries a charge-dependent piece $\propto Q^2$ that cancels only within a charge sector. It is this electromagnetic cancellation that makes same-charge ratios (m_c/m_u, m_s/m_d) the scheme-stable targets, while cross-charge comparisons (up versus down, quark versus lepton) retain a Q^2-dependent scale sensitivity and are therefore **[Conditional — scheme-sensitive]** and weaker than the within-sector ratios.

Second-generation extraction. Using current-quark masses and the localization reference $e^{(16/3)} \approx 207$:

$m_c / m_u \approx 589$, so $\Theta_c / \Theta_u \approx 589 / 207 \approx 2.8$. **[Conditional — assumes C1]** $m_s / m_d \approx 20$, so $\Theta_s / \Theta_d \approx 20 / 207 \approx 0.10$. **[Conditional — assumes C1]**

The two residuals fall on opposite sides of unity. Completion amplifies the up-type second-generation ratio and suppresses the down-type second-generation ratio.

Result 4.1 [Conditional — assumes C1]. Conditional on sector-universal κ, the completion factor is non-monotonic across sectors: $\Theta_c/\Theta_u > 1$ while $\Theta_s/\Theta_d < 1$. No single monotonic completion cost can reproduce both.

This is the central observation of Part I and the hardest constraint any account of completion must meet.

What variable the sign reversal demands. The up/down contrast is a *same-generation, same-closure-class* phenomenon: up-type and down-type quarks are both partial closures, both second-generation. Closure class therefore does no work here — it distinguishes quarks from leptons, not up from down. Generation depth does no work either, since the contrast holds at fixed generation. The sign reversal must be carried by a variable that distinguishes the two charge sectors and enters with opposite sign. An isospin-like or charge-sector degree of freedom is therefore mandatory in any viable completion structure.

Result 4.2 [Conditional — assumes C1]. The completion structure must contain a charge-sector (isospin-like) variable acting with opposite sign on up-type and down-type closures. Closure class and generation depth are individually insufficient to produce the up/down sign reversal.

The completion factor must accordingly distinguish at least three things: generation depth, charge sector, and closure class — the first to carry the tau suppression and the generation ladder within a sector, the second to carry the up/down sign reversal, and the third to separate quarks from leptons.

---

5. The Load-Bearing Assumption: Is Localization Sector-Universal?

Result 4.1 is stated conditional on Import C1 because the reference e^(16/3) was calibrated on the lepton sector. Three points follow.

First, the lepton programme supports localization *within the lepton sector*. It provides no independent evidence that the quark sectors share the same compression rate κ. Reusing κ = 8/3 for quarks is an extrapolation, not an inheritance.

Second, this means the disjunction sometimes used to motivate structured completion — "either localization is non-universal, or completion is non-trivial" — cannot be resolved in favour of structured completion without *assuming* universal κ, which is one horn of the disjunction itself. Stated baldly, the inference "localization is universal, therefore completion is structured" is close to circular. The honest statement is conditional:

Result 5.1 [Conditional — assumes C1]. Conditional on sector-universal κ, completion carries non-trivial, sign-changing structure. Absent C1, the observed hierarchy differences are shared between localization-rate variation and completion variation, and cannot be assigned to completion alone.

Third, and constructively, this clarifies what the §4 numbers actually measure. They are *deviations from lepton-calibrated localization scaling*. The quantity 2.8 is unambiguous as a deviation; its decomposition into "completion" versus "a different κ in the up sector" is not fixed by the data. The next derivation must either (a) derive κ for each sector and confirm universality, discharging C1, or (b) derive a sector-dependent κ, in which case the §4 residuals are partially reabsorbed into localization and the completion target is correspondingly revised.

Gate G-κ [Open]. Sector-universality of the compression rate κ is undischarged. All Θ-ratios in §4 and §7 are gated on G-κ. Discharging G-κ is a prerequisite for promoting Result 4.1 / 5.1 above [Conditional].

---

6. Eliminating a Class of Candidates

The candidate survey is not a neutral list. Result 4.2 already forecloses an entire class of mechanisms, and stating this explicitly converts the survey from a catalogue of open items into an active narrowing of the search space.

Proposition 6.1 (Completion Structure Proposition) [Conditional — assumes C1]. The charge-sector variable demanded by Result 4.2 cannot arise from closure class alone.

Reason. A quantity that is a function of closure class takes a single value across all members of a given closure class. The up-type and down-type quarks of a fixed generation share closure class — all are partial closures. Yet the second-generation completion ratio reverses sign between them:

$\Theta_c / \Theta_u > 1$ (same closure class) versus $\Theta_s / \Theta_d < 1$ (same closure class).

By Result 3.1 a closure-class-only quantity is constant across the u, c, s, d set and cancels from both ratios, contributing nothing to either and hence nothing to their difference. A sign reversal between two same-closure-class comparisons therefore cannot originate in closure class. The discriminating variable must distinguish charge sectors *within* a common closure class. ■

The inference from the datum to the elimination is deductive; only the datum (the sign reversal of §4) is conditional, which is why the proposition inherits **[Conditional — assumes C1]** rather than **[Proven]**.

The proposition does real work: it removes closure-class-level mechanisms from contention as the *source* of the up/down reversal, before any model is built. The candidate assessment below is organized by the single sharpened question — *does the candidate supply a charge-sector variable acting within a closure class?*

Confinement completion. Closure-class-level by construction: it separates quarks (partial) from leptons (complete). By Proposition 6.1 it is **eliminated** as the source of the up/down sign reversal. It may still contribute a uniform quark-versus-lepton offset, but that offset cancels from every intra-quark ratio and cannot carry the charge-sector constraint. **[Eliminated as source — closure-class-level]**

Saturation. Carries generation-depth dependence and is already tied to tau suppression. It addresses the depth axis, not the charge-sector axis. Admissible as a contributor to the generation ladder and the tau datum; **not a candidate** for the up/down variable unless shown to depend on charge sector. **[Open — wrong axis for the reversal]**

Persistent distinguishability load. A substrate-stiffness quantity and a candidate amplifier. Survives Proposition 6.1 *only if* it can be shown to vary across charge sectors within the partial-closure class with opposite sign. Its charge-sector dependence is undemonstrated; this is the precise question to settle. **[Open — charge-sector dependence undemonstrated]**

Interface transport complexity. Likewise a substrate-stiffness quantity and candidate amplifier, subject to the identical test: opposite-sign variation across charge sectors within a closure class. Undemonstrated. **[Open — charge-sector dependence undemonstrated]**

The survey now has a definite shape. One candidate is eliminated as a source, one is confined to the wrong axis, and exactly two — the two substrate-stiffness loads — remain live, each reduced to a single sharp question: does it carry an opposite-sign charge-sector variable within the partial-closure class? Proposition 6.1 has converted "find the variable" into "decide between two named candidates on one criterion."

Gate G-Θ [Open]. No surviving candidate has yet been shown to supply the opposite-sign charge-sector variable of Result 4.2 from closure geometry without fitting. The completion factor is underived, but the search space is narrowed by Proposition 6.1 to charge-sector-resolving substrate-stiffness structure.

Status at the close of Part I. Completion is an independent object (Result 3.1), carries sign-changing structure (Result 4.1, conditional on C1), demands a charge-sector variable (Result 4.2), and cannot source that variable from closure class (Proposition 6.1). What completion *is* remains open. Part II proposes an answer.

PART II — THE ADMISSIBILITY-MAINTENANCE HYPOTHESIS

The remainder of this paper is graded **[Conjectural]** except where a step is pure logic on Part I's data. It proposes a physical identity for the completion factor and is to be read as a hypothesis that meets — and could fail to meet — the constraints Part I established.

The factor is renamed to carry its proposed meaning:

$M = M_{loc} \cdot \Theta_{adm}$,

where Θ_{adm} is the same mathematical object as Part I's Θ, now interpreted as an **admissibility-maintenance coupling**: the signed energetic price a closure pays to remain admissible.

One claim governs everything that follows. It is stated here, at the head of Part II, because each later section is a consequence of it, and because it is the single proposition on which the entire interpretation stands or falls.

Central Conjecture [Conjectural]. The tau suppression and the quark maintenance couplings are two limits of a single admissibility-maintenance functional. Self-maintenance (complete closures, internal environment) and bath-coupling (partial closures, shared confinement environment) are not two mechanisms but one functional evaluated in two environments.

What hangs on it. If the conjecture holds, the signed factorization (§7), the self/bath unification (§9), the charge-sector sign structure (§10), and the ontological reading of mass (§13) form one picture. If it fails — if the tau cost and the quark couplings have unrelated origins — then Part II reduces to the modest claim that the completion factor has been renamed and shown to carry a sign, and the leptons and quarks remain separate stories. The whole gamble of Part II is the *singleness* of the functional. The reduction requirement R-limit (§12) is precisely the test of this conjecture.

The sections below develop the conjecture's consequences and convert it into a falsifiable derivation target.

---

7. The Maintenance Term Must Be Signed

The first task is to show the maintenance contribution cannot be a one-signed cost. This is what licenses a multiplicative log-form over a positive additive one.

Consider a candidate additive form $M = M_{loc} + M_{maint}$ with M_{maint} interpreted as a maintenance *cost*, hence $M_{maint} \geq 0$. Take the down sector, second over first generation. With $M_{loc}(g) \propto e^{16g/3}$,

$$M_{loc}(s) / M_{loc}(d) = e^{16/3} \approx 207, \text{ while } m_s / m_d \approx 20.$$

The observed ratio sits *below* the localization ladder. Writing $M_{loc}(d) = m$ and $M_{loc}(s) = 207m$,

$$m_s / m_d = (207m + M_{maint,s}) / (m + M_{maint,d}).$$

For this to fall below 207 requires $M_{maint,s} < 207 \cdot M_{maint,d}$. A maintenance cost that does not decrease with generation depth ($M_{maint,s} \geq M_{maint,d}$, the natural expectation that a deeper, more strongly localized closure requires at least as much support) gives a ratio at or above 207 — amplification, not the observed suppression.

Result 7.1 [Conditional — assumes C1]. A maintenance term that is sign-fixed (constrained to add mass) and non-decreasing in generation depth cannot reproduce the down-sector suppression. It necessarily amplifies the ratio, as it does in the up sector. *The logical step is [Proven]; it inherits [Conditional] from the down-sector datum.*

The down sector therefore forces one of two structures: a maintenance contribution that *decreases* with depth (anti-natural for a "cost"), or a maintenance contribution that can take the opposite sign — *screening*, in which the environment returns energy and lowers the mass. The screening reading is the physical one and is adopted; the depth-decreasing-cost alternative is noted and not excluded by the data alone.

Result 7.2 [Conditional — assumes C1]. The admissibility-maintenance contribution is signed. It amplifies mass in some sectors and suppresses it in others. No single-signed cost reproduces the spectrum.

This justifies writing the contribution multiplicatively about unity, where the two signs are symmetric excursions:

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

with $\ln \Theta_{adm} > 0 \longrightarrow$ burden (constructive coupling, amplification), $\ln \Theta_{adm} < 0 \longrightarrow$ screening (destructive coupling, suppression), $\ln \Theta_{adm} = 0 \longrightarrow$ self-supporting (e, μ to current accuracy).

The log form is not cosmetic: it makes the observable ln(Θ_adm,2 / Θ_adm,1) a *signed difference of maintenance couplings*, so Part I's "variation of completion" becomes literally a coupling difference, and the cancellation result (Result 3.1) is the statement that an additive constant in ln Θ_adm is unobservable.

---

8. The Admissibility-Maintenance Factor

Definition. Θ_adm is the multiplicative factor by which a closure's mass departs from its localization value on account of the energetic price of remaining admissible. Equivalently, ln Θ_adm is the signed maintenance coupling: the work the closure (internally or via its environment) must perform per unit localization energy to sustain its admissible state, with screening counted negative.

The naming problem is resolved by reservation rather than overload. "Burden," "cost," and "support cost" name the **positive branch** ln Θ_adm > 0 only. The umbrella quantity is sign-neutral: *admissibility-maintenance coupling*. A screening contribution is not a negative cost — it is a stabilization, energetically distinct, and the vocabulary now reflects that.

Read as signed couplings, the inherited residuals become:

Mode	Comparison	ln(Θ ratio)	Reading
up-type	charm / up (gen 2 / 1)	ln 2.8 ≈ +1.03	burden, grows with depth
down-type	strange / down (gen 2 / 1)	ln 0.10 ≈ −2.30	screening, grows with depth
lepton	tau / muon (gen 3 / 2)	ln(1/12) ≈ −2.48	self-suppression at saturation

[Conditional — assumes C1; lepton row inherited]

Two structural observations follow, both [Conditional]:

Observation 8.1. The sign of the maintenance gradient with depth is sector-specific: positive for up-type, negative for down-type and for the lepton third generation. The amplifying sector is the lone outlier; suppression with depth is the more common behaviour.

Observation 8.2 [Conjectural — possibly spurious]. The down-type screening (−2.30) and the lepton self-suppression (≈ −2.5, from the inherited Θ_τ ≈ 1/12, or −2.51 by direct extraction in §11) are close in log magnitude. *If* self-maintenance and bath screening are one mechanism (§9), a comparable magnitude is the kind of agreement one would hope to see; but with two data points this may be coincidence and is recorded only as a flag for the derivation, not as evidence. The systematic extraction of §11 records this as one entry in a wider branch pattern.

---

9. One Mechanism, Two Regimes: The Self/Bath Unification

A clean dichotomy — leptons self-support, quarks require support — is too simple, and the tau shows why: the tau is a complete closure that nonetheless pays a large maintenance price ($\Theta_\tau \approx 1/12$). A complete closure paying maintenance is not an exception to explain away; it is the key to the Central Conjecture. This section unpacks that conjecture into the structure a derivation must realize.

Unification structure (elaboration of the Central Conjecture) [Conjectural]. Every closure pays an admissibility-maintenance coupling. The coupling differs not in *kind* but in the *environment* through which it is paid:

- **Complete closures** (leptons) maintain themselves. The environment is trivial; the coupling is *internal*. Saturation is the internal limit — a closure spending its structure to hold its own admissible state. The tau row is this regime.
- **Partial closures** (quarks) cannot remain admissible alone. The environment is the shared confinement bath; the coupling is *external* and bath-mediated. The quark rows are this regime.

Self-maintenance is the limiting case of bath-coupling in which the bath reduces to the closure itself.

This is stronger and more falsifiable than the dichotomy it replaces. It does not assert tau and quarks are the *same* case — their environments, signs, and magnitudes differ — only that they are governed by the *same maintenance functional* evaluated in different environments, which is exactly the content of the Central Conjecture. The clean separation "complete vs partial closure" is preserved exactly where it belongs: as the variable that selects internal versus bath-mediated maintenance. This is consistent with Proposition 6.1 — closure class selects the *regime*, not the *sign within the quark regime*, which §10 assigns to charge sector.

The conjecture imposes a hard reduction requirement on any derivation — its direct operational test:

Requirement R-limit [Open]. A derived maintenance functional must reduce, in the complete-closure (trivial-environment) limit, to the saturation self-cost, reproducing $\Theta_\tau \approx 1/12$; and in the partial-closure regime, to a bath-coupling cost reproducing the quark rows. A functional that yields the quark sign structure but cannot recover the tau in its self-limit fails the unification and is, at best, a quark-only model.

---

10. The Sign Structure: Charge Sector Sets the Phase

Proposition 6.1 requires a variable that distinguishes charge sectors within the partial-closure class. The maintenance picture identifies it physically:

Sign Identification [Conjectural]. The sign of the bath coupling is set by the charge sector. Up-type partial closures couple to the shared confinement bath *constructively* ($\ln \Theta_{adm} > 0$, burden growing with depth); down-type partial closures couple *screening-wise* ($\ln \Theta_{adm} < 0$, screening growing with depth).

This is the first concrete, falsifiable content of the reframing. It is not a fitted sign — it is read directly from the data via the decomposition (up amplifies, down suppresses) — but as a *derivation target* it is open: the mechanism must produce these signs from closure geometry, not assume them.

The robust statement concerns the first generation step. The full extraction of §11 shows the up-type sector does not amplify monotonically — it turns over by the third generation — so the prediction is stated at the gen1→gen2 step, where the charge-sector reversal is sharpest, with the higher-generation behaviour recorded separately as a feature the functional must reproduce. The C1-dependence of the *sign* (as opposed to the C1-free raw splitting ratio) is treated below.

Prediction P-sign [Open — falsification target]. At the first generation step (gen 1 → gen 2), the maintenance departure has opposite sign in the two charge sectors:

$\ln(\Theta_{adm,c} / \Theta_{adm,u}) > 0$ (up-type amplifies), $\ln(\Theta_{adm,s} / \Theta_{adm,d}) < 0$ (down-type screens).

A derivation that yields the wrong relative sign between the charge sectors at this step falsifies the bath-coupling hypothesis, independently of any magnitude fit — conditional on C1, since the sign is read against the common localization reference (see the C1-gating discussion below; the C1-free residue is the raw splitting ratio).

Feature F-turnover [Open]. The down-type sector screens at both available steps (gen 1 → 2 and gen 2 → 3), while the up-type sector amplifies at the first step and screens mildly at the second (see §11). The functional must reproduce this asymmetry: monotone screening down-type, sign-changing up-type. The third-generation up-type datum carries the top quark and is the least reliable entry, the top lying outside the confinement picture.

P-sign concerns the sign of the *maintenance coupling*, and that sign is C1-dependent — a point that must be stated carefully, because it is easy to over-claim. The sign of the within-sector coupling is sign $\ln(\Theta\_c/\Theta\_u) = \text{sign}[\ln(m\_c/m\_u) - 2\kappa\_up]$, sign $\ln(\Theta\_s/\Theta\_d) = \text{sign}[\ln(m\_s/m\_d) - 2\kappa\_down]$.

Renormalization-group invariance protects the *first* term in each bracket — the mass ratios m_c/m_u and m_s/m_d are scheme-stable. It does *not* protect the second term, the localization reference, which is exactly what C1 fixes. Under ¬C1, with independent κ_up and κ_down, the up-type bracket is positive only if $\kappa\_up < \frac{1}{2}\ln(m\_c/m\_u)$ and the down-type negative only if $\kappa\_down > \frac{1}{2}\ln(m\_s/m\_d)$; nothing forces both at once. So even the *relative* (opposite-sign) statement is C1-dependent. This is precisely Result 5.1: absent C1, the hierarchy differences are shared between localization-rate variation and completion variation and cannot be assigned to completion alone. P-sign is therefore gated on G-κ, like everything else built on the localization reference.

The genuinely C1-free, scheme-stable invariant is weaker and more primitive: the raw fact that the up-type second-generation splitting vastly exceeds the down-type one,

$m\_c/m\_u \approx 505$ versus $m\_s/m\_d \approx 20$, a factor ≈ 25,

a ratio of like-charge ratios that survives ¬C1 entirely. The figure ≈ 505 is the co-scaled RG-invariant value (charm run to 2 GeV, $m\_c(2\text{ GeV}) \approx 1.09$ GeV); the scale-mixed PDG division $m\_c(m\_c)/m\_u(2\text{ GeV}) \approx 589$ appears only in the §11 table, where it is flagged as convention-dependent, and is *not* the invariant to quote here. The *interpretation* of this disparity as an opposite-sign maintenance coupling requires the common localization reference, i.e. C1. The durable claim is thus: the up- and down-type gen-2/1 splittings differ by a factor ≈ 25 (scheme-stable, C1-free); reading this as an opposite-sign maintenance coupling requires sector-universal κ (C1, Gate G-κ).

---

11. The Extracted Maintenance Hierarchy

The analysis admits one further step, and it is the step that turns the maintenance factor from a quantity to be derived into a quantity already visible in the data. Every observed within-charge-sector mass ratio can be divided through by the localization ladder to expose the maintenance coupling it implies. The result is not a scatter of residuals but a structured hierarchy of signs.

With $M = M\_loc \cdot \Theta\_adm$, the observable for any two same-sector modes is

$\Theta\_adm,2 / \Theta\_adm,1 = (M\_2 / M\_1) / (M\_loc,2 / M\_loc,1)$,

and with the adjacent-generation localization reference $M\_loc,2 / M\_loc,1 = e^{(16/3)} \approx 207$, the observed spectra imply the following. Each within-sector ratio is RG-invariant *in principle* — same-charge mass ratios are scale-independent for the electromagnetic reason set out in §4 (QCD running is flavour-universal and so leaves all ratios invariant; the residual scale-dependence is the charge-dependent $\propto Q^2$ piece, which cancels only within a charge sector). In practice the

PDG values used here are quoted at differing scales (u, d, s at μ = 2 GeV; c and b at their own running masses m_c(m_c), m_b(m_b); top by direct measurement), so as tabulated only the s / d entry and the lepton ratios sit at a common scale. The c / u, b / s, and t / c entries mix scales and are therefore convention-dependent at the 10–20% level; the table is a PDG-convention extraction, not yet a fully co-scaled one. See the convention note below.

Step	Mass ratio	Θ_adm ratio	ln Θ_adm	Branch
up-type, gen 2 / 1 (c / u)	≈ 589	≈ 2.8 (2.4–2.8)	+1.03	amplifying
up-type, gen 3 / 2 (t / c)	≈ 136	≈ 0.66	−0.42	mild screening *
down-type, gen 2 / 1 (s / d)	≈ 20	≈ 0.10	−2.30	strong screening
down-type, gen 3 / 2 (b / s)	≈ 45	≈ 0.21	−1.56	screening
lepton, gen 2 / 1 (μ / e)	≈ 207	≈ 1.00	0.00	neutral (anchor)
lepton, gen 3 / 2 (τ / μ)	≈ 16.8	≈ 0.081	−2.51	strong screening

[Conditional — assumes C1; only s / d and the lepton rows are at a common scale as quoted; the top entry (*) lies outside the confinement picture and is the least reliable. RG-invariance does not protect t / c: the top does not hadronize, m_t is a pole/threshold mass, and pairing it with MS-bar charm is scheme-mixed, so "same charge" does not confer scheme-stability on this row.]

Convention. The table uses PDG 2025 central values: m_u = 2.16 MeV, m_d = 4.70 MeV, m_s = 93.5 MeV (MS-bar at μ = 2 GeV); m_c = 1.273 GeV, m_b = 4.183 GeV (running masses at their own scale, m_c(m_c), m_b(m_b)); m_t = 172.56 GeV (direct top-mass measurement). Charged-lepton pole masses: m_e = 0.511 MeV, m_μ = 105.66 MeV, m_τ = 1776.86 MeV.

Provisional status. Because c, b, and t are quoted at scales different from the light quarks, the c / u, b / s, and t / c residuals are PDG-convention extractions, not fully RG-unified ones. A future version should rerun the table with all quark masses co-scaled to a common reference (e.g. 2 GeV) or built from published scale-invariant ratios. The *sign pattern* — one entry above unity, the rest below — is robust against this; the *magnitudes* of the scale-mixed entries are not. The table should be read accordingly: illustrative for magnitudes, reliable for signs.

The hierarchy does not produce a random collection of residuals. Read as signed couplings, the entries organize into three behaviours.

The neutral anchor. The lepton gen 2 / 1 step gives Θ_adm ≈ 1.00 (ln ≈ 0). The electron-to-muon step pays essentially no maintenance. This is why the earlier picture "mass = localization cost" survived as a working description: it was read off precisely the sector where maintenance vanishes. The anchor is not a separate phenomenon to explain — it is the zero against which the other entries are amplification or screening.

Amplification is the exception. Exactly one entry lies above the anchor: the up-type gen 2 / 1 step, Θ_adm ≈ 2.8 (ln ≈ +1.03). No other observed step amplifies. This is the single most load-bearing number in the section — the entire sign opposition rests on it — and so its convention sensitivity must be stated plainly: co-scaling c and u to a common reference lowers the ratio to ≈

2.4 (ln ≈ +0.9), so the magnitude ranges ≈ 2.4–2.8 across reasonable conventions. What is convention-*independent* is the sign: $\Theta_c/\Theta_u > 1$ regardless. Nothing downstream depends on the precise value, only on the fact that this entry sits above unity while the others sit below. The amplifying branch is thus a single, localized feature of the up-type second generation, not a general tendency.

Screening is the rule. Four of the six entries screen. Among them, two land remarkably close together:

down-type gen 2 / 1 (s / d): $\ln \Theta_{adm} \approx -2.30$, lepton gen 3 / 2 (τ / μ): $\ln \Theta_{adm} \approx -2.51$.

A down-type quark step and a lepton step coincide to within $\Delta(\ln \Theta_{adm}) \approx 0.21$ — the log-couplings agree to within ~10%, though the Θ factors themselves then differ by $e^{0.21} \approx 1.23$, i.e. ~23%. If self-maintenance (the lepton step) and bath screening (the down-type step) are two limits of one functional, as the Central Conjecture asserts, a near-coincidence of this kind is exactly the signature one would hope to find. Two cautions attach to it, not one. First, it is two data points and may be accidental. Second, it is cross-scheme as well as cross-sector: the s/d entry is an MS-bar current-quark coupling while the τ/μ entry uses pole masses, so the adjacency compares couplings extracted under different conventions. It is recorded as a flag for the derivation, not as evidence (cf. Observation 8.2). The remaining screening entries — b / s at −1.56 and the caveated t / c at −0.42 — share the screening *direction* but are milder; they are family by sign, not by magnitude.

The down/up asymmetry. The down-type sector screens at both available steps (−2.30, then −1.56), and the lepton sector screens at its high step (−2.51) after the neutral anchor. The up-type sector alone changes sign with depth: strongly amplifying at gen 1 → 2 (+1.03), then mildly screening at gen 2 → 3 (−0.42). The up-type maintenance coupling is therefore non-monotonic, while the down-type coupling is monotone screening. This is Feature F-turnover of §10, now read directly from the table. The third-generation up-type entry carries the top quark, which does not hadronize and may lie outside the maintenance picture altogether; the apparent turnover is consequently either a real feature the functional must reproduce or a signal that the picture does not extend to the non-confining top. The extraction does not decide between these, and says so.

Maintenance-Branch Hypothesis [Conjectural]. After localization is removed, the residuals *exhibit a pattern consistent with* two branches — an *amplifying* branch lying above unity (the single entry c / u) and a *screening* branch lying below it (τ / μ, s / d, b / s), with t / c intermediate. Whether this pattern reflects real structure — distinct maintenance regimes the functional would have to produce — or is numerical coincidence is precisely the open question; it is not settled here. The only thing genuinely established is the sign opposition itself: c / u lies above the localization expectation while the other observed steps lie below it. The clustering of the three sub-unity residuals into a *family* is a further, separable conjecture (see below). A viable maintenance functional, if the branch pattern is real, must reproduce not merely the magnitudes but which sectors amplify, which screen, and where the up-type sign change falls.

The value of this section is not a discovery. It is that, after removing localization, the residuals show a pattern rather than a scatter — and *if* the pattern is real, the maintenance functional must reproduce it. That is a weaker and safer claim than announcing a structure, and it is the right one: the residuals are arithmetic (conditional on C1), the sign opposition of c / u against the rest is a direct reading of them, and everything beyond that — branches, families, regimes — is hypothesis to be tested. What the section adds is a second empirical hook beyond the tau suppression: after localization is removed, the tau sits alongside the down-type residuals rather than apart from them. That is worth recording even if it later proves coincidental, because a single underlying functional, were it real, would be expected to produce exactly such an adjacency. The residuals are built entirely from scheme-stable, within-charge-sector ratios, so the pattern, whatever its ultimate status, is not an artefact of mass conventions.

---

12. The Target and the Maintenance Functional

The acceptance test combines Part I's parameter-free standard with Part II's unification, sign, and branch-pattern requirements. A viable derivation of the admissibility-maintenance functional must, with no parameters fitted to the masses:

1. **Reproduce the extracted hierarchy** of §11 — the full set of signed couplings, including the neutral anchor, the lone up-type amplification, and the screening entries (whose clustering into a *family* §11 holds back as a separable conjecture) — from closure geometry. **[Open — Gate G-Θ]**
2. **Satisfy the unification reduction** R-limit: recover the saturation self-cost in the complete-closure limit and the bath-coupling cost in the partial-closure regime from one functional. **[Open]**
3. **Predict the charge-sector sign** P-sign at the first generation step and the up-type turnover F-turnover, from closure geometry rather than assumption. **[Open — falsification target]**

The derivation reduces to constructing one object:

The maintenance functional — a map from closure-geometric data (and the environment: trivial for complete closures, the shared bath for partial ones) to the signed coupling $\ln \Theta_{adm}$. **[Open — Gate G-Θ]**

It must exhibit, in order of decreasing robustness of the constraint:

1. the C1-free floor: the up- and down-type gen-2/1 splittings differ by a factor ≈ 25 (the co-scaled, RG-invariant $m_c/m_u \approx 505$ vs $m_s/m_d \approx 20$), scheme-stable and independent of sector-universal κ;
2. the first-step charge-sector sign reversal, opposite on up-type and down-type (P-sign — C1-gated, since it converts the splitting ratio into signed couplings against a common localization reference);

3. a self-limit reproducing tau saturation, numerically adjacent to down-type screening (R-limit and the §11 near-coincidence — the unification test);
4. the branch pattern and magnitudes of §11, *if* that pattern reflects real structure rather than coincidence (Gate G-Θ — conditional on C1 and on resolving Gate G-κ; the top entry conditional on the picture extending to non-confining modes).

The bar is parameter-freeness. The six extracted couplings are trivially fit by a sufficiently flexible functional; the content lies entirely in deriving the branch pattern, the unification limit, and the signs together, without adjustable inputs — as achieved for the electron–muon ratio, which the table now identifies as the maintenance-neutral anchor of the whole hierarchy.

---

13. What Mass Measures

This paper began as a question about quarks. Its deepest consequence is not about quarks.

Earlier in the programme the operative picture was

mass = localization cost,

with quark hierarchy anomalies treated as departures to be explained. The factorization established here, *read under the Central Conjecture*, replaces that picture with

mass = localization × admissibility-maintenance.

Localization fixes *where* a structure sits on the refinement ladder. Maintenance measures *how hard reality must work* to keep that structure admissible — internally for a complete closure, through the shared bath for a partial one, with a sign set by whether the environment burdens or screens the mode.

On this reading the quark sector was never special in kind. It is the sector in which the maintenance term cannot be neglected, because there it runs negative and large (down-type screening) or large and positive (up-type burden). The charged leptons were always paying a maintenance coupling too; the electron and muon simply pay $\ln \Theta_{adm} \approx 0$ in these units, which is why "mass = localization cost" survived as a working picture for as long as it did. The tau is the first lepton where the term becomes visible, and — if the Central Conjecture holds — it is visible there for the *same reason* it dominates the quark sector, not a different one.

The grade discipline here is essential, and the section is placed in Part II for exactly this reason. The two-piece *factorization* $M = M_{loc} \cdot \Theta_{adm}$ is not at issue — it is the algebra of Result 3.1. What is conjectural is the *interpretation* of the second factor as a single physical quantity, "the work to remain admissible," spanning leptons and quarks alike. That interpretation stands or falls with the Central Conjecture.

Ontological reading [Conjectural — conditional on the Central Conjecture]. Mass is not localization cost with corrections. It is localization composed with a signed admissibility-maintenance coupling throughout, for every matter sector. The lepton sector is the regime in which maintenance is nearly trivial; the quark sector is the regime in which it is not. Should the Central Conjecture fail, this reading fails with it, and the factorization retains only its established status as an observable decomposition (Result 3.1) without a unified physical content.

Stated as a target rather than a claim: the prize of the maintenance functional is not a quark mass formula. It is a single account of what the second factor of mass *is*, valid from the electron to the bottom quark, with the sector differences emerging as differences of environment rather than of mechanism.

Conclusion

Part I isolated completion as an independent object and proved the constraints any account of it must meet: only its variation is observable (Result 3.1); that variation reverses sign across charge sectors (Result 4.1, conditional on sector-universal localization); the reversal demands a charge-sector variable (Result 4.2); and that variable cannot come from closure class (Proposition 6.1), which eliminates confinement completion as its source and narrows the search to two named candidates.

Part II proposed what the object is, under one governing claim — the Central Conjecture, that the tau suppression and the quark maintenance couplings are two limits of a single admissibility-maintenance functional. From that conjecture the paper derived a picture that is forced, not chosen (the down-sector suppression rules out any single-signed cost, Result 7.2, making the signed log-form minimal); unifying (self-maintenance and bath-coupling as two regimes of one functional, dissolving the soft lepton/quark dichotomy into a statement about environment); falsifiable (the charge-sector sign, up-type constructive and down-type screening — a prediction a derivation can get wrong, gated on C1, with the C1-free residue being the raw factor-≈25 splitting disparity); and, at its deepest, a reframing of what mass measures — localization composed with signed admissibility maintenance, throughout.

The hierarchy is not thereby solved, and the whole interpretive edifice rests on the singleness of the functional, which is conjectural. The open question has sharpened into something a derivation can attack and fail at:

Construct the admissibility-maintenance functional: derive, from closure geometry, why up-type partial closures couple constructively and down-type screening-wise, and show — this is the test of the Central Conjecture — that the same functional reduces to saturation self-maintenance in the complete-closure limit.

That construction, gated on G-κ and G-Θ, is the next computational target of the programme.