A Three-Gate W_7 Audit of the χ-Halving Law

A Three-Gate W_7 Audit of the χ-Halving Law in VERSF — Gate-1 Fold Realisation, Gate-2 Twin–Census Closure, and Gate-3 Log-Access Intertwining

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

Companion to "Uniform Binary Admissibility and the Half-Lazy Closure Operator: A Structural Reduction of the χ-Halving Law in VERSF" (the ***reduction****), whose conditional theorem and Gate-1 specification this paper inherits.*

Summary for the General Reader

The previous paper gave a recipe for a "gate" — a small step that, if it has the right shape, forces the up/down mass pattern in quarks to halve from one generation to the next. The recipe has three ingredients, called admissibility, transport, and completion, but it stopped short of saying what those ingredients concretely are. It ended by pointing at a single calculation that decides everything downstream: does a genuine *two-sided* up/down structure exist in the gate at all? That is the gate's first checkpoint.

This paper supplies concrete ingredients and runs that checkpoint. The ingredients are not picked arbitrarily — they are built from a shape that already lives in the theory: a small wheel-shaped network called W_7. The symmetry of that wheel naturally singles out a pair of opposite "directions," and the construction takes two of the three ingredients to be those two mirror-image directions, and the third to be the mirror that swaps them. With these in place, the first checkpoint passes cleanly, and a companion reading that extracts the up/down quantity does so without ever consulting the measured quark masses.

But the paper is careful about what that actually shows, and this is its honest centre. The checkpoint does not merely pass — it *cannot fail*, for a reason that turns out to be a small theorem rather than a lucky outcome. Any wheel-symmetric pair of opposite directions, paired with the mirror that swaps them, produces exactly this clean result; the algebra is identical no matter which pair you start from. So the checkpoint is not testing whether *this particular* wheel and *this particular* pair are the physically right ones. It is confirming that the recipe, once you have committed to a wheel and a mirror-pair at all, has the shape the gate needs.

The real question therefore moves onto the commitments, and we state four of them plainly. Is the wheel the right shape to be working on in the first place? Is it *this* pair of directions and not the other, equally good pair — a choice the checkpoint provably cannot make for us, because the two pairs are mirror copies that the checkpoint reads as identical, and the only thing that tells

them apart (how fast the wheel's rotation turns each pair — a sixth of a turn versus a third) lives in a part of the structure the checkpoint never looks at? Are the two ingredients truly mirror images, as the construction in fact requires (pairing them any other way makes the recipe collapse to nothing)? And is the up/down reading drawn from the theory's own pre-existing notion of "access," rather than quietly tuned to give the answer we were hoping for?

The candidate is a real advance — it is grown from the theory's own structure instead of from placeholders, it is mathematically clean, and its central claims are now closed-form facts rather than numbers that merely came out small on a computer. But its credibility rests squarely on those four commitments, which this checkpoint confirms rather than independently corroborates. The lasting prize is more practical: once the shape is fixed, the harder checkpoints still ahead stop being abstract conditions and become concrete calculations on an actual network.

This version carries those harder checkpoints through. There are three in all. The first (does a genuine two-sided structure exist?) is the one above. The second asks whether the two sides are true *twins* — indistinguishable by any fair measurement, separated by no shortcut, and sorted into exactly two bins — and the answer here is yes, except for one physical assumption about how "cost" is counted. The third asks for an explicit rule that actually halves the up/down quantity from one generation to the next; the paper writes that rule down, shows it is the only rule consistent with a short list of natural requirements, and shows the halving factor of one-half is forced by the two-bin structure of the second checkpoint — not chosen. Three of the four original commitments have since become *principled requirements* rather than free choices: which shape to read the structure on, which pair of directions, and which "access" reading are now pinned by natural demands (a faithful representation of the wheel's rotation, the meaning of reversible transport, and a short list of access axioms). What remains is one structural identification — that the physical carrier really is this wheel's cycle space — together with two finite checks that can be run on a computer: that the two bins cost the same under the theory's own ledger, and that the real dynamics actually follow the written-down halving rule. So the honest claim is conditional: grant one identification, three natural operational requirements, and two checkable audits, and the one-half halving law follows by calculation. Nothing here is asserted as physically proved. And one boundary should be stated plainly up front: the "up/down quantity" this paper halves is an abstract internal measure defined on the wheel; that it is *the actual quark susceptibility* is taken over from the earlier paper and is not checked anywhere here. So what is shown is that an internal contrast halves — its tie to real quark masses is assumed, not tested.

Abstract

The reduction established that the χ-halving law $\rho = \Delta\chi_2/\Delta\chi_1 = \frac{1}{2}$ follows from a nondegeneracy prerequisite plus six structural conditions on an up/down "fold" gate built from primitive closure operations A (admissibility), B (transport / fold-exchange), C (completion), with oriented closure words U = ABC, D = CBA. It specified a two-headed Gate-1 audit — selected by $r = \mathrm{rank}(U)$ — that decides whether the fold sector is genuine, and flagged r as uncomputed pending concrete A, B, C. This paper supplies concrete operators and runs the audit.

The operators are **generated**, not chosen: from the W_7 closure cell. On $H^1(W_7; \mathbb{C})$ (dimension $|E| - |V| + 1 = 12 - 7 + 1 = 6$) the rim rotation acts as the regular representation of C_6, with spectrum the six sixth-roots of unity. The reflection-conjugate doublet at $e^{\pm i\pi/3}$ (the E_1 sector) yields rank-one orthogonal projectors P_+, P_- exchanged by a D_6 reflection $S = S^\dagger = S^{-1}$. Setting **$A = P_+$, $B = S$, $C = P_-$** (self-adjoint), the closure words satisfy — and we prove this as a closed-form **Dihedral Automatic-Pass Lemma** (§3.1), so the witnesses below are exact, not numerical: $U = \alpha|v_+\rangle\langle v_-|$ and $D = \bar{\alpha}|v_-\rangle\langle v_+|$ with $\alpha = \langle v_+|S|v_-\rangle$, $|\alpha| = 1$; $\mathrm{rank}(U) = 1 < \dim H^1 = 6$ (Route A); $D = U^\dagger$ identically; $UU^\dagger = A$, $U^\dagger U = C$ (distinct rank-1 supports); $\|U - D\| = \sqrt{2}$ exactly; $\det \mathrm{Gram}(U, D) = 1$ exactly (U, D Hilbert-Schmidt-orthonormal). A **mass-blind** candidate access map $E_m(R) = \mathrm{Tr}(P_+R)\cdot U + \mathrm{Tr}(P_-R)\cdot D$, induced from the pre-existing trace-accessibility readings, satisfies the encoder tests (image-containment automatic by construction; reversal-equivariance exact — and, we note, itself entailed by the mirror identity $SAS = C$), with odd readout $\chi(m) = \ln \mathrm{Tr}(P_+R_m) - \ln \mathrm{Tr}(P_-R_m)$.

The honest grade. Gate 1 is **automatically satisfied** by this self-adjoint, symmetry-canonical candidate — the Automatic-Pass Lemma shows the satisfaction is a property of *any* two-line-plus-reflection fold structure, not specific to W_7 — so the audit is *reduced to* four identification bridges it does not itself test: (a) the carrier; (b) the doublet is E_1 not E_2; (c) admissibility and completion are reflection-conjugate; (d) the access reading is $\mathrm{Tr}(P\pm R)$. Gate 1 cannot decide (b): E_1 and E_2 are **distinct, non-isomorphic D_6-irreps** (their rotation characters differ), but they induce **isomorphic two-line-plus-reflection data on the exchange quotient Gate 1 actually reads**, so the construction is bit-identical on either — the distinguishing datum (ord R = 6 on E_1 versus 3 on E_2) lives in the ambient action the quotient forgets. §8 then **upgrades three of the four bridges from free choices to lemmas conditional on explicit operational axioms**: (b) E_1 follows from the **Primitive-Faithfulness Lemma** (no non-identity element of the six-step rim group C_6 acts trivially — E_2 fails since $R^3 = I$); (c) the reflection-conjugate pairing follows from the **Source–Target Support Theorem** (reversible transport, $P_{comp} = S\, P_{adm}\, S$); (d) the trace reading follows from the **Trace-Access Uniqueness Theorem** (four rank-one access axioms). Only (a), the identification of the physical carrier with the W_7 cycle space, remains a bare structural assumption. **Gate 2** (Part III) is a conditional pass: operational twinhood (dim-4 invariant algebra, with the external/internal observable split), an *executed* nonadjacency rewrite audit on $\mathrm{Cay}(S_3,\{(12),(23)\})$ giving d(ABC,CBA) = 3, and a binary census *classified* not posited (Terminal-History Classification, Thm 2.4); its open content is a finite audit — the rewrite-generator and terminal-history premises, and the two equalities $\omega_I = \omega_Q$ (uniform class weight) and $\Delta\ell_I = \Delta\ell_Q$ (ledger cost) that fix $w = ½$. **Gate 3** (Part IV) derives the exact contrast-transport **family** $\chi(\Phi R) = w\cdot\chi(R)$, pins the inner mean by a regularised Geometric-Mean Selection Lemma, and takes **$w = ½$ from Gate 2** — its open content is the **executable generator-level descent audit** $E_{m+1}S_m = W E_m$. The resulting **Three-Gate Conditional Realisation Theorem**: granting the carrier (a), three operational axioms (faithfulness, source–target, access), and the two finite audits (Gate-2 census/cost, Gate-3 descent), the candidate passes all three gates and $\Delta\chi_{g+1} = ½\Delta\chi_g$. This is **conditional derivation, not physical proof**. The full dependency ledger is stated flatly below: **one structural assumption (the carrier), three natural-but-unverified operational premises (faithfulness, source–target, access), and two finite audits (the Gate-2 audit comprising four sub-conditions, including the genuinely independent equalities $\omega_I = \omega_Q$ and $\Delta\ell_I = \Delta\ell_Q$; and the Gate-3 descent check)**, plus one upstream interpretive bridge. The reduction from arbitrary choices is major, but "natural" does visible work in three places, and the headline is the *enumerated* ledger, not a smaller "one + two." **Scope:** what is derived, granting all of these, is the halving of an **abstract internal log-access contrast**, $\Delta\chi_{g+1} = \frac{1}{2}\Delta\chi_g$ for the functional χ defined here; that this functional *is* the physical quark susceptibility is assumed upstream (the reduction's interpretive bridge plus the carrier identification) and is **tested nowhere in this paper**. A correction of record is folded in: an earlier draft conflated these χ gates with a separate transport/locality sequence (orientability; the refinement-loop trace τ), which supports the wider programme but does not by itself close χ-halving and is demarcated as §13.

The Three-Gate Conditional Realisation Theorem

This paper is organised around a single conditional theorem. The reduction reduced the χ-halving law to a gate audit; this paper executes all three gates on the W_7 candidate and isolates every remaining dependency as a named premise.

The three χ-halving gates (as the reduction defines them — and distinct from the separate transport/locality sequence of §13):

- **χ Gate 1 — fold realisation.** The up/down fold sector is genuine: $U = ABC$ and $D = CBA$ are linearly independent, so the fold space $\mathcal{F} = \mathrm{span}\{U, D\}$ is real and the access state is encoded into it.
- **χ Gate 2 — twin / census / cost.** The two folds are operational twins (indistinguishable by invariant observables), the orientation words are non-adjacent, the operational census is binary, and there is no branch-cost tilt between the folds.
- **χ Gate 3 — log-access intertwining.** A mass-blind access update Φ intertwines with the doubling/halving map W ($E \circ \Phi = W \circ E$), halving the odd log-access readout.

Three-Gate Conditional Realisation Theorem. Assume:

Carrier and operational semantics. (1) the physical persistent carrier is the W_7 cycle space $H_1(W_7) \cong H^1(W_7)$; (2) the fold representation is faithful to the full C_6 rim action (no non-identity rim element acts trivially); (3) completion is the reversible transport of the admissible input support ($P_{comp} = S\, P_{adm}\, S$); (4) accessibility satisfies the rank-one trace-access axioms (positivity, support-locality, P-invariance, normalisation).

Gate-2 finite audit. (5) the primitive rewrite generators are exactly $\{(12),(23)\}$ with no direct reversal generator; (6) every admissible terminal history satisfies $T J = J T$, $T^2 = T$, $T|_{\mathcal{F}_e} = I$; (7) the class base weights are uniform, $\omega_I = \omega_Q$ (a premise about the census *measure*, not a corollary of binarity); (8) the ledger costs are equal, $\Delta\ell_I = \Delta\ell_Q$.

Gate-3 finite audit. (9) the substrate refinement passes the generator-level descent check $E_{m+1} S_m = W E_m$.

Then the W_7 candidate passes χ Gate 1, χ Gate 2 and χ Gate 3, and the χ increment halves across generations:

$\Delta\chi_{g+1} = \frac{1}{2} \Delta\chi_g$ ($\rho = \frac{1}{2}$),

with the ½ supplied by Gate 2's two equalities $\omega_I = \omega_Q$ and $\Delta\ell_I = \Delta\ell_Q$ (giving $w = \frac{1}{2}$, from (7)–(8)) and transported by Gate 3's family $\chi' = w\chi$. Assumptions (2)–(4) are the §8 lemma hypotheses that upgrade Gate-1 bridges (b)–(d) from free choices to theorems; (1) is the one bare structural identification; (5)–(9) are finite, executable audits. (This numbering matches the flat ledger below item-for-item.)

The complete dependency ledger (stated flatly, so nothing reads smaller than it is). §8 turns three Gate-1 bridges into lemmas, which is real progress — but the lemmas do not eliminate their assumptions, they convert free choices into *named requirements*, and "natural" then does quiet work in several places. The honest count is **one structural assumption + three natural-but-unverified operational premises + two finite audits (the Gate-2 audit itself comprising four sub-conditions)**, plus one upstream interpretive bridge inherited from the reduction. In full:

- **(1) Carrier** *(structural assumption)* — the physical persistent carrier is the W_7 cycle space $H_1(W_7)$.
- **(2) Faithfulness** *(operational premise; see the post-hoc caveat below)* — the fold representation is faithful to the C_6 rim action. Selects E_1 (Lemma 8.1).
- **(3) Source–target transport** *(operational premise)* — completion is the reversible transport of admissible input support (Theorem 8.2).
- **(4) Access axioms** *(operational premise)* — accessibility is positive, linear, support-local, P-invariant, normalised (Theorem 8.3).
- **(5) Rewrite generators** *(audit)* — the primitive rewrite generators are exactly {(12),(23)}, with no direct reversal generator (hypothesis of the Prop. 2.2 audit).
- **(6) Terminal-history axioms** *(audit)* — every admissible terminal history satisfies $T J = J T$, $T^2 = T$, $T|_{\mathcal{F}_e} = I$ (hypotheses of Thm 2.4).
- **(7) Uniform class weight** *(audit/premise)* — $\omega_I = \omega_Q$. **This is a third Gate-2 premise, on the same footing as the ledger cost, not a corollary of binarity** (binarity fixes the *number* of classes, not the *measure* on them; see §10).
- **(8) Ledger cost** *(audit)* — $\Delta\ell_I = \Delta\ell_Q$ ($\kappa_I = \kappa_Q$).
- **(9) Descent** *(audit)* — the substrate refinement passes $E_{m+1} S_m = W E_m$ on generators.

Grouped: (1)–(4) carrier and operational semantics; (5)–(8) the Gate-2 finite audit; (9) the Gate-3 finite audit. The advance over the prior framing is that **all of these are now named, and (5)–(9) are finite and executable** — not that there are only three. No input is hidden, none is claimed proved here, and the theorem asserts **conditional realisation on the candidate, not physical proof:** granting (1)–(9), the χ-halving ratio ½ follows by finite, checked computation.

Scope: what is, and is not, about quarks (read this before the gate machinery). The headline object is the quark ratio $\rho = \Delta\chi_2/\Delta\chi_1$, but the χ this paper actually manipulates is an **abstract odd log-access readout** on the W_7 access cone. Granting all of (1)–(9), what has been derived is the

halving of that internal contrast functional: $\Delta\chi_{g+1} = \frac{1}{2} \Delta\chi_g$ for the functional χ defined here. That this functional *is* the physical quark susceptibility is **assumed upstream** — it lives in the reduction's interpretive bridge together with the carrier identification (1), and it is **examined nowhere in this paper**. This is a deliberate scope decision, but it inverts the apparent weight of the open questions: pages of gate machinery sit downstream of a single, least-scrutinised identification ("this χ = the physical χ"). The physically substantive claim is that identification; the gate audits are conditional structure built on top of it. Nothing here tests the link to quark masses.

Structure. *Part I* (§§1–2) inherits the reduction and fixes the three gate definitions. *Part II* (§§3–9) audits Gate 1 — the Dihedral Automatic-Pass Lemma proves the fold sector real once the W_7 mirror-pair is granted, and §8 (Lemmas 8.1–8.3) reduces bridges (b)–(d) to natural requirements, leaving (a) the carrier as the one foundational input. *Part III* (§10) audits Gate 2 — twinning (with the observable split), an executed nonadjacency rewrite audit, and a *classified* binary census (Thm 2.4) — isolating Cost Neutrality as the ledger equality $\Delta\ell_I = \Delta\ell_Q$. *Part IV* (§11) derives the contrast-transport family $\chi' = w\chi$, pins the inner mean by the regularised Geometric-Mean Selection Lemma, takes $w = \frac{1}{2}$ from Gate 2, and isolates Descent as a generator audit. §12 assembles the conditional realisation; §13 records the separate transport/locality sequence; §14–§15 give the ledger and conclusion.

Table of Contents

Part I — Inheritance and the Three Gates

1. Inheritance from the Reduction

The reduction reduced $\rho = \frac{1}{2}$ to a nondegeneracy prerequisite (0) and six conditions on an up/down fold gate, and located the entire remaining burden in a structural audit of whether the physical gate satisfies them. Three inherited facts frame this paper; they are stated, not re-derived (see the reduction for proofs).

The fold sector and its two readings. From closure primitives A, B, C, the oriented words U = ABC and D = CBA span a candidate fold space $\mathcal{F} = \mathrm{span}\{U, D\}$, with reversal involution J(U) = D and even/odd lines $\mathcal{F}_e = \mathrm{span}\{U + D\}$, $\mathcal{F}_o = \mathrm{span}\{U - D\}$; the odd line U − D is what the susceptibility χ reads. The sector is **genuine** iff U, D are linearly independent (operator Gram det ≠ 0; order-sensitivity K = U − D ≠ 0 is necessary but not sufficient). The reduction gives two non-interchangeable realisations: the **word-span** (Route B) and the operator **supports** $\mathcal{F}_u = \mathrm{Ran}(U)$, $\mathcal{F}$_d = Ran(U†) (Route A).

The two-headed Gate-1 audit, selected by rank(U). Whether the fold sector is real is decided by one upstream computation, r = rank(U): if r < dim $\mathcal{H}$ (closure loses rank) Route A is available and the strongest configuration (Option 3) with it; if r = dim $\mathcal{H}$ (invertible closure) the supports collapse (P_u = P_d = I) and Route B is forced, with the entire physical burden on the access-row identification. The full pass/fail specification is the reduction's Gate-1 spec, recapped compactly in Appendix A.

What the reduction left open, and this paper supplies. The reduction flagged r = rank(U) as uncomputed pending concrete A, B, C, and noted that supplying them — together with an independently specified physical access map E_m — converts its reference run into a physical verdict, *provided* E_m is not read off the observed masses (the circularity trap). This paper supplies concrete operators generated from the W_7 closure cell, computes the verdict in closed form, and is explicit about what the verdict does and does not establish. It does **not** re-open the reduction's conditional theorem, the six conditions, or the §6L analysis; those stand.

Two scope fences carried from the reduction hold throughout. **Gate 1 establishes only that the fold sector is real and is encoded into** — it is silent on twinning, census granularity, and branch-cost (Gate 2), and on full log-access intertwining (Gate 3). And the deliverable is a *verdict-or-a-named-dependency*: where a canonical substrate input is not in hand, the honest output is a flagged dependency, graded as such.

---

2. The Instantiation Fork — and Two Grades of Instance

The audit needs (A, B, C) as concrete operators. The reduction's fork stands: either the substrate fixes (A, B, C) canonically (case I), or it fixes only the structural roles — admissibility excludes inadmissible distinctions, transport carries the fold-exchange, completion closes — and the matrices are not yet uniquely pinned (case II). This paper proceeds under (II), the honest weaker reading. But within (II) there are **two grades of instance**, and the distinction is this paper's substance:

- **Generic instance** — three role-faithful but otherwise arbitrary operators, chosen to make the protocol executable and to demonstrate that its witnesses *can fail* (a generic encoder leaves the word-span; distinct supports are a genuine fact about specific operators). This certifies the protocol; it makes no claim of canonicity. The reduction's reference run was of this grade; it is recapped here only where it still does work the candidate cannot (Appendix C).

- **Symmetry-canonical instance** — operators *generated* from a structure already in the VERSF corpus (the W_7 closure cell and its D_6 symmetry), so the choice of (A, B, C) is dictated by symmetry rather than picked. This is far less arbitrary, and it is the primary instance audited here (§3–§8).

The upgrade from generic to symmetry-canonical is real progress, but it carries a cost that §3.1 makes a *theorem* and §8 makes explicit, and that the reader should hold from the outset: **as the instance becomes less arbitrary, the audit becomes less discriminating**, because the structural facts Gate 1 tests migrate from the operators into the symmetry identification that generates them. By the time the instance is fully symmetry-canonical, Gate 1 passes automatically — the Automatic-Pass Lemma shows it passes for an *entire class* of instances at once — and its content lives entirely in the identification. The candidate is therefore stronger as a *construction* and weaker as a *test* — and both must be stated.

Part II — Gate 1: the W_7 Fold-Sector Realisation

3. The W_7 Symmetry-Canonical Candidate

Carrier. The wheel W_7 has one hub and six rim vertices (7 vertices, 12 edges: 6 spokes, 6 rim edges). The construction lives on its **cycle space** — the kernel of the vertex–edge boundary map,

$\ker \partial = Z_1(W_7; \mathbb{C}) = H_1(W_7; \mathbb{C})$, $\dim = |E| - |V| + 1 = 12 - 7 + 1 = 6$.

Since W_7 has no 2-cells, the cycle space is the graph homology H_1; the finite-graph (Hodge) inner product identifies it canonically with the cohomology $H^1(W_7; \mathbb{C}) \cong H_1(W_7; \mathbb{C})$, and we write H^1 for the carrier throughout *under that identification* (the operators below are self-adjoint with respect to that inner product). Nothing depends on the distinction beyond fixing the inner product.

The six triangular bounded faces form a basis, and the 60° rim rotation R permutes them cyclically. So R acts on H^1 as the regular representation of C_6, with spectrum the six sixth-roots of unity, each once:

$\mathrm{spec}(R) = \{1, e^{+i\pi/3}, e^{-i\pi/3}, e^{+2i\pi/3}, e^{-2i\pi/3}, -1\}$.

Under the full rim symmetry D_6 this decomposes into the two rotation-fixed modes (eigenvalues ±1) and two two-dimensional sectors: $\mathbf{E_1}$ = the $e^{\pm i\pi/3}$ pair and $\mathbf{E_2}$ = the $e^{\pm 2i\pi/3}$ pair. Let S be a D_6 reflection, $S^2 = I$, $SRS = R^{-1}$, so S exchanges each conjugate pair. A structural fact we will use in §8: R restricted to E_1 has **order 6** ($e^{i\pi/3}$ is a primitive sixth root), while R restricted to E_2 has **order 3** ($e^{2i\pi/3}$ is a primitive cube root). This is the one invariant that tells the two doublets apart, and the construction below does not see it.

Projectors and primitives. Define the spectral projectors onto the rotation eigenlines,

$$P_\pm = (1/6) \Sigma_{n=0}^{5} e^{\mp i n \pi/3} R^n,$$

which are orthogonal rank-one projectors with $S P_+ S = P_-$. On the E_1 doublet set the **self-adjoint** closure primitives

$A = P_+$ (admissibility), $B = S$ (transport), $C = P_-$ (completion),

and form the ordered closure words $U = A B C$, $D = C B A$.

3.1 The Dihedral Automatic-Pass Lemma

The W_7 witnesses are not a numerical coincidence at machine precision; they are the closed-form output of a small representation-theoretic lemma whose only hypothesis is "a two-dimensional dihedral irrep." Stating it once, in generality, does three things: it makes every §5 witness exact, it explains in advance why E_1 and E_2 pass identically (§8b), and it pins down precisely what Gate 1 is and is not detecting.

Lemma (Dihedral fold-sector automatic pass). Let a finite-dimensional Hilbert space $\mathscr{H}$ carry a unitary rotation R and a unitary reflection S with $S^2 = I$ and $SRS = R^{-1}$. Suppose $\mathscr{H}$ contains a two-dimensional R-invariant subspace on which R has *distinct* conjugate eigenvalues $e^{\pm i\theta}$ ($\theta \not\equiv 0 \mod \pi$), with orthonormal eigenvectors $|v_+\rangle$, $|v_-\rangle$, and on which S exchanges the two eigenlines: $S|v_-\rangle = e^{i\psi}|v_+\rangle$ for some phase ψ. Let $P_\pm = |v_\pm\rangle\langle v_\pm|$ be the rank-one spectral projectors. Set $A = P_+$, $B = S$, $C = P_-$, $U = ABC$, $D = CBA$. Then, **exactly**:

1. **(Words are conjugate outer products.)** $U = \alpha|v_+\rangle\langle v_-|$ and $D = \bar{\alpha}|v_-\rangle\langle v_+|$, where $\alpha = \langle v_+|S|v_-\rangle$ and $|\alpha| = 1$.
2. **(Self-adjoint closure.)** $D = U^\dagger$. *Proof:* A, C are self-adjoint projectors and B = S is self-adjoint ($S = S^{-1}$ unitary involution), so $U^\dagger = C^\dagger B^\dagger A^\dagger = CBA = D$.
3. **(Rank one, distinct supports.)** rank $U = 1$, $UU^\dagger = P_+ = A$, $U^\dagger U = P_- = C$, and $A \neq C$.
4. **(Order-sensitivity, exact.)** $\|U - D\|_{HS} = \sqrt{2}$; $\langle U, D\rangle_{HS} = 0$; $\|U\|_{HS} = \|D\|_{HS} = 1$, so $\det \mathrm{Gram}(U, D) = 1$.
5. **(Exchange involution realised by S.)** $SAS = C$ and $S^2 = I$, so the reduction's exchange involution J ($J P_u J = P_d$) is realised on the support pair by the candidate's own transport operator.

Proof. Since R is unitary with distinct eigenvalues on the doublet, $\langle v_+|v_-\rangle = 0$. Then $U = P_+ S P_- = |v_+\rangle\langle v_+| S |v_-\rangle\langle v_-| = \langle v_+|S|v_-\rangle \cdot |v_+\rangle\langle v_-| = \alpha|v_+\rangle\langle v_-|$, giving (1); S unitary with $S|v_-\rangle = e^{i\psi}|v_+\rangle$ forces $|\alpha| = |\langle v_+|e^{i\psi} v_+\rangle| = 1$. (2) is the displayed two-line computation. For (3), $UU^\dagger = \alpha|v_+\rangle\langle v_-| \cdot \bar{\alpha}|v_-\rangle\langle v_+| = |\alpha|^2 |v_+\rangle\langle v_+| = P_+$, and symmetrically $U^\dagger U = P_-$; $A \neq C$ because $P_+ \neq P_-$. For (4), $\langle U, D\rangle_{HS} = \mathrm{Tr}(U^\dagger D) = |\alpha|^2 \langle v_+|v_-\rangle \mathrm{Tr}(|v_-\rangle\langle v_+|) = 0$, and $\|U\|^2_{HS} = \mathrm{Tr}(U^\dagger U) = |\alpha|^2 = 1$, so $\|U - D\|^2 = 1 + 1 - 0 = 2$ and the Gram matrix is the identity. (5) is $SAS = S P_+ S = P_- = C$. ■

Three consequences. First, **every §5 witness is exact** — √2, det Gram = 1, D = U†, the distinct supports — with no appeal to a 10^{-16} residual; the floating-point run of Appendix D merely confirms the lemma's arithmetic in a particular basis. Second, the hypotheses are satisfied by **both** E_1 and E_2 (each is a 2D dihedral irrep), so Route-A Gate 1 passes **bit-identically** on the two doublets — a corollary, not a separate numerical observation (this is bridge (b), §8). Third, and most sharply: the lemma's hypotheses make **no reference to θ**. Gate 1 reads only the internal exchange-quotient structure (two orthogonal lines + a reflection that swaps them), and *that data* is isomorphic across all such doublets — even though E_1 and E_2 are distinct, non-isomorphic D_6-irreps. The datum that distinguishes them — the *order* of R on the doublet, $\mathrm{ord}(R|E_1) = 6$ versus $\mathrm{ord}(R|E_2) = 3$ — is a property of how the doublet sits inside the *ambient* H^1, and is structurally invisible to a test built from P±, S alone. (This is exactly what Lemma 8.1 exploits: the ambient action distinguishes the sectors; the gate quotient forgets the distinction.)

Why this is a candidate, not the canonical primitives. The construction is *generated* by W_7/D_6 symmetry, which is genuine progress over a generic instance — there is no fine-tuning, and the operators are forced once the doublet and the reflection are fixed. But the Automatic-Pass Lemma is exactly the precise statement of the residual: the symmetry produces *two* eligible doublets, both satisfy the lemma's hypotheses, and the identification of admissibility/completion with the two reflection-halves of *one* of them is an input, not an output. §8 isolates these as named bridges. Here we record only that the primitives are self-adjoint and symmetry-canonical — which, on its own, favourably settles the self-adjointness question a generic instance can only assume (lemma item 2).

---

4. The Rank-Selecting Test and the Route Verdict

The selecting computation is r = rank(U) via SVD. On the W_7 candidate, lemma item (1) gives the answer in closed form:

$U = P_+ S P_- = \alpha |v_+\rangle\langle v_-|$, $\alpha = \langle v_+|S|v_-\rangle$, $\mathrm{rank}(U) = 1 < \dim H^1 = 6$.

Closure loses rank — maximally, to rank one. The verdict, by the Gate-1 spec (Appendix A, C.0):

$r < n \Rightarrow$ **Route A available** (support subspaces), and **Option 3 available** (audit in A, exposit in B, certified by the induction bridge).

Two remarks on what rank-one means here. First, it is the *cleanest* possible rank-deficient case: the two fold orientations are single lines, $\mathcal{F}_u = \mathrm{Ran}(U) = \mathrm{span}\{|v_+\rangle\}$ and $\mathcal{F}$_d = Ran(U†) = $\mathrm{span}\{|v_-\rangle\}$, and the support story is correspondingly transparent. Second, and less comfortably, the Automatic-Pass Lemma now makes precise *how little* room the witnesses have to fail: at rank one with a reflection-conjugate pair, distinct supports and the exchange involution are not "near-immediate consequences" but **theorems** of the construction class. The rank verdict is therefore not an independent fact about opaque operators but a direct readout of "A and C are the two lines of one reflection-doublet." It is recorded honestly in the ledger as *entailed by the identification*, not as a discriminating measurement.

---

5. Route-A Gate-1 Audit — the W_7 Witnesses (Now Closed-Form)

With Route A selected and the self-adjoint preconditions met ($D = U^\dagger$ by lemma item 2), the three Route-A witnesses are read off the lemma. Each is exact; the Appendix-D residuals ($\approx 10^{-16}$) are confirmation, not the source of the value.

A.1 — Order sensitivity ($U \neq D$).

$\|U - D\| = \sqrt{2} = 1.41421356\ldots > \text{tol} \Rightarrow$ **PASS.**

(By lemma item 4: $U = \alpha|v_+\rangle\langle v_-|$ and $D = \bar{\alpha}|v_-\rangle\langle v_+|$ are HS-orthogonal unit-norm outer products, so $\|U - D\|^2 = 1 + 1 - 0 = 2$.) Necessary, not sufficient.

A.2 — Distinct support subspaces. The support projectors are the closure words' own products (lemma item 3):

$P_u = U U^\dagger = P_+ = A$, $P\_d = U^\dagger U = P_- = C$, rank P_u = rank $P\_d = 1$, $P_u \neq P\_d$ (since $P_+ \neq P_-$) $\Rightarrow$ **PASS (distinct range, no collapse).**

The forward and reverse supports are the two distinct rotation eigenlines. (The collapse failure $P_u = P\_d$ is what invertible closure would force; it does not occur.)

A.3 — Exchange involution. The reflection S *is* the exchange (lemma item 5): $S P_+ S = P_-$ realises $J P_u J = P\_d$ on the support pair, with $S^2 = I$. Equivalently, the polar phase V of U ($U = V|U|$, $V = |v_+\rangle\langle v_-|$ the partial isometry $\mathcal{F}\_d \to \mathcal{F}_u$) coincides with S restricted to the E_1 sector up to phase, so the doubled-carrier involution of the spec is realised by the candidate's own transport operator $\Rightarrow$ **PASS.** The even/odd projectors $\Pi_e = \frac{1}{2}(I + J)$, $\Pi_o = \frac{1}{2}(I - J)$ act on span{U, D}.

Route-A Gate 1: PASS (A.1 ∧ A.2 ∧ A.3) on the W_7 candidate — two distinct fold lines, exchanged by the D_6 reflection, with order-sensitivity confirmed. **[Proven on candidate, conditional on bridges (a)–(c).]**

What this establishes: two distinct fold subspaces exchanged by S. What it does **not** establish: that physical residues live in them (Gate 3), twinning (Gate 2), weights (Gate 2 census), or that W_7/E_1 is the *unique* substrate choice (§8). §9 holds that line.

---

6. The Induction Bridge — the Word-Span Reading Agrees

Because $r < n$, Option 3 is available: run the support test (§5) as the audit, present the fold space as the word-span $\mathcal{F} = \text{span}\{U, D\}$, and certify the two readings agree via the polar phase. The word-span genuineness condition is linear independence of U, D, tested by the operator Gram determinant under $\langle X, Y\rangle = \text{tr}(X^\dagger Y)$:

$G = [[\langle U,U\rangle, \langle U,D\rangle], [\langle D,U\rangle, \langle D,D\rangle]] = [[1, 0], [0, 1]]$, $\det G = 1 \neq 0 \Rightarrow$ U, D linearly independent.

This is the strongest possible word-span pass, and lemma item 4 makes it analytic: U and D are not merely independent but **Hilbert-Schmidt-orthonormal** (the orthogonal outer products $\alpha|v_+\rangle\langle v_-|$ and $\bar{\alpha}|v_-\rangle\langle v_+|$), so the Gram matrix is the identity and $\det G = 1$ exactly. It is also strictly stronger than order-sensitivity: $K = U - D \neq 0$ ($\|K\| = \sqrt{2}$) is necessary but not sufficient, while $\det G = 1$ certifies genuineness. The agreement is explicit:

($P_u \neq P\_d$, Route A) $\Leftrightarrow$ ($\det G \neq 0$, Route B): **consistent.**

The word-span even/odd lines are nontrivial and complementary ($\|U + D\| = \|U - D\| = \sqrt{2}$), so the orientation-distinguishing content $U - D$ the readout reads is present and bounded away from zero. The candidate sits in the strongest configuration: a finite support audit, an agreeing word-span exposition, $\det G = 1$, and a certified bridge.

---

7. The Access-Row Identification — and Why It No Longer Discriminates Here

The reduction located the real Route-B Gate-1 content in the **access-row identification** (spec test B.3, Appendix A): exhibiting the physical encoder $E\_m : X\_m \to \mathcal{F}$ mapping the substrate's access state into the fold space, with (i) $\text{Im}(E\_m) \subseteq \text{span}\{U, D\}$ and (ii) $E\_m \hat{J}\_m = J E\_m$, on a generating set, for all m, with E_m **independently specified** — not read off the masses.

The candidate encoder. The W_7 candidate supplies a mass-blind encoder built from the pre-existing trace-accessibility readings $a\pm(R) = \text{Tr}(P\pm R)$:

$E\_m(R) = a_+(R)\cdot U + a_-(R)\cdot D$, with odd readout $\chi(m) = \ln a_+(R\_m) - \ln a_-(R\_m)$.

It satisfies both B.3 conditions, and the second is now seen to be a corollary of the lemma's mirror identity. Condition (ii) holds because

$a_+(S R S) = \text{Tr}(P_+ S R S) = \text{Tr}(S P_+ S \cdot R) = \text{Tr}(P_- R) = a_-(R)$

(using SAS = C, i.e. $S P_+ S = P_-$, and cyclicity of the trace), and symmetrically $a_-(SRS) = a_+(R)$, so

$E\_m(S R S) = a_-(R) U + a_+(R) D = J E\_m(R)$.

The numerical run confirms this (reversal-equivariance error ≈ 8×10^{-15}), but the equality is exact.

Three honesty corrections, all load-bearing. First, condition (i) is **automatic, hence vacuous here**: the encoder is *defined* as a combination of U and D, so its image lies in span{U, D} by construction and cannot fail. The genuine content is in (ii), reversal-equivariance.

Second — and this is the new tightening — (ii) is **not independent content either**: the computation above shows it follows entirely from SAS = C, which is bridge (c). So the encoder, on this candidate, contributes exactly **one** genuinely new dependency, the non-circularity of a±, and its structural test (ii) is downstream of a bridge already on the books. A discriminating B.3 — one in which (i) is a real constraint — needs a *generic* encoder against which to fail; that demonstration is exactly what the generic reference model still provides (Appendix C), where an arbitrary encoder leaves the 2-dimensional word-span with image-residual ≈ 5.2 in a 16-dimensional operator space and correctly fails (i). The candidate cannot demonstrate that B.3 discriminates, because its encoder is constructed to pass (i); the reference can and does. Both are needed, for different jobs.

Third, the *non-circularity* of the candidate encoder rests **entirely on a corpus citation**: that a±(R) = Tr(P± R) is the independently-motivated accessibility functional of the prior programme, not a reading chosen to make $\chi = \ln a_+ - \ln a_-$ come out odd. If that citation holds, the encoder is genuinely mass-blind and the construction is clean. If a± were reverse-engineered from the desired χ, it is precisely the circularity trap the reduction named. This is gradeable only against the accessibility corpus, and it is flagged as the load-bearing bridge (d) of §8 — not settled here.

So under Route A (the candidate's verdict) Gate 1 is already cleared by §5 without the encoder; the encoder enters as the Gate-3 half (B.3 is "the encoder half of Gate 3"), and its status is: (ii) passes, entailed by (c); (i) vacuous; non-circularity **now derived, not cited** — §8's Trace-Access Uniqueness Theorem (8.3) shows the four access axioms force a±(R) = Tr(P±R), so mass-blindness is a theorem conditional on those axioms rather than a corpus citation.

---

8. Where the Verdict Now Lives — Four Identification Bridges

The W_7 candidate passes Gate 1, but — per §2, §3.1 and §4 — **the pass is entailed by the symmetry identification, not an independent corroboration of it.** The Automatic-Pass Lemma makes this a theorem: once one fixes "the carrier is $H^1(W_7)$, the fold sector is a reflection-conjugate doublet, and admissibility/completion are its two reflection-halves," every Gate-1 witness follows. So the honest content is "Gate 1 has been *reduced to* four identification bridges." This section now goes further: **three of the four bridges are upgraded from assumptions to lemmas**, each conditional only on a natural requirement, leaving (a) the carrier identification as the single foundational input.

(a) The persistent carrier is the W_7 cycle space. The whole construction lives on $H_1(W_7) \cong H^1(W_7)$ (the Hodge-identified cycle carrier of §3). That the physical persistent carrier *is* this cycle space — not merely an abstract six-dimensional space, or another closure cell's carrier — is assumed, not derived here. It is the foundational bridge, and the one this section does not eliminate; the lemma and the upgrades below all take the carrier as given.

(b) The fold doublet is E_1 — now a lemma (Primitive Faithfulness), not a free selection. E_1 and E_2 are **distinct, non-isomorphic irreps of the full D_6 action** — their rotation characters differ (tr $R = 2\cos(\pi/3) = 1$ on E_1, $2\cos(2\pi/3) = -1$ on E_2). But Gate 1 does not see the full action: it reads only the **exchange-pair quotient** (two eigenlines of R, swapped by the reflection S), on which E_1 and E_2 induce *isomorphic* two-line-plus-reflection data. So the construction is **bit-identical at Gate 1** on either (verified, Appendix B), and — the point an earlier draft missed — the same blindness extends downstream: the complete Gate-2 and Gate-3 audits of Parts III–IV (twinning, $d = 3$, the binary census, and $\chi(\Phi R) = \frac{1}{2}\chi(R)$) also transfer bit-identically to E_2, since each reads only that quotient. **No gate calculation discriminates the doublet** — the gates forget precisely the ambient datum (the rotation order) that distinguishes the sectors. That ambient datum is the discriminator:

Lemma 8.1 (Primitive Faithfulness). Require that the fold representation be **faithful to the full six-step rim group C_6** — the physically natural no-aliasing requirement that *no non-identity element of C_6 acts trivially*. (This is sharper than "a rim move cannot act trivially": on E_2 the one-step R is non-trivial, acting as e^{±2iπ/3}; the failure is at the order-2 element.) On E_1, R has order 6 and C_6 is represented faithfully; on E_2, R has order 3, so $R^3 = I$ — the physical half-turn, a non-identity element of C_6, acts trivially, and the representation is **unfaithful** (verified: R^3 acts as -1 on E_1, $+1$ on E_2). Only E_1 realises the full six-step closure order. Hence faithfulness selects E_1 uniquely.

Honest grade of the lemma. Faithfulness is the **exact invariant that separates E_1 from E_2** — and it is introduced here precisely to break the tie the three gates cannot. That is suspicious on its face: a tie-breaker chosen for its outcome is post-hoc, however natural it sounds. The lemma escapes that charge **only** if the corpus motivates a faithful rim action *independently* — for a reason stated before anyone knew it selected E_1 (e.g. that the substrate's six-step closure order is physically primitive and must be represented without aliasing). If such a motivation exists it should be cited at this point of use; this paper does not have it in hand. So, graded honestly, **faithfulness is a natural-*looking* tie-breaker whose independent justification is not established here** — better than a free choice (it is a single, sharply-stated requirement), but not yet "physically natural" in the sense of being motivated prior to its selecting role. **[Premise — tie-breaker; independent corpus motivation outstanding.]**

So bridge (b) is no longer a free selection between two identical-passing options; it follows from no-aliasing. It remains conditional on the faithfulness requirement, but that is a single, physically natural premise about the representation — and the honest point stands that faithfulness, not any gate computation, is what separates the doublets. **[Lemma — conditional on the no-aliasing/faithfulness requirement.]**

(c) Admissibility and completion are reflection-conjugate — now derived from reversible transport, not from non-degeneracy. The bare necessity argument (a non-conjugate pairing gives rank U = 0, $\langle v_+|S|v\_other\rangle = 0$; verified, Appendix B) is real but is not a physical derivation — it shows the pairing must be conjugate for the word not to vanish, not why the operational roles make it so. Operational meaning supplies the why:

Theorem 8.2 (Source–Target Support). For a reversible fold transport S ($S^2 = I$), the admissible *input* sector is the initial support P_adm, and the completed *output* sector is its transported final support; hence P_comp = S P_adm S. With A = P_adm = P_+ and S the D_6 reflection, C = P_comp = S P_+ S = P_-.

So A = P_+ and C = P_- follow from the operational meanings — admissibility = input support, transport = the reversible map S, completion = the S-transported input support — not from the requirement that ABC not vanish. Bridge (c) is thereby reduced to the single identification "completion is the S-transport of admissibility," which is what reversible fold transport means. **[Theorem — conditional on that operational reading of transport.]**

(d) The access reading Tr(P±R) is the accessibility functional — now a uniqueness theorem, not a citation. The encoder (§7) is mass-blind iff a±(R) = Tr(P±R) is the independently-motivated accessibility functional. A short uniqueness theorem replaces the citation:

Theorem 8.3 (Trace-Access Uniqueness). Let a_P be an access functional that is (1) positive and linear in R, (2) support-local — a_P(R) = 0 whenever PRP = 0 — (3) invariant under unitaries preserving P, and (4) normalised, a_P(P) = 1. Then a_P(R) = Tr(PR). *Proof.* Positivity and linearity give a_P(R) = Tr(FR) with F ≥ 0; support-locality forces F to be HS-orthogonal to every R with PRP = 0, hence F = PFP; for rank-one P this is F = cP, and normalisation gives c = 1. ■

Applied to P±, this *derives* a±(R) = Tr(P±R) from four access axioms, none of which mentions χ — the circularity trap is closed by uniqueness, not citation. **[Theorem — conditional on the four access axioms being the operative accessibility law.]**

The shape of the residual — reduced. Three of the four bridges are now lemmas: (b) follows from faithfulness (Lemma 8.1), (c) from reversible transport (Theorem 8.2), (d) from the access axioms (Theorem 8.3). What remains is **(a)** — the identification of the physical persistent carrier with the W_7 cycle space — the single foundational structural input this audit does not derive, together with the two physical premises of Parts III–IV (Cost Neutrality and Descent). So the candidate's standing is the conjunction of (a) + Cost Neutrality + Descent, with (b), (c), (d) discharged to natural requirements rather than left as free choices. The original six inputs are thereby reduced toward three — one carrier identification and two audits — and, should the corpus derive the W_7 carrier itself, to the two audits alone. That is the precise, bounded, and now substantially smaller standing of the W_7 candidate.

---

9. What Gate 1 Does and Does Not Decide

The cleanest statement of the verdict is as a **conditional with a discharged side-condition**:

(a) ∧ (b) ∧ (c) ∧ (d) ⟹ the fold sector is real and is encoded into, and the W_7 candidate makes the antecedent concrete, finite, and named, while discharging the reduction's self-adjointness side-condition ($D = U^\dagger$) as a theorem rather than an assumption.

Unpacking the three registers:

Decided (on the W_7 candidate, conditional on bridges (a)–(c)). Granting the carrier, the E_1 choice, and the mirror-pair identification, the fold sector is real: closure loses rank ($r = 1 < 6$), the two fold lines are distinct, exchanged by the D_6 reflection, the word-span reading agrees ($\det G = 1$), and the odd sector $U - D$ is nontrivial. The audit is a finite computation passing exactly (closed form) and to machine precision (Appendix D). **[Proven on candidate, conditional on (a)–(c).]**

Favourably resolved. Self-adjointness — which a generic instance can only *assume* to make Route A well-posed — is here a *property* of the symmetry-canonical primitives (lemma item 2): A, C are orthogonal projectors and B is a real reflection, so $D = U^\dagger$ holds by construction, a two-line consequence. The candidate exhibits a genuinely self-adjoint realisation, removing the well-posedness caveat for this instance. **[Proven.]**

Reduced to named inputs, three now lemmas (for the physical gate). §8 supplies three of the four bridges as theorems conditional on natural operational axioms: (b) E_1 by Primitive Faithfulness (Lemma 8.1), (c) the reflection-conjugate pairing by Source–Target Support (Theorem 8.2), (d) the trace reading by Trace-Access Uniqueness (Theorem 8.3). What this audit does **not** supply is (a), the carrier identification, which remains a bare structural assumption; the physical Gate-1 verdict rests on (a) plus the three operational axioms. **[(a) Open — structural; (b),(c),(d) Lemma/Theorem, conditional on faithfulness / transport / access axioms.]**

A genuine upside. Pinning the carrier to a concrete graph turns the *remaining* gates into finite computations rather than abstract conditions. On W_7/D_6 one can now pose, as explicit checks: **twinning** (are the two E_1 lines reversal-fixed twins with equal external closure neighbourhoods on the W_7 closure-order graph?); **nonadjacency** (the graph distance between the orientation words, the reduction's $d(U, D) = 3$ analogue); **census granularity** (the D_6 orbit structure on the cell, and whether the operational quotient is class-level); and the **intertwining naturalities** of Gate 3. These were verbal conditions in the reduction; on the W_7 candidate they are computations — and §10 (Gate 2) and §11 (Gate 3) now carry them out, with Gate 2 a conditional pass and Gate 3 a constructive exact-½ pass, assembled into the conditional realisation of §12. That, not the Gate-1 pass, is the candidate's real prize, and it is where the test *discriminates* again because it constrains structure the symmetry identification does not already fix. (A separate transport/locality sequence — orientability and the refinement-loop trace τ — is demarcated in §13 and does not bear on χ closure.) (In particular, twinning and census granularity *do* depend on the ambient embedding — and so are exactly the place where the E_1/E_2 ambiguity of bridge (b) may finally be broken.)

Out of scope, by construction. Gate 1 remains silent on twinning, census granularity, and branch-cost (Gate 2), and on full log-access intertwining (Gate 3, of which B.3 is the encoder half; the S_m circularity trap stands). Passing Gate 1 means the fold sector is real and encoded into. It does not mean the gate halves. Halving is Gates 2 and 3.

Part III — Gate 2: Twin, Census, and the Cost Premise

10. Gate 2 on W_7/E_1 — Conditional Pass

Gate 2 asks four things of the fold pair: that the two folds are operational twins, that the orientation words are non-adjacent, that the operational census is binary, and that there is no branch-cost tilt between the folds. On the W_7/E_1 candidate the first three are theorems; the fourth reduces to a single named physical premise. This is the second link of the conditional-closure chain — main content, not forward status.

Proposition 2.1 (Twinning — invariant observables cannot distinguish the folds). The algebra of D_6-invariant observables on $H^1(W_7)$ has dimension four: under D_6, $H^1 \cong \text{triv} \oplus \text{sign} \oplus E_1 \oplus E_2$ with each irrep at multiplicity one, so the invariant commutant has dimension $1^2 \cdot 4 = 4$. Every invariant observable O satisfies

$\mathrm{Tr}(P_+ O) = \mathrm{Tr}(P_- O)$ (exact analytically; $\approx 6 \times 10^{-16}$ numerically),

because $S P_+ S = P_-$ and $[O, S] = 0$ give $\mathrm{Tr}(P_- O) = \mathrm{Tr}(S P_+ S \cdot O) = \mathrm{Tr}(P_+ \cdot S O S) = \mathrm{Tr}(P_+ O)$. Equivalently, each invariant O acts as a single scalar on E_1 and so cannot separate its two eigenlines. The folds are therefore **operational twins** — indistinguishable by any D_6-invariant measurement, which is strictly stronger than being mere mirror images. **[Proven on candidate]**

The observable split (to forestall the obvious objection). If no observable distinguishes the lines, how can the mass ratio? The resolution is that twinhood and the χ-readout live in **two different algebras**, and the distinction must be formal. Twinhood and the census use the external closure algebra $O\_ext \subseteq \mathrm{Comm}(D_6)$ — the D_6-invariant observables, which are blind to the two eigenlines. The χ-readout is *not* in O_ext: it is the odd internal access functional ℓ_χ with $\ell\_\chi \circ J = -\ell\_\chi$, defined post-selection on the access cone. There is no contradiction: external closure labels are D_6-invariant (hence twin), while the sector-specific access *contrast* is odd under reversal and therefore lies outside O_ext. The mass ratio reads ℓ_χ, not O_ext.

Proposition 2.2 (Nonadjacency — an executed rewrite audit). Rather than *assume* the elementary rewrites are adjacent transpositions, build the one-step closure-order graph explicitly and audit it. With the closure letters at three positions and the primitive rewrites the adjacent swaps (12) and (23), the rewrite graph is the Cayley graph

$G\_ord = \mathrm{Cay}(S_3, \{(12), (23)\})$.

A breadth-first computation on G_ord (Appendix D) returns: all six arrangements reachable; diameter 3; and **d(ABC, CBA) = 3** — the reversal CBA = (13)·ABC sits at the antipode, since (13) = (12)(23)(12) has word-length 3. Crucially, (13) — the direct U↔D fold-flip — is **not** a generator (verified: it is not in {(12),(23)}, and its distance from the identity is 3, not 1), and no adjacent swap implements reversal in one step. So nonadjacency is not an assumption but a finite, exhaustive audit of the rewrite graph, conditional only on the primitive generator set being the adjacent swaps {(12),(23)}. **[Proven on the stated generator set; the generator set itself is flagged — see below.]** *(The one residual: that the substrate's primitive rewrites are exactly the adjacent swaps, with no extra primitive implementing (13), reversal, or a direct U↔D flip. The audit enumerates the consequences of that generator set exhaustively; deriving or enumerating the generator set from the corpus is the remaining check.)*

Proposition 2.3 (Binary census). The even projector $Q = \Pi_e = ½(I + J)$ is exactly idempotent ($Q^2 = Q$), reversal-equivariant ($[Q, J] = 0$), and annihilates the odd contrast ($Q \Pi_o = 0$). It collapses the closure tower to the two-element operational quotient

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

a binary class census of the required granularity. But idempotence of Q alone collapses a Q^n tower without yet excluding a *third* admissible history; the binarity should be a classification, not a definition:

Theorem 2.4 (Binary Terminal-History Classification). Let T be an admissible terminal one-step history on $\mathcal{F} = \mathcal{F}_e \oplus \mathcal{F}_o$ satisfying (i) reversal-equivariance $T J = J T$, (ii) terminality/idempotence $T^2 = T$, and (iii) even-sector preservation $T|_{\mathcal{F}_e} = I$. Then T commutes with J, hence is block-diagonal in $\mathcal{F}_e \oplus \mathcal{F}_o$; on the (one-dimensional) odd line it acts as a scalar μ with $\mu^2 = \mu$, so $\mu \in \{0, 1\}$; therefore $T = I$ ($\mu = 1$) or $T = \Pi_e = Q$ ($\mu = 0$). There are **at most two** admissible terminal histories, and since both I and Q are realised, exactly two:

$$H_\chi / \sim_{op} = \{[I], [Q]\}.$$

So binarity is derived from (i)–(iii), not posited. **[Proven, given (i)–(iii)]** *The cost of honesty:* this relocates the burden onto justifying that **all** admissible gate histories satisfy reversal-equivariance, terminal idempotence, and even-sector preservation — a stronger and more natural claim than simply declaring the history set to be {I, Q}, and one the commitment corpus must supply.

The one remaining premise — and what closing it actually requires. The branch weight that Gate 3 will need (Part IV) is

$$w = W_I / (W_I + W_Q),\quad W_{[c]} = \omega_{[c]} \cdot \kappa_{[c]},$$

where ω_I, ω_Q are class-level base weights and κ_I, κ_Q physical cost factors for the two classes [I], [Q]. The target $w = ½$ holds iff $W_I = W_Q$, i.e. iff $\boldsymbol{\omega_I = \omega_Q}$ **and** $\boldsymbol{\kappa_I = \kappa_Q}$ — and *both* equalities are premises, not just the second. **Binarity does not by itself give $\boldsymbol{\omega_I = \omega_Q}$ — this conflates the *cardinality* of the census with its *measure*.** Theorem 2.4 establishes that there are exactly two classes; it says nothing about their relative weight. The point is sharp: either ω is a **count of micro-histories per class**, in which case $\omega_I = \omega_Q$ requires the two classes to have equal size — which is *not shown* — or ω is **one-unit-per-class by definition of the census measure**, in which case uniformity holds by construction of that measure, not by the number of classes. Either way, $\omega_I = \omega_Q$ is a premise about *how weight attaches to classes*, on exactly the same footing as Cost Neutrality, not a corollary of *how many* classes there are. An earlier draft treated it as immediate from binarity; that was a genuine third Gate-2 premise hiding in plain sight, and it is promoted here to the visible ledger alongside $\Delta\ell_I = \Delta\ell_Q$. Closing it needs a **uniform-admissibility principle** (an equal-class-weighting theorem, cited with an applicability proof, or a direct audit of class sizes). The κ-equality is the ledger content: a J-invariant structural cost gives $\kappa_I = \kappa_Q$ automatically ($\ell(Q x) = \ell(I x)$), but symmetry does not select the functional — the operator norm makes I and Q equal ($\|Q\|_{op} = \|I\|_{op} = 1$) while Hilbert–Schmidt does not ($\|Q\|_{HS} = 1 \neq \sqrt{2} = \|I\|_{HS}$) — so $\kappa_I = \kappa_Q$ is a definite ledger computation, not a norm choice. Hence $w = ½$ rests on two equalities:

Cost Neutrality, as two equalities to audit. (i) **$\omega_I = \omega_Q$** — equal class base weights, from a uniform-admissibility principle (cited/derived, not from binarity alone). (ii) **$\Delta\ell_I = \Delta\ell_Q$** ($\kappa_I = \kappa_Q$) — equal incremental cost of the canonical representatives of [I], [Q] under the VERSF closure-state ledger ℓ, motivated by "one fold, one record" and becoming a calculation once ℓ is evaluated. Together they give $w = ½$.

Neither equality is closed here (no ledger values, no equal-weighting citation in hand); both are stated as the finite checks to run, not asserted. **[Conjectural — $w = ½$ rests on $\omega_I = \omega_Q$ (uniform admissibility) and $\Delta\ell_I = \Delta\ell_Q$ (ledger).]**

Gate 2 verdict: conditional pass. Twinhood (Prop. 2.1, with the observable split) holds outright. The remaining items are finite premises, honestly itemised: (G2-i) the primitive rewrite generators are exactly {(12),(23)} with no extra reversal generator (the hypothesis of the Prop. 2.2 audit); (G2-ii) every admissible terminal history satisfies $T J = J T$, $T^2 = T$, $T|_{\mathcal{F}_e} = I$ (the hypotheses of Thm 2.4) — to be derived from the closure-state machine or carried as premises; and (G2-iii) the two weight equalities $\omega_I = \omega_Q$ and $\Delta\ell_I = \Delta\ell_Q$ that fix $w = ½$. **[Conditional pass — modulo (G2-i) generators, (G2-ii) terminal-history axioms, (G2-iii) the two weight equalities.]**

Part IV — Gate 3: the Log-Access Intertwiner and the Descent Premise

11. Gate 3 on W_7/E_1 — Constructive Pass

Gate 3 asks for a mass-blind access update that intertwines with the doubling/halving map. The candidate supplies one explicitly, on the positive fold-access cone. The construction is deliberately staged so that the two gates' roles are visibly separated: **Gate 3 derives the**

contrast-transport family $\chi' = w\chi$ for a single weight w, and **Gate 2 supplies $w = ½$**. That separation is the conceptual core of this part.

The inner fold mean (exchange-symmetric). The inner pair a_+, a_- (with $a_\pm = \mathrm{Tr}(P_\pm R)$ on the cone $R = a_+ P_+ + a_- P_-$, $a_\pm > 0$) *is* exchanged by the fold symmetry J, so closing it by a symmetric mean is justified — and is forced:

Lemma 3.2 (Geometric-Mean Selection — the symmetric closure is the geometric mean). Let $m : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ close an unresolved *exchange-symmetric* access pair, with m (i) positive, (ii) degree-1 homogeneous ($m(\lambda a_+, \lambda a_-) = \lambda\, m(a_+, a_-)$), (iii) exchange-symmetric ($m(a_+, a_-) = m(a_-, a_+)$), (iv) multiplicative under independent access composition ($m(a_+b_+, a_-b_-) = m(a_+, a_-)\, m(b_+, b_-)$), and (v) **regular** (continuous — equivalently Borel-measurable, or monotone, or locally bounded — in each argument). Then $m(a_+, a_-) = \sqrt{(a_+ a_-)}$.

Proof. Set $f(x, y) = \log m(e^x, e^y)$. By (iv), f is additive: $f(x + x', y + y') = f(x, y) + f(x', y')$. Additivity plus the regularity (v) makes f **linear**, $f(x, y) = \alpha x + \beta y$ (regularity is essential — additive-but-discontinuous Cauchy solutions are otherwise possible). Exchange-symmetry (iii) gives $\alpha = \beta$; homogeneity (ii) gives $\alpha + \beta = 1$; hence $\alpha = \beta = ½$, i.e. $m = \sqrt{(a_+ a_-)}$. (Condition (v)-fixed-points $m(a,a) = a$ is then automatic.) ■

So the inner closure is $g = \sqrt{(a_+ a_-)}$, uniquely. *(Note: regularity here replaces the earlier loose "no additive odd offset"; the affine-offset condition belongs to the log-intertwining statement, not to this classification of multiplicative means.)*

The outer gate mean (weighted, because persistence ≠ closure). The outer pair is *not* the inner pair. It is (branch, closure) = (a_±, g) — persistence versus closure — and these are **not** exchanged by the fold symmetry: I and Q are functionally distinct classes (Part III). Applying a *symmetric* mean a second time would silently impose equal weight. So use a weighted multiplicative mean with a free weight $w \in (0, 1)$:

$M\_w(x, y) = x^w\, y^{\{1-w\}}$, $\Phi_+ = M\_w(a_+, g)$, $\Phi_- = M\_w(a_-, g)$.

Proposition 3.1 (Gate 3 derives the family $\chi' = w\chi$). With $g = \sqrt{(a_+a_-)}$,

$\Phi_+ = a_+^{\{(1+w)/2\}}\, a_-^{\{(1-w)/2\}}$, $\Phi_- = a_+^{\{(1-w)/2\}}\, a_-^{\{(1+w)/2\}}$, $\log \Phi(a) = [[(1+w)/2, (1-w)/2], [(1-w)/2, (1+w)/2]] \cdot \log a$,

a matrix fixing the even (sum) line at eigenvalue 1 and scaling the odd (contrast) line by w. Hence with E the access encoder, W_w the corresponding intertwiner and χ the odd readout,

$E \circ \Phi = W\_w \circ E$, $\chi(\Phi R) = w \cdot \chi(R)$, exactly (verified symbolically, Appendix D).

Reversal-equivariance is exact and the even (persistence) component is preserved for every w. Gate 3 thus delivers an exact contrast-transport **family**, parameterised by the single outer weight w. **[Proven on candidate — exact family $\chi' = w\chi$]**

Proposition 3.2 (Gate 2 supplies $w = ½$, hence $\chi' = ½\chi$). The outer weight is the branch weight of Part III: $w = W_I/(W_I + W_Q)$ with $W_{[c]} = \omega_{[c]} \kappa_{[c]}$. The binary census (Thm 2.4) is uniform, so $\omega_I = \omega_Q$; cost neutrality (the ledger equality $\Delta\ell_I = \Delta\ell_Q$) gives $\kappa_I = \kappa_Q$; together $W_I = W_Q$ and therefore

$$w = ½, \chi(\Phi R) = ½ \chi(R), \Delta\chi_{g+1} = ½ \Delta\chi_g.$$

With $w = ½$, $M_{1/2}(a_\pm, g) = \sqrt{a_\pm g}$ recovers the explicit $\Phi = (a_+^{3/4}a_-^{1/4}, a_+^{1/4}a_-^{3/4})$, $\log \Phi = ½(I + \Pi_e) \log a$ (numerical intertwining error 1.4×10^{-16}). The ½ is **not fitted**: it is the equal-weight binary persistence/closure census of Gate 2, transported by Gate 3's family. Gate 2 supplies the coefficient; Gate 3 transports it. **[Exact ½ — conditional on the Gate-2 ledger equality]**

The one remaining premise — sharpened to an executable descent audit. Propositions 3.1–3.2 exhibit an exact $w = ½$ map on the access cone; they do not yet prove the substrate's actual refinement *induces* it. That is the decisive remaining Gate-3 job, and it is now a definite generator-level computation rather than a slogan. Construct the closure state machine $S_m : X_m \to X_{m+1}$ and the access encoder $E_m : X_m \to \mathbb{R}^2$, and test

$$E_{m+1} \circ S_m = W \circ E_m$$

on a finite generating set. A complete execution must specify four things the present paper can name but not evaluate (the closure state machine is not reproduced here): **(1) the finite generating set** of X_m — the primitive refinement moves that generate the relevant state algebra; **(2) the concrete substrate update** S_m and encoder E_m as explicit maps on those generators; **(3) an algebra-closure argument** — why $E_{m+1}S_m = W E_m$ on the generators implies it on the whole state algebra (here the multiplicativity of the access law does the work: a multiplicative map agreeing with W on a generating set agrees on all products, so generator-level equality propagates); and **(4) pass/fail tolerance and the exact symbolic identities** where available (the access checks — positivity, reversal-equivariance, independent-composition multiplicativity, equal-access fixed points, generation independence — are exact symbolic identities, tolerance 0; only any floating realisation carries a residual). The content to confirm is that the microscopic history census delivers the two class exponents $(1\pm w)/2$ with $w = ½$.

Descent Premise (as an audit). The substrate's S_m, E_m satisfy $E_{m+1}S_m = W E_m$ on generators — equivalently, the microscopic access dynamics satisfy axioms (i)–(v) of Lemma 3.2 (forcing the inner geometric mean) and the history census supplies $w = ½$ (forcing the outer weight). Then descent realises Φ.

This is **not executed here** — building S_m, E_m requires the closure state machine from the corpus, which the present paper does not reproduce — so descent is stated as the named generator-level audit to be run, not a result. But it is reproducible and finite: a generator check, not an appeal to "real refinement realises Φ." **[Open — reduced to a finite generator-level descent audit]**

Gate 3 verdict: constructive pass. An exact contrast-transport family $\chi' = w\chi$ is derived and pinned to the geometric mean by Lemma 3.2 (regularised); Gate 2 fixes $w = ½$; the only open item is the executable Descent audit. **[Constructive pass — modulo the Descent audit]**

---

12. The Three-Gate Conditional Realisation

The three parts now combine into the theorem stated at the front. Gate 1 (Part II) makes the fold sector real and encoded; §8 reduces its load from four free bridges to one — the carrier identification (a) — with (b), (c), (d) now lemmas (faithfulness, reversible transport, the access axioms). Gate 2 (Part III) establishes operational twinhood (with the external/internal observable split), an executed-audit nonadjacency, and a *classified* binary census (Thm 2.4), isolating the Cost Neutrality ledger equality $\Delta\ell\_I = \Delta\ell\_Q$. Gate 3 (Part IV) derives the exact contrast-transport family $\chi' = w\chi$ and pins the inner mean (regularised Lemma 3.2); Gate 2's uniform binary census with cost neutrality fixes $w = ½$, isolating the Descent generator audit. Assembling — the inputs now being the carrier identification (a), the three operational axioms of §8 (faithfulness, source–target transport, the access axioms), and two finite audits (the Gate-2 census/cost audit, the Gate-3 descent audit):

Granting the carrier identification, the three operational axioms (2)–(4), and the Gate-2 and Gate-3 audits — the Gate-2 audit being the four sub-conditions (5)–(8), i.e. the rewrite-generator and terminal-history premises **and the two equalities $\omega\_I = \omega\_Q$ and $\Delta\ell\_I = \Delta\ell\_Q$** — together with the Gate-3 descent (9), the W_7 candidate passes χ Gate 1, χ Gate 2 and χ Gate 3, and the generational χ increment halves:

$\Delta\chi\_\{g+1\} = ½ \Delta\chi\_g$ $(\rho = ½)$,

the ½ being the binary-census weight of Gate 2 carried by the geometric-mean update of Gate 3 (Remark 3.3).

This is **conditional derivation**, not physical proof — and the conditional set has been reduced. After §8, the structural inputs are **one** carrier identification (a) plus the §8 lemma hypotheses (faithfulness, reversible transport, the access axioms), and the physical content is **two finite audits**: the Cost Neutrality ledger equality $\Delta\ell\_I = \Delta\ell\_Q$ (Part III) and the Descent generator check $E\_\{m+1\}S\_m = W E\_m$ (Part IV). Nothing is hidden or assumed proved. The reduction's residual conditions (3)/(5) and the magnitude $\Delta\chi_1$ remain its own scope; what this paper delivers is that, modulo the carrier identification and the two audits, the χ-halving *ratio* ½ follows from the W_7 structure by finite, checked computation rather than by assertion — with Gate 2 supplying $w = ½$ and Gate 3 transporting it. The honest headline: **χ-halving is conditionally derived from one carrier identification, three operational axioms, and two executable audits** — a major reduction from arbitrary choices, not waved at, and not over-claimed as physically closed.

---

13. A Separate Transport/Locality Sequence (Supporting, Not χ-Closing)

For completeness, the orientation-transport material that an earlier draft mislabelled as "Gate 2 / Gate 3" is recorded here as what it is — a **separate transport/locality sequence**, supporting the wider commitment programme but not part of χ closure. Its full treatment belongs to a dedicated transport paper; the status, conservatively graded, is:

- **Global orientability (boundary-cancellation step).** A completed vacuum interface as an integral cellular cycle Z_vac with ∂Z_vac = 0 forces opposite induced orientations on shared faces, so the matching signs are a coboundary s = $\delta\varepsilon$ and [s]_vac = 0 — strictly stronger than any K = 7 orientation argument, but valid only over an integral, oriented Z_vac with free coefficients (it fails over $\mathbb{Z}/2$). **PASS-pending** the integrality/freeness source checks; the rank of the group [s]_vac inhabits to be stated. **[Conditional — separate sequence]**
- **Refinement-loop transport (collar / amplitude-lift / τ step).** The minimal six-sector collar lifts to the forced alternating germ Λ_a = $((-1)^i a)$, which passes the full K = 7 ledger (I_inc, I_hub, I_circ, I_comp, anti-alignment, T_L, non-vacuity) — admissible, not merely evaluable. The corpus types the transport trace τ on full refinement motion (τ : Loop_rev(Γ_MS) $\to$ Z_1(C_cl)), with ρ : T_L $\to$ $\mathcal{D}$ a separate forgetful map and no factorisation $\tau = \hat{\tau} \circ \rho$, so the forced constant germ does not collapse the loop. **OPEN**, residual narrowed to one corpus-flagged item: the generator-level rule for τ on σ_j, μ_j, deciding *records* (τ(w_D) = sd(γ_D)) versus *quotients* (0). **[Open — separate sequence]**

Neither result bears on χ-halving; both are carried forward to the transport sequence's own treatment, and neither is invoked by the Three-Gate Conditional Realisation Theorem.

14. Status Ledger

Object	Status after this paper	Grade
Gate-1 audit protocol (the reduction's spec)	**Executed** — run end-to-end via the scripts of Appendix D, and now backed by a closed-form lemma	[Proven]
Instance grade	**Symmetry-canonical W_7 candidate** is primary (§3–§8); generic reference retained only as the discrimination vehicle (§7, Appendix C)	—
Dihedral Automatic-Pass Lemma (§3.1)	Closed-form proof that *any* 2D dihedral irrep yields rank-1 U, D = U†, ‖U−D‖ = √2, det Gram = 1, SAS = C	[Proven]
Carrier $H^1(W_7; \mathbb{C})$	dim = |E|−|V|+1 = 6; rotation spectrum = six sixth-roots of unity; E_1, E_2 the two reflection-conjugate doublets; ord($R\|E_1$) = 6, ord($R\|E_2$) = 3	[Proven]

Object	Status after this paper	Grade
Candidate primitives $A = P_+$, $B = S$, $C = P_-$	Self-adjoint, symmetry-canonical; rank A = rank C = 1, rank B = 6; S A S = C (mirror-conjugate pair)	[Proven]
Self-adjointness / $D = U^\dagger$	**Property** here (lemma item 2; $\lVert D - U^\dagger \rVert \approx 3 \times 10^{-16}$ confirms). Favourably resolves the reduction's well-posedness caveat for this instance	[Proven]
Rank verdict r = rank(U)	**r = 1 < 6** ⇒ Route A + Option 3. **Entailed by the identification** (lemma item 1), not an independent measurement (§4)	[Proven, entailed]
A.1 $\lVert U - D \rVert$	PASS ($\sqrt{2}$, exact) — necessary, not sufficient	[Proven]
A.2 distinct supports	PASS — $P_u = UU^\dagger = P_+$, $P\_d = U^\dagger U = P_-$, rank 1 each, distinct	[Proven]
A.3 exchange involution	PASS — realised by the D_6 reflection S ($S P_+ S = P_-$); polar phase coincides with S on E_1	[Proven]
Route-A Gate 1	**PASS on W_7 candidate** — but entailed by bridges (a)–(c), not corroborating them	[Proven, conditional on (a)–(c)]
Induction bridge	det Gram(U, D) = **1 exactly** (U, D HS-orthonormal); support-distinct ⇔ word-independent, consistent	[Proven]
Candidate encoder $E\_m(R) = Tr(P_+R)U + Tr(P_-R)D$	(ii) reversal-equivariance PASS (exact; **entailed by SAS = C**); **(i) image-containment automatic/vacuous**; non-circularity **derived by Thm 8.3** (Trace-Access Uniqueness), not cited	[Mixed: Proven / vacuous / Thm 8.3]
Discrimination of B.3	Shown only by the **generic reference** (Appendix C), where an arbitrary encoder fails (i) (residual ≈ 5.2). The candidate cannot show it	[Proven on reference]
(a) carrier = W_7 cycle space $H_1 \cong H^1$	The single foundational identification; not derived here (carrier-notation fixed via the graph Hodge inner product, §3)	[Open — structural assumption]
(b) doublet = E_1 → **Lemma 8.1**	Now a lemma: Primitive Faithfulness (no-aliasing) selects E_1, since $ord(R\|E_1)=6$ but $ord(R\|E_2)=3$ (R^3 aliases to I on E_2; verified). All three gates are blind to the doublet — faithfulness is the discriminator	[Lemma — modulo faithfulness]
(c) A,C reflection-conjugate → **Thm 8.2**	Now a theorem: Source–Target Support — $P\_comp = S\ P\_adm\ S$ from reversible transport, so $A=P_+$, $C=P_-$ follow from operational meaning, not from "else U=0"	[Theorem — modulo the transport reading]
(d) $Tr(P\pm R)$ access reading → **Thm 8.3**	Now a theorem: Trace-Access Uniqueness — positivity+linearity+support-locality+normalisation force $a\_P(R)=Tr(PR)$; circularity closed by uniqueness, not citation	[Theorem — modulo the access axioms]

Object	Status after this paper	Grade
Physical Gate-1 verdict	**Reduced to (a) plus the §8 lemma hypotheses**; three of four bridges discharged to natural requirements	[Conditional on (a)]
Gates 2, 3 made concrete	Now **run** (§§10–12): twinning / nonadjacency / census (Gate 2) and the log-access intertwiner (Gate 3) are finite computations on W_7/D_6, assembled into the Three-Gate Conditional Realisation Theorem	[Conditional realisation]
χ Gate 2 (twinning, nonadjacency, census, cost)	§10 — **conditional pass**: twinning (dim-4 invariant algebra + external/internal observable split), nonadjacency *executed* on Cay(S_3,{(12),(23)}) d=3, binary census *classified* (Thm 2.4). Open items: rewrite-generator and terminal-history premises, plus **two** weight equalities ω_I = ω_Q (uniform-admissibility measure) and Δℓ_I = Δℓ_Q (ledger), which together fix w = ½	[Conditional — two equalities + census premises]
χ Gate 3 (log-access intertwining)	§11 — **constructive pass**: derives family χ(ΦR)=w·χ(R); inner mean pinned by regularised Lemma 3.2; **w=½ from Gate 2's two equalities** (ω_I=ω_Q, Δℓ_I=Δℓ_Q). Open item sharpened to the executable generator audit E_{m+1}S_m=W E_m	[Constructive; descent audit Open]
Transport/locality sequence (separate)	§13 — orientability PASS-pending (integrality); τ on refinement loops OPEN (generator rule). **Does not bear on χ closure**	[Conditional / Open — separate]
Whether ρ = ½	Gate 3's Φ gives χ(ΦR) = ½χ(R) exactly on the cone, the ½ inherited from Gate 2's uniform binary census (Remark 3.3); full physical ρ = ½ is conditional on the two Gate-2 equalities (ω_I=ω_Q, Δℓ_I=Δℓ_Q), Descent, the χ↔physical-susceptibility identification inherited from the reduction, and the reduction's own conditions (3)/(5) and magnitude $\Delta\chi_1$	[Conditional]

Epistemic grades: **[Proven]** (the lemma, the witnesses as closed-form facts, self-adjointness); **[Proven, conditional on (a)–(c)]** (the Route-A pass); **automatic/vacuous** (encoder condition i); **derived by Thm 8.3** (encoder non-circularity, bridge d — no longer citation-dependent); **(a) Open — structural; (b)–(d) Lemma/Theorem** conditional on the operational axioms (faithfulness, transport, access); **entailed-not-corroborated** (the Gate-1 pass as a whole); **conditional-realisation** (Gates 2 and 3 — §§10–12 — pass on the candidate modulo the Gate-2 audit and the Descent audit; the Three-Gate Conditional Realisation Theorem assembles them).

---

15. Conclusion

This paper began as a Gate-1 audit and closes as a three-gate conditional derivation: all three χ-halving gates are run on the W_7 candidate, and what remains is a finite, named list of premises. The Gate-1 narrative below is Part II of that arc; §§10–12 carry the rest.

The reduction ended at a computation it could not run: r = rank(U). This paper runs it — not on an arbitrary reference but on a candidate generated from the W_7 closure cell, which is the right kind of instance: corpus-internal, self-adjoint, forced by symmetry rather than chosen.

On $H^1(W_7; \mathbb{C})$, the rim rotation's regular-representation spectrum yields a reflection-conjugate doublet whose two rank-one projectors, taken as admissibility and completion with the reflection as transport, give a closure word that loses rank to one. The Route-A audit passes — and the **Dihedral Automatic-Pass Lemma** upgrades each witness from a 10^{-16} residual to a closed-form fact: distinct fold lines ($P_u = P_+$, $P\_d = P_-$), an exchange realised by the D_6 reflection, $\|U - D\| = \sqrt{2}$ exactly, and the strongest possible word-span certificate, det Gram = 1 exactly. A mass-blind candidate access map induced from the pre-existing trace-accessibility readings satisfies the encoder's reversal-equivariance exactly, with odd readout $\chi(m) = \ln \mathrm{Tr}(P_+R\_m) - \ln \mathrm{Tr}(P_-R\_m)$. Self-adjointness, which a generic instance can only assume, is here a two-line property of the symmetry-canonical primitives.

But the honest centre is that **the Gate-1 pass has migrated upstream — and the lemma proves the migration is total.** Once the W_7/E_1/mirror-pair identification is fixed, every witness follows by a theorem about 2D dihedral irreps; the audit no longer discriminates, because there is no residual operator freedom left to test, and the very generality of the lemma shows an *entire class* of instances would pass identically. The verdict's whole load rests on four identification bridges — and §8 then upgrades three of them to lemmas: (a) the carrier is $H_1(W_7) \cong H^1(W_7)$; (b) the fold doublet is E_1 rather than E_2 — which Gate 1 *provably cannot* decide, since E_1 and E_2 (distinct, non-isomorphic D_6-irreps) induce isomorphic data on the exchange quotient Gate 1 reads, the separating invariant (ord R = 6 vs 3) sitting in the ambient action the audit forgets — but which the **Primitive-Faithfulness Lemma (8.1)** decides, since faithfulness of the C_6 action excludes E_2 ($R^3 = I$ there); (c) admissibility and completion are reflection-conjugate — derived by the **Source–Target Support Theorem (8.2)** from reversible transport, not merely forced by non-degeneracy; and (d) the access reading $\mathrm{Tr}(P\pm R)$ — derived by the **Trace-Access Uniqueness Theorem (8.3)** from four access axioms, not cited. Only (a) remains a bare assumption. The candidate is a genuine, self-adjoint, symmetry-generated *construction*; its Gate-1 pass certifies internal consistency, not substrate-uniqueness, and §8 reduces its standing to (a) plus three operational axioms.

So the standing is bounded and exact. The protocol is executed; the W_7 candidate clears Gate 1 conditional on (a)–(c) with (d) deciding encoder circularity; self-adjointness is resolved for this instance; and the physical verdict is reduced to the conjunction of four named bridges that Gate 1 entails rather than corroborates. The paper does not show the substrate selects this construction, and does not pretend to.

The next burden is sharp, and §8 has made it sharper than a Gate-1 audit alone could. Three of the four old bridges are now lemmas: (b) the E_1 selection from Primitive Faithfulness, (c) the mirror pairing from Source–Target Support, (d) the trace-access reading from Trace-Access

Uniqueness. Only the carrier identification (a) remains a structural assumption. On top of it sit exactly **two physical audits**, each now a finite executable check rather than a slogan: the **Cost Neutrality ledger equality** $\Delta\ell\_I = \Delta\ell\_Q$ (Gate 2), which fixes the census weight w = ½; and the **Descent generator audit** $E\_\{m+1\}S\_m = W E\_m$ (Gate 3), which is what would show the substrate realises the $\chi' = w\chi$ family. Gate 2 supplies the coefficient, Gate 3 transports it, and $\Delta\chi\_\{g+1\} = ½\Delta\chi\_g$ follows. So the result is not a Gate-1 audit with gestures at what follows: it is a conditional derivation of the χ-halving ratio ½ from **one carrier identification, the three operational axioms of §8, and two finite audits**, with a separate transport/locality sequence (§13) demarcated as not bearing on χ closure. The honest standing: χ-halving is conditionally derived, not physically proved; the remaining work is the carrier identification, confirming the three operational axioms (one of which — faithfulness — is currently a post-hoc tie-breaker awaiting independent corpus motivation), and running the two finite audits (whose four Gate-2 sub-conditions include the independent equalities $\omega\_I = \omega\_Q$ and $\Delta\ell\_I = \Delta\ell\_Q$). And one caveat dominates all of them: **what halves is an internal contrast functional; its identification with the physical quark susceptibility is inherited from the reduction and is examined nowhere here.** That single upstream identification, not the gate machinery, is where the physical weight of the claim actually rests.

Appendix A — The Gate-1 Specification (Recapped)

The full two-headed specification is the reduction's Gate-1 audit; it is recapped here compactly so the witnesses of §4–§7 are self-contained. Inputs: concrete A, B, C on $\mathscr{H}$ (dim n); U = ABC, D = CBA.

C.0 — selecting test. r = rank(U) by SVD. **r < n** → Route A available + Option 3 (audit in A, exposit in B). **r = n** → supports collapse (P_u = P_d = I); Route B forced, B.3 carries Gate 1. ("r < n requires admissibility or completion to genuinely lose rank" — a physics commitment.)

Route A (requires r < n; A, B, C self-adjoint so D = U†).

- **A.1** order sensitivity: ‖U − D‖ > tol (necessary, not sufficient).
- **A.2** distinct supports: P_u = projector onto Ran(U), P_d onto Ran(U†); pass iff $P_u \neq$ P_d (equal-rank, distinct-range). Collapse P_u = P_d fails.
- **A.3** exchange involution J ($J^2 = I$, J P_u J = P_d) from the polar phase of U on the doubled carrier $\mathcal{F}$_d ⊕ $\mathcal{F}_u$.
- Route-A Gate 1 passes iff A.1 ∧ A.2 ∧ A.3. Establishes two distinct fold subspaces exchanged by J; does not establish Gate 2/Gate 3.

Route B (no rank requirement; self-adjointness not required).

- **B.1** nondegeneracy: |det Gram(U, D)| > tol (U ≠ λD any λ; strictly stronger than K = U − D ≠ 0).
- **B.2** involution/split well-defined on $\mathcal{F}$ = span{U, D}.

- **B.3** access-row identification: an **independently specified** encoder $E_g : X_g \to \mathcal{F}$ with (i) $\mathrm{Im}(E_g) \subseteq \mathrm{span}\{U, D\}$ and (ii) $E_g \hat{J}_g = J E_g$, on a generating set, all g. Defining E_g from observed masses voids the test (circularity).
- Route-B Gate 1 passes iff B.1 ∧ B.2 ∧ B.3.

Option 3 (r < n). Run Route A (finite SVD, decisive witness) as audit, present $\mathcal{F} = \mathrm{span}\{U, D\}$ for clarity, certify agreement via the polar phase ($P_u \neq P_d \Leftrightarrow$ U, D independent).

What Gate 1 does not decide. Twinning and census granularity and branch-cost (Gate 2); full log-access intertwining (Gate 3, of which B.3 is the encoder half; defining S_g as the census average $\frac{1}{2}(P_g + C_g)$ makes the Gate-3 intertwining test vacuous — the second circularity trap).

Appendix B — The W_7 Symmetry-Canonical Candidate, Explicitly (with Non-Uniqueness Probe)

Carrier. W_7 = one hub (vertex 0) and six rim vertices (1…6); 12 oriented edges (6 spokes 0→rim, 6 rim edges cyclic). Boundary matrix $\partial \in \mathbb{R}^{7\times12}$; $H^1 = \ker \partial$, dim = 12 − 7 + 1 = 6, in an orthonormal basis.

Symmetry. Rim rotation R: rim $r \mapsto 1 + (r \bmod 6)$; reflection S: rim $r \mapsto 1 + ((-(r-1)) \bmod 6)$ (fixes 1, 4; swaps (2,6),(3,5)); both lifted to signed permutations on oriented edges and restricted to H^1. On H^1: $\mathrm{spec}(R) = \{1, e^{\pm i\pi/3}, e^{\pm 2i\pi/3}, -1\}$ (each simple); $S^2 = I$, $S R S = R^{-1}$. Rotation orders on the doublets: $\mathrm{ord}(R|E_1) = 6$, $\mathrm{ord}(R|E_2) = 3$.

Projectors and primitives. $P_\pm = (1/6) \sum_{n=0}^{5} e^{\mp in\pi/3} R^n$ (rank-one, orthogonal, $S P_+ S = P_-$). Candidate primitives on the E_1 doublet: $A = P_+$, $B = S$, $C = P_-$; $U = ABC$, $D = CBA$.

Closed-form witnesses (Dihedral Automatic-Pass Lemma, §3.1), with confirming residuals (Appendix D, machine precision):

rank A = rank C = 1, rank B = 6; $A^2 = A = A^\dagger$, $C^2 = C = C^\dagger$ ($\|\cdot\| \approx 3 \times 10^{-16}$), $B^2 = I$, $B = B^\dagger$ ($\|B^2 - I\| \approx 6 \times 10^{-16}$), $S A S = C$ ($\|\cdot\| \approx 6 \times 10^{-16}$); $U = \alpha|v_+\rangle\langle v_-|$, $D = \bar{\alpha}|v_-\rangle\langle v_+|$, $|\alpha| = 1$; rank(U) = 1; $D = U^\dagger$ ($\|D - U^\dagger\| \approx 3 \times 10^{-16}$); $UU^\dagger = A$, $U^\dagger U = C$ ($\|\cdot\| \approx 4 \times 10^{-16}$); $\|U - D\| = \sqrt{2} = 1.41421356\ldots$; det Gram(U, D) = 1.000000000000; encoder image residual ≈ 5×10^{-15}; encoder reversal-equivariance error ≈ 8×10^{-15}.

Candidate access encoder (mass-blind). $a_\pm(R) = \mathrm{Tr}(P_\pm R)$; $E_m(R) = a_+ U + a_- D$; $\mathrm{Im}\, E_m \subseteq \mathrm{span}\{U, D\}$ automatically; $E_m(S R S) = a_- U + a_+ D = J E_m(R)$ (since $S A S = C$, exact); odd readout $\chi(m) = \ln a_+(R_m) - \ln a_-(R_m)$. No quark mass enters.

Non-uniqueness probe (the bridges (b), (c) computations). Running the identical construction on the other doublet, and on a non-conjugate pairing, confirms the lemma's predictions:

E_1 ($e^{\pm i\pi/3}$): rank(U) = 1, ‖U − D‖ = √2, det Gram = 1, S A S = C → Route-A PASS; E_2 ($e^{\pm 2i\pi/3}$): rank(U) = 1, ‖U − D‖ = √2, det Gram = 1, S A S = C → Route-A PASS, **bit-identical**; non-conjugate (k = +1, +2): rank(U) = **0**, ‖S A S − C‖ ≈ 1.4 (≠ 0) → word vanishes; not a mirror pair.

So Gate 1 does **not** select E_1 over E_2 (bridge b — and the lemma shows this is structural, since both are 2D dihedral irreps), and the construction **requires** a reflection-conjugate pair (bridge c). Both are computations *predicted by the lemma*, not assertions. The probe also exhibits the one invariant Gate 1 ignores: $\mathrm{ord}(R|E_1) = 6 \neq 3 = \mathrm{ord}(R|E_2)$.

Appendix C — The Generic Reference Certification Model

A generic instance is retained for the one job the W_7 candidate cannot do: **demonstrating that B.3 discriminates** — that the image-containment condition (i) is a real constraint capable of failing. The candidate's encoder satisfies (i) by construction (§7), so it cannot exhibit discrimination; the generic model can.

Carrier $\mathscr{H} = \mathbb{C}^4$. Generating unit vectors |a⟩ ∝ (1, 0.3i, −0.2, 0.1), |c⟩ ∝ (0.2, 1, 0.4i, −0.3), |b⟩ ∝ (0.1, −0.5, 1, 0.6i). Primitives A = I − |a⟩⟨a| (rank 3), C = I − |c⟩⟨c| (rank 3), B = I − 2|b⟩⟨b| (involution). Then U = ABC, D = CBA = U†, rank(U) = 3 < 4 (Route A), with ‖U − D‖ = 1.275, distinct rank-3 supports $\|P_u - P\_d\| = 1.328$, polar-phase $J^2 = I$ and $J P_u J = P\_d$ to 10^{-15}, and |det Gram| = 2.689 ≠ 0.

Discrimination demonstration. On a two-generator toy access space: a *generic* encoder (generators → arbitrary operators) leaves the 2-dimensional word-span with worst image-residual ≈ 5.2 in 16-dimensional operator space and **correctly fails (i)**; a *structurally specified* encoder (generators → U, D) lands in the span (residual ≈ 10^{-16}) and respects J-equivariance exactly. So B.3's condition (i) is non-vacuous — a fact the W_7 candidate, whose encoder is built to pass (i), cannot itself establish.

The reference model is role-faithful but otherwise arbitrary; it makes no canonicity claim. Note that it does *not* satisfy the Automatic-Pass Lemma's hypotheses (A, C are rank-3 reflections, not rank-1 conjugate projectors), which is precisely why its witnesses (‖U−D‖ = 1.275, det Gram = 2.689) differ from the candidate's clean √2 and 1 — and why it has room to fail where the candidate does not. It is the protocol-certification and discrimination vehicle, nothing more.

Appendix D — The Audit Scripts

Four companion scripts (fixed seeds; numpy and sympy) accompany this paper. Two cover Gate 1; two cover the Gate-2/Gate-3 audits added in Parts III–IV.

`gate2_nonadjacency.py` — builds the one-step closure-order graph G_ord = Cay(S_3, {(12),(23)}) and does a breadth-first audit: all six arrangements reachable, diameter 3, d(ABC, CBA) = 3, and (13) (the direct U↔D flip) absent from the generators and at distance 3 (Prop 2.2). `gate3_weight.py` — verifies symbolically that the weighted outer mean M_w gives $\chi(\Phi R) = w \cdot \chi(R)$ for general w, and the inner closure is the geometric mean (Lemma 3.2), recovering the w = ½ map $\log \Phi = ½(I + \Pi\_e) \log a$ (Props 3.1–3.2).

The original two:

Two companion scripts (fixed seeds, numpy) accompany this paper.

`w7_gate1_candidate.py` — builds W_7, computes H^1 and the D_6 action, selects the E_1 projectors, forms $A = P_+$, $B = S$, $C = P_-$, U = ABC, D = CBA, and audits rank, adjointness, support projectors, the Gram determinant, and the candidate encoder's image-containment and reversal-equivariance. The accompanying probe runs the E_2 and non-conjugate cases and reports $\mathrm{ord}(R|E_1)$, $\mathrm{ord}(R|E_2)$. No quark mass enters; output residuals are those quoted in §3–§8 and Appendix B, and confirm the closed-form values of the Automatic-Pass Lemma.

`gate1_audit.py` — the generic reference model (Appendix C): constructs the role-faithful primitives, runs the Route-A witnesses, the induction-bridge Gram determinant, and the B.3 discrimination demonstration (generic encoder fails, structural encoder passes). Retained for protocol certification.

Both are structured so that replacing their operators with canonical substrate primitives — and the toy/candidate encoder with a physical, independently specified E_m — converts the run into the physical audit with no change to the protocol. The two circularity guards are explicit in both: the encoder is not constructed from masses, and the refinement step is not defined as its own census average.

Appendix E — Referee Objections and Replies

Objection 1 — The W_7 candidate is just a prettier rigged instance: you chose A, C to be projectors so closure would lose rank. Sharpened, not dismissed. Yes, $A = P_+$ and $C = P_-$ are rank-one projectors, so rank-deficiency is built in — and the Automatic-Pass Lemma (§3.1) now states *exactly* how built-in: rank-one is a theorem for any 2D dihedral irrep. §8 is the admission that this, and more, is built in: the Gate-1 pass is *entailed by* the W_7/E_1/mirror-pair identification, not a corroboration of it. The candidate's value is not that it passes (it must) but that it passes *from corpus-internal symmetry* rather than from arbitrary operators, and that pinning the carrier makes Gates 2/3 concrete (§9). The grade is "construction, not test"; nothing is hidden.

Objection 2 — Then the W_7 pass establishes nothing a generic instance did not. It establishes three things a generic instance could not: (a) a *self-adjoint* realisation in which $D = U^\dagger$ is a two-line property, not an assumption (§3.1 item 2, §9); (b) operators *generated* by W_7/D_6 symmetry rather than chosen, so the residual freedom is named and finite — bridges (a)–(d) —

rather than open-ended; (c) a concrete carrier on which the remaining gates become finite computations. What it does *not* establish — substrate-uniqueness — it explicitly disclaims.

Objection 3 — "The unique E_1 sector" oversells: E_2 would do as well. Sustained *as a statement about Gate 1*, and now answered by a lemma. E_2 passes Gate 1 — and the whole Gate-2/Gate-3 audit — bit-identically, because all three gates read only the exchange quotient on which E_1 and E_2 (distinct, non-isomorphic D_6-irreps) induce isomorphic data. No gate selects the doublet. But the **Primitive-Faithfulness Lemma (8.1)** does: faithfulness of the C_6 rim action excludes E_2, since $R^3 = I$ there (a non-identity rim element acting trivially). So E_1 is selected by a natural no-aliasing requirement, not imported as a free choice — and the honest point is preserved that faithfulness, not any gate calculation, is the discriminator.

Objection 4 — Admissibility and completion are forced to be mirror images; that is a strong, unargued physical claim. Sustained, and it is bridge (c). The construction sets C = S A S; a non-conjugate pairing gives rank(U) = 0 (Appendix B probe), so the mirror relation is *required*, not incidental. In the reduction, admissibility and completion are distinct operations with no a priori reason to be reflection-conjugate. The candidate identifies them as the two halves of one doublet; §8 isolates this as its own bridge rather than letting it ride for free on the symmetry. §7 adds the observation that the encoder's only non-trivial test (reversal-equivariance) is itself downstream of this same identity, so (c) carries more of the load than first appears.

Objection 5 — The candidate encoder passes B.3 (i) trivially; that is not a real check. Sustained (§7). $E\_m = a_+U + a_-D$ is defined inside span{U, D}, so (i) cannot fail — it is vacuous here. The genuine content is (ii) reversal-equivariance, which holds *exactly* — but, as §7 now shows, it holds *because* S A S = C, i.e. it is entailed by bridge (c) and not independent new content. Discrimination of (i) — that a generic encoder *does* fail it — is shown only by the generic reference (Appendix C, residual ≈ 5.2). The grade is "(ii) passes, entailed by (c); (i) automatic."

Objection 6 — Tr(P±R) was chosen to make χ come out odd; that is the circularity trap. This was bridge (d), and it is now met by a theorem rather than a citation. The **Trace-Access Uniqueness Theorem (8.3)** shows that any access functional that is positive, linear, support-local, P-invariant, and normalised *must* be $a_\pm(R) = \mathrm{Tr}(P_\pm R)$ — none of those axioms mentions χ. So mass-blindness is forced by the access axioms, not reverse-engineered from the desired readout. The residual is only whether those four axioms are the operative accessibility law (an operational-semantics question), not whether Tr(P±R) was hand-picked.

Objection 7 — rank(U) = 1 is degenerate; does the downstream two-dimensional fold story survive? Yes. Although U, D are rank-one *operators*, the word-span $\mathcal{F}$ = span{U, D} is two-dimensional with U, D Hilbert-Schmidt-orthonormal (det Gram = 1, lemma item 4), and the even/odd lines U ± D are complementary and nontrivial ($\|U \pm D\| = \sqrt{2}$). So the reduction's 2-dimensional fold structure is realised exactly; rank-one is the *cleanest* such realisation (two orientations = two lines), not a degeneration of it.

Objection 8 — Does the W_7 candidate move ρ toward ½? No, and §9 is explicit. Gate 1 establishes the sector is real; ½ is fixed by Gate 2 (census, granularity, cost) and Gate 3 (intertwining). A real sector that does not halve is entirely possible (the reduction's failure knobs). The candidate's contribution to ½ is indirect: by pinning the carrier it makes the Gate-2/3 computations finite, where the actual constraint on ½ lives.

Objection 9 — Why retain the generic reference at all? Because it does one job the candidate cannot: demonstrate that B.3's image-containment condition (i) is a real, falsifiable constraint (a generic encoder fails it). The candidate's encoder is built to pass (i), so it cannot show (i) discriminates. Protocol certification and substrate candidacy are different roles; the reference holds the first, the W_7 candidate the second. The lemma makes the division of labour exact: the reference deliberately violates the lemma's hypotheses (rank-3 reflections, not rank-1 conjugate projectors), which is what gives it room to fail.

Objection 10 — Self-adjointness was a flagged caveat in the reduction; is it now resolved? For this instance, yes, and now by proof rather than computation. A, C are orthogonal projectors and B is a real reflection, so $D = U^\dagger = CBA$ is the two-line consequence of §3.1 item 2 (the residual $\lVert D - U^\dagger \rVert \approx 10^{-16}$ merely confirms the basis arithmetic). The general caveat (canonical primitives might not be self-adjoint, forcing Route B) stands for the substrate at large, but the W_7 candidate exhibits a self-adjoint realisation, removing the caveat for itself.

Objection 11 — If Gate 1 is entailed by the identification, the audit is empty. Not empty — *reduced*. The audit's value was never to pass a pre-decided instance but to convert "is the fold sector real?" into a finite, named set of conditions. On the W_7 candidate it has done exactly that: the verdict is reduced to bridges (a)–(d), and §8 then upgrades three to lemmas — (a) foundational (carrier); (b) decided by Primitive Faithfulness; (c) by Source–Target Support; (d) by Trace-Access Uniqueness. That reduction is the deliverable. The discriminating tests now live one gate downstream — Gate 2 on W_7 — which §9 names as the next audit and as the place where bridge (b)'s ambient invariant finally enters a test.

Objection 12 — §10 conflates two different gate sequences, and earlier even mislabelled them. Conceded, and corrected in §10 itself. The reduction's χ-halving gates are χ Gate 2 (twinning, nonadjacency, binary census, branch-cost) and χ Gate 3 (log-access intertwining); a *separate* transport/locality sequence (global orientability; the refinement-loop trace τ) is a different object that supports the wider programme but does not by itself close χ-halving. An earlier draft audited the transport sequence under the labels "Gate 2 / Gate 3," which was a genuine numbering collision. §§10–12 now run the actual χ gates (Gate 2 a conditional pass with Cost Neutrality the lone open premise; Gate 3 a constructive pass with $\chi(\Phi R)=\tfrac{1}{2}\chi(R)$ exact and Φ pinned by the Selection Lemma, Descent open) and assemble them into the Three-Gate Conditional Realisation Theorem, while §13 demarcates the transport material as the separate sequence it is (orientability PASS-pending; τ OPEN on the generator rule). The conditional-realisation chain is explicit about its named inputs — one carrier identification, three operational axioms, two finite audits — and claims conditional derivation, not physical proof. A forward-status ledger that catches and corrects its own mis-attribution is behaving as intended, not failing.

Objection 13 — The Automatic-Pass Lemma proves the construction is generic to all 2D dihedral irreps; doesn't that *weaken* the W_7 claim? It clarifies it, in the direction of honesty. The lemma does not weaken what was actually being claimed — it makes precise that Gate 1 detects "a 2D dihedral fold structure is present," nothing finer. What remains specific to W_7 is supplied by the bridges, not by the audit: which carrier (a), which of its two doublets (b), the physical role-assignment (c), the accessibility citation (d). The lemma's service is to draw the line between "what the audit proves" (the dihedral structure, for free, for a whole class) and "what the corpus must supply" (the four bridges) exactly where it belongs, so neither side is over- or under-credited. A construction whose passing is a theorem is more trustworthy as a construction, not less — provided one says plainly, as §8 does, that the theorem is not a corroboration.

Objection 14 — χ Gate 3 "constructive pass" overstates: Φ is one map among many. Partly sustained, and now partly answered. The selection half of the worry is met by the regularised Geometric-Mean Selection Lemma (3.2): under positivity, homogeneity, exchange-symmetry, multiplicativity, and a regularity condition (continuity/measurability/monotonicity — which is what rules out discontinuous Cauchy solutions), the closure of the *exchange-symmetric* inner pair is forced to be the geometric mean. The outer (persistence/closure) weight is deliberately left free, giving the family $\chi' = w\chi$, with $w = ½$ supplied separately by Gate 2 — so Φ is forced, not one map among many, and the ½ is earned rather than imposed. What remains genuinely open is the **descent** half — that real refinement descends to the positive fold-access cone and that its dynamics satisfy those axioms (equivalently, realises Φ). That is the Descent Premise, named and graded. "Constructive pass" means an exact, axiom-forced realiser is exhibited — a real advance over a verbal condition, and strictly less than "the physical map halves." The same discipline applies to χ Gate 2: every structural check passes, but Cost Neutrality — that physics uses a branch-common cost under which $\ell(Qx) = \ell(Ix)$ — is a physical premise symmetry does not fix (operator norm neutral, Hilbert–Schmidt not), graded conjectural. Neither gate is unconditionally closed; both pass conditional on one named premise each.

Objection 15 — re-titling this a "three-gate conditional realisation" overclaims; it is still mostly open. It is not a closure of the χ-halving law, and the title and theorem say *conditional realisation* for exactly that reason. The claim is bounded and exact: §8 reduces Gate 1's load to the carrier identification (a) plus three operational axioms (faithfulness, transport, access), turning (b)–(d) into lemmas; Gate 2 and Gate 3 then pass conditional on two finite audits — the Gate-2 census/cost audit (generators, terminal-history axioms, ω_I = ω_Q, Δℓ_I = Δℓ_Q) and the Gate-3 descent check E_{m+1}S_m = W E_m — with the ½ derived from Gate 2's binarity and transported by Gate 3's family $\chi' = w\chi$. That is a conditional derivation of the ratio ½ from **one carrier identification, three operational axioms, and two executable audits**, every dependency on the table, none claimed proved. Calling that a "three-gate conditional realisation" is accurate; calling it a proof of χ-halving would not be, and the paper does not. The advance is real: the gates are executed, not gestured at, and the residual is reduced from four free bridges plus two vague premises to one identification, three principled axioms, and two checks a subsequent paper can pass or fail on a computation.

Objection 16 — the Primitive-Faithfulness Lemma (8.1) just renames the E_1 assumption "faithfulness." Partly: it relocates it, and the relocation is the point. The bare bridge (b) was a free choice between two options Gate 1 reads as identical — and, as §8 now shows, *all three* gate audits read them as identical, so no calculation in this paper favours E_1. Lemma 8.1 replaces that free choice with a single, physically natural requirement: a primitive 60° rim move must not act trivially in the fold representation (no aliasing). On E_2 the half-turn R^3 aliases to the identity (verified: $R^3 = +1$ on E_2, -1 on E_1); on E_1 it does not. So the residual is no longer "pick a doublet" but "require faithful representation of the rim group" — a single, sharply-stated requirement rather than a free binary choice. But the concession must go further than relocation: faithfulness is *the* invariant separating the doublets and is introduced to break the tie, so it is **post-hoc unless the corpus motivates a faithful rim action independently of this selection**. Absent that citation (which this paper does not have), it is graded as a tie-breaker chosen for its outcome — better than a coin-flip between E_1 and E_2, because it is one principled-looking condition, but not yet "physically natural." The honest status is: a much better premise than a free choice, still a premise, and currently without the independent motivation that would make it more than a tie-breaker.

Objection 17 — the Source–Target Support Theorem (8.2) assumes its conclusion: that completion is transported admissibility. It assumes the *operational meaning* of the three roles — admissibility gates the input support, transport is the reversible map S, completion gates the output support — and derives $C = S\,P\_adm\,S = P_-$ from them. That is genuinely more than the prior necessity argument (which only said a non-conjugate pairing makes the word vanish): it grounds the mirror relation in what "reversible fold transport" means, rather than in the algebraic accident that other pairings give rank 0. The residual is the operational reading itself, which the reduction's definitions of admissibility/transport/completion should either license or refute — a corpus check, not a free assignment.

Objection 18 — separating an exchange-symmetric inner mean from a weighted outer mean is a contrivance to land on ½. The opposite: collapsing them was the contrivance. The inner pair (a_+, a_-) *is* exchanged by the fold symmetry, so a symmetric mean is forced (Lemma 3.2). The outer pair (branch, closure) is *not* — I and Q are functionally distinct classes — so imposing symmetry there would smuggle in equal weight unexamined. Keeping the outer weight free yields the honest family $\chi' = w\chi$ (verified symbolically), and $w = ½$ then has to be *earned* from Gate 2's uniform binary census plus cost neutrality. That makes the ½ a derived coefficient with a stated source, not a number chosen by applying the same mean twice. The asymmetry here is **structural, not stipulated**: J exchanges $a_+ \leftrightarrow a_-$ but does *not* exchange persistence with closure (I and Q are functionally distinct — one preserves, one collapses the odd sector), so symmetry is *available* on the inner pair and *unavailable* on the outer one. A construction that respected the symmetry it has and declined the symmetry it lacks is the disciplined choice; collapsing both into one symmetric mean is what would have smuggled ½ in unexamined. This is the paper's cleanest single step and deserves to be read as such. The regularised Lemma 3.2 also fixes a real gap: multiplicativity alone admits discontinuous Cauchy solutions, so a regularity hypothesis (continuity/measurability/monotonicity) is now explicit, and the affine-offset condition has been moved out of the mean-classification where it did not belong.

Objection 19 — Cost Neutrality and Descent are still unproved, so "conditional realisation" is the same open status with nicer words. The status is conditional, and the words are more than nicer — both are now *finite, executable checks* rather than slogans, which is what makes them dischargeable. The Gate-2 audit reduces to explicit equalities — $\omega\_I = \omega\_Q$ (uniform admissibility) and $\Delta\ell\_I = \Delta\ell\_Q$ (ledger cost) between two named representatives under ℓ — neither evaluated here, because the present paper does not carry ℓ's values or an equal-weighting citation. Descent reduces to a generator-level test $E_{m+1}S_m = W E_m$ on the closure state machine; again not executed here, because that machine is not reproduced. Neither is claimed proved. The advance is in the *shape* of the residual: from "four free bridges plus two vague premises" to "one carrier identification, three operational axioms, and two finite audits," each sharply enough stated that a subsequent paper can pass or fail it on a computation.

Objection 20 — for a paper about quark masses, nothing here touches a quark. Conceded, and now flagged prominently (the Scope note after the theorem, the abstract, and the conclusion). The χ this paper halves is an abstract odd log-access readout on the W_7 cone; the four-axiom uniqueness theorem (8.3) characterises its *form* as $\mathrm{Tr}(P \pm R)$ but does not show that functional computes a physical susceptibility. The identification "this χ = the quark χ" lives entirely upstream — in the reduction's interpretive bridge and the carrier identification (1) — and is the single least-scrutinised assumption in the whole construction, precisely because it is inherited rather than re-derived. So the centre of gravity is honestly inverted: pages of gate machinery are conditional structure sitting on top of one interpretive identification that this paper assumes and does not test. The deliverable is a conditional halving of an internal contrast; the bridge to quark masses is upstream scope, named and not claimed.