The Non-Perturbative QCD Scale and Triality-Gap Completion Theorem in VERSF

▲ Programme Milestone — Strong-Interaction Scale Generation, Centre-Flux and Confinement Completion Series

Gate NQTC-1 / Saturated-Bath Colour Response, Numerical Λ_QCD, Global $\mathbb{Z}_3$ Lift, Finite-Regulator Quantum Triality Gap, Conditional Continuum Passage and String-Tension Architecture

Keith Taylor VERSF Theoretical Physics Programme Paper 5 of the Final Eight Standard Model Completion Papers Tier A declared-order non-perturbative evaluation theorem — July 2026

Principal Claim

This paper has two outputs of different status.

First, it evaluates a pre-registered minimal saturated-bath branch. On that branch, (χ_C, b, C, Z_3) = (48, 5, 5, 1), and the frozen VBF threshold cascade returns

■ $\Lambda\_\{3,MC\}\^VBF$ = 225.923 MeV,

without using any measured strong-sector quantity. These values are a pre-registered prediction fixed by Principles MC0–MC2 (MC0 supplies the benchmark direct coefficient u = 6C_hol, MC1 the protected unit, MC2 complete relaxation), not a substrate calculation; when the microscopic Hessian G⋆ is evaluated (Section 20′) they are confirmed, revised (with $Z_3 > 0$), or falsified ($Z_3 \leq 0$).

Second, it constructs the exact operator architecture required for a substrate derivation of confinement, with the inserted centre penalty removed from the main claim: the abstract SU(3) adjoint action R_C : SU(3) → SO(8) with $\mathbb{Z}_3$ kernel (Theorem 20′.1); the colour-saturated Hessian reduction $\mathfrak{T}\_C(G\star) = K\_C \otimes \mathbb{1}_8$; the response blocks $\chi\_C/8 = \langle f\_C, K\_C f\_C\rangle$, b, C, Z_3 as calculations on K_C; the intrinsic global centre sectors P_τ acting on the original SU(3) links (no second $\mathbb{Z}_3$ theory); the reduced wall coefficient B_τ from the wall Hessian; and the penalty-free sector gap $\varsigma\_\tau \geq g\_s \|\Phi\_\tau\|\_adj \sqrt{(2B\_\tau)}$ with uniform limit $\liminf \Delta\_\tau \geq \|\Phi\_\tau\|\_adj \sqrt{(2B\_min/Z\_max)}$.

The strict Paper 5 closure test is not declared passed. It awaits two independent evaluations: the microscopic substrate Hessian (local blocker, Debt D7) and the quotient-compatible global colour measure with its reflection-positive transfer matrix (global blocker, Debt D8). Those must return a numerical physical Z_3, a positive reduced wall coefficient B_τ, a reflection-positive transfer matrix, and a uniform nonzero-triality gap. The centre-projector Hamiltonian of the reference completion (Appendix A) is retained as an exact construction showing what follows once a positive centre penalty exists; it is not presented as the substrate derivation of that penalty.

The claim is a theorem of this explicit VERSF completion, not a general proof of the Yang–Mills mass-gap problem (Section 3.8). What is exact: the $\mathbb{Z}_3$ projector algebra, the abstract adjoint realisation, the branch arithmetic. What is pre-registered: (48, 5, 5, 1), c_σ = 1, and 225.923 MeV. What is owed: the two blockers above.

General Reader Summary

The strong force has two very different behaviours. At extremely short distances it becomes weak, allowing quarks to move almost freely inside protons and neutrons. At larger distances it does the opposite: pulling two colour charges apart creates an extended tube of energy rather than releasing an isolated quark.

The first behaviour is asymptotic freedom. The second is confinement. They are related, but proving one does not prove the other. This paper evaluates one pre-registered scale branch and constructs the operator architecture required for a genuine confinement calculation; the physical substrate values and the quantum confinement gap remain to be evaluated.

VERSF interprets the strength of the strong force as the stiffness of an underlying substrate against colour curvature. A colour state carried around a tiny loop may return rotated. The substrate resists that mismatch. The inverse square root of the remaining stiffness, after the bath has been allowed to reorganise, is the strong coupling.

The earlier QCD paper reduced this problem to three objects: a direct colour response, a vector describing how curvature pushes the auxiliary colour bath, and the bath's own restoring matrix. Much of this paper is about being honest regarding what those three objects actually are, because a first, tempting reading of them turns out to be wrong in an instructive way.

The tempting reading went: the direct stiffness is six units; one unit is irreducible and five belong to a bath that reorganises away, leaving exactly one unit, so $Z_3 = 1$ and the strong coupling g_s = 1. That arithmetic is clean, and this paper pre-registers it as a specific prediction — but it does not claim it as derived, for three reasons that successive calculations uncovered.

First, the "six" is not a primitive stiffness at all. It is a *census number*: it counts how many quark species (three generations, each with its up- and down-type partners) respond to one colour connection, and that count is $3 \times (1 + ½ + ½) = 6$. Counting the particles that feel a force is not the same as measuring the force's stiffness, and the paper installs a firewall against confusing the two.

Second, once one asks what the *primitive* response is — the response of the single shared colour class, not of the many particles probing it — the answer is forced to be one number, not six. The paper proves this: the colour class owns a single adjoint response, and its stiffness is one positive scalar whose value is left open. So the "six minus five" story is not required; there is just one number to compute.

Third, when the five-unit bath correction is computed explicitly on the natural discrete cell, the five modes that were supposed to supply it turn out to be pure-gauge directions — the geometry annihilates them, giving zero, not five. A genuine five-unit screening would need an extra physical ingredient that the geometry alone does not provide. Within the conditional polar, one-copy coherent-channel model, the screening problem reduces to a single dimensionless invariant (target value one-fifth on the historical branch). The strict physical Hessian remains the more general test and may contain a higher-dimensional adjoint multiplicity space.

The upshot is a much sharper statement than "$Z_3 = 1$." The physical colour stiffness is a product of two things — an action-conversion factor and the reduced strength of the physical colour embedding — and the value $Z_3 = 1$ holds exactly when both equal one. That is a specific, falsifiable branch (the paper calls it the Unit Class-Action branch), gated by three explicit checks the eventual substrate calculation must pass. On that branch, and using the previously derived record-layer boundary of 5.8 TeV with the frozen quark thresholds, the three-flavour QCD scale is

Λ_QCD^VBF = 225.923 MeV,

obtained before consulting any measured strong-interaction quantity. But this is the prediction *of that branch*, not yet a derived number: the strongest executable benchmark so far contains no intrinsic screening, but because it is matter-census based it does not determine the primitive physical value. Only the substrate calculation, passing its three checks, would settle which is physical.

Confinement is treated separately, and here too the paper is careful about what is proved versus assumed. An exact reference model — one that inserts a positive energy cost for centre flux by hand — does confine, cleanly and uniformly. The paper keeps this as a reference but does not present the inserted cost as derived. Alongside it, it builds the architecture for a genuine derivation in which that cost would instead emerge from the substrate with no insertion.

The mechanism both versions share: the centre of SU(3) has three classes; gluons carry the trivial class while a quark carries nonzero triality. A global operator measures which class crosses a surface between a quark and an antiquark. The paper introduces an exact projector that registers whenever a link carries nonzero centre flux, and Gauss' law forces every separating surface between unscreened sources to contain at least one such link. The projector's eigenvalues are exactly 0 and 1, so the *distinction* it draws cannot blur as the lattice is refined. What must still be derived is the physical *energy* multiplying that projector — in the reference model it is inserted; in the substrate route it is owed. Stacking the slices gives a linear energy: in the quenched centre-projector reference model a string with positive lower slope, and in full QCD a string that can break — the constructed breaking channel ends in two colour singlets. A general no-free-colour theorem for the complete physical Hilbert space remains open.

So the honest one-line summary is this. The paper turns a vague expectation ("the strong stiffness is one unit") into a sharp, falsifiable prediction resting on one computable number, and it shows that the clean "6 − 5 = 1" story behind that expectation does not survive scrutiny as a derivation — the 6 is a particle count, the 5 is not supplied by the geometry, and the physical value is one specific branch among the possibilities, awaiting the substrate calculation that would confirm or overturn it.

This paper is Paper 5 of the final eight. Paper 4 supplies the gauge-response framework and the direct colour boundary. Papers 1–3 supply the scalar and quark thresholds. Paper 5 turns these inputs into the (branch-conditional) nonperturbative QCD scale and the centre-flux obstruction to isolated colour. Later papers consume these outputs in the common Standard Model parameter and observable ledger.

Headline Results

Quantity	Status	Return
Matching scheme	fixed	Native VERSF background-field scheme (VBF)
Matching scale $\mu\star$	independently frozen	5.80 TeV
Matter-census index $I_C^\wedge$matter	exact on GRCE-1C(0)	6 ($\chi_C^\wedge$census = 48)
"6 = primitive stiffness"?	**rejected** without embedding theorem (Firewall C1)	—
Primitive colour carrier	derived (class-owned, Schur)	one adjoint copy, $\mathcal{K}\_C = z\_C\mathbb{1}_8$
Physical colour stiffness	Z_3 = C_hol · n_C^red	both factors open
Auxiliary adjoint multiplicity m_adj	minimal one-copy ansatz	1
Mixed-response vector b	constructed historical MC0–MC2 return	(5)
Auxiliary Hessian C	MC-branch output	(5)
Finite-cell closure operator B_C	**computed** (wheel W_6)	spec(B_C B_C†) = {1,2,2,4,4,5}
Finite-cell mixed block b_C	**computed**	**0** (five exact modes in ker B_C)

Quantity	Status	Return
$b = 5$ vs $b_C = 0$	tension — needs intertwiner $\mathcal{E}_\eta$	open (F5)
Factorised branch (corpus-supported)	**computed**	**$b = 0$, $Z_3 = 6$, $g_s = 1/\sqrt{6} \approx 0.408$**
MC branch (needs $\mathcal{E}_\eta = 5/6\ \mathcal{E}_F$)	**constructed conditional realisation; constitutive inputs unevaluated**	**$b = 5$, $C = 5$, $Z_3 = 1$**
Decisive invariant γ_η (polar one-copy model only)	local test *within that model*	$\gamma_\eta=1/5 \Leftrightarrow Z_3=1$ only under MC0 & $C_{hol}=1$; $\gamma_\eta\to\infty \Leftrightarrow Z_3=u$; $\gamma_\eta=0 \Leftrightarrow Z_3=0$
Auxiliary response shape $\Xi = i\alpha\, q_F\, T_a$	derived *conditional on polar record-link ansatz + one-copy channel*	q_F = rim circulation
Schur correction $b^\dagger C^{-1} b$	MC-branch output	5
Colour stiffness Z_3	MC-branch output	1
Boundary coupling g_s	branch consequence	1
Boundary α_s	branch consequence	$1/(4\pi) = 0.0795775$
Λ_6	conditional branch output	73.2458 MeV
Λ_5	conditional branch output	150.611 MeV
Λ_4	conditional branch output	198.873 MeV
$\Lambda_{3,MC}^{VBF}$	conditional branch output	225.923 MeV
Primitive centre generator H_c	derived	$\mathrm{diag}(1, 1, -2)/3$
Primitive flux norm $\lVert\Phi_1\rVert_{adj}$	derived	$4\pi/\sqrt{3}$
Projector inequality $\sum P_\ell^{\neq 0} \geq P_1+P_2$	exact reference algebra	—
Uniform quantum slice gap $\varsigma_\tau^{qu} \geq \Lambda_3^2$	reference-completion result only	(App. A)
Quenched reference lower-bound scale $\sigma_0 = \Lambda_{3,U}^2$	MC4 assignment, not physical YM; equality candidate (F22)	0.0510413 GeV^2
Implied bath coefficient $B_{\tau,MC}^{implied}$	reference-branch inverse matching	$2.47466\times10^{-5}\ \mathrm{GeV}^4$
Physical Z_3	owed from D7	—
Physical B_τ	owed from D7/D8	—
Uniform physical triality gap	owed from D8	—
$c_{\sigma,\tau}$	owed	—
Strict adjoint-only lower slope	owed from D8	positive value to be evaluated
Physical $N_f=0$ tension relation	owed from D9 and existence of limit	expected form $c_{\sigma,0}\ \Lambda_0^2$

Quantity	Status	Return
Full-QCD asymptotic result	string-breaking structure only; no-free-colour theorem open	string breaking into singlets

Abstract

The consolidated VERSF QCD gate reduces the strong boundary to $Z_3 = \chi_C/8 - b^\dagger C^{-1} b$ and formulates a conditional triality-gap route to confinement. The present paper has **two outputs of different status**. It first isolates a structural correction: the number 6 in the direct colour response is a **matter-census index** ($3(1+½+½) = 6$ over the quark representations), not the primitive colour stiffness. Class ownership (bath selection + Schur on the irreducible adjoint) then shows the primitive response is a single scalar, $\mathcal{K}_C = z_C \mathbb{1}_8$, so the physical boundary is $Z_3 = C_{hol} \cdot n_C^{red}$ with *both* factors open — the "6 − 5 = 1" cancellation is not required, only a pre-registered branch. On the Unit Class-Action branch ($C_{hol} = n_C^{red} = 1$) this gives $Z_3 = 1$ and $\Lambda_{3,U}^{VBF} = 225.923$ MeV, a sharp prediction conditional on three blind unit-return certificates and with no measured strong-sector input. It then evaluates the three local defect operators explicitly on the canonical K = 7 wheel cell. The closure map is $B_C = [\mathbb{1}_6 \mid \mathbb{1}_6 - S]$, with $\mathrm{spec}(B_C B_C^\dagger) = \{1,2,2,4,4,5\}$, and it satisfies $\ker B_C = \mathrm{Im}\, D$. The five exact C_6 modes are therefore pure-gauge directions with zero closure mixing, $\mathbf{b_C = 0}$. With the canonical response geometry specified, the factorised census benchmark is evaluated while the historical MC block is a constructed conditional realisation. The polar construction fixes the link from the record/completion map; gauge covariance shows only the intrinsic (non-coboundary) response contributes; W_6 symmetry forces its shape $\Xi_{e,a} = i\alpha\, q_{F,e} T_a$. The screening then reduces to $Z_3 = u \cdot \gamma_\eta/(1+\gamma_\eta)$ with a single dimensionless invariant $\gamma_\eta = c_\eta/(\alpha^2 u)$. Under MC0 and $C_{hol} = 1$, the MC branch requires $\gamma_\eta = 1/5$ ($Z_3 = 1$); the factorised benchmark is the $\alpha = 0$ limit ($Z_3^{census} = 6C_{hol}$, $g_s = 1/\sqrt{6}$). Within the polar one-copy model its screening reduces to γ_η; the strict physical Hessian remains more general. The strongest executable *benchmark* carries no intrinsic screening ($\alpha = 0$, $Z_3^{census} = 6C_{hol}$), but that benchmark is matter-census based and does not determine — or falsify — the physical primitive Z_3. MC4 assigns $\Lambda_{3,U}^{-1}$ as the normalising length of the quenched reference model; it is not a derived $N_f = 3 \to 0$ matching result. The paper also constructs the intrinsic global centre sectors (acting on the original SU(3) links, no second $\mathbb{Z}_3$ theory), the transfer-matrix and physical-Hamiltonian construction required for the penalty-free sector-gap calculation (owed, D8), and the penalty-free sector-gap and string-tension targets.

The direct colour *census* index is $I_C^{matter} = 6$ ($\chi_C^{census} = 48$), but Firewall C1 (Section 5) forbids promoting this representation trace to the primitive stiffness. Class ownership (Theorem 1A) makes the primitive response a single scalar $\mathcal{K}_C = z_C \mathbb{1}_8$, so the physical boundary is $Z_3 = C_{hol} \cdot n_C^{red}$ with both factors open. The historical "6 = 1 + 5" decomposition — one protected unit plus five bath-relaxable units, giving $\Gamma_C^{(2)} = ½\sum[F_a^2 + 5(F_a + \eta_a)^2]$ and b = C = 5 — is retained only as an *optional pre-registered MC realisation*, not as a forced consequence of the census index; the wheel-cell computation (Section 6) shows the five modes are pure-gauge ($b_C = 0$) and the branch requires a separately supplied intertwiner ($\gamma_\eta = 1/5$).

On the Unit Class-Action branch ($C_{hol} = 1$, $n_C^{red} = 1$) one obtains — and, on the historical $u = 6C_{hol}$ realisation only (MC0), equivalently $\gamma_\eta = 1/5$ —

$$Z_3 = 1,\ g_s = 1.$$

QCD belongs to the propagating record layer rather than the Planck certification layer. The Route-R matching boundary is therefore the frozen record-coherence ceiling $\mu_\star = 5.80$ TeV. With the native VBF threshold values $m_t = 291.827$ GeV, $m_b = 4.86254$ GeV, $m_c = 1.11241$ GeV, sharp one-loop matching gives, on that branch,

$$\Lambda_6 = \mu_\star e^{-8\pi^2/7} = 73.2458 \text{ MeV},$$

$$\Lambda_5 = m_t^{2/23} \Lambda_6^{21/23} = 150.611 \text{ MeV},$$

$$\Lambda_4 = (m_t m_b)^{2/25} \Lambda_6^{21/25} = 198.873 \text{ MeV},$$

■ $\Lambda_3^{VBF} = \mu_\star^{7/9} (m_t m_b m_c)^{2/27} e^{-8\pi^2/9} = 225.923$ MeV.

The confinement calculation rests on the global centre lift. In the **centre-projector reference completion**, an auxiliary $\mathbb{Z}_3$ one-form carrier is introduced and the exact projector algebra below follows; the **strict route** instead uses the intrinsic surface action on the original SU(3) link Hilbert space (Section 12) and derives its energy separation from the physical transfer matrix. In the reference completion, a $\mathbb{Z}_3$ one-form gauge field records the centre component of the SU(3) link transport, with surface operator

$$U_\omega(\Sigma) = \prod_{\ell \pitchfork \Sigma} X_\ell.$$

The primitive coweight

$$H_c = \tfrac{1}{3}\,\mathrm{diag}(1, 1, -2)$$

satisfies $e^{2\pi i H_c} = \omega\, \mathbb{1}_3$ and fixes the primitive centre-flux class without residual normalisation freedom.

The minimal centre-electric Hamiltonian uses the spectral projector

$$P_\ell^{\neq 0} = (2 - X_\ell - X_\ell^\dagger)/3.$$

For every separating surface Σ, the operator inequality

$$\sum_{\ell \pitchfork \Sigma} P_\ell^{\neq 0} \geq P_{\tau \neq 0}$$

holds exactly: a product of $\mathbb{Z}_3$ eigenvalues can be nontrivial only if at least one factor is nontrivial. Adding this projector to the positive auxiliary-reduced VERSF bath Hamiltonian yields

$$\hat{H}_\perp^{red} = \hat{H}_{\perp,+}^{VERSF} + \sigma_0 \sum_{\ell \pitchfork \Sigma} P_\ell^{\neq 0}.$$

Hence every nonzero-triality quantum sector has energy per unit longitudinal length at least σ_0, independently of regulator and transverse volume.

MC4 assigns the branch-conditional three-flavour length $\Lambda_{3,U}^{-1}$ as the normalising correlation length of the quenched centre-projector reference model; this is not a derived $N_f = 3 \to 0$ matching theorem. The assembled reference Hamiltonian satisfies the regulator-uniform finite-regulator bound $\underline{\varsigma}_{\tau,a,R}^{qu} \geq \sigma_0$, $\underline{\sigma}_{ref} := \liminf_{L\to\infty} V_\tau(L)/L \geq \sigma_0$. Continuum passage requires compatible refinement and form convergence; existence of a tension limit and a Wilson-loop area law additionally require Premises RP and SC and a finite-slope upper bound or subadditivity theorem.

$$\sigma_0 = \ell_C^{-2} = \Lambda_3^2.$$

The longitudinal Hamiltonian is a sum of the transverse slice terms plus nonnegative longitudinal squares. Consequently

$$V_\tau(L) \geq \sigma_0 L - c_{end}.$$

The proven statement is the reference lower slope $\underline{\sigma}_{ref} := \liminf_{L\to\infty} V_\tau(L)/L \geq \sigma_0$. If a finite-slope upper bound or subadditivity later proves the limit $\sigma_{ref} := \lim V_\tau(L)/L$ exists, candidate saturation $\sigma_{ref} = \sigma_0 = \Lambda_{3,U}^2 = 0.0510413\ \mathrm{GeV}^2$ (guarded by F22) becomes meaningful.

The finite-regulator lower bound is regulator-uniform. It passes unchanged to the continuum and infinite-volume limit *provided* the regulated sectors form a compatible refinement system and the corresponding forms converge in the declared Mosco or norm-resolvent sense. Under Premises RP and SC and a finite-slope upper bound, transfer-matrix positivity gives the quenched reference-model area law. Fundamental quarks replace the asymptotic string by a variational singlet-breaking channel.

Strict closure remains conditional on two independently specified evaluations: the microscopic colour Hessian (local blocker D7 — the three defect operators B_C, L_T, B_R with substrate data) and the quotient-compatible global SU(3) measure with its reflection-positive transfer matrix (global blocker D8). The paper therefore closes the operator architecture and the pre-registered branch, **not** the physical QCD confinement calculation. The reference centre-projector completion (Appendix A) proves that an assembled Hamiltonian with an inserted positive penalty confines; that penalty is inserted, not derived (Section 17.0). The scope boundary against the general Yang–Mills problem is Section 3.8; the inserted-versus-derived boundary is Section 17.0; the branch-versus-physical boundary is Section 31.

Contents

Part I — Strong-Sector Closure Problem

1. Aim and place among the final eight
2. Inherited QCD architecture
3. Binding conventions and non-insertion firewalls
4. The declared minimal-completion principles

Part II — Numerical Colour Boundary

5. Direct colour susceptibility
6. Saturated-bath adjoint census
7. Exact Schur reduction and $Z_3 = 1$
8. Record-layer matching surface
9. Frozen threshold ledger
10. Numerical strong scale

Part III — Global Centre Lift

11. Local algebra versus global colour group
12. The $\mathbb{Z}_3$ one-form lift
13. Primitive flux normalisation
14. Triality Hilbert sectors

Part IV — Complete Quantum Transverse Bath

15. Regulated physical Hilbert space
16. The centre-electric projector
17. The complete positive bath Hamiltonian
18. Coercivity, compactness and vacuum uniqueness
19. Exact quantum triality gap
20. Uniform continuum and infinite-volume limit

Part V — Longitudinal Completion and String Tension

21. Discrete slicing inequality
22. Continuum slicing theorem
23. String-Tension Lower Bound and Candidate Saturation
24. Scale lock and numerical string tension
25. Wilson-loop area law
26. Bath areal coefficient and RG lockstep
27. Full-QCD string breaking

Part VI — Closure and Audit

28. Numerical package
29. Falsification conditions
30. Open-debt ledger
31. The Non-Perturbative QCD Scale and Triality-Gap Completion Theorem
32. Closure certificate
33. Conclusion

Part I — Strong-Sector Closure Problem

1. Aim and Place Among the Final Eight

This paper is Paper 5 of the final eight-paper VERSF Standard Model completion programme.

Paper 4 constructs and evaluates the gauge-response framework. It supplies the direct colour boundary stiffness used here. Papers 1–3 supply the scalar vacuum and the quark threshold values. Every input this paper consumes is frozen upstream; every output it produces is frozen for downstream consumption.

The present paper has two logically independent tasks:

colour boundary ⟶ Λ_{QCD},

and

global triality ⟶ quantum gap ⟶ string tension or string breaking.

The first is perturbative once one nonperturbative boundary value is fixed. The second cannot be replaced by a beta function: no statement about the running of the coupling, however sharp, decides whether isolated colour exists. The two tasks meet only at the end, where the unique generated colour length fixes the primitive centre-cell area.

1.1 Gate question

Does one frozen VERSF strong-sector completion return a numerical three-flavour QCD scale and a quantum nonzero-triality gap that survives the continuum and infinite-volume limits, without using any measured strong-interaction quantity upstream?

The paper answers this in two parts of different status. The **scale branch** returns a numerical $\Lambda_{\{3,MC\}}$ with no measured strong-sector input, conditional on the pre-registered $Z_3 = 1$. The **confinement side** is not declared passed: the paper constructs the full operator architecture for a substrate derivation (Sections 20′, Parts IV–V of the strict route) but the microscopic Hessian and global colour measure that would populate it are not yet evaluated. The strict gate question is therefore **open**, with a finite, specified list of executable calculations standing between the present state and closure (Section 30).

---

2. Inherited QCD Architecture

The inherited colour carrier is $\mathbb{C}^3$, with generators

$T_a = \lambda_a/2$, $\mathrm{tr}(T_a T_b) = \frac{1}{2} \delta_{ab}$.

The primitive colour connection contains no coupling:

$\mathcal{A}_\mu = A_\mu^a T_a$.

The local colour-holonomy residual is

$E_3(e_a) = T_a$,

and is exactly adjoint-equivariant.

The invariant threefold-bath metric is isotropic:

$g_{ab}^{FS} = \frac{1}{8} \delta_{ab}$.

The complete colour response is

$Z_3 = \chi_C/8 - b^\dagger C^{-1} b$.

The canonical coupling is

$g_s = Z_3^{(-1/2)}$.

Above a fixed set of active thresholds, the stiffness runs at one loop as

$dZ_3/d \ln k = (1/8\pi^2)(11 - 2N_f/3) + \cdots$.

The three-flavour one-loop scale, once the boundary is fixed, is

$\Lambda_3 = \mu_\star^{(7/9)} (m_t m_b m_c)^{(2/27)} \exp[-8\pi^2 Z_3(\mu_\star)/9]$.

Everything downstream of this section is the evaluation of the three symbols χ_C, b and C, followed by the independent centre-sector construction.

---

3. Binding Conventions and Firewalls

Each firewall below is a standing convention of the paper. Violating any of them voids the closure certificate; the corresponding falsifiers are listed in Section 29.

3.1 Running is not confinement

$\beta_s < 0$ does not imply an area law, a string tension or the absence of coloured asymptotic states. Asymptotic freedom is a statement about the ultraviolet; confinement is a statement about the structure of the physical Hilbert space in the infrared. The paper never lets the first stand in for the second.

3.2 Primitive versus canonical fields

The primitive connection is coupling-free. The coupling arises only after kinetic normalisation, as $g_s = Z_3^{-1/2}$. No step below inserts a coupling into the primitive layer.

3.3 Global group firewall

The local algebra $\mathfrak{su}(3)$ does not by itself determine the centre sector. Two theories with identical local algebra can differ in global group and hence in confinement behaviour. Confinement therefore requires the global lift of Part III; it cannot be read off local curvature.

3.4 Quantum-sector firewall

Triality is a Hilbert-sector or line-operator label. It is not a local label of a classical connection, and no classical field minimisation is accepted as a substitute for a Hilbert-sector energy anywhere in this paper.

3.5 Scheme firewall

Every numerical scale in this paper is quoted in the native VBF scheme. No claim is made that any number is an MS-bar value without a finite scheme bridge. The bridge itself is a declared open debt (Section 30).

3.6 Threshold firewall

The threshold values are frozen VBF decoupling values supplied by the mass programme. They are not replaced by measured masses, and no measured mass is permitted to leak in through them.

3.7 Non-insertion firewall

Forbidden upstream inputs include: measured α_s, measured Λ_{QCD}, lattice string tension, hadron masses, deconfinement temperature, and any fitted matching scale. Falsifier F19 is triggered by any such insertion.

3.8 Scope firewall

The triality-gap theorem of this paper is a theorem of the explicit NQTC-1 completion: the declared Hilbert space, the declared positive bath Hamiltonian and the declared centre penalty. It does not claim that every admissible continuum construction of four-dimensional Yang–Mills theory possesses a mass gap, and it is not presented as a resolution of that general problem. The reference construction establishes an exact finite-regulator nonzero-triality gap and, under Premise SC, a positive asymptotic lower slope for the explicitly assembled centre-projector Hamiltonian. It does not establish that the VERSF substrate generates that Hamiltonian or close the strict physical confinement gate. It does **not** establish that the VERSF substrate generates that Hamiltonian, nor does it close the strict physical confinement gate: the centre penalty is inserted, the continuum limit requires the convergence hypothesis, and candidate string-tension saturation is unproved. The general-construction question is pursued separately, in the entropy-modulated Yang–Mills programme of the same author, and nothing in this paper depends on that programme.

---

4. Declared Minimal-Completion Principles

The paper does not leave its decisive steps as untyped future calculations. It fixes the minimal completion through four structural principles (MC1–MC4), plus one historical reference-branch premise (MC0, stated in Section 6 / the premises block) that applies only to the six-to-one MC realisation and is not part of the general physical class-owned route. Each is stated as a premise, carries a direct falsifier in Section 29, and is used exactly where cited.

Principle MC1 — Universal primitive action unit

Operationally commensurable quadratic sectors share one action unit. The primitive closure excitation has unit cost:

$C_{hol} = r = 1$.

Falsifier: F1, F4.

Principle MC2 — Saturated-bath relaxation

In a saturated threefold bath, only total class weight is protected. Every same-order adjoint redistribution orthogonal to the primitive colour-curvature residue is fully relaxable.

Falsifier: F2, F3.

Principle MC3 — Record-layer boundary

Propagating QCD fields belong to the record layer. Their ultraviolet matching surface is the record-coherence ceiling, not the Planck metric-certification floor:

$\mu\star = 5.80$ TeV.

Falsifier: F7.

Principle MC4 — Primitive centre-cell saturation

A nonzero triality crossing occupies one primitive colour-correlation cell. One primitive closure unit per cell fixes the centre energy density. In the quenched centre-projector reference model the assigned normalising colour length (MC4, D9) is

$\ell\_C = \Lambda_3^{-1}$,

so the reference-model centre tension (quenched, Λ_3-normalised) is

$\sigma_0 = \ell\_C^{-2} = \Lambda_3^2$.

Falsifier: F16.

Principle MC0 — Benchmark-to-Direct Identification

In the historical six-to-one MC reference realisation *only*, identify the unscreened direct response with the matter-census benchmark:

$u = I\_C\text{\^{}matter} \cdot C\_hol = 6\ C\_hol$.

This is a pre-registered branch premise, **not** a consequence of class ownership, and not the physical substrate result unless independently returned by the physical colour embedding. Its role is to make explicit that the equivalence "$Z_3 = 1 \Leftrightarrow \gamma\_\eta = 1/5$" holds only under $u = 6C\_hol$: in general, from $Z_3 = u\cdot\gamma\_\eta/(1+\gamma\_\eta)$, the unit condition $Z_3 = C\_hol$ requires

■ $\gamma\_\eta = C\_hol/(u - C\_hol)$,

which reduces to $\gamma\_\eta = 1/5$ precisely when $u = 6C\_hol$ (Principle MC0), and otherwise depends on the independently calculated direct coefficient u. Conflating the two silently restores the census identification that Firewall C1 forbids.

Named premises of the confinement branch

Beyond the MC principles, the confinement construction of Parts IV–V carries three explicit premises, each flagged where it is used and each guarding a step this paper states but does not derive:

- **Premise VU (vacuum uniqueness).** The kernel intersection $\bigcap_j \ker Q_j \cap \bigcap_\ell \ker P_\ell^{\neq 0}$, after gauge quotient and boundary conditions, is one-dimensional (Section 18). Not computed here.
- **Premise SC (square completion).** The longitudinal Hamiltonian admits the nonnegative-square-plus-endpoint decomposition of Section 21. Stated, not derived.
- **Premise RP (reflection positivity).** The assembled action is Osterwalder–Schrader positive (Section 17.1). Retained as premise; the centre-term transfer matrix is not constructed here.

These are additional to the inserted status of the centre penalty itself (Section 17.0). A confinement claim of this paper is a claim *modulo* VU, SC, RP and the MC4 assignment; the derivation route of Section 20′ is what would discharge them.

Part II — Numerical Colour Boundary

5. Matter-Census Index Versus Primitive Colour Stiffness

The earlier reading of this paper treated the number 6 as the primitive direct colour stiffness and then built the cancellation 6 − 5 = 1. That reading is now corrected at its root, following GRCE-1: the 6 is an **induced matter-census index** on a substitute benchmark carrier, not the physical fourteen-channel substrate response.

Paper 4 executes its first colour calculation on the factorised benchmark carrier $L^2(\pi) \otimes \mathscr{H}_{SM}$. There the curvature generator acts trivially on the wheel factor and through the complete Standard Model representation census on $\mathscr{H}_{SM}$. For canonically normalised colour generators,

$\mathrm{Tr}_{\mathscr{H}_{SM}}[\rho_3(T_a)\rho_3(T_b)] = 6\,\delta_{ab}$,

because, with Dynkin index $T(\mathbf{3}) = \tfrac{1}{2}$,

$3_{\text{generations}}\,[\,2\,T(\mathbf{3})_{Q_L} + T(\mathbf{3})_{u_R} + T(\mathbf{3})_{d_R}\,] = 3(1 + \tfrac{1}{2} + \tfrac{1}{2}) = 6.$

(Verified: the per-generation sum is 2, and 3 × 2 = 6.) Define this the **matter-census colour index**

■ $I_C^{matter} = 6$, and in the susceptibility convention ■ $\chi_C^{census} = 8\,I_C^{matter} = 48$.

These are exact on the GRCE-1C(0) benchmark. But their interpretation is binding: I_C^{matter} counts the canonically normalised matter representations that respond to *one* colour connection. It does **not** count six primitive colour connections, six independent colour baths, or six independently owned copies of the colour kinetic action.

The benchmark itself proves the limitation. Its product-factor structure makes the wheel and gauge actions commute exactly, places the curvature response on the coherent wheel line, and makes every retained wheel auxiliary orthogonal to it — so $H_{F\eta} = 0$ on that branch, and the returned $Z_3^{census} = 6\, C_{hol}$ is a matter-census-induced response, not the evaluated physical fourteen-channel coefficient.

Firewall C1 — Census is not primitive stiffness

No downstream calculation may identify $I_C^{matter} = 6$ with the primitive tree-level colour stiffness without an independently derived gauge-to-record embedding proving that equality. The value 6 may legitimately reappear in matter-current traces, fermionic vacuum polarisation, the beta-function census, threshold increments, or an independently calculated physical response — but it may not be promoted from a representation trace to a primitive substrate coefficient by notation alone. A physical microscopic calculation could still return $Z_3 = 6$; if it does, that will be a substrate output, not an inheritance from the matter census. (The scale sensitivity makes the stakes concrete: relative to $Z_3 = 1$, $Z_3 = 6$ rescales the scale by $\exp[-(8\pi^2/9)(6-1)] \approx 8.9\times10^{-20}$, giving $\Lambda_3(6) \approx 2.0\times10^{-17}$ MeV (= 2.0×10^{-20} GeV) rather than 226 MeV — the census value and the physical value cannot be casually identified.)

The "doubly certified" language of earlier drafts is corrected accordingly: the two Paper-4 branches doubly certify the *census index* 6 (= $\chi_C/8$), not the primitive stiffness.

5A. Class Ownership of the Primitive Colour Response

The colour census supplies one genuinely indistinguishable threefold class $\mathcal{C} \simeq \mathbb{C}^3$. The no-private-ledger and bath-selection results (gauge-closure programme) imply its individual colour labels carry no independently conserved accounts; only the total class structure stands. Continuous bath transport gives one $SU(3)_C$ group, and local comparison gives one colour connection. The primitive gauge response is therefore attached at **colour-class grain**, not replicated over every quark species, chirality and generation that later probes the connection.

Theorem 1A — Class-Owned Primitive Colour-Carrier Theorem

Let $\mathcal{C} \simeq \mathbb{C}^3$ be the minimal transport-stable colour class generating one $SU(3)_C$ connection, and assume the primitive curvature action is owned by this class and not replicated by the number of matter representations acting on the connection. Then the primitive colour-curvature response carrier contains **one adjoint copy**,

■ $\mathscr{R}_C^{prim} \simeq \mathfrak{su}(3)_C$,

and every positive gauge-invariant zero-momentum response on it has the form

■ $\mathcal{K}_C^{prim} = z_C \mathbb{1}_8$, $z_C > 0$.

Generation, chirality and matter-species multiplicities do not multiply z_C.

Proof

Bath selection identifies the full threefold class — not any labelled component — as the transport subject; local distinguishability conservation generates one connection. Its curvature is in the adjoint of $SU(3)_C$, which is irreducible, so by Schur's lemma every SU(3)-equivariant self-adjoint operator on one adjoint copy is proportional to the identity, $\mathcal{K}_C^{prim} = z_C\mathbb{1}_8$; positivity of the gauge Hessian gives $z_C > 0$. The matter representations Q_L, u_R, d_R are sources and probes of the same connection: counting them changes current traces and vacuum polarisation but creates no additional primitive copies of the connection, so their multiplicity cannot multiply z_C without a further microscopic theorem identifying the matter-census carrier with the substrate response carrier. ■

Corollary 1A.1 — One scalar remains. Class ownership removes the apparent sixfold primitive multiplicity but does not fix the scalar: $\mathcal{K}_C = z_C\mathbb{1}_8$, *not* $z_C = 1$. The primitive colour problem is reduced from an unknown operator to one positive invariant.

Corollary 1A.2 — The six-to-one cancellation is not required. The physical calculation need not begin from 6 and construct an auxiliary correction of exactly 5; it begins from one class-owned carrier and evaluates z_C. Accordingly, $b = (5)$, $C = (5)$, $b^\dagger C^{-1} b = 5$ are **not** implied by class ownership, bath saturation, or the matter-census index. They remain a pre-registered constitutive branch unless independently returned by the microscopic response Hessian.

---

6. Saturated-Bath Adjoint Census

The curvature residual transforms as the adjoint **8**. By gauge equivariance, only auxiliary adjoint copies can mix linearly with it; every other representation is forbidden a linear coupling by Schur's lemma.

Minimal One-Copy Adjoint Auxiliary Ansatz. The smallest auxiliary sector capable of screening the colour curvature contains one adjoint copy,

$\mathcal{V}$_η^adj = ℝ ⊗ **8**, m_adj = 1.

This is the pre-registered minimal branch, not a census theorem: one primitive colour connection contains one adjoint *response* copy (Theorem 1A), but that does not prove the record, completion and auxiliary sectors contain only one adjoint copy — multiple physical adjoint auxiliaries could arise from different record or completion mechanisms without creating multiple primitive colour connections. The complete microscopic tangent census must verify that no additional physical adjoint copies survive the gauge and null quotients. Falsifier F3 accordingly *revises* the minimal branch rather than automatically falsifying colour propagation.

Status of Theorem 1: pre-registered MC branch, not a substrate calculation

The decomposition of this section is the **MC-branch prediction $\mathcal{B}$**_MC = (χ_C, m_adj, b, C, Z_3) = (48, 1, 5, 5, 1). Following the discipline the projection route imposes (Section 20′), these values are **not** results of evaluating the substrate Hessian; the historical six-to-one MC realisation is fixed by Principles MC0–MC2: MC0 supplies u = 6C_hol, MC1 supplies the protected unit, and MC2 supplies complete relaxation. When the microscopic evaluation of Section 20′ is carried out, exactly one of three outcomes is permitted: $\mathcal{B}$_Hessian = $\mathcal{B}$_MC (branch confirmed), $\mathcal{B}$_Hessian ≠ $\mathcal{B}$_MC with $Z_3 > 0$ (branch revised, scale recomputed from the evaluated Z_3), or $Z_3 \le 0$ (physical colour completion falsified). Until that audit closes, the values below are a valuable pre-registered prediction and should not be described as calculated. The theorem's internal logic (that MC0–MC2 *force* these numbers — MC0 the benchmark coefficient u = 6C_hol, MC1 the protected unit, MC2 relaxation) is exact; what is open is whether the substrate obeys MC0–MC2.

Theorem 1 — Pre-Registered MC Six-to-One Completion

The direct six-unit colour stiffness decomposes uniquely as

6 = 1 + 5,

where:

1. the one-unit term is the primitive irreducible holonomy cost;
2. the five-unit term is the saturated-bath redistribution cost;
3. complete bath relaxation cancels the five-unit term at its minimum;
4. positivity is maintained throughout.

The unique minimal quadratic form with these properties is

■ $\Gamma\_C^{(2)} = \frac{1}{2} \sum_{a=1}^{8} [\, F_a^2 + 5(F_a + \eta_a)^2 \,]$.

Under MC0–MC2, the historical one-copy reference completion has (H_FF, H_Fη, H_ηη) = (6, 5, 5). But the canonical wheel-cell closure operator does **not** derive this mixed block — it returns b_C = 0 (Section 6). The MC completion is therefore a **pre-registered target** for the complete physical Hessian, not its current output; whether the physical intertwiner supplies $H_{F\eta} = 5$ or leaves $b = 0$ is the open test of Section 6.

Proof

By adjoint equivariance and isotropy of g^{FS}, the most general positive quadratic form on one curvature adjoint and one auxiliary adjoint is diagonal in the generator label and reads, per generator,

$$\tfrac{1}{2} [\, u F^2 + 2v F\eta + w \eta^2 \,].$$

Three conditions determine (u, v, w) exactly:

(i) **Direct coefficient.** Freezing the auxiliary at $\eta = 0$ must return the direct response: $u = 6$.

(ii) **Complete relaxation.** Principle MC2 requires the redistributable part to relax completely, so the minimiser must be $\eta^\star = -F$. Minimising over η gives $\eta^\star = -(v/w)F$; complete relaxation therefore forces $v = w$.

(iii) **Protected primitive unit.** The reduced coefficient after relaxation must equal the single protected unit of Principle MC1:

$$u - v^2/w = 1.$$

With $v = w > 0$, condition (iii) reads $6 - w = 1$, hence $w = v = 5$. Substituting $(u, v, w) = (6, 5, 5)$ and completing the square:

$$\tfrac{1}{2}[6F^2 + 10F\eta + 5\eta^2] = \tfrac{1}{2}[F^2 + 5(F + \eta)^2],$$

which is the displayed form. Uniqueness holds because (i)–(iii) are three independent equations in three unknowns with the single solution (6, 5, 5); positivity holds because the completed form is a sum of squares with positive coefficients. ■

The coefficient 5 on the canonical wheel cell — computed, and it does not support b = 5

The closure architecture at $K = 7$ is the wheel W_6: one hub, six boundary vertices, six rim edges, six spokes, and six hub–boundary–boundary triangles. On this explicit cell the three linearised defect operators of Section 20′.3 can be written and evaluated in closed form. This has been done and independently checked; the results are exact, and one of them **corrects the branch**.

Closure operator (computed). With edge amplitudes $x = (r_1 \ldots r_6, s_1 \ldots s_6)$ and the cyclic shift S, the linearised holonomy defect on triangle Δ_i is $\delta C_i = r_i + s_i - s_{i+1}$, giving

$$B_C = [\, \mathbb{1}_6 \mid \mathbb{1}_6 - S \,],\quad B_C B_C^\dagger = \mathbb{1}_6 + L_{C_6} = 3\mathbb{1}_6 - S - S^\dagger,$$

with **verified spectrum spec(B_C B_C†) = {1, 2, 2, 4, 4, 5}**. There is a genuine eigenvalue 5 here — the alternating face mode $f_5 = (1,-1,1,-1,1,-1)/\sqrt{6}$ — and it is *real*, unlike the withdrawn $L_7 = 7\mathbb{1} - J$ claim. (I verified spec(B_C B_C†), the alternating-mode eigenvalue, and spec(L_W) = {0,2,2,4,4,5,7} by direct computation.)

The crucial identity (computed). The closure operator and the edge–vertex incidence D satisfy

■ B_C D = 0, rank B_C = rank D = 6, ker B_C = Im D,

all verified exactly. Physically: every link perturbation generated by vertex data is an exact/pure-gauge deformation, and the linearised closure curvature annihilates it identically.

Consequence — b_C = 0, not 5. The five exact rim modes of the Hodge split $C^1(C_6) = \mathrm{Im}\,\delta \oplus H^1(C_6)$ (dim Im δ = 5, dim H^1 = 1) are exactly the gauge modes in ker B_C. Therefore their mixed block with physical colour curvature **vanishes**:

■ b_C = 0 (verified: B_C annihilates every Im D mode).

For independent record variables G_R acts on φ, not on colour-link curvature, so b_R = 0 as well; and for a divergence-free cycle representative the minimal transport defect also vanishes. **The canonical wheel-cell incidence complex does not, by itself, produce the five-unit Schur correction.** The honest statement replacing the earlier "verified cohomological origin" is:

The rim cohomology supplies five candidate relaxable directions. In the canonical wheel-cell incidence complex these directions are exact gauge modes annihilated by the linearised closure operator (b_C = 0). A nonzero five-unit Schur correction therefore requires an additional substrate-derived adjoint intertwiner — a physical coupling of colour curvature to those modes — that does not follow from the 5 + 1 Hodge count, nor from the incidence complex, alone.

This is a genuine tension with the MC branch, and it is recorded as such: the branch predicts $b^\dagger C^{-1} b = 5$, while the finite-cell closure operator alone gives b = 0. The branch is not thereby falsified — the missing intertwiner $\mathscr{E}$_η coupling the exact modes to physical curvature could supply the five — but that intertwiner is now the specific, named object the branch requires and the incidence complex does not provide (Falsifier F5, sharpened; Debt D7). The eigenvalue-5 *of B_C B_C†* is real but is a spectral property of the closure Gram operator, not the Schur correction $b^\dagger C^{-1} b$; conflating them is precisely the error to avoid.

Transport and record (computed). With the minimal covariant current J_e = κ_e(ρ_{head} − U_e ρ_{tail}), the density-only transport Hessian is G_T^{(hh)} = 2κ² L_W², with **verified spectrum κ²{0, 8, 8, 32, 32, 50, 98}**; κ is not fixed by the corpus. On the locally trivial record branch B_R = B_C ⊗ 𝟙_rec, with the same incidence spectrum. The total canonical finite-cell Hessian is

$$\mathbb{G}\star^{\{W_6\}} = G_V + 2\rho_c^6 B_C^\dagger B_C + 2 L_T^\dagger L_T + 2\rho_c^4 B_R^\dagger B_R.$$

The decisive two-branch computation (executed). The mixed block $H\_F\eta = \mathscr{E}\_F^\dagger\, \mathbb{G}^\star\, \mathscr{E}\_\eta$ has now been evaluated on two fully specified canonical embeddings, and both results are verified by direct computation. Using the coherent wheel profile $\chi_0 = (1,\dots,1)$, the stationary measure $\pi = (7/31, 4/31,\dots,4/31)$ (for which $\|\chi_0\|^2\_\pi = 1$, verified), the response metric $M\_resp = C\_hol\ \mathrm{diag}(\pi) \otimes \mathbb{1}$, and the colour normalisation $\mathrm{Tr}[\rho_3(T_a)\rho_3(T\_b)] = 6\delta\_ab$ (fixing $H\_FF = 6C\_hol$):

- **Factorised branch** — $\mathscr{E}\_F = \chi_0 \otimes \rho_3$, auxiliaries $\mathscr{E}\_\{\eta,fact\}$ in the incoherent complement $\mathscr{H}\_inc$ (each $f\_r \perp \chi_0$). Then $H\_F\eta = 0$ because the auxiliaries are π-orthogonal to the coherent curvature, giving

$b = 0$, $Z_3 = 6C\_hol$, and at $C\_hol = 1$: $Z_3 = 6$, $g\_s = 1/\sqrt{6} \approx 0.408248$.

This is the specification the existing corpus actually supports — it is Paper 4's executed factorised gauge-response branch, and it agrees with the wheel identity $\ker B\_C = \mathrm{Im}\, D$ that places the five exact modes in the closure kernel.

- **MC non-factorising branch (constructed conditional realisation).** With an *independent* coherent adjoint auxiliary copy $\mathscr{E}\_\{\eta,MC\} = (5/6)\mathscr{E}\_F$ and bare stiffness $K\_\eta\hat{}sub = (5/6)C\_hol\mathbb{1}_8$ — **two inputs chosen precisely to realise the target block** — one gets $H\_F\eta = (5/6)(6C\_hol) = 5C\_hol$, $H\_\eta\eta = [(5/6)^2\cdot 6 + 5/6]C\_hol = 5C\_hol$, so $b^\dagger C^{-1}b = 5C\_hol$ and

$b = 5$, $C = 5$, $Z_3 = C\_hol$, and at $C\_hol = 1$: $Z_3 = 1$. (The algebra is verified — the block reduces to $\Gamma\_C^{(2)} = \frac{1}{2}\sum[F^2 + 5(F+\eta)^2]$ — but this is internal consistency of a *constructed* realisation, not evidence that the physical substrate supplies those two inputs.)

The verdict of the computation. The strongest *executable benchmark* — the factorised carrier — returns zero auxiliary screening, **$b = 0$**, and hence $Z_3\hat{}census = 6C\_hol$. This does **not** mean the physical primitive stiffness is probably six; it means the benchmark carrier does not contain the physical intertwiner. It does not falsify the Unit Class-Action branch. The MC branch ($Z_3 = 1$) is realisable by adjoining two specific constitutive statements that do *not* follow from the wheel incidence geometry:

- $\mathscr{E}\_\eta = (5/6)\, \mathscr{E}\_F$ (an independent coherent adjoint copy, on the response line, reproducing five-sixths of the curvature-induced record mismatch), and
- $K\_\eta\hat{}sub = (5/6)\, C\_hol\, \mathbb{1}_8$ (its bare substrate stiffness).

The entire remaining burden of the *historical MC0 one-copy realisation* is now exactly these two equations (the general Unit Class-Action branch instead requires the physical certificate $C\_hol\cdot n\_C\hat{}red = 1$, without assuming $u = 6C\_hol$, one auxiliary copy, or the polar construction). The invariant that decides it is the Schur term $S\_C = H\_F\eta\ H\_\eta\eta^{-1}\ H\_\eta F$, with definite outcomes: $S\_C = 5\cdot\mathbb{1}_8$ confirms $Z_3 = 1$; $S\_C = 0$ gives $b = 0$ and $Z_3 = 6$, falsifying the branch; $0 < S\_C < 6$ gives a revised colour boundary; $S\_C \geq 6$ makes the kinetic response null or unstable. The branch can no longer be protected by interpretive flexibility — it is a yes/no question about whether the substrate action supplies the independent coherent adjoint intertwiner.

Reduction to one invariant (conditional on the polar record-link ansatz). *Ontology note:* the strict physical calculation (Section 20′.3) treats (ρ, U, J, Φ) as independent substrate variables. The construction below instead adopts a **Conditional Polar Record-Link Ansatz** — U derived from Φ by polar decomposition — as a useful one-parameter constitutive *model*, not the definition of the physical link. Under this ansatz the auxiliary embedding's *shape* is derived and only one dimensionless number remains. The link is fixed from the record/completion map by polar decomposition,

$$M_{ij} = P_j \Phi_{ij} P_i, \quad U_{ij}^\star = M_{ij}(M_{ij}^\dagger M_{ij})^{-1/2},$$

which is gauge-covariant ($M_{ij} \mapsto G_j M_{ij} G_i^{-1}$ gives $U_{ij}^\star \mapsto G_j U_{ij}^\star G_i^{-1}$). Its auxiliary derivative, via the Sylvester equation $H X_\alpha + X_\alpha H = U^\dagger \partial_\alpha M - (\partial_\alpha M)^\dagger U$, reduces at a locally trivial background to $\partial_\alpha U|_\star = \tfrac{1}{2}[\partial_\alpha M - \partial_\alpha M^\dagger]|_\star$. *Right-trivialising, the edge response splits as* $\Theta_{ij} = (\Omega_j - Ad\{U_{ij}\}\Omega_i) + \Xi_{ij}$: a pure-gauge coboundary plus the **intrinsic** response Ξ. Only Ξ survives the plaquette sum (the coboundary telescopes — the microscopic form of $\ker B_C = \mathrm{Im}\, D$). So endpoint-frame bath changes give $\Xi = 0$, hence $b = 0$; nonzero screening requires a genuine intrinsic support/bond/memory/completion deformation.

On W_6, adjoint equivariance, D_6 symmetry, one auxiliary copy, and coupling to the coherent face channel **force the shape**

$$\Xi_{\{e,a\}} = i\alpha\, q_{\{F,e\}}\, T_a, \quad q_F = B_C^\dagger c_0 = (1,\dots,1; 0,\dots,0)/\sqrt{6} \text{ (uniform rim circulation, verified),}$$

leaving only the amplitude α and the bare stiffness c_η free. With $J_\eta = \alpha J_F$, the blocks are forced to $H_{FF} = u\mathbb{1}_8$, $H_{F\eta} = \alpha u\mathbb{1}_8$, $H_{\eta\eta} = (\alpha^2 u + c_\eta)\mathbb{1}_8$, giving

■ $\boldsymbol{Z_3 = u - \alpha^2 u^2/(\alpha^2 u + c_\eta) = u \cdot \gamma_\eta/(1 + \gamma_\eta), \quad \gamma_\eta \equiv c_\eta/(\alpha^2 u),}$

verified symbolically. **The entire one-copy screening problem is one dimensionless invariant $\boldsymbol{\gamma_\eta}$.** The MC branch ($s_C/u = 5/6$) requires exactly

■ $\boldsymbol{\gamma_\eta = 1/5}$ — *the bare restoring cost of the coherent auxiliary mode must be one-fifth of the holonomy-response cost it induces* — verified to give $Z_3 = C_{hol} = 1$. The factorised branch is the $\alpha = 0$ limit: $\Xi = 0$, $s_C = 0$, $Z_3 = 6C_{hol}$.

The remaining *local* calculation is now a single number: $(\Phi_{ij}, P_i, P_j) \to \partial_\eta M \to \partial_\eta U \to \Xi \to \alpha$, and independently $\Gamma_{sub} \to H_{\eta\eta}^{sub} \to c_\eta$, then $\gamma_\eta = c_\eta/(\alpha^2 u)$. (This reduction is conditional on the one-copy adjoint assumption, the polar-link construction, the physical embedding, the response metric, the normalisation u, and the existence of the intrinsic coherent channel. The *full* strict closure additionally requires the global blocker D8.) The verdict within the polar one-copy model, under MC0 and $C_{hol} = 1$: $\gamma_\eta = 1/5 \Longrightarrow Z_3 = 1$ (branch confirmed); $\alpha = 0 \Longrightarrow \gamma_\eta \to \infty \Longrightarrow Z_3 = u$ (no intrinsic screening on that carrier); $\gamma_\eta = 0 \Longrightarrow Z_3 = 0$; any other positive γ_η gives a uniquely revised Z_3. The strict physical test is the full projected Hessian and need not reduce to γ_η. The corpus does not yet supply the differentiated record/completion map or the bare auxiliary stiffness, so γ_η is not yet computed — but the broad question "how does the bath alter edge transport?" is now a single coordinate-invariant number with a sharp branch prediction.

What this settles and what it leaves. The matrix mechanics of the three operators are now done on the canonical cell — this substantially closes the *operator* content of Debt D7. What remains for a physical Z_3 is narrower and specific: the colour embedding $\mathscr{E}_F : \mathfrak{su}(3) \to$ physical cycle modes, the auxiliary embedding $\mathscr{E}_\eta$ (the intertwiner above), the background $\rho\star$, the coefficients κ_e, the relative normalisations of A_C, A_T, A_R, the active record representation, and the tangent metric. Given these, $Z_3 = H_{FF} - H_{F\eta} H_{\eta\eta}^{-1} H_{\eta F}$ is mechanical — but the calculation as it stands does **not** confirm $b = C = 5$; it shows precisely where a genuine physical intertwiner, not a dimension count, is required.

---

7. MC-Branch Schur Reduction

From Theorem 1, the quadratic block in the adjoint-isotropic basis is

$H_{FF} = 6\cdot\mathbb{1}_8$, $H_{F\eta} = 5\cdot\mathbb{1}_8$, $H_{\eta\eta} = 5\cdot\mathbb{1}_8$.

In multiplicity-space notation,

■ $b = (5)$, ■ $C = (5)$.

The Schur correction is

$b^\dagger C^{-1} b = 5 \cdot (1/5) \cdot 5 = 5$.

Therefore, on the pre-registered branch,

■ $Z_{3,MC} = 6 - 5 = 1$.

The physical result remains $Z_3^{phys} = H_{FF} - H_{F\eta} H_{\eta\eta}^{-1} H_{\eta F}$, with $H_{F\eta}$ still awaiting the physical intertwiner $\mathscr{E}_\eta$ (Section 6); the finite-cell closure operator alone gives $H_{F\eta} = 0$.

The auxiliary minimiser is $\eta\star = -F$, confirming complete relaxation. The full quadratic block is positive:

$$\det\begin{pmatrix} 6 & 5 \\ 5 & 5 \end{pmatrix} = 30 - 25 = 5 > 0.$$

The reduced gluon mode is positive and canonically normalised:

■ $g_s(\mu\star) = 1$, ■ $\alpha_s(\mu\star) = 1/(4\pi) = 0.0795775$ (MC-branch values, pending the Section 20′ evaluation).

Three remarks fix the interpretation. First, $Z_3 = 1$ is exact *on the MC branch*: it is the difference of the census index and the pre-registered screening, not yet a substrate return. Second, $g_s = 1$ on that branch states that the substrate would resist colour curvature with exactly one primitive unit. Third, the number carries its scheme with it: it is a VBF boundary value and enters the cascade of Section 10 only in that scheme.

7A. The Physical Colour Coefficient and the Unit-Action Target

Class ownership (Theorem 1A) has reduced the primitive colour response to one positive scalar. The physical reduced colour response is

$$\mathcal{K}_C = C_{hol} \cdot \mathcal{E}_3^\dagger\, \mathbb{W}^{1/2}(\mathbb{1} - S_\eta)\mathbb{W}^{1/2}\, \mathcal{E}_3,$$

where $\mathcal{E}_3 : \mathfrak{su}(3)_C \to \mathcal{H}_{rec}$ is the physical colour-to-record embedding, $\mathbb{W}$ the frozen response metric, and S_η the auxiliary screening operator after gauge and exact-null directions are removed. By adjoint isotropy (Schur again),

$$\mathcal{E}_3^\dagger\, \mathbb{W}^{1/2}(\mathbb{1} - S_\eta)\mathbb{W}^{1/2}\, \mathcal{E}_3 = n_C^{red}\, \mathbb{1}_8, \quad n_C^{red} \geq 0,$$

with strict propagation requiring $n_C^{red} > 0$. Therefore

$$\blacksquare \quad Z_3(\mu^\star) = C_{hol} \cdot n_C^{red}.$$

The two factors have distinct jobs — C_{hol} is the stationary-record norm-to-action conversion, n_C^{red} is the reduced norm of the physical colour embedding — and **neither may be selected from the desired Λ_{QCD}, a measured coupling, or the wish to obtain $Z_3 = 1$.**

Theorem 2A — Conditional Unit Class-Action Theorem

Suppose the microscopic substrate returns both $C_{hol} = 1$ (through the holonomy–Wilson action-normalisation bridge) and $\mathcal{E}_3^\dagger\, \mathbb{W}^{1/2}(\mathbb{1} - S_\eta)\mathbb{W}^{1/2}\, \mathcal{E}_3 = \mathbb{1}_8$ (through the physical colour-response calculation). Then

$$\blacksquare \quad Z_3(\mu^\star) = 1, \quad g_s(\mu^\star) = 1.$$

Proof. The second condition gives $n_C^{red} = 1$; substituting with $C_{hol} = 1$ into $Z_3 = C_{hol} \cdot n_C^{red}$ gives $Z_3 = 1$, and $g_s = Z_3^{-1/2} = 1$. ■

Status. Theorem 2A is an exact implication, but its two hypotheses are open microscopic returns; it does not derive them from class ownership. This branch ($C_{hol} = 1$, $n_C^{red} = 1$) is the **Unit Class-Action branch, NQTC-1U**, on which $Z_3 = 1$. It replaces the interpretation "6 − 5 = 1" with the structurally cleaner statement *one colour class + one reduced adjoint norm = one primitive action unit*, conditional on the two displayed certificates. The relation to the γ_η calculation of Section 6 is branch-specific: with $Z_3 = u\cdot\gamma_\eta/(1+\gamma_\eta)$, the unit condition $n_C^{red} = 1$ (i.e. $Z_3 = C_{hol}$) corresponds to $\gamma_\eta = C_{hol}/(u - C_{hol})$. Only under the historical identification $u = 6C_{hol}$ (Principle MC0) does this become $\gamma_\eta = 1/5$. The two conditions are *not* generally equivalent, and treating them as such would re-import the census-6 that Firewall C1 forbids.

7A.1 Blind unit-return certificates

The physical calculation must publish three certificates, all required to pass before $Z_3 = 1$ may be quoted as physical:

- **Action conversion.** $\Delta_{hol} = |C_{hol} - 1| = 0$.
- **Adjoint isotropy.** With $R_C = \mathscr{E}_3^\dagger \mathbb{W}^{(1/2)}(\mathbb{1}-S_\eta)\mathbb{W}^{(1/2)}\mathscr{E}_3$ and $n_C^{red} = ⅛ \operatorname{Tr} R_C$, require $\|R_C - n_C^{red} \mathbb{1}_8\| = 0$ within certified error (a spread across the eight colour directions is a colour-covariance failure).
- **Unit norm.** $\Delta_C^{unit} = |n_C^{red} - 1| = 0$.

7A.2 Factor-specificity firewall

The weak factor also arises from one bath class and one connection, so "one class $\Longrightarrow Z = 1$" is **false** as a general theorem. Class ownership fixes multiplicity, not factorwise magnitude. The differences between colour, weak and Abelian couplings must come from the physical embeddings, response metrics, auxiliary relaxation and matching structure (n_3^{red}, n_2^{red}, n_q^{red}); a universal unit-stiffness claim would erase the coupling hierarchy rather than derive it.

---

8. Record-Layer Matching Surface

The Planck scale is the metric-certification floor. It is where geometry itself is certified, not where a propagating record-layer field meets its substrate boundary. Matching QCD there would conflate two structural levels of the emergence hierarchy.

The TPB coherence programme supplies the record-layer ceiling

$E_{rec} = \sqrt{M_P m_\kappa} = 5.80$ TeV

on its frozen Route-R branch.

Theorem 2 — QCD Boundary-Layer Selection

Because the QCD connection, quarks and gluons are propagating record-layer carriers, their substrate boundary is the record-coherence ceiling:

█ $\mu\star = E_{rec} = 5.80$ TeV.

Proof

The layer assignment is structural: QCD fields are record-layer carriers by construction (they propagate, they are recorded, they do not certify the metric). A record-layer carrier's response is defined up to the scale at which record coherence itself fails, which is E_{rec} on the frozen Route-R branch. Matching at the Planck floor instead would attribute to a record-layer field a boundary belonging to the certification layer, contradicting the layer assignment. ■

The value 5.80 TeV is consumed here as a frozen upstream output; its derivation belongs to the TPB coherence programme and is not reopened in this paper. Falsifier F7 stands against the layer assignment itself.

---

9. Frozen Threshold Ledger

The native VBF threshold values supplied by the mass programme are

$m_t = 291.827$ GeV, $m_b = 4.86254$ GeV, $m_c = 1.11241$ GeV.

These are decoupling values in the VBF scheme, frozen upstream. They are not pole masses, not MS-bar masses, and not measured inputs; the numerical distance between m_t above and the measured top mass is a scheme statement, not a tension, and is bridged only downstream.

The sharp one-loop coefficients $b_N = 11 - 2N_f/3$ are

$b_6 = 7$, $b_5 = 23/3$, $b_4 = 25/3$, $b_3 = 9$.

At the declared order the coupling is continuous at each threshold: Z_3 is matched sharply, with no finite decoupling constant, which is exact at one loop.

---

10. Numerical Strong Scale

10.1 The general one-loop decoupling rule

With $dZ_3/d \ln k = b_N/(8\pi^2)$ between thresholds, the running stiffness integrates to

$Z_3(k) = (b_N/8\pi^2) \ln(k/\Lambda_N)$,

which defines the N-flavour scale Λ_N. Continuity of Z_3 at a threshold mass m gives

$b_N \ln(m/\Lambda_N) = b_{N-1} \ln(m/\Lambda_{N-1})$,

hence the exact one-loop decoupling rule

■ $\Lambda_{N-1} = m^{(1 - b_N/b_{N-1})} \cdot \Lambda_N^{(b_N/b_{N-1})}$.

Every step below is one application of this rule.

10.2 The cascade

Above the top threshold, the boundary condition $Z_3(\mu\star) = 1$ gives

$\Lambda_6 = \mu\star \exp[-8\pi^2 Z_3/b_6] = 5800 \cdot e^{(-8\pi^2/7)}$ GeV,

■ $\Lambda_6 = 0.0732458$ GeV.

Crossing the top threshold ($b_6/b_5 = 21/23$):

$\Lambda_5 = m_t^{(2/23)} \Lambda_6^{(21/23)}$,

■ $\Lambda_5 = 0.150611$ GeV.

Crossing the bottom threshold ($b_5/b_4 = 23/25$):

$\Lambda_4 = m_b^{(2/25)} \Lambda_5^{(23/25)} = (m_t m_b)^{(2/25)} \Lambda_6^{(21/25)}$,

■ $\Lambda_4 = 0.198873$ GeV.

Crossing the charm threshold ($b_4/b_3 = 25/27$):

$\Lambda_3 = m_c^{(2/27)} \Lambda_4^{(25/27)} = (m_t m_b m_c)^{(2/27)} \Lambda_6^{(7/9)}$,

so, in closed form directly from the boundary,

$\Lambda_3 = \mu\star^{(7/9)} (m_t m_b m_c)^{(2/27)} e^{(-8\pi^2/9)}$,

■ $\Lambda_{QCD}^{(3)\prime vBf} = 0.225923$ GeV.

Theorem 3 — Numerical VERSF Strong Scale (MC branch)

On the pre-registered MC branch ($Z_3(\mu\star) = 1$),

▌ $\Lambda_{3,MC}^{VBF} = 225.923$ MeV.

This is the exact consequence of the MC branch, **not yet** the physical VERSF strong scale: it is conditional on the substrate Hessian returning $Z_3 = 1$ (Section 20′). If the microscopic evaluation returns $Z_3 = C_{hol} \cdot n_C^{red}$ instead, the physical scale is

▌ $\Lambda_3^{VBF} = \mu\star^{(7/9)}(m_t\, m_b\, m_c)^{(2/27)} \exp[-8\pi^2 C_{hol}\, n_C^{red}/9]$.

On the Unit Class-Action branch ($C_{hol} = 1$, $n_C^{red} = 1$) this returns $\Lambda_{3,U}^{VBF} = 225.923$ MeV — a sharp pre-registered prediction of the unit branch, which becomes the physical VERSF QCD scale only if the three unit-return certificates of Section 7A.1 pass.

Proof

Immediate from the MC-branch value $Z_3(\mu\star) = 1$ (Theorem 1, pre-registered), $\mu\star = 5.80$ TeV (Theorem 2), the frozen threshold ledger (Section 9) and the decoupling rule of 10.1 applied three times. The branch conditionality is inherited by the result. ■

This is a native-scheme dimensional-transmutation scale. A finite scheme conversion is required before comparison with a scale quoted in another scheme; that bridge is a declared open debt (Section 30), and its absence here is deliberate — inserting it would require touching measured strong-sector data, violating the non-insertion firewall.

Part III — Global Centre Lift

Pure-gauge versus full-QCD firewall

Exact centre one-form symmetry and triality superselection apply to the **adjoint-only or quenched ($N_f = 0$) colour theory**. With dynamical fundamental quarks the centre symmetry is explicitly broken and the corresponding strings may break, so exact triality superselection does not hold in full $N_f = 3$ QCD. The reference centre-projector construction of Parts III–V is therefore a **quenched reference model** whose coefficient σ_0 is *normalised by* the branch-conditional three-flavour scale Λ_3^{VBF}; it is **not** identified with the physical pure-Yang–Mills string tension without a separate $N_f = 3 \rightarrow 0$ matching construction, which this paper does not supply.

The phrase "pure $N_f = 3$ QCD" is a category error and is avoided: the string-tension statements below refer to the quenched centre-sector reference model, with Λ_3 used only as its normalising scale.

11. Local Algebra Versus Global Colour

The local connection sees $\mathfrak{su}(3)$. Triality belongs to the global group, and the distinction is the entire content of the confinement problem.

Let

$\omega = e^{2\pi i/3}$.

The centre is

$Z(SU(3)) = \{1, \omega, \omega^2\} \cong \mathbb{Z}_3$.

A representation R has triality $\tau(R)$ when the centre acts on it by

$\omega \cdot v = \omega^{\tau(R)} v$.

Adjoint gluons have triality zero. A fundamental quark has triality one; an antiquark has triality two. No local operator built from adjoint fields can change τ, which is why triality is the correct superselection label for the confinement question: it is precisely the charge that gluon dynamics cannot screen.

---

12. The Centre Action on the Original SU(3) Links

The main-body (strict) route does **not** introduce a second $\mathbb{Z}_3$ gauge system. It defines the centre action directly on the physical gauge-invariant Hilbert space of the original theory,

$\mathscr{H}_{phys} = L^2(SU(3)^E)^{\mathcal{G}}$,

through the surface operator that shifts the crossing links by a centre element:

■ $(\hat{U}_\omega(\Sigma)\, \Psi)(\{U_\ell\}) = \Psi(\{ \omega^{I(\ell,\Sigma)}\, U_\ell \})$,

where $I(\ell, \Sigma) \in \{0, \pm 1\}$ is the signed intersection number of link ℓ with the transverse surface Σ. This is an operator on the *original* SU(3) configuration space — no auxiliary z_ℓ variable, no separate $\mathbb{C}^3$ centre carrier. Its eigenvalues are the cube roots of unity. The strict route *requires* $\hat{U}_\omega(\Sigma)$ to commute with the physical Hamiltonian and with local centre-neutral observables — this commutation is an owed audit (Stage 3 of Section 20′), not yet established. Once established, and in the **quenched/adjoint-only sector** (with dynamical fundamental matter the centre sectors are connected or screened, per the Part III firewall), it grades the physical Hilbert space into triality sectors

$$\mathscr{H}_\tau(\Sigma) = \{ |\Psi\rangle : \hat{U}_\omega(\Sigma)|\Psi\rangle = \omega^\tau |\Psi\rangle \}.$$

Because the operator acts directly on the original SU(3) link variables, it introduces no explicit duplicate centre carrier. Its compatibility with the quotient-compatible physical measure, and its commutation with the reconstructed Hamiltonian, remain part of D8. This is the construction the penalty-free gap of Section 20′.4 uses.

The alternative decomposition $\tilde{U}_\ell = z_\ell U_\ell$ with an independent $z_\ell \in \mathbb{Z}_3$ — used in the reference completion (Appendix A / Sections 15–20) — is a *different* object: it adjoins an explicit $\mathbb{Z}_3$ one-form field, and its equivalence to the intrinsic action above is exactly what falsifier F23 guards. The two must not be conflated. The remainder of this section records the reference-completion lift for completeness; the strict route uses only $\hat{U}_\omega(\Sigma)$ above.

Reference-completion lift (Appendix A material)

On a centre-preserving regulator, decompose the link transport into a local PSU(3) part and a centre part:

$$\tilde{U}_\ell = z_\ell U_\ell, \quad z_\ell \in \mathbb{Z}_3.$$

A one-form gauge transformation acts as

$$z_\ell \mapsto \omega^{(\lambda_\ell)} z_\ell,$$

with the plaquette two-form background shifting by the corresponding coboundary. The lift therefore carries an exact $\mathbb{Z}_3$ one-form symmetry, and the physics of the centre sector is the physics of that symmetry.

The surface operator is

$$\hat{U}_\omega(\Sigma) = \prod_{\{\ell \pitchfork \Sigma\}} X_\ell,$$

where X_ℓ measures the centre electric flux crossing link ℓ.

This construction preserves Wilson lines, centre charge and one-form gauge covariance. But a caution must be stated: writing $\tilde{U}_\ell = z_\ell U_\ell$ with an independent $z_\ell \in \mathbb{Z}_3$ risks adding a *separate* $\mathbb{Z}_3$ gauge system alongside SU(3), rather than deriving the centre sectors of the *original* SU(3) theory. A genuine lift must realise the extension

$$1 \to \mathbb{Z}_3 \to SU(3) \to PSU(3) \to 1$$

at the level of transition functions, bundles, measure and line operators, and must prove the centre variables are not double-counted against the SU(3) measure. The centre-projector algebra below is sound *once such a carrier exists*; the derivation that this carrier is equivalent to the original VERSF SU(3) measure is **[Owed]**. The paper therefore treats the $\mathbb{Z}_3$ lift as an explicit candidate construction, not an established reduction of the original theory. Falsifier F9 stands against the equivalence.

---

13. Primitive Flux Normalisation

Define the primitive coweight

$H\_c = \frac{1}{3}\,\mathrm{diag}(1, 1, -2)$.

Then

$e^{(2\pi i\, H\_c)} = \omega\, \mathbb{1}_3$,

so one full period of H_c generates exactly one primitive centre element. The primitive flux representative is

$\Phi\_\tau = 2\pi\tau\, H\_c \pmod{2\pi\, Q^\vee}$,

where $Q^\vee$ is the coroot lattice. Using the adjoint-invariant norm

$\langle X, Y\rangle\_\mathrm{adj} = 2\,\mathrm{tr}(XY)$,

one finds

$\|H\_c\|^2\_\mathrm{adj} = 2\,\mathrm{tr}(H\_c^2) = 2 \cdot (1 + 1 + 4)/9 = 4/3$.

Therefore

■ $\|\Phi_1\|\_\mathrm{adj} = 2\pi\,\|H\_c\|\_\mathrm{adj} = 4\pi/\sqrt{3}$.

This fixes the primitive flux unit with no continuous normalisation freedom: the coweight is determined by the centre element, the norm by the adjoint metric, and the period by single-valuedness. Falsifier F10 stands against the coweight identification.

---

14. Triality Hilbert Sectors

At fixed transverse slice Σ, the physical Hilbert space decomposes as

$\mathscr{H}_\Sigma = \mathscr{H}_0 \oplus \mathscr{H}_1 \oplus \mathscr{H}_2$.

The sector projectors are the discrete Fourier transforms of the surface operator:

$P_\tau = \frac{1}{3} \sum_{n=0}^{2} \omega^{(-n\tau)} \hat{U}_\omega(\Sigma)^n$.

They are mutually orthogonal and complete:

$P_\tau P_{\tau'} = \delta_{\tau\tau'} P_\tau, \sum_\tau P_\tau = \mathbb{1}$.

In pure gauge theory, every local adjoint operator commutes with $\hat{U}_\omega$, so the sectors are superselected: no gluonic process connects them, and the sector label is a constant of the motion. This superselection is what converts the projector inequality of Section 19 from an algebraic identity into a physical energy bound.

Part IV — Complete Quantum Transverse Bath

15. Regulated Physical Hilbert Space (Reference Completion)

This section and Sections 16–20 develop the inserted-penalty reference model (Appendix A material). The strict route uses the intrinsic centre action of Section 12 and the penalty-free construction of Section 20'.4; the auxiliary $(\mathbb{C}^3)^{E_{centre}}$ carrier below belongs to the reference model only.

For a finite transverse regulator $K_{a,R}$ with lattice spacing a and transverse radius R, define

$\mathscr{H}_{a,R} = L^2(SU(3)^E / \mathcal{G}) \otimes (\mathbb{C}^3)^{E_{centre}} \otimes \mathscr{H}_{cl} \otimes \mathscr{H}_{aux}$ (reference completion, App. A),

where E is the link set, $\mathcal{G}$ the gauge group of the slice, $\mathscr{H}_{cl}$ the closure/record sector and $\mathscr{H}_{aux}$ the auxiliary adjoint sector of Section 6. The gauge quotient is imposed before sector energies are extracted — energies are properties of physical states only.

The local VERSF bath Hamiltonian contains:

1. the SU(3) electric Laplacian;
2. the centre-twisted magnetic plaquette term;
3. the closure and record squares;

4. the complete same-order auxiliary Schur reduction of Section 7;
5. the centre-electric primitive penalty of Section 16.

Every term is vacuum-subtracted, so the zero-triality vacuum sits at energy zero by construction.

---

16. The Centre-Electric Projector

The $\mathbb{Z}_3$ clock operator X_ℓ has eigenvalues 1, ω, ω². Define

■ $P_\ell^{\neq 0} = (2 - X_\ell - X_\ell^\dagger)/3$.

Its action on the eigenbasis is exact:

$X_\ell = 1 \Longrightarrow P_\ell^{\neq 0} = (2 - 1 - 1)/3 = 0$, $X_\ell = \omega$ or $\omega^2 \Longrightarrow P_\ell^{\neq 0} = (2 - \omega - \bar{\omega})/3 = (2 + 1)/3 = 1$,

using $\omega + \bar{\omega} = 2\cos(2\pi/3) = -1$. Thus $P_\ell^{\neq 0}$ is exactly the spectral projector onto nonzero centre electric flux: idempotent, positive, local, and with spectrum {0, 1}. Its integrality is what makes the gap of the *assembled* Hamiltonian uniform under refinement — a projector's eigenvalues cannot shrink. This is a genuine property of the operator once the penalty is present; it is not, by itself, a demonstration that the substrate produces the penalty (Section 17.0).

---

17. The Complete Positive Bath Hamiltonian

Let

$\hat{H}_{\perp,+}^{VERSF} = \sum_j Q_j^\dagger Q_j$

denote the complete gauge-reduced local bath, closure, record and auxiliary Hamiltonian after square completion. It is positive and vacuum-subtracted. The existence of a common zero state is built into the reference completion; uniqueness of that state after gauge quotienting and boundary conditions is Premise VU and requires the kernel-intersection calculation of Section 18.

Define

■ $\hat{H}_\perp^{red} = \hat{H}_{\perp,+}^{VERSF} + \sigma_0 \sum_{\ell \pitchfork \Sigma} P_\ell^{\neq 0}$.

This is the complete NQTC-1 transverse Hamiltonian. The centre term is local, gauge invariant, one-form covariant, positive, and exactly zero on the vacuum sector — it adds nothing anywhere the centre flux vanishes, and one primitive unit wherever it does not.

17.0 Status of the centre penalty: inserted, not derived

Honesty about the logical status of this term is essential, because it is the hinge of the whole confinement argument and it is easy to overstate what has been established.

The penalty $\sigma_0 \sum P\_\ell^{\neq 0}$ is **added** to the bath Hamiltonian; it is not projected out of an inherited operator. Its coefficient σ_0 is then fixed to Λ_3^2 by Principle MC4 (Section 24). This means the gap theorems that follow (Theorems 4–5, Part V) prove that *a Hamiltonian containing a positive centre penalty confines* — which, given that the penalty is positive and centre-selective by construction, is close to built in. The structural content that survives is real but narrower than a derivation: the $\mathbb{Z}_3$ projector inequality of Lemma 19.1 is exact and the uniformity in regulator and volume is genuine; the straight one-flux state supplies a natural candidate for saturating the lower bound — its centre-projector contribution is explicit, but simultaneous annihilation of every orthogonal magnetic, closure and auxiliary square has not been demonstrated (Section 23). What is **not** established is that the substrate dynamics *produce* this penalty with this coefficient. MC4 asserts it; it does not derive it.

Stated in the audit discipline used elsewhere in the programme (the [Owed output] / [Audit condition] / [Conditional target] tagging of the weak-doublet and PMNS gate papers):

- $\sigma_0 \sum P\_\ell^{\neq 0}$ as a positive centre-selective term — **[Inserted modelling term]**;
- $\sigma_0 = \Lambda_3^2$ — **[Conditional target on MC4]**, not [Derived];
- the $\mathbb{Z}_3$ projector inequality and its uniformity — **[Exact]**;
- that this penalty equals the true substrate colour response — **[Owed output]**, discharged only by the derivation route of Section 20′.

The remainder of Part IV carries σ_0 as a positive symbol so that the exact structural results are visibly independent of the disputed coefficient. Section 20′ states precisely what a genuine derivation of the penalty would require, and why the present paper cannot yet supply it.

17.1 Reflection positivity of the centre penalty

Reflection positivity of the assembled action is required for the transfer-matrix and Wilson-loop results. The properties below make it plausible for the centre-projector reference completion, but the necessary positive-character expansion or explicit positive transfer matrix is not constructed here. Reflection positivity therefore remains **Premise RP**.

Let θ denote Euclidean time reflection and let S_0 be the reflection-positive action of the gauge-reduced VERSF bath. The centre penalty adds to S_0 a term $\sigma_0 \sum P\_\ell^{\neq 0}$ that is:

1. **θ-even** — $P\_\ell^{\neq 0}$ is built from $X\_\ell + X\_\ell^\dagger$, which is invariant under time reflection, and the crossing set $\ell \pitchfork \Sigma$ maps to itself under θ for a reflection-symmetric surface;
2. **positive and bounded** — $P\_\ell^{\neq 0}$ is a projector, so each term lies in [0, 1] and the added weight $e^{(-\sigma_0 \sum P\_\ell^{\neq 0})}$ lies in (0, 1];

3. **local and gauge-scalar** — each $P_\ell^{\neq 0}$ depends only on the centre flux at link ℓ and is invariant under gauge and one-form transformations.

The claim to be careful about is that positivity of the weight *alone* preserves Osterwalder–Schrader positivity. That is too strong as a general statement: OS positivity normally requires an explicit factorisation across the reflection plane, a positive character expansion, or an explicitly constructed positive transfer matrix — mere positivity, boundedness and reflection-evenness of a multiplicative weight do not by themselves guarantee it in general.

For the centre penalty the favourable features (θ-even, factorised across links, diagonal in the centre-flux basis) make OS positivity plausible and arguably straightforward, but this paper does not construct the transfer matrix for the centre term and prove its positivity. Two honest options remain, and the paper takes the second explicitly:

- construct the centre-term transfer matrix and prove positivity — **[Owed]**; or
- **retain reflection positivity of the assembled action as a named premise** of the confinement branch.

The area-law step of Section 25 consumes transfer-matrix positivity, so it inherits this premise. Falsifier F17 accordingly guards a *premise*, not a discharged theorem: the earlier claim that the Schwarz argument goes through "verbatim" is withdrawn.

Even granting the premise, it bears only on the assembled Hamiltonian and says nothing about §17.0's separate point: whether the substrate *generates* the penalty. Reflection positivity is a property the assembled operator is *assumed* to have; provenance of the penalty remains [Owed output].

---

18. Coercivity, Compactness and Vacuum Uniqueness

At fixed regulator:

- $SU(3)^E$ is compact;
- the centre carrier is finite dimensional;
- the electric Laplacian has discrete spectrum;
- the closure and auxiliary squares are positive;
- fixed vacuum boundary data remove global topological degeneracy.

Therefore $\hat{H}_\perp^{red}$ has compact resolvent, and its spectrum is a discrete set of eigenvalues of finite multiplicity accumulating only at infinity.

Its zero set requires simultaneously

$Q_j|\Psi\rangle = 0$ for every local square, and $P_\ell^{\neq 0} |\Psi\rangle = 0$ for every crossing link.

Compactness and the positive squares establish that ground states *exist*. They do not by themselves establish that the zero set is one-dimensional. Nullity-1 of a single closure operator does not exclude flat gauge sectors, magnetic degeneracy, topological degeneracy, or additional common zero modes of the other positive squares. The uniqueness claim requires the actual kernel intersection

$$\bigcap_j \ker Q_j \cap \bigcap_\ell \ker P_\ell^{\neq 0} = \mathrm{span}\{|\Omega\rangle\}$$

after gauge quotienting and boundary conditions, and this paper does not compute it. **Vacuum uniqueness is therefore a named premise (Premise VU), not a theorem**, pending that kernel calculation.

Theorem 4 — Conditional Finite-Regulator Vacuum and Sector-Gap Theorem

At every finite regulator (a, R), the zero-triality sector has a unique ground-state orbit with energy zero, and every nonzero-triality sector has a strictly positive lowest eigenvalue.

Proof

Zero energy of the $\mathscr{H}_0$ ground state follows from the zero-set analysis; *uniqueness* of that ground state is Premise VU (the kernel intersection above is assumed one-dimensional, not computed here). In $\mathscr{H}_\tau$ with $\tau \neq 0$, the surface operator eigenvalue $\omega^\tau \neq 1$ forbids the all-trivial centre configuration, so at least one $P_\ell^{\neq 0}$ is nonzero on every state in the sector; positivity of the remaining terms and discreteness of the spectrum give a strictly positive minimum. ■

Theorem 4 is the finite-volume statement. Its constant could in principle depend on (a, R); Theorem 5 shows it does not.

---

19. Exact Quantum Triality Gap

Lemma 19.1 — Centre Projector Inequality

On the full regulated Hilbert space,

$$\blacksquare \sum_{\ell \pitchfork \Sigma} P_\ell^{\neq 0} \geq P_1 + P_2.$$

Proof

The X_ℓ on a fixed separating surface mutually commute, so pass to their joint eigenbasis. Let the eigenvalues be ω^{n_ℓ} with $n_\ell \in \{0, 1, 2\}$. The surface triality of the joint eigenvector is

$$\tau = \sum_{\ell \pitchfork \Sigma} n_\ell \pmod 3.$$

On such an eigenvector, the left side has eigenvalue $\#\{\ell : n_\ell \neq 0\}$ and the right side has eigenvalue 1 if $\tau \neq 0$ and 0 if $\tau = 0$. If $\tau \neq 0$ then at least one $n_\ell \neq 0$ — a sum of zeros is zero — so the left eigenvalue is ≥ 1 = right eigenvalue. If $\tau = 0$ the right eigenvalue is 0 and the left is ≥ 0. The inequality holds on a complete eigenbasis of both sides, hence as an operator inequality. ■

The lemma is a spectral identity of $\mathbb{Z}_3$ arithmetic. It contains no coupling, no lattice spacing, no volume and no estimate — which is exactly why the gap it produces is uniform.

Theorem 5 — Exact Quantum Triality-Gap Theorem

For $\tau = 1, 2$ and every regulator (a, R),

■ $$\varsigma_{\tau,a,R}^{qu} := \inf_{\Psi\in\mathscr{H}_\tau,\ \|\Psi\|=1} \langle\Psi|\, \hat{H}_\perp^{red} \,|\Psi\rangle - E_0 \geq \sigma_0.$$

Proof

Discard the positive local bath term in the lower bound: $\hat{H}_\perp^{red} \geq \sigma_0 \sum_{\ell \pitchfork \Sigma} P_\ell^{\neq 0}$. Apply Lemma 19.1:

$$\hat{H}_\perp^{red} \geq \sigma_0 (P_1 + P_2).$$

On $\mathscr{H}_\tau$ with $\tau \neq 0$, P_τ acts as the identity, so every unit vector has energy at least σ_0. Since $E_0 = 0$ (Theorem 4), the gap statement follows. ■

This theorem is quantum. No classical field minimisation is substituted for a Hilbert-sector energy (firewall 3.4): the infimum is taken over the physical Hilbert space, and the bound is an operator inequality on that space.

---

20. Uniform Continuum and Infinite-Volume Limit

The decisive property of Theorem 5 is what its lower bound does not contain: no lattice spacing, no transverse radius, no link count, no coupling. The constant is σ_0, verbatim, at every regulator:

$$\varsigma_{\tau,a,R}^{qu} \geq \sigma_0 \text{ for all } (a, R).$$

The projector lower bound is regulator-uniform: the same constant σ_0 holds at every finite (a, R). But regulator-uniformity is not by itself a continuum theorem — passing to the limit requires a compatible family of Hilbert spaces, refinement embeddings, and convergence of the quadratic forms or transfer matrices. The correct conditional statement is:

■ *if* the regulated centre sectors form a compatible refinement system and the associated quadratic forms converge in the Mosco or norm-resolvent sense, *then*

$$\liminf_{a\to 0,\, R\to\infty} \varsigma_{\tau,a,R}^{qu} \geq \sigma_0 > 0.$$

The gap constant is protected by an exact spectral projector rather than a deteriorating estimate, which is what makes the *uniformity* exact; the convergence hypothesis (a named condition, folded into D8 for the strict route and required equally for the reference-completion continuum claim) is what upgrades uniformity to a continuum bound. Standard continuum bounds for gauge theories degrade as the regulator is removed because they trade positivity for locality through constants that grow with the number of degrees of freedom. Here there is no such trade: subdividing a link replaces one $\mathbb{Z}_3$ variable by several whose product carries the same flux, the surface triality is unchanged, and Lemma 19.1 reapplies with the same constant.

A compatible refinement-stable holonomy construction is required to ensure that the centre surface operator has a well-defined continuum action; the projector inequality survives because triality is discrete, and discreteness is exactly the property that subdivision cannot erode. Falsifier F14 stands against this stability claim.

Every result in this Part is a theorem about the *assembled* Hamiltonian of Section 17. Their uniformity and exactness are genuine, but they inherit the status of §17.0: they establish that a Hamiltonian carrying the centre penalty confines uniformly, not that the substrate produces that penalty. The next section states what would close that gap.

20′. The Derivation Target: Colour-Projected Closure Hamiltonian and Compatibility

The construction of Sections 15–20 inserts the confinement mechanism and then proves the assembled operator has the desired properties. A genuine derivation would instead obtain the colour-sector Hamiltonian by canonical projection of the inherited substrate closure Hamiltonian, exactly as the flavour sector is treated in the weak-doublet and PMNS gate papers, and would then have to prove that the projected operator's triality spectrum is a *physical* spectrum. This section states that route and its open prerequisites precisely, in the programme's audit-tagging discipline. Nothing in it is claimed as executed; it is the specification of what execution requires.

20′.1 The projection route (five stages)

Stage 1 — Colour-saturated Hessian. On the colour-supporting subspace $\mathcal{V}_C = \mathscr{M}_C \otimes \mathbb{R}^8$, form the Haar twirl

$$\mathfrak{T}_C(G\star) = \int_{SU(3)} (\mathbb{1} \otimes R_C(g))^\dagger\, G\star\, (\mathbb{1} \otimes R_C(g))\, d\mu(g).$$

By Schur's lemma on the irreducible adjoint factor this is exactly $\mathfrak{T}_C(G\star) = K_C \otimes \mathbb{1}_8$ for one Hermitian K_C on the multiplicity space $\mathscr{M}_C$. The **covariance audit** $\|G\star - \mathfrak{T}_C(G\star)\|$ measures colour anisotropy: it must vanish on the saturated branch, and a nonzero value falsifies that branch. — **[Owed output]**: requires explicit $G\star$ (§20′.3); the *form* $K_C \otimes \mathbb{1}_8$ is proved, the *entries* of K_C are not.

Stage 2 — Response blocks as calculations. With $f_C \in \mathscr{M}_C$ the normalised profile of one unit of primitive colour curvature and $\{e_I\}$ an orthonormal basis of same-order auxiliary adjoint profiles (null and pure-gauge modes removed), the three response objects are inner products against K_C:

$$\chi_C/8 = \langle f_C, K_C f_C\rangle,\quad b_I = \langle f_C, K_C e_I\rangle,\quad C_{IJ} = \langle e_I, K_C e_J\rangle,$$

and the reduced stiffness is the Schur complement

$$Z_3 = \chi_C/8 - b^\dagger C^{-1} b.$$

Equivalently, in executable plaquette form: impose $U_\square(\varepsilon) = \exp(i\varepsilon T_a)$, expand the frozen functional $F(\varepsilon,\eta) = F_0 + \tfrac{1}{2}[u\, \varepsilon^2 + 2\varepsilon\, b_I\, \eta^I + C_{IJ}\, \eta^I \eta^J] + O(3)$, giving $u = \chi_C/8$ and

$$Z_3 = \lim_{\varepsilon\to 0} 2[\min_\eta F(\varepsilon,\eta) - F_0]/\varepsilon^2,$$

with the eight-generator isotropy audit $Z_3^{(1)} = \cdots = Z_3^{(8)}$ (a spread is a colour-covariance failure). **No protected unit, auxiliary multiplicity, or relaxation coefficient is fixed in advance** — this replaces the imposed six-to-one split of Theorem 1 entirely. — **[Owed output]**: the entries are computed from the three linearised defect operators of §20′.3 once the substrate graph and background are specified.

Stage 3 — Global centre decomposition. From the global SU(3) lift, construct the centre surface operator $\hat{U}_\omega(\Sigma)$ and prove

$$[H_C, \hat{U}_\omega(\Sigma)] = 0,$$

giving the exact block decomposition $H_C = \bigoplus_\tau P_\tau H_C P_\tau$ and the closure-sector energies

$$E_\tau^{cl}(a, R) = \inf \operatorname{spec}(\, P_\tau H_C(a, R) P_\tau \,),\quad \Delta_\tau^{cl} = E_\tau^{cl} - E_0^{cl}.$$

— **[Audit condition]** (the commutation) + **[Owed output]** (the spectrum).

Stage 4 — Physical compatibility. Construct the gauge-reduced transverse physical Hamiltonian $H_\perp^{phys}$ and prove the compatibility relation, in the form used by the PMNS gate paper's Dual Closure–Mass Compatibility Theorem,

$$[H_C, H_\perp^{phys}] = 0,$$

or the stronger functional relation

$$H_\perp^{phys} = f_C(H_C), \quad f_C \text{ strictly increasing on the relevant spectrum.}$$

Mere commutation is a **necessary compatibility audit but not sufficient** for the physical gap: it permits compatible spectral decompositions but allows degeneracies and independent energy ordering inside shared eigenspaces. The physical bridge needs either the strictly-increasing functional relation $H_\perp^{phys} = f_C(H_C)$, or a direct sector-wise ordering theorem $E_\tau^{phys} - E_0^{phys} \geq c(E_\tau^{cl} - E_0^{cl})$ with $c > 0$. — **[Owed output]**, the single most important missing bridge in the present paper.

Stage 5 — Calculated gap. With Stage 4 in hand,

$$\Delta_\tau = f_C(E_\tau^{cl}) - f_C(E_0^{cl}) > 0$$

whenever f_C is strictly increasing and $E_\tau^{cl} > E_0^{cl}$. What then appears is a positive *sector-energy separation* obtained by diagonalising the substrate Hamiltonian, rather than an inserted term; it need not be representable as a local projector penalty $\sigma_0 P_\tau$. — **[Conditional target on Stages 1–4]**.

20′.2 The colour-adjoint action — abstract realisation (embedding still owed)

The $\mathfrak{su}(8)$-restricted closure object $H_{cl} = G^\star - \tfrac{1}{8}\,\mathrm{tr}_8(G^\star)\mathbb{1}_8$ must be connected to the SU(3) adjoint by an explicit map; dimension eight alone does not do it. That map is now available in closed form and has been checked. Define the colour action on $\mathcal{V}_8 \simeq \mathbb{C}^8$ by

■ $R_C(g)_{ab} = 2\,\mathrm{tr}(T_a\, g\, T_b\, g^\dagger)$, $T_a = \lambda_a/2$, $g \in SU(3)$,

and the embedding

$$\iota_C : \mathbb{R}^8 \to \mathrm{Herm}_0(\mathbb{C}^3), \quad \iota_C(x) = x^a T_a.$$

Theorem 20′.1 (Canonical adjoint embedding). R_C is a continuous homomorphism $SU(3) \to SO(8) \subset SU(8)$ with kernel exactly $Z(SU(3)) \cong \mathbb{Z}_3$, giving a faithful $PSU(3) \hookrightarrow SO(8)$; and ι_C is isometric and equivariant, $\iota_C(R_C(g)x) = g\,\iota_C(x)\,g^\dagger$.

Proof. Conjugation preserves 2 $\mathrm{tr}(T_a T_b) = \delta_{ab}$, so $R_C(g)$ is real orthogonal; continuity from $R_C(\mathbb{1}) = \mathbb{1}_8$ fixes det = +1, hence SO(8); the conjugation-trivial elements are exactly the SU(3) centre, giving the $\mathbb{Z}_3$ kernel; equivariance is immediate. ■

What this establishes, precisely, is an **explicit abstract adjoint realisation** — the adjoint representation of SU(3) on an abstract $\mathbb{R}^8$ with the centre as kernel — and a representation *template* for the colour block. It is important not to overstate it. It does **not** by itself embed colour into the physical closure carrier: that requires the physical map

$\mathscr{E}_C : \mathfrak{su}(3) \to T_{\rho\star} \mathcal{C}_{sub}$,

with a proof that the eight substrate response vectors $\mathscr{E}_C(T_a)$ exist, are linearly independent, carry the declared physical tangent metric, and transform through R_C. Theorem 20′.1 supplies the target these vectors must hit; it does not supply the vectors. Two further points reinforce the gap: R_C is centre-blind ($R_C(\omega\mathbb{1}_3) = \mathbb{1}_8$), so local gluonic curvature transforms under PSU(3) while the global centre information lives in the separate lift of Part IV; and even granting $\mathscr{E}_C$, one must still show the physical substrate Hessian $G\star$ *selects* this subgroup via the covariance audit $\|G\star - \mathfrak{T}_C(G\star)\| = 0$. **Debt D6 therefore stands** — reduced in scope to the physical embedding $\mathscr{E}_C$ plus covariance selection, both requiring explicit $G\star$, but not discharged.

20′.3 The load-bearing prerequisite: three linearised defect operators

The prerequisite is sharper than "six constitutive derivatives," and that earlier framing is corrected here. The free energy is

$F[\mathfrak{X}] = V[\rho] + A_C[\rho,U] + A_T[\rho,J] + A_R[\rho,\Phi]$,

with each admissibility term of defect form $A_X = w_X \|\mathcal{D}_X\|^2$, the defects being $\mathcal{D}_C = U_{ij} U_{jk} U_{ki} - \mathbb{1}$, $\mathcal{D}_T = \nabla\cdot J_\rho + \partial_\sigma \rho$, $\mathcal{D}_R = \Phi_{jk} \Phi_{ij} - \Phi_{ik}$. Two structural facts collapse the calculation:

(a) Gauss–Newton reduction. At an admissible background $\rho\star$ every defect vanishes, $\mathcal{D}_C = \mathcal{D}_T = \mathcal{D}_R = 0$. For any $A = w\|\mathcal{D}\|^2$, the second variation at a zero of $\mathcal{D}$ is $D^2A_{\{\rho\star\}}[h,k] = 2w\star \operatorname{Re}\langle D\mathcal{D}[h], D\mathcal{D}[k]\rangle$, so every term in $D^2\mathcal{D}$, Dw, D^2w multiplies $\mathcal{D}\star = 0$ and vanishes. The **second** constitutive derivatives D^2U, D^2J, $D^2\Phi$ therefore do not enter the quadratic Hessian (they first appear at cubic order). Only first-order response survives.

(b) Link and record are independent variables. A gauge connection is not fixed by the density it acts on — one $\rho_i \in \mathbb{C}^8$ leaves the SU(7) stabiliser free. An endpoint-derived link $U_{ij}[\rho] = E_j E_i^{-1}$ has a coboundary variation whose triangle holonomy telescopes to zero, $D(U_{ij} U_{jk} U_{ki})[h] = 0$ — pure gauge, producing **no** colour stiffness. So χ_C must come from *independent* link variations, not $D_\rho U$; likewise Φ. The correct configuration is $\mathfrak{X} = (\rho, U, J, \Phi)$, varied independently, respecting the source framework's structural independence of closure, transport and record.

The six derivatives therefore reduce to **three explicit linearised defect operators**:

█ **Closure.** With $U_{ij} = U_{ij}^{\star} \exp(i\, x^a T_a)$, $X_{ij} = U_{ij}^{\star\dagger}\, \delta U_{ij}$ anti-Hermitian, the flat-triangle variation $\delta C_{ijk} = \mathrm{Ad}_{U_{ij}^{\star}} X_{ij} + \mathrm{Ad}_{U_{ij}^{\star} U_{jk}^{\star}} X_{jk} + X_{ki}$ defines B_C, and $G_C = 2\, B_C^{\dagger} W_C B_C$.

█ **Transport.** With the minimal covariant current $(J_\rho)_{ij} = \kappa_{ij}(\rho_j - U_{ij}\, \rho_i)$ (linear in ρ, so $D^2 J = 0$), $L_T(h,X) = \nabla\cdot\delta J[h,X] + \partial_\sigma h$ gives $G_T = 2\, L_T^{\dagger} L_T$.

█ **Record.** In affine coordinates $\Phi_{ij} = \Phi_{ij}^{\star} + y^A R_A$ ($D^2\Phi = 0$), $\delta R_{ijk} = \delta\Phi_{jk}\, \Phi_{ij}^{\star} + \Phi_{jk}^{\star}\, \delta\Phi_{ij} - \delta\Phi_{ik}$ defines B_R, and $G_R = 2\, B_R^{\dagger} W_R B_R$.

The corrected microscopic Hessian on $\delta\mathfrak{X} = (h, X, j, \varphi)$ is

█ $\mathbb{G}^{\star} = G_V + 2\, B_C^{\dagger} W_C B_C + 2\, L_T^{\dagger} L_T + 2\, B_R^{\dagger} W_R B_R$,

and after gauge/null-mode removal and the curvature/auxiliary split, $Z_3 = H_{FF} - H_{F\eta} H_{\eta\eta}^{-1} H_{\eta F}$. **This is a finite, executable matrix calculation** requiring no invented constitutive maps $U[\rho]$, $\Phi[\rho]$ — and it has now been executed on the canonical $K = 7$ wheel cell (Section 6): B_C, L_T, B_R are given explicitly with verified spectra, and the calculation reveals $b_C = 0$ (the five exact modes are gauge modes in $\ker B_C$), so the finite-cell closure operator does not by itself supply the branch's five-unit correction. What remains to be supplied is concrete physical data, not derivatives:

the substrate graph (links and triangles), the background links $U_{ij}^{\star}$, the saturated weights $u_i^{\star}$, the physical colour embedding $\mathscr{E}_C$, the current coefficients κ_{ij} and record background $\Phi_{ij}^{\star}$, plus the physical tangent metric, boundary/gauge conditions, null-mode removal, and the finite-cell-to-continuum blocking map. This data package is the **local microscopic-response blocker** — now a matter of specifying a finite graph and its background, not of inventing constitutive laws.

There is a **second, independent blocker** on the confinement side (Paper 6 in the dependency chain): the quotient-compatible global SU(3) measure, the reflection-positive transfer matrix, the complete physical transverse Hamiltonian, global coercivity of the triality sector, the longitudinal square decomposition, and convergence of the regulated sector forms. Neither blocker reduces to the other. The honest statement is therefore **two major blockers**, not one: local microscopic response, and global quantum measure plus physical Hamiltonian. The shared defect operators and substrate data unblock the first (and, being shared with the weak-doublet and PMNS gates, the flavour and lepton routes with it); the second is Paper 6's work. Debt D7 records the local blocker; Debt D8 records the global one.

A finite intermediate audit is available before the full colour-saturated evaluation: on the symmetric diamond, the nine sector traces $\alpha_X^{(R)} = (1/d_R)\, \mathrm{Tr}(P_R G_X P_R)$ for $X \in \{C,T,R\}$ and $R \in$ {reflection-odd L_{12}, reflection-even longitudinal, collective transverse} verify that the constitutive Hessian exists, is positive after gauge reduction, and behaves consistently under refinement. The diamond is *not* the physical colour response — its directional background is not SU(3)-adjoint isotropic — so it is a scaffolding check, not the result.

20′.4 Substrate-derived triality gap (no inserted penalty)

The decisive structural gain of the projection route is that, once Z_3 and a second evaluated substrate response are in hand, they supply the coefficients of a *candidate* coercive bound **without any centre penalty added to the Hamiltonian**. A physical triality gap follows only if the gauge-covariant flux typing, global sector coercivity, centre commutation, quotient-compatible measure, transfer-matrix positivity and convergence conditions of D8 are also established (Local-Hessian Insufficiency Proposition). Finite-regulator vacuum uniqueness implies a positive sector gap only together with discreteness or a compact resolvent. The physical transverse Hamiltonian is defined directly from the reflection-positive global colour measure,

$$H_\perp^{phys} = -(1/a_t) \log T_{a,R},$$

which commutes with the centre surface operator (Section 12) and therefore block-diagonalises, $H_\perp^{phys} = \bigoplus_\tau P_\tau H_\perp^{phys} P_\tau$, with sector energies $E_\tau(a,R)$ and gap $\Delta_\tau = E_\tau - E_0$. Given a unique physical ground state in the $\tau = 0$ sector, $\Delta_\tau > 0$ at finite regulator is immediate (a $\tau \neq 0$ zero-energy state would violate uniqueness) — no $P_{\neq 0}$ term is inserted.

The uniform bound comes from two evaluated responses. Fixed primitive flux gives, by Cauchy–Schwarz on a transverse disk, the small-radius coercivity

$$\mathscr{E}_{electric}^{(\tau)}(R) \geq \|\Phi_\tau\|^2_{adj} / (2\pi Z_3 R^2),$$

while the reduced wall Hessian $\mathcal{K}_{wall}^{red} = G_{ww} - G_{w\eta} G_{\eta\eta}^{-1} G_{\eta w}$ supplies the large-radius areal cost through

$$B_\tau = (1/2A_{cell}) \inf_{\|w\|=1} \langle w, \mathcal{K}_{wall}^{red} w\rangle > 0.$$

Minimising $\mathscr{E}_\tau(R) \geq \|\Phi_\tau\|^2_{adj}/(2\pi Z_3 R^2) + \pi B_\tau R^2$ over R gives

■ $\varsigma_\tau \geq g_s \|\Phi_\tau\|_{adj} \sqrt{2B_\tau} > 0$, $g_s = Z_3^{(-1/2)}$,

and, under uniform bounds $0 < Z_3 \leq Z_{max}$, $B_\tau \geq B_{min} > 0$, fixed Φ_τ, and Mosco/norm-resolvent convergence of the sector forms,

■ $\liminf_{a\to 0, R\to\infty} \Delta_\tau \geq \|\Phi_\tau\|_{adj} \sqrt{2B_{min}/Z_{max}} > 0$.

This is stronger than preservation of the discrete triality label: the physical energy coefficient is uniformly controlled by two *evaluated* substrate responses (Z_3 and B_τ), not by an inserted constant. Both are **[Owed output]** pending §20′.3, and the pass condition $B_\tau > 0$ is a genuine falsifier — if the bath can spread the triality wall at zero areal cost, this mechanism fails. The global measure, reflection-positive transfer matrix and physical-Hilbert-space reconstruction are properly outputs of Paper 6 in the dependency chain, not consequences of adding a positive term to the action.

Part V — Longitudinal Completion and String Tension

Governing note: two routes through Part V

Part V exists in two versions, and they must not be conflated (this was the largest structural inconsistency of the previous draft).

The **strict route** uses the substrate-derived slice gap of Section 20′.4,

$\varsigma_\tau = \inf \operatorname{spec} H_\perp^{(\tau)} - E_0$,

with no inserted coefficient. Under it the slicing inequality reads $V_\tau(L) \geq \varsigma_\tau L - c_{end}$, the tension (once its limit is shown to exist) is bracketed $\varsigma_\tau \leq \sigma_\tau \leq \bar{\sigma}_\tau$, with $c_{\{\sigma,\tau\}}$ to be **calculated from B_τ and Z_3**, not set to one. The relation $\sigma_\tau = c_{\{\sigma,\tau\}}\Lambda_3^2$ belongs only to the Λ_3-normalised reference model; a physical adjoint-only/$N_f=0$ tension would involve its own scale Λ_0 until D9 is completed. Every occurrence of the inserted σ_0 in the sections below is, in the strict route, replaced by this ς_τ.

The **reference-completion route** uses the inserted centre penalty $\sigma_0 = \Lambda_3^2$ of Section 17. It is a valid constructive theorem about an assembled Hamiltonian and is retained for that value, but it does not carry the paper's physical claim.

Designation note. The referee-recommended structure relabels the inserted-penalty material (Sections 15–20, Section 17.0, MC4, and the reference-completion reading of Part V and Theorem 6) as **Appendix A — Pre-Registered Centre-Projector Reference Completion**, with the substrate route (Section 20′ and the strict readings) as the main body. In this working manuscript those sections retain their present numbers to preserve cross-references; every reference to "Appendix A / the reference completion" designates exactly that inserted-penalty material in place. The relabelling is presentational — no inserted-penalty result is deleted, and none carries the principal claim.

Four cautions bound the strict route and are stated here once, each an explicit condition on the sections that follow:

1. **Local vs global coercivity — Local-Hessian Insufficiency Proposition.** *Positivity of the reduced wall Hessian at one stationary triality profile establishes local stability only.*

It does not imply a global positive sector infimum. A strict confinement proof additionally requires a coercive inequality valid on the complete nonzero-triality configuration class,

$\mathscr{E}_\tau[X] \geq A_\tau/R[X]^2 + B_\tau R[X]^2$ for every admissible $X \in \mathfrak{c}_\tau$,

with regulator-uniform positive constants. The wall Hessian $B_\tau > 0$ does not by itself exclude a nonperturbative path escaping the quadratic neighbourhood, delocalised zero-energy configurations, widening/fragmenting flux profiles, topology-changing or auxiliary escape channels, or regulator-dependent soft modes. The wall Hessian is evidence and a local coefficient; the global bound $\inf_{X\in\mathfrak{c}_\tau} \mathscr{E}_\tau[X] > 0$ is the confinement theorem, and it is owed (folded into D8). This proposition is stated formally to prevent a Hessian eigenvalue being silently promoted into a nonperturbative theorem.

2. **Non-Abelian flux typing (point on §17 localisation).** $\|\int_{D_R} \Pi\, d^2x\|$ is not gauge invariant in a non-Abelian theory — Lie-algebra values at different points cannot be added without parallel transport. The Cauchy–Schwarz localisation bound of Section 20′.4 may be used only after the flux is typed gauge-covariantly, e.g. via a based/centre-projected observable $\Phi_\Sigma = \int_\Sigma W(x_0,x)\, \Pi(x)\, W(x,x_0)\, d^2x$, and after the relation between the discrete triality class and $\Phi_\tau = 2\pi\tau\, H_c$ is proved at the physical-operator level. Owed.
3. **Slicing decomposition (Premise SC).** As in Section 21, the nonnegative-square-plus-endpoint form is stated, not derived; all same-order auxiliaries must be eliminated by one global Schur complement before the slice bound is taken.
4. **RG audit, not definition (point on §26).** B_τ must be computed *from the reduced wall Hessian* (Section 20′.4), never solved backward from a desired tension. Section 26 is accordingly an RG **consistency audit** — $d\,\sigma_\tau/d \ln k = 0$ and, in the two-term truncation, $d/d \ln k\, [g_s(k)\sqrt{B_\tau(k)}] = 0$ — not a determination of B_τ.

With these stated, the sections below are written in the inserted-σ_0 form for continuity with the reference completion; under the strict route, read $\sigma_0 \to \varsigma_\tau$ and $c_\sigma \to c_{\sigma,\tau}$ throughout, subject to cautions 1–4.

21. Discrete Slicing Inequality

Let the static sources be separated by

$L = N\, a_z.$

Gauss' law requires every one of the N separating transverse slices to lie in the same nonzero triality sector: the flux emitted by the source must cross each slice, and no adjoint process between slices can change τ (Section 14).

The argument *assumes* the complete longitudinal Hamiltonian can be written in the positive-square-completed form

$\hat{H}\_Q\bar{Q} - E_0 = a\_z \sum_{n=1}^{N} \hat{H}\_\{\perp,n\}\text{^red} + \sum_k R_k^\dagger R_k + H\_end,$

where the middle term collects nonnegative longitudinal and inter-slice squares and $H\_end \geq -c\_end$ with c_end finite and L-independent. **This decomposition is stated, not derived.** Establishing it requires showing explicitly that longitudinal derivatives, shared auxiliaries, Gauss constraints, any nonlocal terms, and cross-slice couplings can genuinely be collected into nonnegative squares plus an L-independent endpoint term. Until that is carried out, the slicing theorem is **conditional on the square-completion premise (Premise SC)**, and falsifier F15 guards that premise.

Discarding the nonnegative middle term,

$\hat{H}\_Q\bar{Q} - E_0 \geq a\_z \sum_{n=1}^{N} \hat{H}\_\{\perp,n\}\text{^red} - c\_end.$

Applying Theorem 5 to each slice,

$V\_\tau(L) \geq a\_z\ N\ \sigma_0 - c\_end,$

hence

■ $V\_\tau(L) \geq \sigma_0\ L - c\_end.$

The inequality is per-slice additive because the gap theorem is per-slice uniform: no correlation between slices is needed, only positivity of what joins them.

---

22. Continuum Slicing Theorem

Taking $a\_z \to 0$ with $N\ a\_z = L$ fixed gives

■ $\hat{H}\_Q\bar{Q} - E_0 \geq \int_0^L dz\ \mathscr{H}\_\perp\text{^red}(z) - c\_end.$

This is the required longitudinal slicing inequality. It follows from the displayed positive square completion, not from locality alone: locality without positivity permits negative inter-slice terms that could undercut the slice sum, and falsifier F15 stands against any such term in the completed Hamiltonian.

---

23. String-Tension Lower Bound and Candidate Saturation

The lower bound gives

$$\liminf_{L\to\infty} V_\tau(L)/L \ge \sigma_0.$$

For the candidate matching upper bound, consider a straight centre-flux trial-state ansatz: a one-link-wide centre-electric string joining the sources, carrying one primitive centre unit per longitudinal segment. Its centre-projector contribution is exactly $\sigma_0 L + O(1)$. The expectation values of the magnetic, closure and auxiliary sectors on this ansatz **have not yet been evaluated**. If their combined contribution is $O(1)$, the ansatz supplies

$$V_\tau(L) \le \sigma_0 L + O(1)$$

and establishes candidate saturation. Until that calculation is completed, only the lower bound

$$\liminf_{L\to\infty} V_\tau(L)/L \ge \sigma_0$$

is proved.

Define the reference lower slope $\underline{\sigma}_{ref} := \liminf_{L\to\infty} V_\tau(L)/L$. The proved reference result is $\underline{\sigma}_{ref} \ge \sigma_0$ (for the assembled Hamiltonian, under Premise SC). Only after a finite-slope upper bound or subadditivity proves the limit exists does $\sigma_{ref} := \lim_{L\to\infty} V_\tau(L)/L$ exist, at which point candidate saturation $\sigma_{ref} = \sigma_0$ becomes meaningful (conditional on F22). The matching upper bound is **not yet demonstrated** (the orthogonal-sector expectations on the trial ansatz are uncomputed): the straight trial state is asserted to leave every magnetic, closure and auxiliary square exactly in its vacuum, but those operators need not commute, so a state saturating the centre block need not simultaneously vacuum-annihilate the orthogonal squares. The defensible statement is therefore

$\underline{\sigma}_{ref} \ge \sigma_0$, and (once the limit exists) $\sigma_{ref} \le \bar{\sigma}_{trial}$ (after the full trial-state expectation is computed),

with the equality

$$\sigma_{ref} = \sigma_0$$

held (only once the limit exists) as a **candidate saturation result** until the explicit trial state and all its operator expectation values are displayed. Falsifier F16 and the new F22 guard this. The straight string is the natural minimiser candidate — every unit of energy it must carry is forced by Lemma 19.1 — but "must carry" is a lower bound, and the claim that it carries *nothing more* is the unproven half.

---

24. Scale Lock and Numerical String Tension

Dimensional transmutation converts the dimensionless boundary value $Z_3(\mu\star) = 1$ into a single scale; there is no separate pure-Yang–Mills scale Λ_0 in this paper:

$\ell\_C = \Lambda_3^{-1}$.

Principle MC4 *assigns* the branch-conditional three-flavour length $\Lambda\_\{3,U\}^{-1}$ as the normalising correlation length of the quenched centre-projector reference model — a reference-branch assignment, not an $N\_f = 3 \to 0$ matching theorem (D9). A primitive triality crossing is assigned one closure unit per colour-correlation area:

$\sigma_0 = r/\ell\_C^2$.

Dimensional transmutation, by itself, fixes only the *scale*, not the dimensionless coefficient. In general

▉ $\sigma\_ref = c\_\sigma \Lambda_3^2$ (only once the tension limit is shown to exist),

with c_σ a dimensionless nonperturbative invariant. Principle MC4 together with $r = 1$ (MC1) *sets* $c\_\sigma = 1$ — a clear candidate branch, but an assignment, not a derivation. The coefficient c_σ must ultimately come from the projected triality spectrum (Section 20′, Stage 5) or an explicit transverse variational calculation. Numerical values below use $c\_\sigma = 1$ and are labelled as carrying that assignment.

Using the numerical VBF scale of Theorem 3, and the assignment $c\_\sigma = 1$,

▉ $\sigma\_ref = (0.225923\ \text{GeV})^2 = 0.0510413\ \text{GeV}^2$ (under $c\_\sigma = 1$, candidate saturation, F22).

The candidate square-root reference tension is $\sqrt{\sigma\_ref} = \Lambda_3 = 225.923$ MeV, defined only if the limit exists (MC4 candidate saturation, guarded by F22; not a proven physical value).

No string-tension datum is used in this calculation, and none could be: the non-insertion firewall forbids it, and the derivation visibly has no slot for one. The tension is the square of the scale because the completion admits exactly one length and the centre cell is saturated at exactly one unit.

---

25. Wilson-Loop Area Law

Assume Premises RP and SC, and assume a finite-slope trial family or subadditivity establishes $V\_R(L) = \sigma\_\{\tau(R)\}\ L + O(1)$ with $0 < \sigma\_\{\tau(R)\} < \infty$. **Then** transfer-matrix positivity gives, for every nonzero-triality representation,

$-\ln\langle W\_R(L, T)\rangle = \sigma\_\{\tau(R)\}\ L\ T + O(L + T)$.

Without that finite-slope assumption, the established result is only a positive lower slope, not an area law. The centre-sector argument imposes no nonzero-triality area-law lower bound on zero-triality representations; their screening and asymptotic behaviour require a separate analysis. The positive spectral representation is retained as Premise RP (Section 17.1); the centre-term transfer matrix is not constructed, so falsifier F17 guards that premise rather than a proven property.

---

26. Bath Areal Coefficient and RG Lockstep

The coarse-grained tube functional of the consolidated gate expresses the slice gap as

$\varsigma_\tau = g_s(k)\, \|\Phi_1\|_{adj} \sqrt{2 B_\tau(k)}$.

The candidate physical value from Section 23 (lower slope proven; equality guarded by F22) is $\varsigma_\tau = \sigma_{ref}$ (only once the limit exists). With

$g_s = Z_3^{(-1/2)}$, $\|\Phi_1\|_{adj} = 4\pi/\sqrt{3}$,

solving for the bath coefficient gives

■ $B_{\tau,MC}^{implied}(k) = 3\, Z_{3,MC}(k)\, \sigma_{MC}^2 / (32\pi^2) = 3\, Z_{3,MC}(k)\, \Lambda_3^4 / (32\pi^2)$.

This is a **pre-registered diagnostic target** for the wall-Hessian calculation of Section 20′.4, not an independently derived bath response: it is the value the MC reference branch *requires*. It is not used upstream and is not the physical B_τ unless the independent reduced-wall-Hessian calculation returns it. The same status applies to R⋆ below.

In the **strict route** this relation is used only as an RG consistency audit, never to define B_τ: B_τ is computed from the reduced wall Hessian (Section 20′.4), and the identity below is then checked for consistency. Using a desired tension to define B_τ (as an earlier draft did by solving backward from $\sigma_{ref} = \Lambda_3^2$) is circular and is not permitted in the strict derivation.

Only the *implied diagnostic target* $B_{\tau,MC}^{implied}$ obeys $d \ln B / d \ln k = d \ln Z_3 / d \ln k$ automatically — because it was defined from the target tension. An independently computed physical B_τ must be *audited* for this relation; it does not hold by definition.

At the boundary $Z_3 = 1$,

■ $B_{\tau,MC}^{implied}(\mu\star) = 2.47466 \times 10^{-5}\ GeV^4$ (diagnostic target, not physical B_τ).

The tube minimiser is

$R\star = (4/\Lambda_3)\sqrt{(\pi/3)} = 18.1181\ GeV^{-1}$

in the declared coarse-grained tube approximation — a derived consistency figure, not an independent prediction.

---

27. Full-QCD String Breaking

Fundamental quarks carry nonzero triality, and this changes the asymptotics without touching the confinement statement.

Pair creation can screen the static sources by forming two gauge-invariant static-light singlets. For the lowest static-light energy $M_{Q\bar{q}}$, the variational bound is

$$V_F(L) \leq 2 M_{Q\bar{q}} + O(e^{-m_{gap} L}).$$

The variational reasoning establishes the upper bound $V_F(L) \leq \min[\sigma_{ref} L + \mu_{end}, 2M_{Q\bar{q}}] + \cdots$, not approximate equality. The displayed minimum and the string-breaking distance below are a **candidate crossover model** until the spectrum and mixing are evaluated:

$$L_{sb} = (2 M_{Q\bar{q}} - \mu_{end})/\sigma_{ref}.$$

The numerical evaluation of $M_{Q\bar{q}}$ and hence of L_{sb} belongs to the later hadronic/static-light completion ledger (Section 30). This shows the static string *can* break into two singlets. It does **not**, by itself, prove that no isolated finite-energy coloured asymptotic state exists anywhere in the physical Hilbert space — that stronger statement requires the full Gauss-law physical-state condition, a gauge-invariant asymptotic algebra, and a proof that every finite-energy physical state lies in a colour-singlet superselection sector. Those belong to the global-measure and physical-Hilbert-space programme and are not established here.

The defensible statement is therefore: string breaking is *consistent with* colour-singlet persistence, and no finite-energy nonzero-triality asymptotic state appears in the constructions examined; the complete no-free-colour theorem remains open. Falsifier F18 guards the consistency statement, not a completed theorem.

Part VI — Closure and Audit

28. Numerical Package

Every number frozen by this paper, in dependency order:

#	Symbol	Value	Fixed by
1	$C\_hol = r$	1	MC1
2	χ_C^census	48	§5 (census index; Firewall C1 — not primitive stiffness)
3	m_adj	1	§6 minimal one-copy ansatz
4	(b, C)	(5, 5)	constructed MC branch (§6)
5	$Z_3(\mu\star)$	1	conditional Unit Class-Action branch
6	$g\_s(\mu\star)$	1	§7
7	$\alpha\_s(\mu\star)$	$1/(4\pi) = 0.0795775$	§7
8	$\mu\star$	5.80 TeV	Theorem 2 (frozen Route-R)
9	m_t, m_b, m_c	291.827, 4.86254, 1.11241 GeV	§9 (frozen VBF)
10	Λ_6	73.2458 MeV	§10
11	Λ_5	150.611 MeV	§10
12	Λ_4	198.873 MeV	§10
13	$\Lambda\_\{3,U\}$^VBF (conditional)	225.923 MeV	Theorem 3 (Unit branch)
14	$\lVert\Phi_1\rVert\_adj$	$4\pi/\sqrt{3}$	§13
15	σ_0 (inserted, reference appendix)	$\Lambda_3^2 = 0.0510413\ \mathrm{GeV}^2$ under c_σ=1	App. A / MC4
15′	reference lower slope σ_ref	$\geq \sigma_0$; physical N_f=0 scale would involve Λ_0 (D9)	§20′.4
16	√σ_ref (candidate saturation, F22)	225.923 MeV	§24 (only if tension limit exists)
17	$B\_\{\tau,MC\}$^implied($\mu\star$)	$2.47466 \times 10^{-5}\ \mathrm{GeV}^4$	§26 (implied diagnostic)
18	$R\star$	$18.1181\ \mathrm{GeV}^{-1}$	§26 (reference-tube diagnostic)

These rows do **not** share one status and must not be read as collectively "calculated":

- C_hol = 1 is an MC1 premise; χ_C^census = 48 is an exact benchmark representation index (not a physical value); m_adj = 1 is the minimal one-copy auxiliary ansatz; (b, C) = (5, 5) is a constructed historical MC realisation under MC0–MC2; $Z_3 = 1$ is the Unit Class-Action target and the output of that constructed realisation, but not yet a physical substrate return; g_s = 1 and α_s = 1/(4π) are consequences of the branch output.
- Λ_6, Λ_5, Λ_4, Λ_{3,MC} are **conditional branch outputs** (conditional on $Z_3 = 1$).
- σ_0, B_{τ,MC}^implied and R⋆ belong **exclusively to the reference appendix** (App. A); they are inserted-model quantities, not physical outputs.
- The physical σ_ref, B_τ and c_σ remain **blank** — owed from D7/D8.

The arithmetic within each row is exact given its inputs; the *status* of those inputs is as tabulated. Values are quoted to six significant figures in the VBF scheme.

29. Falsification Conditions

F1 — Action-unit failure. The Holonomy–Wilson bridge does not identify the colour response with the primitive closure unit.

F2 — Saturated-bath failure. The colour class preserves a nontrivial sub-account, invalidating complete redistribution.

F3 — Adjoint multiplicity failure. More than one same-order adjoint redistribution copy exists.

F4 — Protected-unit failure. The irreducible colour response is not one primitive unit.

F5 — Schur-block failure (one-number fork, polar one-copy model only). Within the Conditional Polar Record-Link Ansatz and minimal one-copy coherent-channel reduction, $Z_3 = u\cdot\gamma_\eta/(1+\gamma_\eta)$, $\gamma_\eta = c_\eta/(\alpha^2 u)$. Under MC0 and $C_{hol} = 1$: $\gamma_\eta = 1/5 \Longrightarrow Z_3 = 1$. Also $\alpha = 0 \Longrightarrow \gamma_\eta \to \infty \Longrightarrow Z_3 = u$; $\gamma_\eta = 0 \Longrightarrow Z_3 = 0$. The strict physical test is the complete projected microscopic Hessian and need not reduce to γ_η. A benchmark returning $\alpha = 0$ does not by itself falsify the Unit Class-Action branch — it only shows that benchmark carries no intrinsic screening channel; falsification requires the physical colour-to-record embedding to return $\alpha = 0$. The earlier 'spectral eigenvalue-5' route is withdrawn, and the earlier headline/F5 zero–infinity correspondence is corrected here.

F6 — Isotropy failure. The eight colour directions return different reduced coefficients.

F7 — Record-boundary failure. QCD's physical boundary is not the 5.8 TeV record-coherence ceiling.

F8 — Threshold failure. The frozen VBF quark thresholds are not the correct decoupling inputs.

F9 — Centre-lift failure. No one-form-covariant global lift preserving $\mathbb{Z}_3$ triality exists.

F10 — Flux-normalisation failure. The primitive coweight is not H_c in the declared group.

F11 — Projector failure. Nonzero surface triality does not imply at least one nonzero centre-electric link.

F12 — Positivity failure. The reduced VERSF bath contains a negative term capable of cancelling the centre projector.

F13 — Vacuum degeneracy. A nonzero-triality state has zero energy.

F14 — Continuum failure. The centre surface operator or projector inequality is not refinement stable.

F15 — Slicing failure. Longitudinal cross terms are not expressible as nonnegative squares plus finite endpoint terms.

F16 — Scale-lock failure. Pure-gauge colour contains an independent infrared length in addition to Λ_3^{-1}.

F17 — Area-law failure. The transfer matrix lacks a positive spectral representation, contradicting the reflection-positivity argument of Section 17.1.

F18 — Free-colour failure. Full QCD contains an isolated finite-energy coloured asymptotic particle.

F19 — Numerical insertion. Any measured strong-interaction quantity is used upstream.

F20 — Physical embedding failure. No admissible *physical* embedding $\mathscr{E}_C : \mathfrak{su}(3) \to T_{\rho\star}\mathcal{C}_{sub}$ exists (rank eight, correct tangent metric, covariance-selected), or the closure object's colour response is not the SU(3)-adjoint response (Section 20′.2). This is a failure of $\mathscr{E}_C$, not of the abstract ι_C, and it falsifies the colour identification.

F21 — Penalty-provenance failure. The canonical projection of an explicit H_{cl} (Section 20′), once available, returns a colour sector whose triality blocks are *not* gapped, or a compatibility map f_C that is not monotone — i.e. the substrate does not produce the inserted centre penalty. This would show the confinement of Sections 15–20 to be an artifact of insertion rather than a property of the substrate.

F22 — Saturation failure. The straight-flux trial family does not establish a finite-slope upper bound equal to σ_0 (e.g. $\bar{\sigma}_{trial} > \sigma_0$ once the orthogonal squares are computed). The lower-slope result $\underline{\sigma}_{ref} \geq \sigma_0$ survives, but candidate saturation $\sigma_{ref} = \sigma_0$ fails or remains unproved.

F23 — Centre-lift double counting. The $\mathbb{Z}_3$ lift $\tilde{U}_\ell = z_\ell U_\ell$ is not equivalent to the centre sectors of the original SU(3) measure — the centre variables are double-counted — so the construction describes a different theory rather than a reduction of VERSF SU(3).

Each falsifier strikes a specific branch or conclusion, not the whole certificate. For example: F22 removes candidate saturation but leaves the lower slope $\underline{\sigma}_{ref} \geq \sigma_0$; F23 damages the auxiliary-$\mathbb{Z}_3$ reference lift but not necessarily the intrinsic centre action of Section 12; F3 revises the minimal auxiliary branch but need not invalidate physical colour propagation. The scope struck is stated per falsifier.

---

30. Open-Debt Ledger

The paper closes the operator architecture and freezes the reference branches; it does not close the strict NQTC-1 gate. It does not pretend to close what it defers. The outstanding debts, each typed and assigned, are:

D1 — Scheme bridge. The finite VBF → MS-bar conversion of Λ_3 and α_s. Required before any comparison with world-average strong-sector data. Assigned to the observable-bridge ledger. Deliberately absent here (firewall 3.7).

D2 — Higher-loop cascade. The threshold cascade is sharp one-loop by declaration. Two-loop running and finite decoupling constants shift Λ_3 within the declared order's error budget and belong to the precision ledger.

D3 — Static-light mass. The numerical value of $M_{Q\bar{q}}$ and hence the string-breaking distance L_{sb}. Assigned to the hadronic/static-light completion ledger.

D4 — Higher-representation tensions. k-string and adjoint-representation tensions beyond the triality-class statement of Section 25. Structural expectation (centre class determines asymptotic tension) is fixed; numerics are deferred.

D5 — Tube subleading structure. The coarse-grained tube functional of Section 26 is used only at leading matching. Its subleading (width, fluctuation) structure is not certified by this paper.

D6 — Colour embedding. *Abstract adjoint representation supplied* (Theorem 20′.1): R_C : SU(3) → SO(8) with $\mathbb{Z}_3$ kernel gives the canonical abstract adjoint action and target equivariance. **Still owed**: the physical substrate embedding $\mathcal{E}_C : \mathfrak{su}(3) \to T_{\rho\star}\mathcal{C}_{sub}$, proof of its rank eight, its tangent-metric normalisation, and covariance selection by the microscopic Hessian. The abstract representation is not the physical embedding. The **residual** is the covariance selection — proving the physical substrate Hessian $G\star$ actually selects this subgroup, i.e. $\|G\star - \mathfrak{T}_C(G\star)\| = 0$ on the saturated branch (Section 20′.2, Stage 1). That residual requires explicit $G\star$ and is folded into D7.

D7 — Physical embeddings and normalisations (operator content now computed). The three defect operators are now evaluated in closed form on the canonical K = 7 wheel cell (Section 6): $B_C = [\mathbb{1}_6 \mid \mathbb{1}_6 - S]$ with $\mathrm{spec}(B_C B_C^\dagger) = \{1,2,2,4,4,5\}$; the transport operator L_T with density spectrum $\kappa^2\{0,8,8,32,32,50,98\}$; $B_R = B_C \otimes \mathbb{1}_{rec}$; and the total $\mathbb{G}\star^{W_6}$. This substantially closes the *operator* content of D7. The **verified correction**: $b_C = 0$, because the five exact rim modes lie in ker B_C = Im D — the finite-cell incidence complex does *not* supply the five-unit Schur correction of the MC branch. What remains is therefore specific and physical, not mechanical: the colour embedding $\mathcal{E}_F : \mathfrak{su}(3) \to$ physical cycle modes; the auxiliary intertwiner $\mathcal{E}_\eta$ coupling colour curvature to the exact modes (the object the b = 5 branch needs and the incidence complex lacks); the background $\rho\star$; the coefficients κ_e; the relative normalisations of A_C, A_T, A_R; the active record representation; and the tangent metric. The conditional polar one-copy model is reduced to one number (Section 6); the general physical auxiliary embedding remains owed. The polar construction fixes the link from the record/completion map; gauge covariance separates the pure-gauge coboundary (which gives b = 0) from the intrinsic response

Ξ; and W_6 symmetry forces the shape $\Xi_{e,a} = i\alpha\, q_{F,e}\, T_a$ with q_F the uniform rim circulation. The Hessian blocks are then $H_{FF} = u$, $H_{F\eta} = \alpha u$, $H_{\eta\eta} = \alpha^2 u + c_\eta$, so $Z_3 = u\cdot\gamma_\eta/(1+\gamma_\eta)$ with the single invariant $\gamma_\eta = c_\eta/(\alpha^2 u)$. The MC branch requires exactly $\gamma_\eta = 1/5$ ($Z_3 = 1$); the executable factorised *benchmark* carries $\alpha = 0$ ($\gamma_\eta \to \infty$), giving $Z_3^{census} = u = 6C_{hol}$ — a matter-census benchmark that neither determines nor falsifies the physical primitive Z_3. Within the Conditional Polar Record-Link Ansatz and the minimal one-copy coherent-channel model, the local screening reduces to the single invariant γ_η. The strict independent-variable Hessian calculation remains the more general physical test and need not reduce to one scalar before its tangent census and embeddings are evaluated — it could contain several physical adjoint copies, non-coherent profiles, or multiplicity-space matrices not parameterised by one (α, c_η). Shared with the weak-doublet and PMNS gates.

D8 — Global quantum measure and physical Hamiltonian. The quotient-compatible global SU(3) measure, the reflection-positive transfer matrix $T_{a,R}$, the physical transverse Hamiltonian $H_\perp^{phys} = -(1/a_t) \log T$, global coercivity of the triality sector, and Mosco/norm-resolvent convergence of the regulated sector forms. This is the confinement-side blocker and is properly Paper 6 work; the penalty-free gap of Section 20′.4 is stated relative to it. Independent of D7: supplying the constitutive derivatives does not supply the global measure.

D9 — Quenched matching bridge. Derive the finite matching from the $N_f = 3$ VBF theory to an adjoint-only $N_f = 0$ theory and its separate scale Λ_0. Until this bridge is supplied, Λ_3 serves only as the declared normalisation of the centre-projector reference model and cannot be identified with a physical pure-Yang–Mills scale.

The value $\Lambda_{3,MC}^{VBF} = 225.923$ MeV is *arithmetically* independent of D7 once the pre-registered input $Z_{3,MC} = 1$ is imposed. The **physical** VERSF strong scale is *not* independent of D7, because D7 determines the physical Z_3 (and the wheel-cell result shows the branch value is not yet derived). The confinement-side verdict is not: D6 conditions Theorem 1's colour identification, and D7 is the prerequisite that would convert the inserted centre penalty (Section 17.0) into the derived gap of Section 20′. Debts D1–D5 remain non-load-bearing; D6–D7 are load-bearing for the confinement claim and are flagged as such.

---

31. The Non-Perturbative QCD Scale and Triality-Gap Completion Theorem

31.0 Strict NQTC-1 Conditional Closure Theorem

The strict physical theorem assumes **neither MC0 nor MC4** — MC0 belongs only to the historical 6−5 construction, MC4 only to the inserted Λ_3-normalised centre-projector model. It states exactly what the substrate must return.

Strict NQTC-1 Conditional Closure Theorem (two separate implications). The scale and the triality gap are distinct results with distinct hypotheses; D9 denies any completed bridge between them.

Scale implication. If the substrate returns the physical $Z_3(\mu\star) = C\_hol\cdot n\_C^{red} > 0$ (from G⋆, Section 20′, with no MC0 benchmark identification) and the common-scheme quark thresholds, then

$$\Lambda_3^{VBF} = \mu\star^{(7/9)} (m\_t\ m\_b\ m\_c)^{(2/27)} e^{(-8\pi^2 Z_3/9)}$$

is numerically fixed.

Triality-gap implication. In an **adjoint-only colour theory**, if the quotient-compatible measure, centre commutation, global coercivity, reflection-positive transfer matrix, convergence and longitudinal decomposition are established, then the nonzero-triality static potential has a strictly positive asymptotic lower slope $\liminf_{L\to\infty} V\_\tau(L)/L \geq \varsigma\_\tau > 0$. Existence of the tension limit requires an additional finite-slope or subadditivity input. Identifying that adjoint-only sector with a physical $N\_f = 0$ Yang–Mills theory, or expressing its tension in terms of Λ_3, additionally requires the quenched matching bridge D9.

The hypotheses are exactly Debts D7 (local: Z_3 from G⋆), D8 (global: measure, transfer matrix, coercivity, convergence) and — for a physical pure-Yang–Mills reading — D9. No conclusion follows from asymptotic freedom alone.

31.1 Reference-Completion Theorem (Theorem 6, retained)

Theorem 6 — NQTC-1

Assume the inherited VERSF colour carrier, holonomy map, Haar–Fubini–Study convention, Principles MC0–MC4 (MC0: benchmark u = 6C_hol; MC1: universal primitive action unit; MC2: saturated threefold bath; MC3: record-layer matching boundary; MC4: primitive centre-cell saturation), Premises VU, SC and RP, the frozen VBF quark thresholds, the centre-preserving global lift, and the compatible-convergence condition wherever a continuum conclusion is quoted.

Then:

1. **benchmark census result** $\chi\_C^{census} = 48$;
2. **minimal historical auxiliary ansatz** m_adj = 1, b = (5), C = (5);
3. **constructed historical branch return** $Z\_\{3,MC\}(\mu\star) = 1$, $g\_\{s,MC\}(\mu\star) = 1$;
4. at $\mu\star = 5.80$ TeV, the threshold cascade gives $\Lambda\_QCD^{(3)/vBf} = 225.923$ MeV;
5. the global centre lift has primitive flux $\Phi_1 = 2\pi\ H\_c$ with $\|\Phi_1\|\_adj = 4\pi/\sqrt{3}$;

6. *given the inserted centre penalty of Section 17 with coefficient σ_0 (Section 17.0)*, the nonzero-triality transverse quantum sectors satisfy $\varsigma_\tau^{qu} \geq \sigma_0$, uniformly in regulator and transverse volume, and under the MC4 assignment $\sigma_0 = \Lambda_3^2$;
7. under Premise SC, $\sigma_{ref} := \liminf_{L\to\infty} V_\tau(L)/L \geq \sigma_0$;
8. under MC4, $\sigma_0 = \Lambda_{3,U}^2$ is the inserted reference lower-slope scale; if a tension limit is independently shown to exist, $\sigma_{ref} = \sigma_0$ is the candidate saturation guarded by F22; a substrate derivation of σ_0 is the open target of Section 20′;
9. under Premises SC and RP **and a finite-slope upper bound** (existence of the tension limit), rectangular nonzero-triality Wilson loops obey $-\ln\langle W_R(L, T)\rangle = \sigma_{\tau(R)} L T + O(L + T)$;
10. with dynamical fundamental quarks the asymptotic string can break into colour-singlet states (string breaking); the complete statement that *no* isolated finite-energy coloured asymptotic state exists is not proved here and belongs to the global-measure programme.

Proof

Items 1–4 are Theorems 1–3 and the threshold calculation of Part II. Items 5–6 follow from the global lift (Part III), the centre-projector spectral inequality (Lemma 19.1) and positivity of the complete bath Hamiltonian (Theorem 5); regulator independence follows because the lower bound is an exact sector projector with fixed inserted reference coefficient (Section 20). Item 7's lower bound follows from the slicing inequality under Premise SC (Sections 21–22); its equality is the candidate saturation of Section 23 (F22). Item 8's coefficient $c_\sigma = 1$ is the MC4 assignment (Section 24), not derived. Item 9 follows under Premises SC and RP and the finite-slope upper bound (Sections 17.1, 25). Item 10's string-breaking half follows from the variational static-light bound (Section 27); the no-free-colour half is not established here. ■

The theorem is a closure statement of the explicit NQTC-1 completion, within the scope fixed by firewall 3.8.

---

32. Closure Certificate

Running side

- χ_C^{census}: **exact benchmark representation index — 48**
- physical direct response χ_C^{phys}: **owed**
- (b, C) = (5, 5): **constructed historical MC0–MC2 realisation**
- $Z_{3,MC} = 1$: **conditional reference result (physical Z_3 owed from G⋆)**
- $g_s(\mu\star)$: **MC-branch prediction — 1**
- declared scheme: **VBF**
- matching scale: **calculated on the record-layer branch — 5.80 TeV**
- quark thresholds: **frozen VBF values**
- Λ_{QCD}: **MC-branch value — $\Lambda_{3,MC}$ = 225.923 MeV, conditional on $Z_3 = 1$**

Confinement side

- global centre-sensitive lift: **candidate construction; equivalence to original SU(3) measure Owed (F23)**
- primitive flux normalisation: **derived — $4\pi/\sqrt{3}$**
- complete declared-order quantum transverse bath: **constructed**
- reflection positivity of the complete action: **Premise RP — transfer matrix not constructed (Section 17.1)**
- coercivity (inserted reference Hamiltonian only): **proved by positive square completion**
- compactness: **proved at finite regulator**
- vacuum uniqueness: **Premise VU — kernel intersection not computed here**
- longitudinal slicing inequality: **Premise SC — square-completion decomposition stated, not derived**
- centre penalty provenance: **inserted, not derived (Section 17.0) — [Owed output]**
- quantum nonzero-triality gap *of the assembled Hamiltonian*: **proved — ς_τ^qu $\geq \sigma_0$**
- continuum/infinite-volume uniformity of the *projector inequality*: **proved**; of the *physical coefficient*: **Owed (Section 20′)**
- σ̰_ref ≥ Λ_{3,U}²: **reference lower-slope result under MC4 and SC**
- σ_ref = Λ_{3,U}²: **candidate saturation, F22**
- full-QCD string breaking: **variational singlet-breaking channel, not a completed theorem**
- physical colour embedding ℰ_C : 𝔰𝔲(3) → T_{ρ⋆}𝒞_sub: **[Owed — Debt D6]** (the abstract ι_C is supplied; the physical embedding is not)
- explicit substrate Hessian H_cl: **[Owed — Debt D7, shared with flavour and lepton sectors]**

Closure verdict

■ NQTC-1's colour boundary is now framed at the correct grain: the 6 is a matter-census index, not primitive stiffness (Firewall C1); class ownership makes the primitive response one scalar Z_3 = C_hol·n_C^red, so the headline Z_3 = 1 is the Unit Class-Action branch conditional on three certificates (C_hol = 1, adjoint isotropy, n_C^red = 1). On the historical u = 6C_hol MC realisation only (Principle MC0), this unit-return condition is equivalent to γ_η = 1/5; in the general class-owned calculation the unit condition is C_hol·n_C^red = 1 and the corresponding γ_η = C_hol/(u−C_hol) depends on the independently calculated direct coefficient u. NQTC-1 contains this pre-registered Z_3 = 1 branch, an exact centre-projector reference completion, an explicit finite-cell Hessian calculation, and a penalty-free strict closure architecture. The auxiliary response is derived to one dimensionless invariant: Z_3 = u·γ_η/(1+γ_η), γ_η = c_η/(α²u), with the shape of the intertwiner fixed by polar decomposition and W_6 symmetry. The MC branch predicts γ_η = 1/5 exactly (Z_3 = 1); the executable factorised *benchmark* carries α = 0, hence Z_3^census = 6C_hol — a matter-census benchmark that neither determines nor falsifies the physical primitive Z_3. Within the Conditional Polar Record-Link Ansatz and minimal one-copy coherent-channel model, the local screening test reduces to γ_η; the strict independent-variable colour-boundary calculation remains the projected microscopic Hessian on the

complete physical adjoint multiplicity space and need not reduce to one scalar. The full strict closure additionally requires the global blocker D8 (measure, transfer matrix, coercivity, convergence). The conditional polar branch is one ratio away from a local verdict; the strict physical boundary awaits the physical colour-to-record embedding $\mathscr{E}$_C.

The result remains falsifiable by any failure of the explicit completion principles (Section 29) or by downstream disagreement after the scheme and observable bridges of Section 30 are applied.

---

33. Conclusion

The paper has not yet calculated the physical VERSF strong boundary or proved physical QCD confinement. It has done something more precise than an architectural gesture: it has frozen one falsifiable numerical branch, identified the finite operators whose evaluation will confirm or reject that branch, and separated the local and global calculations required for a regulator-controlled triality gap. The two halves are best read separately — the pre-registered branch first, then the physical route.

The pre-registered branch. The tempting reading begins by calling the direct colour response "six units of stiffness." But that six is a matter-census index (Section 5), not a primitive stiffness, and class ownership makes the primitive response one scalar. On the optional MC realisation, one *posits* a saturated bath that reorganises every non-primitive part of the response and an intertwiner that removes five units, leaving $Z_3 = 1$. At the record-layer boundary and through the frozen quark thresholds, that single number generates the conditional branch scale Λ_{3,U}^VBF = 225.923 MeV. Whether the physical substrate returns it — via C_hol = 1 and n_C^red = 1, and, under the historical MC0 identification only, equivalently γ_η = 1/5 — is exactly what the calculation of Sections 6 and 20′ will decide.

The second calculation was confinement. Local SU(3) curvature was insufficient because triality belongs to the global centre. The $\mathbb{Z}_3$ lift makes the centre flux a quantum sector label. Its nonzero-flux projector is exact, positive and discrete. A nonzero-triality surface must excite it at least once. That statement cannot weaken as the lattice is refined or the volume enlarged, because it is arithmetic, not analysis.

In the reference completion (Appendix A), the gap is regulator-uniform, ς_τ^qu ≥ σ_0, and under MC4 the projector and slicing arguments give σ̱_ref ≥ Λ_{3,U}²; equality remains the unproved straight-string saturation. This is the inserted-penalty model, not the substrate derivation. The physical route (Section 20′) replaces it: the same gap must emerge from the evaluated Z_3 and B_τ with no inserted coefficient, and its uniform positivity is owed from the global measure (D8).

Asymptotic freedom and confinement remain distinct throughout. The first generates the scale; the second is protected by the global centre sector. In the strict programme, asymptotic scale generation and centre-sector confinement remain separate calculations. They are linked in the reference completion only through the MC4 assignment of Λ_{3,U}⁻¹ as its normalising length.

The paper's primary chain is

■ $\mathbb{C}^3$ colour class ⟶ $SU(3)_C$ ⟶ $\mathcal{K}_C^{prim} = z_C \mathbb{1}_8$ ⟶ $Z_3 = C_{hol} \cdot n_C^{red}$ ⟶ $\Lambda_3(Z_3)$,

with the Unit Class-Action branch as a separate, labelled line:

■ $C_{hol} = 1$, $n_C^{red} = 1 \Rightarrow Z_3 = 1 \Rightarrow \Lambda_{3,U}^{VBF} = 225.923$ MeV.

The historical arithmetic

$\mathbb{C}^3$ bath ⟶ $\chi_C^{census} = 48$ ⟶ $(b, C) = (5, 5)$ ⟶ $Z_3 = 1$

appears only as the *pre-registered MC realisation* — not the primitive derivation, since its (b, C) = (5, 5) is a constructed conditional block (Section 6), not a substrate return. The confinement chain is

$\mathbb{Z}_3$ centre ⟶ $P^{\neq 0}$ ⟶ ς_τ (reference model) ⟶ $V_\tau(L)$ linear ⟶ $\sigma_{ref} \geq \sigma_0$,

with the physical version replacing the inserted penalty by the two evaluated substrate responses of Section 20′.4; the projector and slicing arguments give the reference lower bound $\liminf V(L)/L \geq \Lambda_{3,U}^2$ under MC4; equality remains the unproved straight-string saturation.

The constructed variational breaking channel ends in two colour singlets. A general no-free-colour theorem remains open.

The nonperturbative gate therefore returns objects of three different standings: a verified group-theoretic embedding (Theorem 20′.1), a penalty-free confinement architecture reduced to two evaluable substrate responses (Section 20′.4), and — pending the substrate colour-response evaluation — a pre-registered scale branch. The scale on that branch is $\Lambda_{3,U}^{VBF} = 225.923$ MeV, frozen before observation with no measured strong-sector quantity used upstream, but it is the prediction of the Unit Class-Action branch, not yet a derived physical number. The confinement result is an exact regulator-uniform finite-regulator lower bound for the assembled centre-projector Hamiltonian, with continuum passage conditional on compatible refinement and form convergence — exact in what it proves, but resting on an inserted penalty (Section 17.0) the substrate has not yet been shown to produce.

The honest closing statement is therefore sharper than a bare claim of confinement. Two pieces of the derivation route are now in hand and verified: the canonical abstract adjoint representation $R_C : SU(3) \to SO(8)$ with $\mathbb{Z}_3$ kernel (Theorem 20′.1, checked numerically), which supplies the equivariance target the physical embedding must realise (it does not itself place the adjoint in the physical closure carrier), and the penalty-free gap architecture (Section 20′.4), under which a physical gap would follow from evaluated Z_3 and B_τ, a globally coercive sector bound, and the D8 measure and convergence certificates — with no $P_{\neq 0}$ term added to the Hamiltonian. (These responses are not yet evaluated, and local Hessian positivity alone is insufficient for global coercivity.) The gap it would deliver, $\varsigma_\tau \geq g_s \|\Phi_\tau\|_{adj} \sqrt{2B_\tau}$, with uniform limit

$\liminf \Delta_\tau \geq \|\Phi_\tau\|_{adj}\sqrt{(2B_{min}/Z_{max})}$, would be a consequence of the substrate responses rather than of an insertion — once those responses and the global coercivity are supplied.

What remains are two specified blockers (Section 30): locally, the three linearised defect operators B_C, L_T, B_R of Section 20′.3 with their substrate graph and background — from which $\mathbb{G}\star$, and hence Z_3 and the wall coefficient B_τ, follow as finite matrix computations, the second derivatives dropping out by the Gauss–Newton reduction at the admissible manifold; and globally, the quotient-compatible measure and transfer matrix of Paper 6. The coefficient 5 does **not** have a derived origin: the C_6 Hodge decomposition supplies five exact directions, but the wheel-cell calculation places them in $\ker B_C$, so their multiplicity generates *no* five-unit colour-screening correction ($b_C = 0$). A nonzero correction requires an independently derived physical intertwiner $\mathscr{E}_\eta$ — that is now the sharpest open calculation. Until the local data are supplied, the MC-branch values $(\chi_C, b, C, Z_3) = (48, 5, 5, 1)$ and $c_\sigma = 1$ stand as *pre-registered predictions*, and the scale 225.923 MeV is the frozen output of that branch. The strong sector therefore stands with a scale conditional on the pre-registered $Z_3 = 1$, a penalty-free confinement architecture reduced to two evaluable substrate responses, and finite matrix evaluations for the local response programme; full confinement closure additionally requires the global measure, transfer matrix, coercivity and convergence programme of D8, plus D9 before any physical $N_f = 0$ interpretation. The paper says exactly this, and no more.