The Strange Confinement Operator in VERSF

Deriving the Balanced Support Module, Complement Breaking, and Unit Self-Weight from Hadronic Closure

Keith Taylor VERSF Theoretical Physics Programme — Quark Mass Hierarchy Series

Successor to *The Strange Current-Mass Operator in VERSF: Resolved Balanced Support Readout, Complement Breaking, and the Twentyfold Down-to-Strange Ratio.* This paper is the attempted closure of that paper's last open problem, **QM-S7**: replace the representation-level support ansatz with an explicit confinement operator.

Summary for the General Reader

The previous strange-current paper proved something precise but conditional. *If* the strange quark is read as twenty balanced support sectors inside a six-channel interface, and *if* the down quark is read as one unresolved baseline, then the strange-to-down current-mass ratio is exactly 20. That paper turned the old observation "six-choose-three equals twenty" into a representation-level operator statement, and it isolated the weak point with full honesty: the balanced support space, the complement breaking, and the unit contribution of each support atom still had to be *derived from confinement* rather than assumed.

This paper attempts that derivation. The aim is to construct a candidate strange confinement operator whose low-readout space is forced to be the balanced three-against-three support shell, rather than chosen to be. The guiding idea is simple. A confined support is not a list of channels; it is a recorded **boundary** between what is occupied and what is exterior. In a six-channel interface, the partial boundary of largest capacity is a split of three occupied channels against three exterior channels. There are twenty ways to choose the occupied side, and if the readout records the occupied side without identifying it with the exterior side, those twenty stay twenty rather than collapsing to ten complement pairs.

The construction makes three moves, and the second and third are the genuinely new content of this version.

1. **Shell selection by one functional.** A single quantity — the boundary capacity $b(T) = |T|\cdot|T^c|$ — selects the central shell. It is maximal, uniquely among partial supports, at the balanced 3+3 split. Crucially it is *not* the binomial coefficient; it is an independent function that happens to peak at the same place, so "k = 3 is largest" and "k = 3 is most balanced" become two readings of one confinement functional rather than a coincidence.
2. **Complement breaking as a role record, with a witness.** Occupied and exterior are physically distinct boundary conditions, so the complement map $T \mapsto T^c$ reverses what is

realised rather than relabelling it. This is certified by an explicit complement-odd observable living in the 10-dimensional antisymmetric sector — the same sign-flipping probe the current-mass paper demanded.

3. **Unit self-weight as a canonical invariant, not a free scalar.** This is the sharpest improvement, but its scope must be stated precisely. The nine occupied/exterior contacts of a 3+3 split form a boundary-incidence module that, under the support's own stabiliser Sym(3)×Sym(3), has a **canonically one-dimensional invariant**. *If* confinement coarse-grains each support onto that invariant, the support atom is that line and the self-weight $a_0 = 1$ follows as a dimension count rather than a chosen number. This does not eliminate the premise; it **relocates it to a more natural one** — "use the canonical invariant projector" replaces "let each atom weigh one." That is real progress (a canonical projector is far better motivated than a free scalar), but the obligation to show VERSF confinement actually uses this projector is moved, not discharged.

The result is not advertised as a closed Standard Model derivation. It is the first explicit version of the missing bridge. Its claim is conditional but sharper than before: if the VERSF hadronic confinement operator carries the balance, orientation, and coherence structure proposed here, then the strange current-mass readout space is *forced* to be the twenty-dimensional balanced support module and the trace is 20. If the explicit confinement theory instead delivers a complement quotient, ordered paths, channel incidences, or a broken channel symmetry, the mechanism fails with a definite wrong number. That is the right kind of progress: the successful strange ratio is no longer merely fitted, but attached to a concrete confinement operator with clean ways to succeed or fail.

Abstract

The previous VERSF strange-current manuscript reduced the strange-to-down current-mass ratio to a trace ratio between an unresolved down baseline and a resolved balanced support module:

$\mathscr{H}\_d = \mathbb{C}[\mathcal{S}_1(I)]^{\mathrm{Sym}(6)}$, dim $\mathscr{H}\_d = 1$; $\mathscr{H}\_s = \mathbb{C}[\,\mathrm{Sym}(6)/(\mathrm{Sym}(3)\times\mathrm{Sym}(3))\,] \cong \mathbb{C}[\mathcal{S}_3(I)]$, dim $\mathscr{H}\_s = C(6,3) = 20$,

with Tr $O\_d = 1$, Tr $O\_s = 20\cdot a_0$, and $a_0 = 1 \Longrightarrow m\_s/m\_d = 20$. That construction solved the representation-level multiplicity problem but left four physical premises open: why confinement selects the balanced 3+3 shell, refuses the complement quotient, ignores internal ordering, and assigns unit self-weight $a_0 = 1$.

This paper gives a first candidate confinement operator that supplies all four. Let *I* be the six primitive interface channels and *G* = Sym(*I*). On the full shell space $\mathscr{H}\_\mathrm{shell} = \bigoplus\_\{k=0\}^\{6\}$ $\mathbb{C}[\mathcal{S}\_k(I)]$, the proposed strange confinement selector is

$K\_s = \alpha\cdot(N - 3)^2 + \beta\cdot K\_\mathrm{role} + \gamma\cdot K\_\mathrm{inc}$,

with N the shell-number operator, and its low-readout (kernel) space is shown to be $\mathbb{C}[\mathcal{S}_3(\mathrm{I})]$. Three results carry the construction:

- **(Shell)** The balance penalty $(N - 3)^2$ is affinely identical to the boundary capacity $b(T) = |T|\cdot|T^c|$ — explicitly K_bal = $4\cdot(9 - b)$ — whose values 0, 5, 8, **9**, 8, 5, 0 single out k = 3 uniquely among partial supports. One Sym(6)-invariant functional, two readings (largest / most balanced).
- **(Orientation)** The complement involution acts *freely* on the twenty triads (no 3-set meets its complement), pairing them into ten pairs; keeping orientation means inducing from Sym(3)×Sym(3) rather than the wreath product Sym(3)≀Sym(2), giving 20 rather than 10. A complement-odd role observable in the 10-dimensional antisymmetric sector witnesses the inequivalence.
- **(Self-weight)** The nine boundary incidences of a 3+3 split form the module $\mathbb{C}^3 \otimes \mathbb{C}^3$ under the support stabiliser Sym(3)×Sym(3), whose invariant subspace is **canonically one-dimensional**. Coarse-graining each support's incidences onto this invariant line yields exactly one trace unit per support — a derivation of $a_0 = 1$ inside the canonical boundary-coherence model, rather than a bare trace-normalisation assumption.

Sym(6)-equivariance then places the induced mass observable in the Johnson-scheme commutant of J(6,3), so its trace is $20\cdot a_0$ for *any* equivariant operator, and the incidence-coherence result fixes $a_0 = 1$. Hence Tr O_s = 20, Tr O_d = 1, and the conditional ratio is m_s/m_d = 20.

One scope line is stated up front rather than buried: the self-weight result derives $a_0 = 1$ *inside the canonical boundary-coherence model*, and the remaining first-principles obligation is to show that VERSF confinement actually uses this canonical invariant projector as its low-readout selector.

The contribution is bounded. The paper does not derive the six primitive channels, the absolute mass scale, the light ratio 6/13, the charm increment, the bottom amplifier, or the five-gate factorisation (QM-S6 remains open). It proposes a confinement operator that would close the strange support-trace gap *if* VERSF confinement geometry supplies the balance functional, oriented boundary record, and incidence coherence. The failure modes stay explicit and numerical: full down multiplicity → 10/3; complement quotient → 10; ordered triads → 120; channel incidences → 60 or 180; noncentral shells → 6 or 15; broken Sym(6) → no constant diagonal.

Table of Contents

---

0. Predictive-Content Ledger

This paper addresses one inherited target:

QM-S7: connect the strange current-mass support ansatz to an explicit hadronic confinement operator.

The previous paper reduced the strange trace to two requirements — choose the correct readout space, and set the self-weight a_0 to one. This paper proposes a confinement operator that attempts to force both. The discipline is unchanged: **matching 20 is not deriving 20.** The task is to show why the confinement operator selects the twenty-dimensional support module rather than a neighbour.

Object	Value or construction	Status in this paper
Primitive interface set	I, $\|I\| = 6$	inherited
Symmetry group	$G = \mathrm{Sym}(I) \cong \mathrm{Sym}(6)$	inherited
Full shell space	$\mathscr{H}$_shell = $\bigoplus$_k $\mathbb{C}$[$\mathcal{S}$_k(I)]	constructed
Down pre-boundary readout	$\mathscr{H}$_d = $\mathbb{C}$[$\mathcal{S}_1$(I)]^G	proposed confinement ground line
Down trace	Tr O_d = 1	exact given $\mathscr{H}$_d
Boundary capacity	$b(T) = \|T\|\cdot\|T^c\|$	proposed closure functional

Object	Value or construction	Status in this paper
Balance penalty	K_bal = $(2\|T\| - 6)^2 = 4(9 - b)$	affinely equivalent to b
Selected strange shell	$\|T\| = 3$	exact unique max of b(k) among partial supports
Strange support space	$\mathscr{H}$_s = $\mathbb{C}[\mathcal{S}_3(I)]$	candidate confinement readout
Stabiliser	Sym(3) × Sym(3)	oriented occupied/exterior support
Complement quotient	absent	role record + complement-odd witness
Incidence module per support	$\mathbb{C}^3 \otimes \mathbb{C}^3$ (dim 9)	boundary contacts of a 3+3 split
Self-weight a_0	1	**dim of canonical Sym(3)×Sym(3)-invariant = 1**
Strange trace	Tr O_s = 20	exact given the above
Five-gate reconciliation	not solved	deferred (QM-S6)

The paper separates three layers, as before: (1) the trace result is exact given the operator; (2) the operator-construction claim is proposed; (3) the dynamical derivation of K_s from a VERSF field action is still owed.

---

1. The Inherited Task

The strange current-mass paper established a clean conditional result. With down assigned to the invariant line of the point module and strange to the full balanced support module,

$\mathscr{H}$_d = $\mathbb{C}[\mathcal{S}_1(I)]$^Sym(6), dim $\mathscr{H}$_d = 1; $\mathscr{H}$_s = $\mathbb{C}$[Sym(6)/(Sym(3)×Sym(3))] ≅ $\mathbb{C}[\mathcal{S}_3(I)]$, dim $\mathscr{H}$_s = C(6,3) = 20.

The mixed rule was the content. If down and strange were both quotient-counted, the ratio is 1; if both raw-counted, it is 20/6 = 10/3; the target 20 appears only when down keeps the trivial invariant baseline while strange keeps the full support module. That shifted the question from "does C(6,3) = 20?" to six structural ones:

1. Why does the down state stay unresolved?
2. Why is strange the first resolved support state?
3. Why is the resolved support rank three?
4. Why is the complement not quotiented?
5. Why is the trace unit one per support atom, not one per channel incidence?
6. Why does Sym(6)-equivariance survive in the mass observable?

This paper answers (1)–(6) by introducing a candidate confinement operator whose readout is not a free choice but the selected stable partial closure of the six-channel interface. (Question 6 was already largely settled by the Johnson-scheme argument; here it is re-derived from a manifestly Sym(6)-invariant selector.)

---

2. What It Means to Derive the Support Module from Confinement

A representation-level ansatz *chooses* a space. A confinement derivation must explain why that space is the physically recorded one. Here a derivation must supply three things.

A selection principle over shells. The shell sizes for a six-channel interface are

C(6,0), …, C(6,6) = 1, 6, 15, **20**, 15, 6, 1.

A derivation cannot merely note that k = 3 is the largest shell; it must connect largeness or balance to confinement. We do this with the boundary-capacity functional of §6, and we show the "largest" and "most balanced" readings are the *same* functional.

An orientation principle. A 3+3 split is either an unordered bipartition $\{T, T^c\}$ (10 states) or an occupied/exterior split (20 states). The strange mechanism needs the latter; §7 supplies the role record and its witness.

A normalization principle. The boundary of a triad against its exterior has nine channel contacts. An incidence-counting theory would give 9×20 = 180, or 3×20 = 60 per occupied channel. The mechanism must instead coarse-grain each stable boundary to one current-mass atom. §8 shows this coarse-graining is *canonical* — projection onto a uniquely one-dimensional invariant — and is therefore the place the self-weight $a_0 = 1$ actually comes from.

The candidate is not a full dynamical derivation from a VERSF field action. It is a minimal **operator-level closure model** that selects exactly the space and normalization the current-mass trace requires, with every premise turned into a checkable spectral condition.

---

3. The Full Six-Channel Shell Space

Let $I = \{1,\dots,6\}$. For each k define the k-support shell $\mathcal{S}$_k(I) = $\{ T \subseteq I : |T| = k \}$, and the full finite shell space

$\mathscr{H}$_shell = ⊕_{k=0}^{6} ℂ[$\mathcal{S}$_k(I)].

A basis vector $|T\rangle$ represents an occupied support T with exterior $T^c = I \setminus T$. The group G = Sym(I) acts by relabelling, $\sigma|T\rangle = |\sigma T\rangle$. Two structural operators:

shell number: $N|T\rangle = |T|\cdot|T\rangle$. **complement involution:** $\mathcal{C}|T\rangle = |T^c\rangle$, $\mathcal{C}^2 = 1$, sending shell k to shell 6 − k.

The central shell k = 3 is *fixed as a shell* by $\mathcal{C}$ (it is the unique fixed point of $k \mapsto 6 - k$), but **no individual triad is fixed**: since $T \cap T^c = \emptyset$, $T = T^c$ is impossible, so $\mathcal{C}$ acts *freely* on the twenty triads, pairing them into ten complementary pairs. This two-level behaviour — fixed shell, free action within the shell — is the whole hinge of the construction. Shell balance is structural; occupation is physical.

---

4. Pre-Boundary Confinement and the Invariant Down Baseline

Before a support boundary is recorded, an individual channel label is not yet a physical support sector; it is a coordinate inside the symmetric interface. This is the proposed reason the down quark reads one baseline, not six singletons.

Let $V_1 = \mathbb{C}[\mathcal{S}_1(I)]$ be the point module, decomposing as

$V_1 \cong S^{(6)} \oplus S^{(5,1)}$, dims 1 + 5 = 6.

The pre-boundary confinement operator K_d is required to commute with G and to place its lowest readout on the invariant line:

$\ker K_d = V_1^G = S^{(6)}$.

A minimal spectral model is $K_d = \lambda_{std}\cdot P_{std}$, where P_{std} projects onto the standard component $S^{(5,1)}$ and $\lambda_{std} > 0$. Then the invariant vector

$\Omega_d = (1/\sqrt{6}) \Sigma_{i\in I} |\{i\}\rangle$

has zero confinement penalty, while any resolved singleton asymmetry costs positive confinement energy. This gives the down denominator

$\mathcal{H}_d = \ker K_d = \mathrm{span}\{\Omega_d\}$, $O_d = I_{\mathcal{H}_d}$, $\mathrm{Tr}\, O_d = 1$.

The physical reading is direct: at depth one the interface has not yet formed a stable occupied/exterior boundary, so the six labels are six redundant coordinates of one symmetric baseline, not six mass sectors. This remains a premise about confinement, but now a precise one — **the depth-one operator must gap $S^{(5,1)}$ and leave $S^{(6)}$** — which is exactly QM-S1.

4.1 The down/strange asymmetry is the load-bearing premise. It must be stated sharply, because the entire ratio is decided here. The strange sector is selected by *capacity + coherence* (§§6–8); the down sector is selected by a *different* operator (K_d gapping $S^{(5,1)}$). These are two mechanisms, and they are not automatically compatible — applying the strange mechanism uniformly to k = 1 does **not** reproduce Tr O_d = 1. Concretely: the boundary capacity is nonzero at k = 1, $b(1) = 1\cdot5 = 5 \neq 0$, so by the strange-style logic a singleton *does* sustain a boundary. Carrying the strange construction through at k = 1 would give the oriented singleton module of dimension C(6,1) = 6, cohere each 1+5 incidence block $\mathbb{C}^1 \otimes \mathbb{C}^5$ onto its Sym(1)×Sym(5)-invariant (again dimension 1·1 = 1), and land on Tr O_d = 6 — hence m_s/m_d = 20/6 = 10/3, which is precisely falsifier 14.1/14.2.

So the construction **requires the down sector to be exempt from the very mechanism that drives the strange sector**, and the only thing currently enforcing that exemption is the narrative "depth one has not formed a stable boundary." Since $b(1) = 5$ is not zero, that narrative cannot be "k = 1 carries no boundary capacity"; it must be the stronger and as-yet-underived claim that nonzero capacity at k = 1 does not *trigger boundary formation* while capacity at k = 3 does. Two sharper observations make the size of this gap precise.

The selector itself contains no threshold. On [0, 3] the capacity $b(k)$ = 0, 5, 8, 9 is strictly increasing — monotone, with no internal feature. So there is nothing *inside b* — the quantity that does all the §6–8 work — that could distinguish "k = 1 doesn't trigger boundary formation" from "k = 3 does." The exempting threshold must be supplied by physics genuinely orthogonal to the capacity functional. The gap is therefore not merely "underived"; it is **"un-derivable from the selector that derives everything else,"** which is a stronger statement and the correct one.

Down and strange are selected by operators of opposite logic. K_d (§4) rewards the *unresolved* symmetric state: its kernel is the invariant line and it gaps S^(5,1) — it says *do not resolve*. K_s (§9) rewards *maximal resolution*: its kernel is the full balanced support module — it says *resolve as much as the boundary allows*. The number 20 is exactly the ratio of these two opposite verdicts (one baseline mode against twenty resolved supports), and at present there is **no single principle from which both operators descend**. So QM-S1/QM-S2 is not a missing lemma to be filled in by routine work; it is a missing *unification* of two operators currently postulated to behave inversely. That is the honest shape of the open problem. (For the record, the k = 2 sector does not create a competing intermediate: within the strange problem it is simply gapped by $(N - 3)^2$, so no spurious shell arises between down and strange. The entire difficulty lives at the down end.)

---

5. Boundary Formation and the Support-Record Map

A support becomes a current-mass sector only when confinement records a boundary between an occupied set T and its exterior T^c. The boundary is not a decorative label; it is the physical event that turns a subset into a support atom.

Define the occupied/exterior incidence set of T,

$\partial T = \{ (i, j) : i \in T, j \in T^c \}, |\partial T| = |T|\cdot|T^c|.$

For $|I| = 6$ this depends only on $k = |T|$: writing $b(k) = k(6 - k)$,

$b(0)=0$, $b(1)=5$, $b(2)=8$, $b(3)=\mathbf{9}$, $b(4)=8$, $b(5)=5$, $b(6)=0$.

So $k = 3$ is the unique partial shell of maximal boundary capacity. The support-record map $\mathscr{R}$ coheres a boundary-incidence pattern into one support atom. For each T define the **uniform boundary vector**

$$|b_T\rangle = (1/\sqrt{|\partial T|}) \Sigma_{(i,j)\in\partial T} |i \to j; T\rangle,$$

and set $\mathscr{R}|b_T\rangle = |T\rangle$. The normalization is not cosmetic: it is precisely what prevents the trace from becoming an incidence count. §8 shows $|b_T\rangle$ is not an arbitrary choice of cohering vector but the *canonical* one — the unique invariant of the boundary module — which is what makes the resulting self-weight forced.

---

6. Central-Shell Selection from One Boundary Functional

The shell-selection operator on $\mathscr{H}$_shell may be written in two ways that look different but are the same functional:

capacity (reward): $B|T\rangle = |T|\cdot|T^c|\cdot|T\rangle = b(k)\cdot|T\rangle$, **imbalance (penalty):** $K\_bal|T\rangle = (|T| - |T^c|)^2\cdot|T\rangle = (2k - 6)^2\cdot|T\rangle$.

These are affinely related: since $b(k) = 9 - (k - 3)^2$ and $(2k - 6)^2 = 4(k - 3)^2$,

$$\mathbf{K\_bal = 4\cdot(9 - B),}$$

so maximizing capacity and minimizing imbalance are literally the same eigenvalue ordering. There is one confinement functional with two readings, not two coincidentally-aligned criteria. Among nontrivial partial supports it is extremal uniquely at $k = 3$.

Because both forms depend only on $|T|$, they commute with Sym(6). Central-shell selection therefore **does not break channel symmetry**: it selects the shell without selecting a preferred triad. The selection is two-stage:

1. balance/capacity selects $k = 3$;
2. symmetry preserves all twenty triads inside that shell with equal status.

Formally, restricted to partial supports, $\ker K\_bal = \mathbb{C}[\mathcal{S}_3(I)]$ and K_bal commutes with Sym(6). This is the first bridge from confinement to the support module.

6.1 Why boundary capacity is not an arbitrary selector. The functional $b(T)$ is not introduced because it happens to peak at $k = 3$. It is the **minimal local scalar a partial confinement boundary can carry** under the assumptions already inherited from the six-channel interface, and that minimality is what answers the obvious objection — that the shell was reverse-engineered to yield twenty.

A partial closure $T \subset I$ is not stabilised by the number of occupied channels alone. A set with no exterior has no boundary; a set with no occupied support has no realised state. What stabilises a confined boundary is the number of occupied/exterior constraints it can sustain across that boundary. If each occupied primitive channel couples uniformly to each exterior primitive channel — the same uniform-allocation symmetry inherited from the interface sector — then the total boundary-contact capacity is the count of occupied/exterior pairs,

$b(T) = |T| \cdot |T^c|$.

This is the **unique** symmetric bilinear contact count under the stated assumptions — dependence only on the occupied/exterior populations, vanishing when either side is absent, and uniform pairwise channel coupling: it is invariant under complement exchange at the shell level, and grows with the number of sustainable occupied/exterior constraints. The honest scope is therefore *uniqueness within the class of bilinear contact counts*: granting that confinement stability is measured by a symmetric bilinear count of occupied/exterior contacts that vanishes on an empty side, the form is forced to a scalar multiple of $|T| \cdot |T^c|$. The bilinearity and contact-count assumptions are themselves the physics input (QM-S3) — they are not derived here, and once they are granted the central maximum is automatic by concavity. So $b(T)$ is not a binomial-counting rule dressed as physics — it is a boundary-stability rule *of an assumed form* — and the binomial coefficient enters only afterwards, to count how many maximal-capacity supports exist.

The logical ordering is the point. If the mechanism merely selected the *largest shell*, it would stay vulnerable to the charge that the shell was chosen to obtain twenty. But confinement here selects the **maximal occupied/exterior boundary**, which fixes k = 3 from $|T|$ alone — *before* any triad is counted. The number twenty is therefore downstream of the confinement selector, not the selector itself.

A clarification on what the uniqueness theorem does and does not buy. The selection of k = 3 is in fact **over-determined** and does not rely on the special form of b at all: *any* increasing symmetric function of the two role-populations peaks at the balanced 3+3 split. For instance $b = |T||T^c|$, $b \cdot (|T|+|T^c|)$, b^2, and $\min(|T|, |T^c|)$ all take their maximum at k = 3 (values 0,5,8,**9**,8,5,0 / ·6 / squared / 0,1,2,**3**,2,1,0 respectively). So the physically robust content — *k = 3 wins* — survives any reasonable choice of stability measure and is a genuine strength: the central-shell outcome is insensitive to functional form. What the uniqueness theorem secures is narrower: *given* that the stability measure is a symmetric bilinear contact count vanishing on an empty side, the form is *forced* to $|T| \cdot |T^c|$. That is a form-fixing convenience (it lets us write K_bal = 4(9 − b) and identify "largest" with "most balanced"), **not** the source of the k = 3 selection. The honest summary is therefore: the k = 3 selection is robust to functional form, and bilinearity-uniqueness merely pins which representative of that robust selection we use. The bilinearity-and-contact-count premise itself is unmotivated input (QM-S3); the §6.1 argument should not be read as deriving k = 3 from it, since k = 3 needs far less.

---

7. Occupied/Exterior Orientation and Complement Breaking

The central shell creates a hazard: since a triad and its complement have equal size, one might quotient T with T^c and obtain ten unordered bipartitions. The strange trace requires twenty, so the operator must record occupied/exterior orientation.

The distinction is between a **partition** and a **support**. An unordered partition says only $I = A \sqcup B$ with $|A| = |B| = 3$ and no physical difference between A and B. An occupied support says $I = T \sqcup T^c$ with T occupied, T^c exterior. The complement map reverses the two roles; it does not rename one state. At the stabiliser level the distinction is exact: internal relabellings inside T and inside T^c are redundant,

$G\_T = \mathrm{Sym}(T) \times \mathrm{Sym}(T^c) \cong \mathrm{Sym}(3) \times \mathrm{Sym}(3)$,

but exchanging the blocks is **not** redundant, so the stabiliser is *not* extended to the wreath product $(\mathrm{Sym}(3) \times \mathrm{Sym}(3)) \rtimes \mathbb{Z}_2 = \mathrm{Sym}(3) \wr \mathrm{Sym}(2)$. The coset counts are

$|\mathrm{Sym}(6)| / (|\mathrm{Sym}(3)|\cdot|\mathrm{Sym}(3)|) = 720/36 =$ **20** (oriented support module), $|\mathrm{Sym}(6)| / (2\cdot|\mathrm{Sym}(3)|\cdot|\mathrm{Sym}(3)|) =$ **10** (unordered-bipartition module).

Inducing the trivial off Sym(3) × Sym(3) gives the 20; inducing off the wreath product gives the 10. Complement breaking is exactly the choice of the smaller subgroup — and because the coset space already distinguishes occupied from exterior, **no extra premise is needed beyond *not* extending by the swap.**

The witness. To make "not quotiented" a property of the readout rather than an assertion, introduce a role observable J on boundary records, acting within each complementary pair $\{T, T^c\}$ by

$J|\text{occupied } T; \text{exterior } T^c\rangle = +|\text{occupied } T; \text{exterior } T^c\rangle$, $J|\text{occupied } T^c; \text{exterior } T\rangle = -|\text{occupied } T^c; \text{exterior } T\rangle$.

J is the involution $\mathcal{C}$ restricted to $\mathbb{C}[\mathcal{S}_3(I)]$ in its −1 eigenbasis; that −1 eigenspace is the **complement-antisymmetric sector, dimension 10** (the +1 symmetric sector is the 10-dimensional unordered-bipartition module). The existence of this complement-odd observable certifies that $\mathcal{C}$ is a physical inequivalence, not a gauge redundancy: if the readout factored through the symmetric sector the trace would collapse to 10. The mass observable need not be J — it may remain the identity on the full 20 — but J's existence is the diagnostic. This is the same complement-odd probe the current-mass paper named as the deliverable of QM-S4; here it is exhibited explicitly.

Complement breaking (operational): the confinement boundary record carries a complement-odd role observable, so the readout space keeps both orientations and the support module is the full 20, not the symmetric 10.

7.1 Explicit form of the role selector. The role term K_{role} in the selector K_s (§9) should be read as a penalty on *unrecorded or quotient-level* boundary states — not as a penalty on either oriented state. The confinement record must contain a two-valued **role register**: for each complementary pair,

$$R_T = \mathrm{span}\{\, |T_{\text{occ}}, T^c_{\text{ext}}\rangle, |T^c_{\text{occ}}, T_{\text{ext}}\rangle \,\},$$

on which the complement involution $\mathcal{C}$ swaps the two basis states and the role observable distinguishes them, $J_T|T\text{occ}, T^c_{\text{ext}}\rangle = +|\cdot\rangle$, $JT|T^c_{\text{occ}}, T_{\text{ext}}\rangle = -|\cdot\rangle$. These satisfy

$J_T \cdot \mathcal{C} = -\mathcal{C} \cdot J_T$ (verified: in the 2-state register, $J = \mathrm{diag}(+1, -1)$, $\mathcal{C}$ = swap, and they anticommute).

The anticommutation expresses that J and $\mathcal{C}$ act *incompatibly* on the register — J distinguishes exactly the two states $\mathcal{C}$ exchanges. This is **consistent with** complement exchange being a physical inequivalence, but it does not by itself **establish** it: $\{J, \mathcal{C}\} = 0$ holds for any diagonal observable and swap on a two-state register, including two states one had artificially sign-labelled that were in fact equivalent. The physical inequivalence is imported separately, by the postulate that the role register is real (encoded in K_{role} / P_{unrec} below); the anticommutation is its diagnostic signature, not its proof. (This is the same point made in the scope note at the end of the section, stated here so the gloss is not over-read.) A precise role selector is therefore

$$K_{\text{role}} = \rho \cdot P_{\text{unrec}}, \quad \rho > 0,$$

where P_{unrec} projects onto boundary descriptions that **lack a role register** or factor through the unordered bipartition $\{T, T^c\}$. Its kernel is the oriented register R_T, not a single complement-symmetric line, so in $\ker K_{\text{role}}$ both orientations survive and the readout retains the full oriented module $\mathbb{C}[\,\mathrm{Sym}(6)/(\mathrm{Sym}(3)\times\mathrm{Sym}(3))\,]$ rather than the bipartition module induced from $\mathrm{Sym}(3)\wr\mathrm{Sym}(2)$. Equivalently: **quotienting by complement is the operation that deletes J_T; strange confinement does not delete it.** The role information carried by the boundary record is exactly what prevents the collapse from twenty oriented supports to ten complement pairs.

A scope note on what J does and does not do. J **certifies consistency and detectability**, not necessity: it shows the oriented (dim-20) readout is internally consistent and that the complement inequivalence is a physical observable *if* the role register is physical. It does not, by itself, *force* orientation — its very existence is a consequence of having taken the oriented kernel, so a reader should not over-read J as a derivation that confinement must break complement. Nothing in the boundary record compels non-quotienting except the assumption that the role register is physical (encoded in K_{role}, §7.1). J characterises the difference between the two candidate readouts and makes complement-breaking falsifiable; the question of whether confinement *must* carry the role register remains QM-S4.

---

8. Support Atoms as Canonical Boundary Invariants (the a_0 = 1 Derivation)

This is the section that most changes the status of the previous paper. There, the self-weight $a_0 = 1$ was the single remaining scalar premise — "let each support atom contribute one trace unit." Here it is derived as the dimension of a canonically determined invariant line.

The wrong counts to be excluded are concrete:

- one unit per occupied channel across twenty triads: $3 \times 20 =$ **60**;
- one unit per occupied/exterior contact: $9 \times 20 =$ **180**;
- one unit per ordered occupied triad: $6 \cdot 5 \cdot 4 =$ **120**.

Each treats internal boundary structure as separate mass sectors. The support-readout claim is that the nine contacts of a single 3+3 split are *internal structure of one closure*, not nine closures.

The canonical coarse-graining (projector lemma). The self-weight $a_0 = 1$ can be stated as a projector theorem.

Lemma (unit self-weight). Fix an occupied triad T. Its boundary-incidence space is $\mathcal{B}_T = \mathbb{C}[\partial T] \cong \mathbb{C}^3 \otimes \mathbb{C}^3$, the first factor labelling occupied channels, the second exterior channels, with stabiliser $G_T = \mathrm{Sym}(T) \times \mathrm{Sym}(T^c) \cong \mathrm{Sym}(3) \times \mathrm{Sym}(3)$. Let $\Pi_T = (1/|G_T|) \Sigma_{g \in G_T} U_g$ be the group-average projector onto the G_T-invariant subspace. Then

$\mathrm{rank}\ \Pi_T = 1$, $\mathrm{Tr}\ \Pi_T = 1$,

with unique invariant vector $|b_T\rangle = (1/3) \Sigma_{i \in T} \Sigma_{j \in T^c} |i \to j; T\rangle$.

Proof. The invariants factorise, $(\mathbb{C}^3 \otimes \mathbb{C}^3)^{\mathrm{Sym}(3)\times\mathrm{Sym}(3)} = (\mathbb{C}^3)^{\mathrm{Sym}(3)} \otimes (\mathbb{C}^3)^{\mathrm{Sym}(3)}$, and each $(\mathbb{C}^3)^{\mathrm{Sym}(3)}$ is the one-dimensional all-channels average, so the invariant subspace is $\mathbb{C} \otimes \mathbb{C} \cong \mathbb{C}$, dimension 1. Hence Π_T has rank 1 and trace 1; the uniform vector $|b_T\rangle$ (unit-norm: $(1/3)^2 \cdot 9 = 1$) is its image. (Verified directly: the group-average projector on the 9-dimensional module has rank 1, trace 1, and fixes $|b_T\rangle$.) ■

Define the support atom as this invariant, $P_T = |b_T\rangle\langle b_T| = \Pi_T$, with $\mathcal{R}\, P_T\, \mathcal{R}^* = |T\rangle\langle T|$. Each stable boundary then contributes exactly one trace unit — not by separate stipulation of a weight, but because the space of G_T-invariant boundary modes is one-dimensional, *given that confinement reads the invariant line*. The nine incidences are not nine current-mass atoms; they are the internal contact structure of a single invariant boundary mode, and the incidence excitations orthogonal to $|b_T\rangle$ are real internal modes that K_{inc} gaps out of the low-readout trace. Summing over the twenty triads,

$\mathrm{Tr}\ O_s = \Sigma_{T \in \mathcal{S}_3(I)} \mathrm{Tr}\ P_T = |\mathcal{S}_3(I)| = 20.$

So $a_0 = 1$ is no longer a free trace-normalisation: it is the rank of the canonical invariant projector associated with each stabilised 3+3 boundary. The premise has not vanished — it has moved from "let $a_0 = 1$" to "confinement coarse-grains onto the canonical invariant" — but the relocated premise is structurally forced rather than numerically chosen, which is the advance. The residual freedom (a coherence selector that kept a higher-dimensional G_*T*-subrepresentation would give $a_0 \neq 1$) is recorded honestly in falsifier 14.7.

Why this is the right a_0. Equivalently: the coherence selector K_inc gaps every non-invariant boundary excitation (the 8-dimensional orthogonal complement of $|b_T\rangle$ inside each $\mathbb{C}^3 \otimes \mathbb{C}^3$), leaving only the invariant line. The self-weight a_0 is then *forced* to be the invariant dimension, which is 1. The previous paper's premise "$a_0 = 1$" is thus replaced by the structural fact "the boundary-incidence module of a balanced support has a one-dimensional stabiliser invariant." An incidence-counting readout (a_0 tied to the 9-dimensional module) gives 180; the canonical-invariant readout gives 20. The choice between them is now a sharp operator question (QM-S5), and the canonical invariant is the natural answer rather than an imposed one.

---

9. The Candidate Strange Confinement Operator

Assemble the pieces. Let $\mathscr{H}$_inc be the boundary-incidence refinement of $\mathscr{H}$_shell (each support resolved into its contact module $\mathbb{C}^3 \otimes \mathbb{C}^3$). The strange readout is produced by the chain

$\mathscr{H}$_inc —(P_3)→ incidence sectors with $|T| = 3$ —(K_inc, $\mathscr{R}$)→ $\mathbb{C}[\mathcal{S}_3(I)]$ —(role record)→ $\mathscr{H}$_s,

with the dimension bookkeeping $20 \cdot 9 = 180$ →(cohere each support to its invariant)→ 20, and orientation kept automatically by using G_*T* = Sym(3)×Sym(3) rather than its wreath extension.

A compact candidate selector is

K_s = $\boldsymbol{\alpha \cdot (N - 3)^2 + \beta}$·K_role + $\boldsymbol{\gamma}$·K_inc, $\alpha, \beta, \gamma > 0$,

with:

1. $\boldsymbol{\alpha \cdot (N - 3)^2}$ selects the balanced central shell (kernel on partial supports: $|T| = 3$), equivalently the maximal-capacity shell via K_bal = $4(9 - B)$.
2. **K_role = ρ·P_unrec** ($\rho > 0$), where P_unrec projects onto boundary descriptions that lack a role register or factor through the unordered bipartition $\{T, T^c\}$ (§7.1); it penalizes unrecorded or quotient-level states, preventing complement quotienting, and its kernel is the oriented register R_*T* keeping both orientations (so the complement-odd observable J of §7 survives).
3. **K_inc** gaps incidence-level excitations orthogonal to the canonical invariant $|b_T\rangle$ inside each support: explicitly, within each *T*-sector,

   K_inc,*T* = $\mu \cdot (I\_T\text{\^{}inc} - |b_T\rangle\langle b_T|)$, $\mu > 0$, so only the one-dimensional invariant survives per support.

The intended low-readout space is the joint kernel,

$\ker K_s \cong \mathbb{C}[\mathcal{S}_3(I)]$, dim 20,

where it should be read that the three terms are assembled on the common incidence space $\mathscr{H}_{inc}$ (the role registers and the shell number N both lift to $\mathscr{H}_{inc}$ via the refinement of §8), and each term is positive semidefinite, so the joint kernel is the intersection $\ker K_s = \ker((N-3)^2) \cap \ker K_{role} \cap \ker K_{inc}$. The "$\ker K_s = \mathbb{C}[\mathcal{S}_3(I)]$" identity is therefore literal on $\mathscr{H}_{inc}$ after the §9 chain of identifications, not merely schematic.

and the induced mass observable is the identity on it,

$O_s = \mathscr{R} \left(\Sigma_{T\in\mathcal{S}_3(I)} |b_T\rangle\langle b_T| \right) \mathscr{R}^* = \Sigma_{T\in\mathcal{S}_3(I)} |T\rangle\langle T| = I_{\mathscr{H}_s}$.

In words: strange confinement first selects a balanced shell, then orients the boundary, then collapses each support's internal incidences onto its canonical invariant mode. Each of the three terms in K_s corresponds to one previously-open premise, and each acts in a manifestly Sym(6)-equivariant way, so the induced observable lies in the Johnson-scheme commutant of J(6,3) and inherits the constant-diagonal property automatically.

---

10. Trace Theorem for the Strange Confinement Readout

Theorem. Let I be a six-element primitive interface-channel set with G = Sym(I). Assume:

1. The down pre-boundary operator has ground readout $\mathscr{H}_d = \mathbb{C}[\mathcal{S}_1(I)]^G$ (it gaps $S^{(5,1)}$, keeps $S^{(6)}$).
2. Partial supports are selected by the Sym(6)-invariant boundary capacity $b(T) = |T|\cdot|T^c|$, whose unique maximum among nontrivial partial supports is $|T| = 3$.
3. The boundary record is occupied/exterior oriented, so T is not identified with T^c.
4. Internal permutations inside T and inside T^c are quotiented.
5. Each support's boundary incidences are cohered onto the canonical Sym(3)×Sym(3)-invariant line, giving one rank-one support atom of trace 1.
6. The induced current-mass observable is Sym(6)-equivariant on the selected shell.

Then the strange confinement readout space is $\mathscr{H}_s \cong \mathbb{C}[\,\mathrm{Sym}(6)/(\mathrm{Sym}(3)\times\mathrm{Sym}(3))\,] \cong \mathbb{C}[\mathcal{S}_3(I)]$, and

$\mathrm{Tr}_{\mathscr{H}_s} O_s = 20$, $\mathrm{Tr}_{\mathscr{H}_d} O_d = 1$, $\mathrm{Tr}\, O_s / \mathrm{Tr}\, O_d = 20$.

Proof. Assumption 1 gives a one-dimensional invariant, so $\mathrm{Tr}\, O_d = 1$. For strange, $b(k) = k(6 - k)$ takes values 0, 5, 8, 9, 8, 5, 0, so the unique nontrivial maximum is k = 3 (assumption 2). Assumption 3 excludes complement-exchange quotienting; assumption 4 quotients only internal relabellings, so the stabiliser of a chosen occupied triad is Sym(3) × Sym(3) and the number of support atoms is |Sym(6)|/(|Sym(3)|·|Sym(3)|) = 720/36 = 20. By assumption 5 each support's boundary module $\mathbb{C}^3 \otimes \mathbb{C}^3$ has one-dimensional stabiliser invariant, contributing trace 1 per support; by assumption 6 the observable lies in the Johnson-scheme commutant, where the diagonal is constant, so the self-weight a_0 is common across supports. Hence

$$\mathrm{Tr}\, O\_s = \Sigma\_\{T\in\mathcal{S}_3(I)\}\ \mathrm{Tr}\, P\_\mathcal{T} = |\mathcal{S}_3(I)| = 20.$$

Dividing by Tr O_d = 1 gives 20. ■

Remark (where a_0 = 1 lives now). In Johnson-scheme language any Sym(6)-equivariant observable on $\mathbb{C}[\mathcal{S}_3(I)]$ has trace $20\cdot a_0$, independent of inter-support couplings (the off-diagonal adjacency operators A_1, A_2, A_3 carry zero diagonal). Assumption 5 is the confinement derivation of $a_0 = 1$: the self-weight equals the dimension of the canonical boundary invariant, which is one. The previous paper isolated $a_0 = 1$ as the single open scalar; this paper proposes that scalar is the dimension of a uniquely determined invariant line, hence forced rather than fitted.

---

11. Relation to the Current-Mass Operator Paper

The previous strange-current paper proved the conditional statement:

if $\mathscr{H}$_d = $\mathbb{C}[\mathcal{S}_1(I)]$^Sym(6), $\mathscr{H}$_s = $\mathbb{C}$[Sym(6)/(Sym(3)×Sym(3))], and a_0 = 1, then m_s/m_d = 20.

This paper attempts to justify the if-clauses. The mapping is one-to-one:

Previous open problem	Attempted answer here
QM-S1 derive invariant down readout	pre-boundary K_d gaps S^(5,1), leaves S^(6) (§4)
QM-S2 derive support-boundary formation	support sectors appear only when an occupied/exterior boundary is recorded (§5)
QM-S3 derive central shell	boundary capacity b(k) = k(6−k), max at k = 3; affine to imbalance penalty (§6)
QM-S4 derive complement breaking	role-sensitive record + complement-odd witness J in the dim-10 antisymmetric sector (§7)
QM-S5 derive a_0 = 1	support atom = canonical 1-dim Sym(3)×Sym(3)-invariant of its $\mathbb{C}^3\otimes\mathbb{C}^3$ boundary module (§8)
QM-S7 connect to confinement	K_s = balance + role + incidence-coherence selector (§9)

What advances is that the support module is no longer introduced only as a representation-theoretic possibility — it is now the low-readout space of a candidate confinement selector, and the formerly-scalar premise $a_0 = 1$ is now a one-dimensional invariant count. The one inherited problem this paper deliberately does **not** touch is **QM-S6**, the five-gate trace-factorisation problem, treated next.

---

12. Relation to the Five-Gate Problem

This paper does not solve the five-gate conflict, and it does not pretend the two pictures agree numerically. The support-trace picture is **additive**, $\mathrm{Tr}\, O_s = \Sigma_T 1 = 20$. The five-gate picture is **multiplicative**,

$$m_s/m_d \approx (G_2/G_1)\cdot \exp(M_{\chi,2} - M_{\chi,1}) \approx 108 \times 0.18 = 19.44.$$

These are not automatically the same kind of object: trace is basis-independent but not multiplicative ($\mathrm{Tr}(AB) \neq \mathrm{Tr}\, A \cdot \mathrm{Tr}\, B$ in general), so an additive count of twenty unit atoms does not become a product of two large factors by a change of coordinates.

12.1 The dichotomy, foregrounded. The bite is sharper than a mismatch of form, and it should lead rather than trail: if 108 and 0.18 are exact structural outputs and 20 is an exact trace, then **at most one of them can be the true current-mass operator** — a single observable cannot carry two exact theoretical values ($20.00 \neq 19.44$). QM-S6 is therefore not a translation problem but a **which-one-is-primitive problem**, and its resolution must *downgrade one of the two "exact" derivations to an effective (approximate) reading*. The four options below are not a menu of equally-likely reconciliations; they are the ways that downgrade can fall out: (i) the five-gate product is an approximate phenomenological decomposition of the exact support trace; (ii) the five-gate product is a different exact operator and the support count is the effective one; (iii) the five-gate product is a downstream effective reading not primitive for current mass; (iv) neither is final. The honest position of this paper is that it supplies a candidate for the *support-trace* operator and takes no stance on which side survives — but it insists the resolution is a primitivity verdict, not a coordinate change.

12.2 Empirical status: a band, not a single anchor. The earlier draft pinned the comparison to one FNAL/MILC pair, $92.47/4.675 = 19.78$. That overstates the precision of the target. The quantity the lattice community actually agrees on at scheme-and-scale-independent level is the ratio $m_s/m_{ud} = 27.23(10)$ (FLAG 2024, $N_f = 2+1+1$); converting to m_s/m_d requires m_u/m_d, which carries the large electromagnetic/isospin uncertainty, $m_u/m_d = 0.47(4)$. Folding,

$$m_s/m_d = (m_s/m_{ud})\cdot(1 + m_u/m_d)/2 = 27.23\cdot(1.47)/2 = \mathbf{20.0(5)},$$

with the ± 0.5 dominated almost entirely by m_u/m_d, giving a 1σ band $\approx$ **19.5 – 20.6**. (Individual determinations scatter consistently across this band: e.g. FLAG $m_{ud} = 3.41$ MeV with $m_u/m_d = 0.47$ gives $m_d \approx 4.64$ MeV and $m_s/m_d \approx 20.0$; other full-lattice sets reach ≈ 20.8.) Against this band the two VERSF readings sit as:

reading	value position in 20.0(5) band
support trace (this programme)	20.00 −0.03σ (at this input)
five-gate residue	19.44 −1.2σ (at this input)

These positions are quoted at the broad-FLAG input m_u/m_d = 0.47(4) and are **not** stable — see the sensitivity table below, which is the actual point.

(For reference, not as a third same-type entry: folding the lattice m_s/m_ud against VERSF's *own* light-ratio input m_u/m_d = 6/13 gives m_s/m_d = 27.23·(19/26) = 19.90. This is a **hybrid** — lattice ratio times a VERSF input — and a consistency check on the 6/13 fold, not an independent VERSF operator prediction of m_s/m_d. It is deliberately kept out of the table so that the only same-type entries are the two genuine VERSF operators; including it would misframe §12.1's "at most one of two can be exact" as "at most one of three.")

Two consequences. First, **neither VERSF reading is empirically excluded** — the band is wide enough to admit both 19.44 and 20.00, so §12.1's dichotomy cannot currently be settled by data; it must be settled by deciding which operator is primitive (QM-S6). Second — and this corrects a temptation the band itself invites — **which reading sits nearer the centroid is not a discriminator, because the centroid is set by the chosen m_u/m_d, and that input is exactly the poorly-determined one.** Across its defensible range the centroid moves straight through both VERSF values:

m_u/m_d input	centroid m_s/m_d	support 20.00	five-gate 19.44
0.47 (broad FLAG average)	20.01	−0.03σ	−1.2σ
0.46	19.88	+0.24σ	−0.9σ
0.455	19.81	+0.38σ	−0.7σ
0.45	19.74	+0.52σ	−0.6σ

At the top of the range the support value sits at the centre and the five-gate value at the edge; a little lower the two become roughly symmetric about the centroid and the apparent preference evaporates. Picking the *loosest* companion input (0.47(4)) is what places the support value at the centre, and that is precisely the kind of single-input flattery §12.2 opened by retiring for the old 19.78 anchor. So the honest reading is stronger for being weaker: the band admits both VERSF values, and **which one looks "central" is an artifact of the m_u/m_d choice across the 0.45–0.47 range — not even a weak empirical discriminator.** The right standard remains structural: **does the closure operator independently force the support module whose trace is 20?** If yes, the agreement (anywhere in the band) becomes meaningful; if no, it stays a striking but underived coincidence, and falsifier 14.10 records that the value cannot presently be adjudicated.

---

13. What Is Proved, Conditional, and Still Open

13.1 Proven here, given the stated operator. If the confinement selector is Sym(6)-invariant, selects the maximal-capacity shell, keeps occupied/exterior orientation, and coheres each support's boundary onto its canonical invariant, then the strange readout space is twenty-dimensional and Tr O_s = 20. The shell values $b(k) = 0,5,8,9,8,5,0$, the affine identity $K\_bal = 4(9 - b)$, the coset count 720/36 = 20, the complement-antisymmetric dimension 10, and the one-dimensional boundary invariant are all exact.

13.2 New structural content. Three genuine advances over the current-mass paper, now in formal rather than verbal form: (i) the central shell k = 3 is selected robustly — *any* increasing symmetric balance measure peaks there — and within the bilinear-contact-count class the stability functional is *forced* to $|T|\cdot|T^c|$ (§6.1), so "largest" and "most balanced" coincide while the selection itself does not depend on that form; (ii) complement breaking is exhibited with an explicit role register R_T and a complement-odd observable satisfying $J\cdot\mathcal{C} = -\mathcal{C}\cdot J$ (§7.1), which is the *diagnostic signature* of orientation (the inequivalence itself is the K_role postulate, not the anticommutation); (iii) the self-weight $a_0 = 1$ is a projector lemma — the rank of the canonical Sym(3)×Sym(3)-invariant of each support's $\mathbb{C}^3\otimes\mathbb{C}^3$ boundary module (§8) — a relocated rather than imposed premise.

13.3 Conditional premises. The single load-bearing premise is the **down/strange exemption** (§4.1): confinement must apply the capacity+coherence mechanism to the strange sector but *not* to the down sector, even though the down singleton carries nonzero boundary capacity $b(1) = 5$. Two facts fix the size of this gap: the capacity functional is monotone on [0, 3], so it contains no internal threshold that could exempt k = 1 (the threshold must come from physics orthogonal to b); and down and strange are selected by operators of *opposite* logic — K_d rewards non-resolution (kernel = invariant line), K_s rewards maximal resolution (kernel = full support module) — so the ratio 20 is the quotient of two inverse verdicts with no single principle yet yielding both. This is the threshold that fixes the ratio at 20 rather than 10/3, and it is a missing unification, not a missing lemma (QM-S1/QM-S2). Beyond it, the candidate assumes VERSF confinement geometry actually supplies the boundary-capacity functional as its stability criterion, treats occupied/exterior role as physical, coheres incidences via the canonical invariant, and preserves Sym(6)-equivariance on the selected shell. These are now spectral conditions on K_s, not free choices, but they are not yet read off a field action.

13.4 Still open. The paper does not derive the six primitive channels, the full hadronic confinement field action, the absolute mass scale, the five-gate trace-factorisation (QM-S6), CKM mixing, or the heavy-quark tensions. The honest status:

candidate confinement-operator derivation of the strange support trace, not a completed first-principles derivation of quark masses.

That status is useful precisely because it makes the next needed derivation — a field-theoretic K_s, and the QM-S6 bridge — much sharper.

---

14. Falsification Conditions

As in the current-mass paper, the conditions fall into three epistemic tiers, and listing them flat would flatter the falsifiability:

- **Tier I — mechanism falsifiers (14.1–14.8).** "If the confinement operator does X, the trace is the wrong number Y." Decidable once the explicit operator is known.
- **Tier II — intra-theory basis competition (14.9).** Not a falsification of VERSF or of the value 20, but of *this account within* VERSF; decidable only by building QM-S6.
- **Tier III — external empirical rejection (14.10).** The only tier that could falsify the value itself.

14.1 Full down multiplicity. If K_d does not gap S^(5,1), the down denominator is 6 and the ratio is 20/6 = 10/3.

14.2 Uniform orbit readout. If confinement treats both down and strange as orbit quotients, both readouts are one-dimensional and r_s = 1.

14.3 Noncentral shell. If the selector lands on k = 2 or 4, the count is 15; on k = 1 or 5, it is 6; on k = 0 or 6, there is no partial boundary.

14.4 Complement quotient. If occupied/exterior role is not physical and $T \sim T^c$, the shell collapses to the symmetric sector and the trace is 10. (Equivalently: if no complement-odd observable J exists, §7's witness is absent.)

14.5 Ordered path readout. If the operator counts ordered triads, the count is 6·5·4 = 120.

14.6 Channel-incidence readout. If current mass is additive over occupied channels, 3 × 20 = 60; if additive over boundary contacts, 9 × 20 = 180. Either contradicts the canonical-invariant support atom of §8.

14.7 Wrong invariant dimension. If the boundary module's relevant invariant is not one-dimensional — e.g. the coherence selector keeps a higher-dimensional G_*T*-subrepresentation — then $a_0 \neq 1$ and the trace is $20 \cdot a_0$. The robustness is that, *given* the canonical invariant, a_0 is forced to 1.

14.8 Broken Sym(6) equivariance. If confinement favours particular channels, the mass observable leaves the Johnson-scheme algebra, the diagonal need not be constant, and the trace need not be 20.

14.9 Five-gate eviction. If the five-gate decomposition is the primitive exact current-mass operator and the support trace has no exact operator home, this support-confinement account is evicted even though m_s/m_d ≈ 20 remains empirically true.

14.10 Empirical rejection. The current scheme-independent record gives m_s/m_d = 20.0(5) (band ≈ 19.5–20.6, dominated by m_u/m_d), which admits both VERSF readings, so the value cannot presently be adjudicated against data (§12.2). If future determinations — with tightened electromagnetic and isospin budgets, principally a sharper m_u/m_d — narrow the band and move its centroid decisively off 20, the value is phenomenologically rejected and the mechanism is wrong however elegant. Equally, a band tight enough to separate 20.00 from 19.44 would, for the first time, let data weigh in on the §12.1 primitivity dichotomy.

---

15. Conclusion

The strange current-mass paper reduced the successful ratio $m_s/m_d \approx 20$ to a sharp representation-level claim — down reads one invariant baseline, strange reads twenty balanced occupied-support atoms, each atom contributing one trace unit — and isolated four premises (balanced shell, complement breaking, support-not-incidence, $a_0 = 1$) as the price of that claim. This paper proposes a confinement selector that supplies all four and, in two cases, upgrades a premise to a derived fact.

The selector has a clear structure. Before a boundary forms, resolved labels are gapped and down stays the invariant all-channel baseline. Once a boundary is recorded, partial supports become physical. The central shell is selected because the boundary capacity $|T|\cdot|T^c|$ is maximal at the balanced 3+3 split — and because capacity and imbalance are affinely the same functional, "largest" and "most balanced" are one criterion. Complement is not quotiented because occupied and exterior are role-distinct boundary conditions, certified by a complement-odd observable in the dimension-10 antisymmetric sector. Incidences are not separate mass sectors because the nine contacts of a 3+3 split cohere onto a **canonically one-dimensional** invariant, which is what forces the self-weight $a_0 = 1$. The readout space is

$$\mathscr{H}\_s = \mathbb{C}[\ \mathrm{Sym}(6)/(\mathrm{Sym}(3)\times\mathrm{Sym}(3))\ ] \cong \mathbb{C}[\mathcal{S}_3(I)],\quad O\_s = \Sigma\_\{T\in\mathcal{S}_3(I)\}\ |T\rangle\langle T|,\quad \mathrm{Tr}\ O\_s = 20,$$

and with the invariant down readout Tr O_d = 1, the conditional ratio is 20.

The claim is bounded but meaningful. It does not prove the full VERSF quark-mass hierarchy, does not solve the five-gate tension (the support count 20.00 and the five-gate residue 19.44 cannot both be exact values of one observable, and which is primitive is unresolved — §12.1), and does not rescue the charm or bottom formulae. On the empirical side the honest statement is a band, $m_s/m_d = 20.0(5)$: both VERSF readings lie inside it, and which one looks central is an artifact of the poorly-determined m_u/m_d input, so data cannot yet adjudicate in either direction. What this paper does is give the first explicit candidate for the confinement operator whose low-readout space is the strange balanced support module, with the self-weight no longer a free scalar but the dimension of a forced invariant — and with the one genuinely load-bearing premise, the down/strange exemption (§4.1), named rather than hidden. If later VERSF confinement geometry derives a K_s of this shape *and* the pre-boundary threshold that exempts down, the strange trace becomes a secured structural result; if it derives a different operator, the mechanism fails cleanly with a definite wrong number — 10, 60, 120, 180, 15, 6, 10/3, or a broken diagonal. That is the right standard for this stage: not just a successful number, but a number attached to a confinement operator that future work can force, modify, or reject.

References

VERSF programme

- Keith Taylor, *Quark Masses from One Anchor in VERSF.* One-anchor calculator and quark-mass grid.
- Keith Taylor, *The Strange Baseline in VERSF: Six-Channel Triad Resolution and the Twentyfold Down-to-Strange Ratio.* Triad-shell proposal and mixed multiplicity obstruction.
- Keith Taylor, *Quark Masses from One Anchor II: The Current-Mass Operator and the Trace Audit of the Diagonal Quark Spectrum.* Operator-trace organisation and empirical grading of open traces.
- Keith Taylor, *The Strange Current-Mass Operator in VERSF: Resolved Balanced Support Readout, Complement Breaking, and the Twentyfold Down-to-Strange Ratio.* Representation-level strange support trace; Johnson-scheme commutant; QM-S1…S7 open problems (this paper closes the operator-construction side of S1–S5, S7).
- Keith Taylor, *From Charge-Sector Maintenance to Quark Masses in VERSF.* Five-gate decomposition and primitive-versus-composite tension for m_s/m_d.
- Keith Taylor, *Entropic Confinement and Effective Quark Masses in the VERSF Framework.* Quarks as confined partial closures; current versus constituent mass.
- Keith Taylor, *Finite Distinguishability and Local Capacity Competition* and related interface/capacity papers. Six primitive local interface channels and uniform primitive allocation.

Mathematical structures used

- Symmetric-group permutation modules and induced modules; Young subgroups Sym(3)×Sym(3) ⊂ Sym(6) and the wreath product Sym(3)≀Sym(2).
- Johnson scheme J(6,3) and the Bose–Mesner commutant of the triad-support action.
- Finite shell spaces, complement involutions, tensor permutation modules $\mathbb{C}^3 \otimes \mathbb{C}^3$, and normalized rank-one projectors onto stabiliser invariants.