

Uniform Binary Admissibility and the Half-Lazy Closure Operator

A Structural Reduction of the χ -Halving Law in VERSF

Keith Taylor *VERSF Theoretical Physics Programme — Quark Mass Hierarchy Series*

Summary for the General Reader

In each generation there is an up-type quark and a down-type quark. Define $\chi(g) = \ln(m_{\text{up}}(g) / m_{\text{down}}(g))$. The observed pattern is that the change in χ from generation 1 to 2 is roughly twice the change from generation 2 to 3: $\Delta\chi_2 \approx \frac{1}{2} \cdot \Delta\chi_1$. The previous paper checked, with all six masses in one consistent convention, that this near-halving is real and not an artefact of mixed bookkeeping — it found $\rho = \Delta\chi_2 / \Delta\chi_1 = 0.50 \pm 0.02$. That confirmed $\frac{1}{2}$ as the right target but did not prove exact equality, and it returned the question to structure.

This paper asks whether VERSF can **derive** exact halving from its own principles. The honest answer it reaches is a precise *conditional reduction*: not "here is a proof that the quarks must halve," but "here is the complete list of structural facts that, taken together, force exact $\frac{1}{2}$ — and each is now a sharp, checkable target rather than a vague hope."

The intuition is simple. In VERSF, when a distinction is real but no admissible fact tells two alternatives apart, they enter symmetrically. Suppose the up/down refinement step is a genuine binary gate with exactly two histories: one in which the up/down difference survives unchanged, and one in which closure averages the pair and erases the difference. Counting two equal histories, each carries weight $\frac{1}{2}$, so after one step exactly half the up/down "difference content" remains — which is the halving law.

The paper is careful about two things that a looser argument would blur. First, **symmetry alone is not enough**. Up/down symmetry says the erasing branch must be a fair average (it cannot prefer up over down), but it does not say the two branches are *equally weighted*, nor that the average is *complete*. Those come from a stronger, separately stated rule — uniform admissibility — which the programme already uses elsewhere. Second, the result is not a single proof but a **reduction to six conditions**. The algebra connecting those conditions to $\frac{1}{2}$ is established and machine-checked; what remains is a structural audit of whether the physical up/down gate actually satisfies them. The two hardest to take for granted are named explicitly: that the gate really has exactly two history classes (and no third "flip" history), and that the two branches carry no hidden difference in cost. Even the fact that the readout is a *logarithm* is no longer assumed — it follows from the up/down mass ratio composing multiplicatively and inverting under up/down exchange.

So the standing is: $\frac{1}{2}$ is not fitted from quark masses, and not merely asserted as "symmetry." It is reduced to one inherited equal-weight principle plus a short list of χ -specific structural claims, which is exactly the right place to put the remaining work.

Abstract

The preceding single-scheme χ audit established that the same-generation up/down susceptibility profile survives consistent $\overline{\text{MS}}$ QCD running, with $\rho = \Delta\chi_2/\Delta\chi_1 = 0.50 \pm 0.02$ — consistent with the structurally targeted $\frac{1}{2}$ but not a measurement of exact equality. This paper addresses the structural question the audit left open: can VERSF derive $\Delta\chi(g+1) = \frac{1}{2} \cdot \Delta\chi(g)$ from admissibility and closure rather than fit it?

The result is a **complete conditional reduction**, not a complete proof. The χ -halving law is reduced to a nondegeneracy prerequisite plus six explicit, individually checkable conditions; the algebra carrying those conditions to exact $\frac{1}{2}$ is established and verified on a non-degenerate model. The construction is positioned against the programme's **Symmetric-Averaging Obstruction (§6L)**, which proves up/down symmetry alone cannot fix the averaging weight — symmetry pins the projection direction Π_e but leaves a free family $a\Pi_e + b\Pi_o$. The reduction does not so much *clear* §6L as *relocate* it: the choice of branch set $\{I, \Pi_e\}$ kills the even-sector freedom (the b direction), the census fixes the weight $a = \frac{1}{2}$, and the branch eigenvalues (survive = 1, close = 0) are a separate single-step-completeness-and-branch-identity postulate — three moves, all contingent on the uniform-admissibility principle, not one principle applied once.

The binary fold sector is built explicitly (Route B) from ordered closure: with primitive closure operations A, B, C , the two oriented composites $U = ABC$ and $D = CBA$ span a two-dimensional word-space $\mathcal{F} = \text{span}\{U, D\}$, with reversal involution J ($J(U) = D$), even/odd split $\mathcal{F}_e = \text{span}\{U+D\}$, $\mathcal{F}_o = \text{span}\{U-D\}$, and a nontrivial odd sector whenever U and D are linearly independent (the operator Gram determinant is nonzero; $K = U - D \neq 0$ alone is necessary but not sufficient). A single upstream computation — the rank of U — selects the route: rank-deficient closure admits the support realisation (Route A), invertible closure forces the word-span (Route B); the full Gate-1 audit for both is Appendix C. The half-lazy operator $W = \frac{1}{2}(I + \Pi_e)$ then has odd eigenvalue exactly $\frac{1}{2}$, and the log-readout — derived, not assumed, from multiplicative composition of access and reversal-inversion of the ratio — gives $\Delta\chi(g+1) = \frac{1}{2} \cdot \Delta\chi(g)$.

The exact $\frac{1}{2}$ is reduced to: (0) a genuine two-fold sector — U, D linearly independent (Gram $\det \neq 0$), with $\text{rank}(U)$ fixing the realisation route; (1) reversal-fixed external signatures; (2) image-even closure, $Q = \Pi_e$; (3) a quotient-binary one-step history $H_{\chi/\sim\text{op}} = \{[I], [\Pi_e]\}$ with no admissible flip primitive (the no-flip resting on local closure-order adjacency, $d(U, D) = 3$), the census taken at one consistent class-level granularity so each class is a single unit; (4) equal scalar weight per primitive class, inherited from the Born-rule equal-weight result; (5) no admissible branch-cost asymmetry between $[I]$ and $[\Pi_e]$; and (6) log-access intertwining. Conditions (2)/(3)/(5) are the structural origins of the failure knobs of §9 — three operator axes (weight w , survive-eigenvalue λ , close-eigenvalue μ ; adjacency β being a negative- μ origin). The

three axes map onto the single scalar ρ , so they are degenerate *in principle*: the audit bounds only their conjunction ($2(\rho - \frac{1}{2}) \approx 0.006$ at the ± 0.05 level, sign uninformative by the audit's own finding), and no attainable ρ precision separates them — that requires structural argument alone. The scalar equal-weight of (4) is relocated to the Born-rule corpus; the **binarity-at-consistent-granularity** claim (3) and the **branch-cost symmetry** claim (5) are χ -specific and are the primary remaining audit targets, together with the access-row identification that, under Route B, is nearly the same task as (3). An earlier draft billed an "equal stabilizer order" sub-clause of (3) as sharpest; that is withdrawn — it is a phantom of inconsistent census granularity (collapsing the Q^n tower while keeping the J-orbit), and under the operative class-level census no orbit/stabilizer factor arises.

Status: Complete conditional reduction. The algebra is established and model-verified; it does **not** prove the physical quark gate realizes exact halving. It proves that if the χ gate satisfies conditions (1)–(6), then $\rho = \frac{1}{2}$ exactly. The remaining work is the χ -specific structural audit of those conditions — above all quotient-binarity, branch-cost symmetry, and access-row identification.

Table of Contents

1. Motivation and Inheritance from the χ Audit
 2. The Structural Question and the Symmetric-Averaging Obstruction
 3. The Binary Fold Sector from Ordered Closure
 4. Same External Neighbourhood and Nonadjacency
 5. Image-Even Closure: $Q = \Pi_e$
 6. Uniform Admissibility and the Exact Weight
 7. Log-Access Intertwining, and Why the Readout Is Logarithmic
 8. The χ -Halving Theorem
 9. Failure Modes — the Knobs and Their Structural Origins
 10. The Twin-Antichain Compression of Conditions (1), (2), (6)
 11. Relation to the Empirical Audit
 12. What This Paper Does and Does Not Establish
 13. Status Ledger
 14. Conclusion
- Appendix A — Minimal Algebraic Skeleton
 - Appendix B — Referee Objections and Replies
 - Appendix C — G-CHI-2 Gate Audit: Gate 1 Specification

1. Motivation and Inheritance from the χ Audit

The VERSF quark hierarchy programme separates the mass structure into a same-charge generation ladder $B(g)$ and a same-generation up/down susceptibility profile measured by

$$\chi(g) = \ln(m_{\text{up}}(g) / m_{\text{down}}(g)), \chi_1 = \ln(m_u/m_d), \chi_2 = \ln(m_c/m_s), \chi_3 = \ln(m_t/m_b),$$

with increments $\Delta\chi_1 = \chi_2 - \chi_1$, $\Delta\chi_2 = \chi_3 - \chi_2$ and profile ratio $\rho = \Delta\chi_2/\Delta\chi_1$.

The preceding single-scheme audit found, under consistent four-loop $\overline{\text{MS}}$ running and threshold matching, $\rho = 0.50 \pm 0.02$ — consistent with the exact-halving target but not a measurement of it; the 0.503-vs-0.500 difference sits below present input precision, dominated by m_u/m_d . The audit's role was narrow and necessary: it showed $\frac{1}{2}$ is not a mixed-scheme artefact, and it returned the burden to structure.

The structural question is whether VERSF can derive $\Delta\chi(g+1) = \frac{1}{2} \cdot \Delta\chi(g)$ from admissibility and closure. This paper answers it as a reduction: it identifies the complete set of structural conditions that force exact $\frac{1}{2}$ and shows the algebra is sound, leaving a sharp, χ -specific audit as the remaining work. Two scope notes carried from the audit hold throughout. The three-generation profile supplies only two increments and hence a single ratio ρ ; the theorem fixes one halving relation, and "profile" should not be read as implying a long geometric sequence.

2. The Structural Question and the Symmetric-Averaging Obstruction

The target law is $\Delta\chi(g+1) = \frac{1}{2} \cdot \Delta\chi(g)$, equivalently $\rho = \frac{1}{2}$. The required operator, acting on the log-access shell, is

$$W = \frac{1}{2}(I + \Pi_e),$$

where I is the survival branch (preserving the up/down odd residue) and Π_e is the closure-evenisation branch (erasing it). If this operator is justified, halving follows immediately because the odd readout sees only the surviving half.

The reason the paper exists is a named obstruction.

The Symmetric-Averaging Obstruction (§6L). Up/down exchange symmetry does *not* fix the averaging weight. Any J -equivariant census is of the form $a \cdot \Pi_e + b \cdot \Pi_o$; symmetry constrains the *direction* of averaging (it forbids preferring up over down) but leaves the *weight* free. A symmetry argument therefore yields only $W_a = a \cdot I + (1 - a) \cdot \Pi_e$, with odd eigenvalue a , hence $\rho = a$ undetermined. Symmetry gives Π_e ; it does not give $\frac{1}{2}$.

This obstruction is the standard against which any discharge of $\frac{1}{2}$ must be measured. §6L isolates three genuinely distinct things that must be pinned, not the two a looser reading suggests. Writing the gate as a mixture of a survive branch and a close branch with odd eigenvalues λ and μ , weighted $(w, 1-w)$, the odd eigenvalue of the realized operator is $\rho = w\lambda + (1-w)\mu$, and the exact point is $(w, \lambda, \mu) = (\frac{1}{2}, 1, 0)$. So three things must be fixed: the **even-sector freedom** (the b direction of $a \cdot \Pi_e + b \cdot \Pi_o$), the **inter-branch weight** w , and the **two branch eigenvalues** λ, μ . The paper's honest claim is *not* that one census principle does all of this. It is that three separate moves do, each contingent on Principle 4.1: the choice of branch set $\{I, \Pi_e\}$ kills the even-sector freedom (both branches act as 1 on the even sector); the census fixes the weight $w = \frac{1}{2}$; and the branch eigenvalues $\lambda = 1$ (the survive history is exactly I — branch identity) and $\mu = 0$ (the close history fully evenises in one step — single-step completeness) are a separate dynamical-and-identity postulate, not an output of the census. What the census does not supply — the branch-multiplicity count (at what granularity the census runs) and the absence of branch-cost asymmetry — becomes a short, explicit, χ -specific list rather than a hidden assumption.

3. The Binary Fold Sector from Ordered Closure

The word-span construction of this section is the Route B realisation, whose validity is conditional on the rank test $r = \text{rank}(U) = n$ (invertible closure); that computation is uncomputed pending concrete A, B, C , and is the single upstream decision of Appendix C.0. Until it is run, §3 is provisional — if U is rank-deficient the support realisation (Route A) governs instead. And invertibility is not "good news": when Route B is forced, the entire physical burden moves to the access-row identification (Appendix C, B.3) — the independently-specified encoder E_g — rather than being discharged, as the abstract notes.

The fold sector is built directly from ordered closure (Route B). Let A (admissibility), B (transport / fold-exchange), C (completion) be the primitive closure operations, and form the two oriented composites

$$U = ABC, D = CBA.$$

Introduce the reversal R on closure words, $R(X_1X_2 \cdots X_n) = X_n \cdots X_2X_1$, so $R(U) = D$ and $R^2 = I$. The genuine-fold condition is that U and D are **linearly independent** as operators — equivalently the Gram determinant

$$\det [[\langle U, U \rangle, \langle U, D \rangle], [\langle D, U \rangle, \langle D, D \rangle]] \neq 0 \text{ under } \langle X, Y \rangle = \text{tr}(X^\dagger Y),$$

equivalently $U \neq \lambda D$ for every scalar λ . Order-sensitivity $K := U - D = ABC - CBA \neq 0$ is *necessary but not sufficient*: it excludes only $\lambda = 1$ and still permits $U = \lambda D$ with $\lambda \neq 1$, which collapses $\text{span}\{U, D\}$ to one dimension and merges the even and odd lines. When U and D are independent the two orientations are genuine; otherwise "up/down" is a relabelling. Define the **fold space** as the word-span

$$\mathcal{F} = \text{span}\{U, D\},$$

a two-dimensional space whose vectors are the oriented closure words, with fold-exchange involution J fixed by $J(U) = D$, $J(D) = U$, so $J^2 = I$. The canonical even/odd decomposition is

$$\mathcal{F} = \mathcal{F}_e \oplus \mathcal{F}_o, \mathcal{F}_e = \text{span}\{U + D\}, \mathcal{F}_o = \text{span}\{U - D\},$$

with projections $\Pi_e = \frac{1}{2}(I + J)$, $\Pi_o = \frac{1}{2}(I - J)$. Because U and D are linearly independent, \mathcal{F}_e and \mathcal{F}_o are genuinely complementary one-dimensional lines and the odd sector \mathcal{F}_o is nontrivial. This is the explicit statement that the two orientations are genuine structural subspaces, not labels: J is the vector-level realisation of the closure-word reversal R , the even sector is the orientation-blind content $U + D$, and the odd sector is the orientation-distinguishing content $U - D$ that χ will read.

Remark (route choice and the rank-selecting test). An alternative realisation (Route A) builds the fold subspaces from the left/right operator supports of U acting on a state space ($\mathcal{F}_u = \text{Ran}(U)$, $\mathcal{F}_d = \text{Ran}(U^\dagger)$), exchanged by the polar-phase involution of the induction bridge. The two routes are *not* interchangeable: on the same operators they can return opposite Gate-1 verdicts. With invertible closure factors U is invertible, $\text{Ran}(U) = \text{Ran}(U^\dagger) = \mathcal{H}$, and Route A collapses ($P_u = P_d = I$) while the word-span Route B remains genuine. The route is therefore selected by a single upstream computation — the rank of U :

$$r = \text{rank}(U).$$

If $r < \dim \mathcal{H}$ (closure loses rank) Route A is available; if $r = \dim \mathcal{H}$ (invertible closure) only Route B applies. This paper adopts Route B for exposition, which makes binarity and the nontrivial odd sector immediate but concentrates the physical burden in the encoder of §7 — the access-row identification. The full Gate-1 pass/fail specification for both routes, and the rank test that selects between them, is **Appendix C**.

4. Same External Neighbourhood and Nonadjacency

For U and D to be **twins** — folds that closure cannot tell apart — two things must hold: they must present the same external closure neighbourhood, and they must not be joined by a direct fold-flip edge. Both follow from reversal structure, but the first carries a condition that must be stated rather than glossed.

Same external neighbourhood. Let $q(h)$ be the externally visible closure signature of a history h , and define the external neighbourhood $N_{\text{ext}}(h) = \{q(c[h]) : c \text{ admissible}\}$ as all signatures reachable by inserting h into an admissible context. Assume (i) admissible contexts are closed under reversal, and (ii) external signatures are **reversal-fixed**, $q(Rh) = q(h)$. Then for $s \in N_{\text{ext}}(U)$ with $s = q(c[U])$,

$$s = q(R(c[U])) = q(R(c)[D]), R(c) \text{ admissible},$$

so $s \in N_{\text{ext}}(D)$; the converse is identical, giving $N_{\text{ext}}(U) = N_{\text{ext}}(D)$.

The proof is valid, but condition (ii) is load-bearing and nearly the conclusion: it says the external signature forgets orientation. Without it, reversal yields only *mirrored* neighbourhoods, $N_{\text{ext}}(U) = R \cdot N_{\text{ext}}(D)$ as sets of reversed signatures — and mirror-equivalence is not twinning, since an observer with orientation access distinguishes the folds through the mirror map. So the honest content is: **U and D have the same external neighbourhood if and only if the completed external signatures are reversal-fixed.** This is condition (1) of the reduction, a real and checkable assumption, not a free consequence of symmetry.

Nonadjacency. Assume an elementary closure-order edge changes the word by one adjacent transposition. The word ABC has inversion number 0 and the reverse CBA has inversion number 3; one adjacent transposition changes the inversion number by exactly ± 1 , so the graph distance is

$$d(U, D) = d(ABC, CBA) = 3,$$

and there is no direct edge $U \leftrightarrow D$. Thus, under local closure-order adjacency, U and D are nonadjacent. This is the structural origin of "no fold-flip primitive": the flip is not reachable in one elementary move, so it is not an admissible one-step history. The result holds *given* the locality assumption — that elementary closure edges are adjacent transpositions. If arbitrary transpositions were elementary, the single swap (1 3) would make $d(U, D) = 1$ and reintroduce a flip edge. Locality is therefore exactly the $\beta = 0$ content, and it is carried inside condition (3) of the reduction, not assumed away.

5. Image-Even Closure: $Q = \Pi_e$

Uniform admissibility (§6) fixes the inter-branch weight once the branches are identified, but the closure branch must be *exactly* Π_e — complete erasure of the odd component (odd eigenvalue 0), not partial. This is the completeness half of the §6L burden, and it is discharged by decomposing it into three named, checkable closure properties rather than asserted from symmetry.

Let Q be the closure operation on the twin sector. Suppose

$Q^2 = Q$ (idempotent: closure completes in one application), $QJ = JQ$ (equivariant: closure prefers no orientation), $\text{Im}(Q) = \mathcal{F}_e$ (closed unresolved states carry only twin-even content).

Lemma 5.1. These three properties force $Q = \Pi_e$.

Proof. For $x_e \in \mathcal{F}_e = \text{Im}(Q)$, idempotency fixes the image, so $Qx_e = x_e$. For $x_o \in \mathcal{F}_o$, equivariance gives $J(Qx_o) = Q(Jx_o) = Q(-x_o) = -Qx_o$, so Qx_o is odd; but $Qx_o \in \text{Im}(Q) = \mathcal{F}_e$, hence $Qx_o \in \mathcal{F}_e \cap \mathcal{F}_o = \{0\}$, so $Qx_o = 0$. Therefore Q acts as identity on \mathcal{F}_e and zero on \mathcal{F}_o : $Q = \frac{1}{2}(I + J) = \Pi_e$. ■

The decomposition is exactly what §6L demands. Equivariance alone admits the whole family of J -equivariant idempotents $\{0, \Pi_e, \Pi_o, I\}$; the condition $\text{Im}(Q) = \mathcal{F}_e$ is the selector that excludes Π_o and I and pins the odd eigenvalue to 0. That image condition *is* the completeness — it is the

structural statement that closure removes the odd part entirely — so Lemma 5.1 is valid given it, with $\text{Im}(Q) = \mathcal{F}_e$ as the load-bearing assumption. This is condition (2) of the reduction.

6. Uniform Admissibility and the Exact Weight

With the branches identified — survival I and closure $Q = \Pi_e$ — the weight is fixed by the census principle the programme uses elsewhere.

Principle 6.1 — Uniform Admissibility. Each distinct primitive admissible history contributes one unit of census weight unless an admissible multiplicity or marker distinguishes it. The realized transition is the normalized census over the admissible histories.

What counts as a distinguisher. Principle 6.1 is not well-posed until "distinguishes" is fixed, and the obvious reading fails: the two histories I and Π_e *do* produce different outcomes by construction — one preserves the odd residue, the other annihilates it — so if "acts differently on the residue" were the disqualifying fact, no binary gate could ever be uniform and the principle would be empty. Distinguishability for census purposes is therefore assessed on the **admissible closure labels** attached to a history — the closure facts an admissible context can read off it — *not* on its dynamical effect. Two histories are census-indistinguishable iff no admissible closure fact (a marker, a multiplicity, a cost label) separates them, even when their operators differ. This is a weaker and more delicate notion than the $u \leftrightarrow d$ case, where indistinguishability is grounded in an actual involution J that literally exchanges the two; there is no involution conjugating I to Π_e (Π_e is not even invertible), so the equal weighting of the $\{I, \Pi_e\}$ gate is a bare census posit, not a symmetry. The whole epistemic weight of the outer weight $\frac{1}{2}$ sits here, and conditions (3) and (5) below are precisely the two kinds of admissible label — multiplicity/stabilizer and cost — that would, if present, break the indistinguishability.

This is the §6L weight-fixing upgrade: from "J-equivariant" (symmetry, which underdetermines the weight) to "uniformly counted over census-indistinguishable admissible histories" (which fixes it). It pins the weight w ; it does not by itself pin the branch eigenvalues. Two parts must be separated, and only the first is inherited.

Scalar equal-weight (inherited). That each primitive admissible history carries equal *base* weight $w(h) = c$ is relocated to the Born-rule corpus: once all admissible path information is contained in the distinguishability geometry, any extra positive weight is an additional distinguishability label unless absorbed into phase or counted as path-class multiplicity. The χ paper inherits equal scalar weight per primitive admissible history. This is condition (4).

Branch multiplicity and census granularity (χ -specific). Equal scalar weight gives $w(h) = c$; it does *not* by itself give $m_0 = m_1 = 1$ until the *unit* of the census is fixed. The census must be taken over **operational equivalence classes**, not raw syntactic words: by idempotency $Q^n = Q$, so Q , QQ , QQQ , ... are not distinct closure histories, and under operational indistinguishability \sim_{op} the one-step history quotient must reduce to exactly two classes,

$$H_{\chi} / \sim_{\text{op}} = \{ [I], [\Pi_e] \},$$

one primitive persistence class and one primitive closure-evenisation class, with **no admissible flip primitive** (the latter resting on the local-adjacency / nonadjacency result of §4, $d(U, D) = 3$). Two classes give two equal weights *provided the census is taken at a consistent granularity* — and this is the real prior condition, sharper than any stabilizer count. But the granularity is **not** something the idempotency identity hands you for free, and the framing must be careful here, because the alternative counts do not bracket $\frac{1}{2}$ — they span the whole interval $[0, \frac{1}{2}]$, all on one side.

Take the sufficiency operator itself, $S = \frac{1}{2}(P + C)$ with $C = \frac{1}{2}(P + \hat{J}P)$, which expands to $\frac{3}{4}P + \frac{1}{4}\hat{J}P$, i.e. $\frac{3}{4}I + \frac{1}{4}J$ on the residue, odd eigenvalue $\frac{1}{2}$. The same physical tokens admit other uniform censuses: a flat census over $\{P, \hat{J}P\}$ gives $\frac{1}{2}(I + J) = \Pi_e$, odd eigenvalue 0 (total evenisation); the idempotency tower $\{I, Q, Q^2, \dots\}$ as $N+1$ equal units gives odd eigenvalue $1/(N+1) \rightarrow 0$; the class census $\{[I], [Q]\}$ gives $\frac{1}{2}(I + Q)$, odd eigenvalue $\frac{1}{2}$. The $Q^n = Q$ collapse kills *only the tower reading*. It does nothing to separate the flat reading ($\rho = 0$) from the class reading ($\rho = \frac{1}{2}$) — those differ by the nesting, not by any power of Q . So class-level granularity is **not settled by idempotency**; it is a summary name for the conjunction of three structural facts — two genuinely open χ -specific gates and one algebraic identity:

(i) **no-flip outer binarity** (§4, open χ -specific gate): no flip primitive $\Rightarrow \hat{J}P$ cannot appear as a third *outer* survive-token \Rightarrow the outer census is forced to $\{\text{survive, close}\}$, blocking the flat $\rho = 0$ reading; (ii) **no inner tower** ($Q^n = Q$, an algebraic identity given idempotency, condition 2 — not an independent gate): the closure branch is one operational history, not a graded tower; (iii) **twins are one operational class** ($(1) \rightarrow (3)$, open χ -specific gate): the J -orbit $\{u, d\}$ inside the close branch is a single outer unit, so "close" counts once.

The open burden is therefore two gates, (i) and (iii); (ii) is free given condition (2).

Class-granularity is the *consequence* of (i) \wedge (ii) \wedge (iii), not an input. Counted that way $[I]$ and $[Q]$ are one unit each, (4) gives equal scalar weight, and $\rho = \frac{1}{2}$ — with no orbit or stabilizer factor (the earlier "equal stabilizer order" sub-clause was an artifact of mixing granularities — collapsing the tower while keeping the J -orbit — and is withdrawn; a microstate census would need a specified (G, X, \cdot) the construction has not built, and is inconsistent with the class-level quotient already in force).

The granularity knob is one-sided, and the sign is a gift. Every deviation from class granularity over-weights the close (evenising) branch — survival is the single token I , which has no internal multiplicity to inflate, while close carries the J -orbit that finer counts re-expand. So every refinement *subtracts* from $\frac{1}{2}$: flat tokens $\rightarrow 0$, tower $\rightarrow 1/(N+1)$, any split of close into k units $\rightarrow 1/(k+1)$, all in $[0, \frac{1}{2}]$, none above. A counting/granularity error is therefore a strictly **subtractive** knob, $\rho \leq \frac{1}{2}$. The one-sidedness is structural, not incidental: it holds *because* the survive branch is single-token — I has no orbit to split — so the only admissible miscounts are refinements of the close branch, which dilute survival. A miscount that manufactured survival multiplicity (counting two persistence sub-histories where there is one) would push ρ above $\frac{1}{2}$; that is excluded only because survival carries no internal structure to split. So the bound is

exactly as general as the claim "the survive branch is structurally single-token," and no more: if survival could ever carry orbit structure, $\rho \leq \frac{1}{2}$ would fail. Granted that, the audit's central value $\rho \approx 0.503$ sits *above* $\frac{1}{2}$ — the wrong sign for any granularity error. So a granularity error, even if the gates (i)/(iii) were to fail, cannot be the source of the observed excess; only $\mu > 0$ (incomplete evenisation, §9) or a cost-driven $w > \frac{1}{2}$ (§5) can lift ρ above $\frac{1}{2}$. The counting content of (3) is sign-restricted against the data in a way μ and the cost side of w are not.

What remains χ -specific and must be checked is therefore the two open gates (i) and (iii) — no-flip outer binarity and twins-as-one-class — with (ii) free given condition (2); and they are sign-constrained to $\rho \leq \frac{1}{2}$. Note that $\sim\text{op}$ is judged by external signature, so condition (3) depends on condition (1): orientation-blind signatures are what make the relevant histories operationally indistinguishable.

No branch-cost asymmetry (χ -specific). Realized weight is base weight times any admissible cost factor. Even with equal base weight (4) and a consistent class-level census (3), unequal realized weight survives unless no admissible entropy/action/realization cost distinguishes $[I]$ from $[\Pi_e]$. This is condition (5), and it is genuinely separate from (3): (3) fixes *how many units* each branch contributes (one each, at class granularity); (5) fixes that those units carry no differential *cost*. It is not routine: persistence (the null operation I) and closure (a genuine averaging onto \mathcal{F}_e) are physically different operations, and a differential realization cost is the generic expectation. If VERSF has a reason the closure branch carries no cost beyond persistence — for instance that closure here is a pure relabelling census with no realization penalty — that reason must be stated. Condition (5) is a primary audit target, on a par with (3), not a footnote.

Lemma 6.2 — Exact Weight. Given (3) quotient-binarity at a consistent class-level granularity, (4) equal scalar weight, and (5) no branch-cost asymmetry, the unnormalized local access operator is $A_{\text{loc}} = I + Q$ and the normalized transition is

$$W = \frac{1}{2}(I + Q) = \frac{1}{2}(I + \Pi_e) \text{ (using } Q = \Pi_e \text{ from Lemma 5.1).}$$

On the odd sector, $\Pi_e x_o = 0$, so $W x_o = \frac{1}{2} x_o$: the odd eigenvalue is exactly $\frac{1}{2}$. With general class multiplicities m_0, m_1 the operator would be $(m_0 I + m_1 \Pi_e)/(m_0 + m_1)$ and $\rho = m_0/(m_0 + m_1)$; at class granularity $m_0 = m_1 = 1$, giving exact $\frac{1}{2}$. (A microstate census would set $m_1 = \text{orbit size}$ and reintroduce a stabilizer dependence — but it is inconsistent with the class-level quotient already in force, §6 above.) ■

The crucial point: $\frac{1}{2}$ is not obtained from an $I \leftrightarrow \Pi_e$ symmetry (those operators are not algebraically interchangeable). It is the normalized count of two primitive admissible history classes at a consistent class-level granularity — the conjunction (3) \wedge (4) \wedge (5).

7. Log-Access Intertwining, and Why the Readout Is Logarithmic

Even with W acting in fold space, the *observed* susceptibility halves only if the physical χ -readout sees W . Let X_g be the physical state space at generation step g , $E_g : X_g \rightarrow \mathcal{F}$ the encoder extracting the fold-access residue, P_g persistence and C_g closure. Assume the encoder naturalities

$E_{g+1}P_g = E_g$ (persistence preserves the residue), $E_g\hat{J}_g = J E_g$ (fold-reversal equivariance, with microscopic involutions \hat{J}_g), $P_g\hat{J}_g = \hat{J}_{g+1}P_g$ (persistence respects reversal),

and define microscopic twin closure as the orbit census and refinement as the outer binary census,

$$C_g = \frac{1}{2}(P_g + \hat{J}_{g+1}P_g), S_g = \frac{1}{2}(P_g + C_g).$$

Here \hat{J} enters only as the **census symmetry of the close operation** — the averaging that defines C_g — not as a dynamical one-step survive-token. This is consistent with §4's prohibition of a flip primitive: §4 forbids $\hat{J}P$ from appearing as a third *outer* history (a flip edge, the β knob, $d(U,D) = 1$), whereas inside C_g the reversal is the symmetry being averaged over to *form* the single close class. Flip-as-census-symmetry-inside-close and flip-as-primitive-survive-token are different objects; only the latter is excluded.

Lemma 7.1 — Intertwining. Under these conditions, $E_{g+1}S_g = W \cdot E_g$ with $W = \frac{1}{2}(I + \Pi_e)$.

Proof. $E_{g+1}C_g = \frac{1}{2}[E_{g+1}P_g + E_{g+1}\hat{J}_{g+1}P_g] = \frac{1}{2}[E_g + J \cdot E_{g+1}P_g] = \frac{1}{2}[E_g + J \cdot E_g] = \Pi_e \cdot E_g$. Then $E_{g+1}S_g = \frac{1}{2}[E_{g+1}P_g + E_{g+1}C_g] = \frac{1}{2}[E_g + \Pi_e \cdot E_g] = \frac{1}{2}(I + \Pi_e) \cdot E_g = W \cdot E_g$. ■

This is a *sufficiency* construction: C_g and S_g are defined as the census averages, so the $\frac{1}{2}$'s are inherited from Principle 6.1 rather than independently re-derived, and the lemma shows a microscopic dynamics realizing W exists and is consistent with the encoder. Its genuinely new, checkable content is the naturality package together with **generation-independence** (the same J and the same local form at every step). This is condition (6).

Why the readout is logarithmic. The log-readout is not an arbitrary choice. Let the physical up/down mass-ratio map be $R(x) = m_{\text{up}}(x)/m_{\text{down}}(x) > 0$, and suppose independent access contributions compose multiplicatively and fold reversal inverts the ratio:

$$R(x + y) = R(x) \cdot R(y), R(Jx) = R(x)^{-1}.$$

Define $\chi(x) = \ln R(x)$. Then $\chi(x + y) = \chi(x) + \chi(y)$ and $\chi(Jx) = -\chi(x)$: χ is an additive, odd functional, so $\chi = \ell_o$ with $\ell_o J = -\ell_o$. Hence

$$\ell_o \Pi_e = \frac{1}{2} \ell_o (I + J) = \frac{1}{2} (\ell_o - \ell_o) = 0,$$

the annihilation the halving proof requires — derived from multiplicative composition and reversal-inversion, not posited. The remaining physical content is exactly the assumption that access composes multiplicatively (mass as an exponential of accumulated odd residue), which is carried in condition (6) along with the absence of any additive generation-dependent odd offset.

8. The χ -Halving Theorem

Theorem 8.1 — χ -Halving (conditional reduction). Assume a **genuine two-fold sector** —

(0) **operator nondegeneracy and route** — U and D linearly independent (Gram $\det \neq 0$, so $\mathcal{F} = \text{span}\{U, D\}$ is genuinely two-dimensional), with the rank of U fixing the realisation route (Route A if rank-deficient, Route B if invertible); §3, Appendix C.0 —

and then:

(1) **reversal-fixed external signatures** — $q(\text{Rh}) = q(h)$, giving $N_{\text{ext}}(U) = N_{\text{ext}}(D)$ (§4); (2) **image-even closure** — Q idempotent, J -equivariant, $\text{Im}(Q) = \mathcal{F}_e$, hence $Q = \Pi_e$ (§5); (3) **quotient-binary one-step history at consistent granularity** — $H_{\chi/\sim\text{op}} = \{[I], [\Pi_e]\}$, exactly two operational classes with no admissible flip primitive (resting on local closure-order adjacency, $d(U, D) = 3$), the census taken at class granularity throughout — which is *not* settled by idempotency alone but follows from the conjunction (i)–(iii) of §6 (no-flip outer binarity §4, the no-inner-tower identity $Q^n = Q$, and twins-as-one-class (1)→(3)) — so each class is one unit and no orbit/stabilizer multiplicity enters (§4, §6); (4) **equal scalar weight** per primitive admissible class, inherited from the Born-rule equal-weight result (§6); (5) **no admissible branch-cost asymmetry** distinguishing $[I]$ from $[\Pi_e]$ (§6); (6) **log-access intertwining** — the encoder is persistence-compatible, reversal-equivariant, generation-independent, and additive/multiplicative so that $\chi = \ell_o$ with no additive odd offset (§7).

Then $W = \frac{1}{2}(I + \Pi_e)$, $\ell_o W = \frac{1}{2}\ell_o$, and $r_{\{g+1\}} = W \cdot r_g$, so

$$\Delta\chi(g+1) = \ell_o(r_{\{g+1\}}) = \ell_o W(r_g) = \frac{1}{2} \cdot \ell_o(r_g) = \frac{1}{2} \cdot \Delta\chi(g),$$

and for the observed profile $\rho = \Delta\chi_2/\Delta\chi_1 = \frac{1}{2}$.

The dependency structure is not flat. Condition (0) is the prerequisite that there is a two-fold sector at all. Condition (1) feeds (3), since operational indistinguishability is judged by external signature. Condition (2) fixes the closure-branch operator. Conditions (3) \wedge (4) \wedge (5) give $m_0 = m_1 = 1$ and hence the exact weight (Lemma 6.2). Condition (6) carries the fold-space identity to the observed χ . The exact $\frac{1}{2}$ follows from the count of two equal-weight primitive classes at a consistent granularity — not from symmetry, which §6L proves insufficient.

The algebra of Lemmas 5.1, 6.2 and 7.1 has been verified on a non-degenerate model (encoder non-invertible, persistence with nontrivial action): $Q = \Pi_e$ is uniquely selected by the three closure properties, and $E_{\{g+1\}}S_g = W \cdot E_g$ holds with W 's odd eigenvalue equal to $\frac{1}{2}$.

9. Failure Modes — the Knobs and Their Structural Origins

Writing the general gate with outer weights $(w, 1 - w)$ on {survive, close}, **survive-branch odd eigenvalue** λ , and close-branch odd eigenvalue μ , the odd eigenvalue of W is

$$\rho = w \cdot \lambda + (1 - w) \cdot \mu,$$

with the exact point at $(w, \lambda, \mu) = (1/2, 1, 0)$. There are exactly **three operator-level axes** — the weight w and the two branch eigenvalues λ, μ — and the earlier draft's "1" silently set $\lambda = 1$, treating the survive branch as rigid while allowing the close branch to deform. That asymmetry is not justified by anything but the branch-identity assumption, and it is now made explicit. Each axis has a definite structural origin in the conditions of §8:

λ — partial survival (condition 3, branch identity). The survive branch contributes $\lambda = 1$ only if it is *exactly* the identity I . If the non-closing history slightly evenises ($\lambda = 1 - \varepsilon$), then at $w = 1/2$, $\mu = 0$ the readout is $\rho = 1/2(1 - \varepsilon) \neq 1/2$. $\lambda = 1$ is therefore not definitional unless "survival" is *defined* as I — which is exactly the branch-identity clause of condition (3) (the survive class is [I]). So $\lambda = 1$ is pinned by (3), on a par with $\mu = 0$ being pinned by (2); the catalogue treats the two branches symmetrically.

μ — incomplete evenisation (condition 2). If one closure step realizes only a partial twin average, the close branch has odd eigenvalue $\mu > 0$ and, at $\lambda = 1$, equal weights, $\rho = 1/2(1 + \mu)$. This is the failure of $\text{Im}(Q) = \mathcal{F}_e / \text{idempotency}$; exact halving needs $\mu = 0$.

β — fold-flip adjacency (condition 3). Adjacency is *not* a fourth operator axis: it is a structural origin of a nonzero close-eigenvalue. A direct fold-flip edge makes the averaging operator act on the contrast with a negative eigenvalue (in §10's twin graph, $-1/(k+1)$ for k shared neighbours), i.e. $\mu < 0$, giving $\rho = 1/2(1 - \beta)$ with $\beta = 1 - 2\rho$. So μ (incompleteness, $\rho > 1/2$) and β (adjacency, $\rho < 1/2$) are opposite-sign values of the *same* operator axis (the close-branch eigenvalue) arising from two different structural failures — image-even completeness (2) and nonadjacency (3) respectively.

w — unequal weights (conditions 3, 4, 5). With both branches at their exact eigenvalues, unequal weight gives $\rho = w = m_0/(m_0 + m_1) \neq 1/2$. At the operative class-level granularity the class multiplicities are $m_0 = m_1 = 1$ by construction, so the live structural sources are two: a base-weight inequality (failure of 4) or a branch-cost asymmetry (failure of 5). A stabilizer/orbit-size mismatch would feed w *only* under a microstate census, which is inconsistent with the class-level quotient in force (§6) — so it is not an operative knob unless the granularity itself is mis-set. Exact halving needs $w = 1/2$, i.e. consistent class granularity (3) \wedge equal scalar weight (4) \wedge no branch-cost asymmetry (5).

Discrete failures. A non-binary gate ($n \geq 3$ operational classes) gives $1/n$, not $1/2$ — the failure of (3)'s binarity; a non-even closure census ($\ell_0 \Pi_e \neq 0$) breaks the central cancellation — the failure

of (2); and a generation-dependent additive odd offset gives $\Delta\chi(g+1) = \frac{1}{2} \cdot \Delta\chi(g) + \ell_o(b_g)$, requiring $\ell_o(b_g) = 0$ — the failure of (6).

The knobs are observationally degenerate, separable only by structure. Three operator axes (w, λ, μ) map onto the single scalar $\rho = w\lambda + (1 - w)\mu$. So the deformations are not merely "unresolved at current precision" — they are unresolvable *in principle* from ρ alone: the pre-image of any $\rho \neq \frac{1}{2}$ is a surface in (w, λ, μ) -space, and sharpening the ρ measurement shrinks that value's error bar without ever shrinking the surface. A measurement of ρ tests only the conjunction $w\lambda + (1 - w)\mu = \frac{1}{2}$ — whether the gate halves exactly — never which axis is responsible if it does not. The knobs are separated only by the structural conditions (binarity, branch identity, single-step completeness, nonadjacency, cost-symmetry), never by the measurement. A complete physical proof must close all three operator axes to the exact point on structural grounds; no attainable ρ precision can do it.

10. The Twin-Antichain Compression of Conditions (1), (2), (6)

The prior χ -profile corpus (the Twin-Antichain Log-Access Route, §6M there) supplies a concrete, falsifiable realization of part of the six-condition reduction. It is worth stating precisely, both for what it compresses and for what it leaves untouched — because the natural temptation is to read it as collapsing the whole audit to "two equations," and that reading would silently re-absorb the weight conditions of §6.

Work on the effective closure-access graph. To avoid collision with the closure primitive A and the down-word $D = CBA$ of §3, write its adjacency matrix as A_G and its degree matrix as D_G , and let $M = D_G^{-1}A_G$ be the one-step neighbour-averaging (random-walk) operator. Let e_+ , e_- be the incidence vectors of the up and down folds. The route rests on two equations.

The closure-order twin incidence equation. The operative condition is that M annihilates the odd contrast as a **right** null vector — the column form —

$$M(e_+ - e_-) = 0,$$

since that is what feeds the half-lazy operator below. This holds iff the two folds have identical *columns* in M , i.e. $A_G(k, u) = A_G(k, d)$ for every external k : the folds are nonadjacent twins with the same neighbours. The corpus writes the condition in **row** form,

$$(e_+ - e_-)^T D_G^{-1}A_G = 0,$$

which states identical neighbour-averaged rows ($e_+ - e_-$ a *left* null vector). For a general $M = D_G^{-1}A_G$ the two are different conditions, and row-null does not imply column-null; they coincide only when A_G is symmetric with equal fold degrees, in which case both reduce to the same equal-neighbourhood statement $N(u) = N(d)$. So the "single incidence equation" collapses

(1)+(2) only under that graph symmetry, which must be assumed or checked; the genuinely operative object is the column condition $M(e_+ - e_-) = 0$.

This still does two jobs at once. It *is* the identical-(not-mirrored)-neighbourhood content of condition (1), stated as equal columns rather than a signature table. And it *delivers* the completeness of condition (2): because the folds share their neighbours, the averaging operator annihilates the odd contrast for any field on those neighbours — the cancellation is topological, not numerical. The half-lazy operator built from M , $W = \frac{1}{2}(I + M)$, then has odd eigenvalue exactly $\frac{1}{2}$, since $W(e_+ - e_-) = \frac{1}{2}(e_+ - e_-) + \frac{1}{2} \cdot 0 = \frac{1}{2}(e_+ - e_-)$. This collapses (1) and (2) into one check and *relaxes* (2): the closure branch need not be the exact even projection Π_e , only an operator that annihilates the odd contrast, which the twin condition guarantees. Note that $W = \frac{1}{2}(I + M)$ is not the same operator as the paper's $\frac{1}{2}(I + \Pi_e)$ — M is stochastic averaging, Π_e a projection — they agree only on the odd contrast mode, which is all the readout sees.

The log-access intertwining equation.

$E_{\{g+1\}} S_g = W E_g$, with W acting on the additive log-access shell, not raw accessibilities.

This is condition (6), with the corpus's sharpening: applying W to raw accessibilities gives only approximate halving near equality, whereas on the additive log shell it is exact — the concrete form of the multiplicative-to-additive transform of §7.

What the compression buys, and what it does not. The genuine prize is real: two verbal conditions (1)+(2) become one falsifiable, field-independent incidence check — the column condition $M(e_+ - e_-) = 0$ — with (2) weakened to odd-contrast annihilation. The route is cleanly refutable: if the folds are adjacent or their neighbour columns mismatch, the check fails, M no longer annihilates the contrast, and the odd eigenvalue deforms off $\frac{1}{2}$. That deformation is precisely the close-eigenvalue knob of §9: for adjacent folds sharing k neighbours the odd eigenvalue is $\frac{1}{2} \cdot k/(k+1)$, tending to $\frac{1}{2}$ only as $k \rightarrow \infty$, *not* to 0 — the route is degraded, not trivially killed, so "adjacent \Rightarrow eigenvalue 0" is not the right falsifier; "adjacent \Rightarrow eigenvalue $\neq \frac{1}{2}$ " is.

But the compression does **not** source the coefficient $\frac{1}{2}$, and the "two equations" summary must not be read as discharging the weight conditions. The $\frac{1}{2}$ in $W = \frac{1}{2}(I + M)$ is the laziness of the walk: the general operator is $(1 - \alpha)I + \alpha M$, with odd eigenvalue $1 - \alpha$ on the twin mode, and exact $\frac{1}{2}$ requires $\alpha = \frac{1}{2}$. That laziness *is* the inter-branch census weight, still carrying condition (4) equal scalar weight, condition (5) no branch-cost asymmetry, and condition (3)'s binarity-at-consistent-granularity. The twin structure secures the *cancellation*; it does not fix the *coefficient*. The incidence check is likewise silent on whether a third admissible history exists and on the census granularity, so the no-third-class and granularity clauses of condition (3) remain separate checks, not consequences of it.

Honest status under the compression. The χ -halving law is conditionally reduced to two finite, falsifiable structural checks from the prior Twin-Antichain route — the closure-order twin incidence check $M(e_+ - e_-) = 0$ (column form; equivalently the row equation $(e_+ - e_-)^T D_G^{-1} A_G = 0$ when A_G is symmetric with equal fold degrees), discharging (1) and (2) at

once and relaxing (2) to odd-contrast annihilation; and the log-access intertwining equation $E_{\{g+1\}}S_g = WE_g$ on the additive log shell, discharging (6) — *together with* the unchanged census conditions that fix the laziness to exactly $\frac{1}{2}$: equal scalar weight (4), no branch-cost asymmetry (5), and binarity at a consistent class granularity with no third class (3). The route secures the cancellation; it does not source $\frac{1}{2}$. The prior corpus is explicit that both equations remain unproved for the actual construction — its OP0.5 (physical identification of the folds as nonadjacent twins, our access-row identification) and OP2 (the log-access readout construction, our condition 6). If the two checks pass *and* the weight conditions hold for the real A, B, C, the physical χ gate satisfies the premises for exact $\rho = \frac{1}{2}$; if the twin check fails, the route is cleanly refuted.

11. Relation to the Empirical Audit

The audit and the reduction play complementary roles. The audit says $\rho = 0.50 \pm 0.02$ is empirically viable and not a mixed-scheme artefact; the reduction says that if conditions (1)–(6) hold, VERSF forces $\rho = \frac{1}{2}$ exactly. Neither proves the other; together they say the observed profile is consistent with the exact value VERSF would derive if the conditions are physically realized.

The §9 knobs make the relation quantitative, and they correspond to the reduction's conditions: $\mu \leftrightarrow$ (2) completeness, $\beta \leftrightarrow$ (3) nonadjacency (the opposite-sign value of the same close-eigenvalue axis), $\lambda \leftrightarrow$ (3) branch identity, $w \leftrightarrow$ (3) consistent class granularity / (4) scalar weight / (5) cost-symmetry, and (6) \supseteq the additive-offset failure. The audit constrains the *net* odd-eigenvalue deviation. The central value $\rho \approx 0.503$ gives $2(\rho - \frac{1}{2}) \approx 0.006$, and the audit bounds this combination at roughly the ± 0.05 level — but, by the audit's own central finding, it cannot separate the knobs from one another or fix the sign, because the central excess is matched in size by perturbative structure, by the m_u/m_d central-value choice, and by an uncomputed same-sign QED differential. The honest joint statement of the two papers is therefore:

the combination $\mu - \beta + 2(w - \frac{1}{2}) + \dots$ is consistent with zero at the ± 0.05 level, unresolved component-wise, with the sign carrying no information.

Exact $\frac{1}{2}$ predicts every knob is individually zero; the audit confirms only that their sum is small. And this is a degeneracy *in principle*, not a precision limit: the three operator axes (w, λ, μ) map onto a single scalar ρ , so no attainable measurement of ρ — however sharp — can attribute a deviation to a particular knob. A sharper m_u/m_d (and a QED treatment, which cleans the same scalar rather than adding an observable) tightens the test of the *conjunction* $w\lambda + (1 - w)\mu = \frac{1}{2}$ — whether the gate halves exactly — but never separates the knobs. If ρ collapses to 0.500 it drives every knob toward its exact value simultaneously and says nothing about which was responsible; if it locks at 0.503 it shows at least one knob is nonzero and still cannot say which. The knobs are separable only by structural argument — binarity, branch identity, single-step completeness, nonadjacency, cost-symmetry — never by the measurement. The empirical side bounds the package; the work that can advance the decomposition is entirely structural.

One asymmetry does partly break the degeneracy *against the data*, at no cost. The counting/granularity content of (3) is sign-restricted: every deviation from class granularity overweights the evenising branch and so subtracts from $\frac{1}{2}$ (§6), giving $\rho \in [0, \frac{1}{2}]$. The observed central value $\rho \approx 0.503$ lies *above* $\frac{1}{2}$. So whatever explains the small excess, it is not a counting error — that knob has the wrong sign. The excess, if real, must come from the two knobs that can lift ρ above $\frac{1}{2}$: incomplete evenisation $\mu > 0$ (§9) or a cost-driven weight $w > \frac{1}{2}$ (§5). This does not separate those two from each other, but it does retire a third candidate on sign grounds alone — a free constraint linking §6 to the audit.

12. What This Paper Does and Does Not Establish

Established here

1. The binary fold sector is constructed explicitly from ordered closure: $\mathcal{F} = \text{span}\{U, D\}$ with $U = ABC$, $D = CBA$, genuine (nontrivial complementary even/odd) whenever U and D are linearly independent (Gram $\det \neq 0$; $K = U - D \neq 0$ alone is insufficient), and J the vector-level reversal.
2. The required operator is $W = \frac{1}{2}(I + \Pi_e)$, with odd eigenvalue exactly $\frac{1}{2}$; $Q = \Pi_e$ is uniquely forced by $\{\text{idempotent, J-equivariant, Im} = \mathcal{F}_e\}$ (Lemma 5.1).
3. The exact weight $\frac{1}{2}$ is the normalized count of two primitive history classes at a consistent class-level granularity (each one unit) — not a symmetry, which §6L proves insufficient (Lemma 6.2).
4. The intertwining $E_{\{g+1\}}S_g = W \cdot E_g$ holds under the encoder naturalities (Lemma 7.1), and the log-readout is derived from multiplicative composition and reversal-inversion, giving $\ell_o \Pi_e = 0$.
5. The whole chain is a **complete conditional reduction**: $\rho = \frac{1}{2}$ follows from the nondegeneracy prerequisite (0) and the six explicit conditions of §8, with the algebra verified on a non-degenerate model. It does not *clear* §6L by a single principle; it *relocates* §6L into three moves — branch-set choice (kills the even-sector freedom), census (fixes the weight), and branch identity + single-step completeness (fix the two branch eigenvalues) — all contingent on Principle 4.1.
6. Each failure knob — the three operator axes w, λ, μ (with β a negative- μ structural origin) and the discrete failures — has a definite structural origin among the conditions; the axes are observationally degenerate on the single scalar ρ , so the audit bounds only their conjunction (± 0.05), and they are separable by structure alone, not by measurement.

Not established here

1. That the physical χ gate realizes conditions (0)–(6). In particular: that the fold sector is genuinely two-dimensional (0); that external signatures are reversal-fixed (1); that closure is an image-even idempotent retraction (2); that the one-step history quotient is exactly $\{[I], [\Pi_e]\}$ with no flip primitive at a consistent class-level granularity (3); that the two branches carry no admissible cost asymmetry (5); and that the encoder is additive, reversal-equivariant and generation-independent (6).

2. That uniform admissibility's *scalar* equal-weight (4) is fully discharged by the Born-rule result — it is relocated there, but the binarity-at-consistent-granularity (3) and branch-cost-symmetry (5) claims are χ -specific and undischarged.
3. The exact empirical identity $\rho = \frac{1}{2}$ (the audit gives 0.50 ± 0.02).
4. The full quark mass hierarchy, the magnitude $\Delta\chi_1$, or the first-generation anchor m_u/m_d as a fresh derivation.

The primary remaining targets, in order of exposure: the **quotient-binary history claim** (3), the **branch-cost symmetry** (5), and — under Route B, nearly the same task as (3) — the **access-row identification** that fixes which physical access rows map into $\text{span}\{U, D\}$, equivariantly and generation-independently.

13. Status Ledger

| Object | Status after this paper |
|--|--|
| Empirical target $\rho \approx \frac{1}{2}$ | Inherited from audit; viable, not an exact measurement |
| Binary fold sector $\mathcal{F} = \text{span}\{U, D\}$ (Route B) | Constructed; genuine iff U, D linearly independent (Gram $\det \neq 0$); $K \neq 0$ necessary, not sufficient |
| Route selection — $\text{rank}(U)$ | $r < \dim \mathcal{H} \rightarrow$ Route A available; $r = \dim \mathcal{H} \rightarrow$ Route B forced; the two can give opposite Gate-1 verdicts (Appendix C) |
| Symmetric-Averaging Obstruction §6L | Relocated , not cleared — branch-set choice kills even freedom, census fixes weight, branch identity + completeness fix eigenvalues; three moves, contingent on 4.1 |
| (0) Operator nondegeneracy + route | Prerequisite: U, D linearly independent (Gram $\det \neq 0$); $\text{rank}(U)$ selects Route A/B; uncomputed pending concrete A, B, C |
| (1) Reversal-fixed external signatures | Condition; = same-neighbourhood iff signatures orientation-blind (open) |
| (2) Image-even closure $Q = \Pi_e$ | Proven from {idempotent, equivariant, $\text{Im} = \mathcal{F}_e$ }; image-even is the load-bearing assumption (open physically) |
| Nonadjacency $d(U, D) = 3$ | Proven given local (adjacent-transposition) closure edges; = $\beta = 0$, folded into (3) |
| (3) Quotient-binary history $\{\{I\}, \{\Pi_e\}\}$, no flip, consistent class granularity | χ -specific; primary open target; depends on (1). Class granularity is not forced by idempotency alone but by the conjunction (i)–(iii) of §6 — no-flip binarity (§4) and twins-as-one-class ((1) \rightarrow (3)), both open, plus the identity $Q^n = Q$; then each class is one unit, $\rho = \frac{1}{2}$, no stabilizer/orbit factor (earlier sub-clause withdrawn as a mixed-granularity phantom). Counting errors are sign-restricted, $\rho \leq \frac{1}{2}$ |
| (4) Equal scalar weight per primitive class | Relocated to Born-rule equal-weight corpus (inherited) |
| (5) No branch-cost asymmetry | χ -specific; primary open target, on a par with (3); a cost (dynamical) factor, distinct from the counting/granularity content of (3) |

| Object | Status after this paper |
|--|---|
| (6) Log-access intertwining + log readout | Sufficiency construction proven; naturalities open; log readout derived from multiplicativity. Generation-independence (one W for all g): assumed, untestable at $N = 3$ — a g -dependent operator giving $\frac{1}{2}$ at the single existing step would reproduce the data |
| Exact weight $W = \frac{1}{2}(I + \Pi_e)$, $\rho = \frac{1}{2}$ | Proven conditional on (3) \wedge (4) \wedge (5); model-verified |
| χ -Halving Theorem | Complete conditional reduction (prerequisite 0 + conditions 1–6) Three axes onto one scalar $\rho = w\lambda + (1-w)\mu$: weight $w \leftrightarrow (3$ granularity, 4, 5), survive-eigenvalue $\lambda \leftrightarrow (3$ branch identity), close-eigenvalue $\mu \leftrightarrow (2)$; adjacency $\beta =$ negative- μ origin $\leftrightarrow (3$ no-flip). |
| Failure knobs (operator axes) | Degenerate in principle — audit bounds only the conjunction (± 0.05); separable by structure, not measurement Column check $M(e_+ - e_-) = 0$ (= row form $(e_+ - e_-)^T D_G^{-1} A_G = 0$ under symmetric A_G , equal degrees) collapses (1)+(2) and relaxes (2) to odd-contrast annihilation; intertwining on the log shell discharges (6). Secures cancellation, not the $\frac{1}{2}$ — laziness $\alpha = \frac{1}{2}$ still carries (3), (4), (5). Falsifier: twin-check failure deforms the eigenvalue off $\frac{1}{2}$, not to 0 |
| Twin-Antichain compression (§10) | |
| Access-row identification | Open; under Route B, nearly the same task as (3) |
| Magnitude $\Delta\chi_1$; full hierarchy; m_u/m_d anchor | Open / not established here |

14. Conclusion

The clean χ audit established $\rho = 0.50 \pm 0.02$, consistent with $\frac{1}{2}$ and no longer vulnerable to the mixed-scheme objection, and returned the burden to structure. This paper does not claim a complete proof; it delivers a **complete conditional reduction**, which is the honest and stronger thing.

The χ -halving law follows if the up/down refinement step realizes the half-lazy operator $W = \frac{1}{2}(I + \Pi_e)$. The decisive point is that this is not a symmetry argument, and not a single principle either. The Symmetric-Averaging Obstruction proves up/down symmetry pins the averaging direction Π_e but leaves the even-sector freedom, the inter-branch weight, and the two branch eigenvalues all free. Pinning them takes three separate moves, each contingent on the uniform-admissibility principle: the branch set $\{I, \Pi_e\}$ kills the even-sector freedom; the census fixes the weight to $\frac{1}{2}$; and branch identity (survive = I) with single-step completeness (close fully evenises) fixes the eigenvalues to 1 and 0. The inner "average over the twin orbit equals Π_e " is the *definition* of Π_e , not a second derivation — the physical content of complete evenisation is the single-step claim, not the algebraic identity. Applying the odd log-readout, derived from multiplicative composition, gives $\ell_o W = \frac{1}{2}\ell_o$ and $\Delta\chi(g+1) = \frac{1}{2}\Delta\chi(g)$.

What the reduction buys is precision about the remaining work. The exact $\frac{1}{2}$ is reduced to a nondegeneracy prerequisite (0) — a genuine two-fold sector — plus six conditions: (1) reversal-fixed external signatures; (2) image-even closure; (3) a quotient-binary one-step history with no flip primitive, counted at one consistent class-level granularity; (4) equal scalar weight per primitive class, inherited from the Born-rule corpus; (5) no admissible branch-cost asymmetry; and (6) log-access intertwining. The $\frac{1}{2}$ is therefore neither fitted from quark masses nor merely asserted as symmetry: it is inherited from the equal-weight principle already used in the Born-rule programme, plus a χ -specific claim that the local gate has exactly two operational history classes, each one census unit. Of the six, (4) is relocated and (2) is proven from its closure axioms, while (1), (3), (5) and (6) are the explicit audit targets — with (3) binarity-at-consistent-granularity and (5) branch-cost symmetry the most exposed, and the access-row identification, under Route B, nearly the same task as (3). An earlier draft elevated an "equal stabilizer order" sub-clause of (3) to sharpest condition; that is withdrawn as a phantom of inconsistent census granularity — under the operative class-level count no orbit/stabilizer factor arises. The genuine prior question is which granularity the inherited census uses, and class-level is *not* handed over by idempotency alone: it follows from the conjunction (i)–(iii) of §6 — no-flip outer binarity (§4) and twins-as-one-class ((1)→(3)), both open χ -specific gates, together with the algebraic identity $Q^n = Q$. And any residual granularity error is sign-restricted ($\rho \leq \frac{1}{2}$), so it cannot account for the audit's central value above $\frac{1}{2}$.

So the standing is exact and bounded. The χ -halving theorem is a complete conditional reduction. The algebra closes. The exact $\frac{1}{2}$ is inherited from equal scalar weight plus a χ -specific quotient-binary, consistent-granularity, no-cost-asymmetry gate. What remains is not a numerical fit but a structural audit of whether the physical χ gate satisfies those conditions — specified as decisive pass/fail tests in Appendix C, with a single upstream decision: the rank of U, which selects whether the support route (A) or the word-span route (B) governs Gate 1. The problem has sharpened: the task is no longer to admire $\frac{1}{2}$, nor to invoke symmetry, but to compute $\text{rank}(U)$ and then audit a short and explicit list.

Appendix A — Minimal Algebraic Skeleton

Closure words and fold space:

$U = ABC, D = CBA = R(U), R^2 = I, U, D$ linearly independent ($\det \text{Gram} \neq 0; K = U - D \neq 0$ necessary, not sufficient), $\mathcal{F} = \text{span}\{U, D\}, J(U) = D, J^2 = I, \Pi_e = \frac{1}{2}(I + J), \Pi_o = \frac{1}{2}(I - J), \mathcal{F}_e = \text{span}\{U + D\}, \mathcal{F}_o = \text{span}\{U - D\}$.

Image-even closure (Lemma 5.1): $Q^2 = Q, QJ = JQ, \text{Im}(Q) = \mathcal{F}_e \implies Q = \Pi_e$ (identity on \mathcal{F}_e , zero on \mathcal{F}_o).

Exact weight (Lemma 6.2): two primitive classes $\{[I], [\Pi_e]\}$ at class granularity, each one unit ($m_0 = m_1 = 1$) $\implies W = \frac{1}{2}(I + Q) = \frac{1}{2}(I + \Pi_e) = \frac{3}{4}I + \frac{1}{4}J$; odd eigenvalue $\frac{1}{2}$, even eigenvalue 1. (A microstate census would set $m_i = \text{orbit size}$ and give $\rho = m_0/(m_0 + m_1)$; inconsistent with the class-level quotient already in force.)

Log readout: $R(x + y) = R(x)R(y)$, $R(Jx) = R(x)^{-1}$, $\chi = \ln R \Rightarrow \chi$ additive, $\chi(Jx) = -\chi(x)$, so $\chi = \ell_o$ with $\ell_o J = -\ell_o$ and $\ell_o \Pi_e = 0$.

Intertwining (Lemma 7.1): $E_{-}\{g+1\}P_{-g} = E_{-g}$, $E_{-g}\hat{J}_{-g} = JE_{-g}$, $P_{-g}\hat{J}_{-g} = \hat{J}_{-}\{g+1\}P_{-g}$, $C_{-g} = \frac{1}{2}(P_{-g} + \hat{J}_{-}\{g+1\}P_{-g})$, $S_{-g} = \frac{1}{2}(P_{-g} + C_{-g}) \Rightarrow E_{-}\{g+1\}S_{-g} = WE_{-g}$.

Halving: $r_{-}\{g+1\} = Wr_{-g}$, $\Delta\chi(g) = \ell_o(r_{-g}) \Rightarrow \Delta\chi(g+1) = \ell_o Wr_{-g} = \frac{1}{2}\ell_o r_{-g} = \frac{1}{2}\Delta\chi(g)$, $\rho = \Delta\chi_2/\Delta\chi_1 = \frac{1}{2}$.

Deformation: survive-branch odd eigenvalue λ , close-branch odd eigenvalue μ , outer weight $w \Rightarrow \rho = w\lambda + (1 - w)\mu$; exact point $(w, \lambda, \mu) = (\frac{1}{2}, 1, 0)$. Three operator axes onto one scalar ρ — degenerate in principle. μ (incomplete evenisation, $\mu > 0$) and adjacency ($\beta = -\mu$ -side, $\mu < 0$) are the two structural origins of a nonzero close-eigenvalue; $\lambda < 1$ is partial survival; β -form $\rho = \frac{1}{2}(1 - \beta)$ at $\lambda = 1$, $w = \frac{1}{2}$.

Appendix B — Referee Objections and Replies

Objection 1 — You have assumed $\frac{1}{2}$. No. $\frac{1}{2}$ is *reduced* to a nondegeneracy prerequisite plus six conditions, of which the equal-weight prior is inherited from the Born-rule corpus and the rest are explicit χ -specific structural claims. The coefficient is the normalized count of two primitive history classes at a consistent class granularity, not an import from quark data.

Objection 2 — Symmetry alone does not imply $\frac{1}{2}$. Correct — this is the Symmetric-Averaging Obstruction §6L. Symmetry gives the family $a\Pi_e + b\Pi_o$. The reduction does not rely on symmetry; it invokes uniform admissibility (which fixes the weight) and the image-even condition $\text{Im}(Q) = \mathcal{F}_e$ (which fixes $b = 0$).

Objection 3 — Same external neighbourhood is nearly assumed. Acknowledged. Reversal symmetry gives only *mirrored* neighbourhoods; identity requires the external signatures to be reversal-fixed, $q(\text{Rh}) = q(h)$. That is stated as condition (1), a real assumption, not a theorem — and it feeds condition (3), since operational indistinguishability is judged by external signature.

Objection 4 — The gate may not be binary. Correct; condition (3). The two-dimensionality of $\mathcal{F} = \text{span}\{U, D\}$ gives two *subspaces*, not two *history classes*. Binarity is the quotient claim $H_{\chi}/\sim_{\text{op}} = \{[I], [\Pi_e]\}$, χ -specific and open, with the no-flip part resting on local closure-order adjacency ($d(U, D) = 3$).

Objection 5 — Evenisation may be incomplete. This is the μ knob. Completeness is $\text{Im}(Q) = \mathcal{F}_e$ together with idempotency (condition 2); incompleteness gives $\rho = \frac{1}{2}(1 + \mu)$.

Objection 6 — The two branches may carry different cost. Correct, and this is condition (5), elevated to a primary target. Equal *base* weight (4) is inherited from the Born-rule result, but a differential entropy/action/realization cost would tilt the realized weights (the w knob, $\rho = w$). Persistence (null) and closure (averaging) are physically different operations, so equal cost is not

free; VERSF must supply the reason, e.g. that closure is a pure relabelling census with no realization penalty.

Objection 7 — The fold operator may not control quark masses. The bridge is condition (6), log-access intertwining, including generation-independence and the multiplicative composition that makes the readout logarithmic. Lemma 7.1 is a sufficiency construction; the naturalities and multiplicativity are the checkable content.

Objection 8 — The empirical audit does not measure $\frac{1}{2}$. Correct; $\rho = 0.50 \pm 0.02$, and the audit established the sign of the central excess carries no information. The reduction claims structural exactness conditional on (1)–(6), not empirical proof of equality.

Objection 9 — Uniform admissibility is an axiom minted for $\frac{1}{2}$. This is the sharpest charge, and the defence is narrower than a confident reading would like. Its scalar equal-weight part is *relocated* to the Born-rule equal-weight / geometric-dependence result — applied here rather than invented — but that relocation is a promissory cross-grounding ("to be cited"), and it is the *only* thing that meets the charge. The principle does give definite answers for other configurations (ternary $\Rightarrow \frac{1}{3}$; a marked branch \Rightarrow unequal weight; an incomplete average $\Rightarrow \mu > 0$), but those are the rule restated on other inputs, not predictive *successes*: they are internal consistency / non-arbitrariness, not independent evidence that the rule is correct. Until one such answer is checked against an independently identified physical gate, 4.1 is constrained only by the one number it was deployed to produce. So the "axiom minted for $\frac{1}{2}$ " charge is met only by the cross-grounding, not by the rule's self-consistency. What is genuinely *not* inherited — binarity at a consistent census granularity (3), and branch-cost symmetry (5) — is stated as explicit χ -specific conditions, so the residual dependency is exposed rather than hidden.

Objection 10 — Does this derive all quark masses? No. Only the conditional χ -increment ratio. It does not derive $\Delta\chi_1$, the first-generation anchor m_u/m_d , the same-charge ladder, or the full hierarchy.

Objection 11 — What makes I and Π_e "indistinguishable" when they act differently? A real well-posedness point. The two histories produce different outcomes by construction, so "acts differently on the residue" cannot be the disqualifying fact — else no binary gate would ever be uniform. Distinguishability for census purposes is assessed on the **admissible closure labels** attached to a history, not on its dynamical effect (§6). And unlike the $u \leftrightarrow d$ case — grounded in an actual involution J — there is no involution conjugating I to Π_e , so the equal weighting of $\{I, \Pi_e\}$ is a bare census posit, not a symmetry. Conditions (3) and (5) are exactly the two admissible labels (a census-granularity/multiplicity fact, and a cost) that would break it. This is upstream of Objection 4, which assumes the branches and asks about hidden markers; Objection 11 asks what would count as a marker for two functionally-asymmetric operators.

Objection 12 — Why is the survive branch rigid ($\lambda = 1$) when the close branch is allowed to deform (μ)? It is not rigid by nature. The general gate is $\rho = w\lambda + (1 - w)\mu$ with three operator axes; $\lambda = 1$ holds only if the survive history is *exactly* I, which is the branch-identity clause of condition (3). §9 now names λ (partial survival) as a knob on a par with μ , pinned by (3) just as μ

= 0 is pinned by (2). (Adjacency β is not a fourth axis: it is a negative value of the same close-eigenvalue axis as μ .)

Objection 13 — Can a sharper ρ ever separate the knobs? No, and §11 now says so. The three operator axes (w, λ, μ) map onto a single scalar ρ , so the degeneracy is in principle, not a precision limit: any measurement tests only the conjunction $w\lambda + (1 - w)\mu = 1/2$, never the decomposition. The knobs are separable by structural argument alone.

Objection 14 — The "equal stabilizer order" sub-clause is a phantom of inconsistent counting. Sustained, and the clause is withdrawn. The orbit-stabilizer mechanics ($|\text{orbit}| = |G|/|\text{Stab}|, \mathbb{Z}_2 \rightarrow 1/3/2/3$) are correct in the abstract but were never made computable — no (G, X, action) was specified, and under the natural conjugation action the direction is indeterminate ($\text{Stab}(I) = G$ gives orbit 1; $\text{Stab}(\Pi_e)$ is unpinned). Worse, the clause is only *live* under microstate counting, while the operative census is class-level — established not by the $Q^n = Q$ identity alone (which kills only the inner tower) but by the conjunction of no-flip outer binarity (§4) and twins-as-one-class ($(1) \rightarrow (3)$), per §6. Counting consistently at class level — the granularity (4)'s "per primitive class" wording also assumes — makes $[I]$ and $[\Pi_e]$ one unit each, $\rho = 1/2$, with no stabilizer factor; the twin orbit $\{u, d\}$ is collapsed by the operational quotient. The genuine prior condition is therefore *granularity consistency*, not stabilizer equality (§6, condition 3). Calling the stabilizer clause "sharpest" was an artifact of collapsing one multiplicity while keeping another.

Objection 15 — §10's incidence equation is row-vs-column. Sustained. The written equation $(e_+ - e_-)^T D_G^{-1} A_G = 0$ is the left-null (row) condition; the annihilation actually used, $M(e_+ - e_-) = 0$, is right-null (column). For non-symmetric M these differ and row does not imply column; they coincide only under symmetric A_G with equal fold degrees. §10 now states the column condition as operative and flags the symmetry assumption under which the single row equation suffices.

Objection 16 — "Six conditions" omits the nondegeneracy prerequisite. Sustained. The two-fold sector requires U, D linearly independent ($\text{Gram det} \neq 0$), which can fail ($U = \lambda D$). It is now condition (0) in §8 — operator nondegeneracy plus the rank fact that selects the route — and the count is stated as "a genuine two-fold sector plus six conditions" in the abstract, ledger, and conclusion.

Appendix C — G-CHI-2 Gate Audit: Gate 1 Specification

This appendix turns the existence claim of §3 — that two genuine fold orientations exist — into a decisive pass/fail computation. Gate 1 alone is specified here; it establishes only that the fold sector is real and is silent on twinning (Gate 2), branch weights (Gate 2 census), and log-access intertwining (Gate 3), which §D below records as the remaining tests. The route fork of §3 makes Gate 1 two-headed, so the appendix opens with the computation that selects the head.

C.0 The Selecting Test — Is VERSF Closure Rank-Deficient?

The two routes test different objects and can return opposite verdicts on the same operators: with invertible self-adjoint closure factors the support route reports collapse ($P_u = P_d = I$) while the word-span route reports a genuine two-fold sector. So Gate 1 cannot be run until one upstream fact is settled.

Input. Concrete operators A (admissibility), B (transport / fold-exchange), C (completion) on a state space \mathcal{H} of dimension n ; $U = ABC$, $D = CBA$.

Computation. $r = \text{rank}(U)$, via SVD (number of singular values $\sigma_i > \text{tol}$).

Decision. $r < n$ (closure loses rank) \rightarrow Route A available, run §C.A; option 3 (audit in A, exposit in B) is then also available via the induction bridge (§C.C). $r = n$ (invertible closure) \rightarrow Route A unavailable, since $\text{Ran}(U) = \text{Ran}(U^\dagger) = \mathcal{H}$ forces $P_u = P_d = I$; commit to Route B and run §C.B. This is a physics commitment: $r < n$ requires admissibility or completion to genuinely lose rank.

C.A Route-A Gate-1 Spec — Support Subspaces (requires $r < n$)

Preconditions. A, B, C self-adjoint, so $D = CBA = U^\dagger$ exactly. (If the factors are not self-adjoint, $D \neq U^\dagger$; the support story decouples from the reversal story and Route A is not well-posed — use Route B.) $r = \text{rank}(U) < n$.

A.1 — Order sensitivity ($U \neq D$). Witness x with $(U - D)x \neq 0$. **Pass:** $\|U - D\| > \text{tol}$. (Necessary, not sufficient.)

A.2 — Distinct support subspaces. From the SVD $U = \sum_i \sigma_i u_i v_i^\dagger$, compute $P_u = \sum_{\{\sigma_i > \text{tol}\}} u_i u_i^\dagger$ (onto $\text{Ran}(U) = \mathcal{F}_u$) and $P_d = \sum_{\{\sigma_i > \text{tol}\}} v_i v_i^\dagger$ (onto $\text{Ran}(U^\dagger) = \mathcal{F}_d$). **Pass:** $P_u \neq P_d$, both nonzero — witness any x with $(P_u - P_d)x \neq 0$. Note $\text{rank}(P_u) = \text{rank}(P_d) = r$ always, so the test is about ranges, not ranks: equal-rank, distinct-range is the target. **Fail (collapse):** $P_u = P_d$.

A.3 — Exchange consistency. An involution J , $J^2 = I$, with $J P_u J = P_d$, built from the polar phase V of U ($U = V|U|$) on the doubled carrier $\mathcal{F}_d \oplus \mathcal{F}_u$; then $\Pi_e = \frac{1}{2}(I + J)$, $\Pi_o = \frac{1}{2}(I - J)$. **Pass:** the constructed J satisfies $J^2 = I$ and $J P_u J = P_d$.

Route-A Gate 1 passes iff $A.1 \wedge A.2 \wedge A.3$. *Establishes* two distinct fold subspaces exchanged by J . *Does not establish* that physical residues live in them (Gate 3), twinning (Gate 2), or weights (Gate 2 census).

C.B Route-B Gate-1 Spec — Word-Span (no rank requirement)

Preconditions. A, B, C operators on \mathcal{H} ; $D = CBA$ defined as the reversal of U (self-adjointness not required). Operator inner product $\langle X, Y \rangle = \text{tr}(X^\dagger Y)$.

B.1 — Nondegeneracy (U, D independent). The correct genuineness condition, strictly stronger than $K \neq 0$ (the case $U = \lambda D$, $\lambda \neq 1$, has $K \neq 0$ yet $\text{span}\{U, D\}$ is one-dimensional). Compute the Gram matrix $G = [[\langle U, U \rangle, \langle U, D \rangle], [\langle D, U \rangle, \langle D, D \rangle]]$. **Pass:** $|\det G| > \text{tol}$, i.e. $U \neq \lambda D$ for every scalar λ . *Do not use $K \neq 0$ as the test.*

B.2 — Involution and split well-defined. Given B.1, $J(U) = D$, $J(D) = U$ extends to a linear involution on $\mathcal{F} = \text{span}\{U, D\}$ with $J^2 = I$, and $\mathcal{F}_e = \text{span}\{U + D\}$, $\mathcal{F}_o = \text{span}\{U - D\}$ are complementary one-dimensional lines. **Pass:** $\|U + D\| > \text{tol}$ and $\|U - D\| > \text{tol}$; the fold-exchange J coincides with the Gate-2 graph automorphism restricted to $\{U, D\}$.

B.3 — Access-row identification (the relocated burden). B.1–B.2 pass cheaply; the real Gate-1 content is here, and under Route B it is nearly the same task as Gate 2's quotient-binarity. Exhibit the physical encoder $E_g : X_g \rightarrow \mathcal{F}$ — the access rows — and verify on a finite generating set B_g of X_g : (i) $\text{Im}(E_g) \subseteq \text{span}\{U, D\}$ (lands in the word-span); (ii) $E_g \hat{J}_g = J E_g$ (J-equivariant). **Pass:** both hold on all generators, all g . **Requires an independently specified E_g** — defining E_g from observed masses voids the test (circularity, as in Gate 3).

Route-B Gate 1 passes iff $B.1 \wedge B.2 \wedge B.3$. *Establishes* a genuine two-dimensional fold word-space with a physical encoding into it. *Does not establish* twinning (Gate 2), weights (Gate 2 census), or full intertwining (Gate 3, of which B.3 is the encoder half).

C.C Option 3 — Audit in A, Exposit in B (only if $r < n$)

If C.0 returns $r < n$, the strongest configuration runs the concrete support test (§C.A, a finite SVD computation with a decisive witness) as the audit, presents $\mathcal{F} = \text{span}\{U, D\}$ in the body for clarity, and certifies the two agree via the induction bridge: the polar phase V of U induces the word-span involution, so $P_u \neq P_d$ (Route A) and U, D independent (Route B) are two readings of one fact. Unavailable when $r = n$, where only §C.B applies and B.3 carries Gate 1.

C.D What Gate 1 Does Not Decide (both routes)

- **Twinning (Gate 2a/2b):** reversal-fixed external signatures and nonadjacency — the closure-graph row comparison $A_G(U, \cdot)|_{\text{ext}} = A_G(D, \cdot)|_{\text{ext}}$ (decisive when the graph is locally finite) and rewrite-generator enumeration $s(U) \neq D$ for every generator s .
- **Equal weights (Gate 2 census):** quotient-binarity $H_{\chi/\sim\text{op}} = \{[I], [\Pi_e]\}$, equal scalar weight, and a **consistent census granularity**. The audit must verify the census is run at class level — which is *not* settled by idempotency alone but is the conjunction of no-flip outer binarity (§4), no inner tower ($Q^n = Q$), and twins-as-one-class ($(1) \rightarrow (3)$), per §6 — under which each of $[I], [\Pi_e]$ is one unit and $\rho = 1/2$. A granularity error is sign-restricted ($\rho \leq 1/2$), so it cannot explain a central value above $1/2$. (An earlier draft demanded "equal branch stabilizer order"; withdrawn as a phantom of mixed granularity — see §6 and Objection 14. A stabilizer/orbit factor is live only under a microstate census, which is inconsistent with the class-level quotient already in force and would need a specified group action; not the inherited census.)
- **Log-access intertwining (Gate 3):** multiplicativity, fold-reversal equivariance, and generation naturality with no odd affine offset, tested on generators with both E_g and S_g independently specified (defining S_g as the census average $1/2(P_g + C_g)$ makes the intertwining test vacuous — the second circularity trap).

Passing Gate 1 means the fold sector is real; it does not mean it halves.