

The W_7 Fold Carrier — A Layer Audit

Transport-Current Shadow, Class-Grain Standing Carrier, and the Demotion of E_1 to a Readout Representation

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

Supersedes the earlier reading in which the physical fold carrier was identified with the E_1 sector of the full graph cycle space $H^1(W_7; \mathbb{C})$. The corpus does not support that identification at the carrier level. This paper resolves the apparent contradiction between the χ -construction's six-dimensional cycle workspace and the sequential-transport strand's one-dimensional outer-cycle homology by splitting "carrier" into three distinct layers, and it relocates E_1 from carrier to candidate readout representation. The dominant open problem is then shown to be not carrier survival but the contrast no-go: a symmetric standing state gives $\chi = 0$.

Summary for the General Reader

Earlier work in this series treated a particular internal pattern of the W_7 object — the " E_1 fold pair" — as the physical thing that carries the difference between the up and down quark. A check against the wider corpus shows that this was using one word, "carrier," for what are really three different things, and the fold pair was placed at the wrong level.

The three levels are these. First, there is the **transport-current shadow**: the bookkeeping of how structure is carried from one step to the next, which the corpus collapses down to a single simple loop. Second, there is the **standing carrier**: not any individual seat or member, but the *class* — the species-level structure that actually owns generation and identity. Third, there is the **registration layer**: individual seats can be probed, but what they read is the class-owned structure, not their own independent content.

The correction is that the fold/species structure lives at the **class** level, not in the loop-bookkeeping and not in any individual seat. The E_1 pair, which the old paper treated as the carrier, is at best a **readout pattern** — a way of representing a binary up/down distinction that sits *over* the class carrier. It is not the carrier itself.

This is partly good news: the old contradiction between the two strands mostly dissolves, because they were describing different layers. But it comes with a hard cost, and this paper is honest about it. The quantities a_+ and a_- — the two readings whose ratio was supposed to *be* the up/down mass ratio — can no longer be two independent seat-level amplitudes. And when they are computed from a symmetric class state, they come out **equal**: the up/down contrast is exactly zero, predicting a mass ratio of one against a measured value near 0.46.

And there is a deeper twist the paper does not soften — but also does not overstate. That zero is not bad luck: it traces back to the very symmetry that made the fold pair "real" in the first place. The earlier work *knew* this and deliberately placed the up/down reading in the "imbalanced" sector by design. What it never showed is that the actual committed state has any imbalance in that sector for the reading to pick up — and a reading aimed at imbalance, applied to a perfectly balanced state, returns zero. So the real unproven step is narrow: *does the committed state carry the imbalance the reading needs?*

Among the **fixed, standing** structures the corpus describes, the answer looks like no — both available levels are either balanced by nature or barred from standing imbalance by another result. That would force zero contrast and, if standing structures were the only option, would falsify the mechanism (it would predict the up and down quark have equal mass, which is false). But there is an escape the corpus itself supplies: the imbalance might live not in a standing structure at all, but in a **transient, changing** quantity that the balanced-state argument simply does not cover. So the whole programme now turns on two questions in order: first, *must* the up/down contrast be a standing quantity, or can it be transient? — and only if it must be standing, *can* a standing structure carry the needed imbalance at all? If the contrast must be standing and cannot carry imbalance, the mechanism is cleanly falsified, and that is a real result worth publishing. If it can be transient, the balanced-state objection was aimed at the wrong target, and a genuine, answerable problem remains with a real shot at the measured 0.46. This paper does not settle which; it shows these are *the* questions, in this order, and that the honest indicators within the standing sector point toward "no" — while the transient escape is live and unexamined.

Abstract

The earlier χ -halving construction identified the physical fold carrier with the E_1 irreducible sector of $H^1(W_7; \mathbb{C})$, the six-dimensional cycle space of the wheel graph. Two corpus facts make that identification untenable at the carrier level. First, the sequential-transport strand treats the $K=7$ architecture as a 2-complex whose physical transport homology collapses to the outer-cycle class, $H_1(W_7) \cong \mathbb{Z}\langle [C] \rangle$, and restricts admissible spoke corrections to $\Delta_i = \lambda_i C$ — a one-dimensional transport-current sector that excludes E_1 (since C is the uniform $k=0$ mode and E_1 is the $k=\pm 1$ wave-mode pair, not a multiple of C). Second, the readout $a_{\pm} = \text{Tr}(P_{\pm} R)$ is carrier-dependent — P_{\pm}, R live on the cycle space — and the corpus supplies no readout carrier defined independently of that homology.

This paper resolves the resulting tension not by inflating the readout carrier to rescue E_1 , but by a **layer split** the later corpus forces. "Carrier" had been used for three distinct objects: the **transport-current shadow** (one-dimensional, $[C]$); the **class-grain standing carrier** (the PFD representation class, where the Carrier Theorem places quantized standing response, since under bath transport the class — not the member — is the minimal transport-stable subject); and the **member-grain registration layer** (seats read class-owned structure under the Realization Theorem's premises, owning no independent standing content). The fold/species structure lives at the class grain. The transport-current sector is its shadow, not the fold carrier. E_1 is neither: it is

demoted to a candidate readout representation — a basis in which a binary fold-access registration over the class carrier might be expressed — graded as not established to be the carrier.

The consequence for the readout is sharp and is the paper's central finding. The Generation Theorem forbids a standing generation invariant at member grain, so a_+ , a_- cannot be two member-owned fold amplitudes; they must be class-grain response components or member-grain registrations of class-owned structure. But the reflection S exchanging the fold projectors, $SP_+S = P_-$, forces $a_+ = a_-$ — hence $\chi = \ln a_+ - \ln a_- = 0$ — for **every** reflection-invariant standing state (verified exactly). The layer split moves this contrast no-go *up* to class grain, where the committed standing state is — under the current corpus reading — symmetric, removing the member-level asymmetry one might have hoped to hide behind. The target is the scheme- and scale-dependent first-generation ratio $m_u/m_d \approx 0.4596$ ($\chi \approx -0.78$), but the no-go does not depend on that value: it requires only the scheme-independent fact $m_u \neq m_d$, against which a symmetric state's prediction of ratio exactly 1 fails qualitatively.

The verdict is therefore a reframing, not a survival check. The old carrier contradiction is **mostly dissolved** by the layer split — at the carrier level, conditional on the adopted layer ontology, with the unverified $H_1 \cong \mathbb{Z}\langle[C]\rangle$ quotient downgraded to Layer-1 bookkeeping rather than formally settled; the old E_1/H^1 fold paper **cannot stand as written**; and the contrast no-go is located precisely. It is *not* "a theorem the construction missed": the three-gate paper anticipated the symmetry and placed χ outside the D_6 -invariant algebra by design, as an odd functional ($\ell_\chi \circ J = -\ell_\chi$, Prop 2.1). The unverified step is tighter — placing an odd readout is not the same as showing the **committed standing state carries odd-sector weight** for that readout to see; an odd functional on a state with zero odd weight returns zero. Among **standing** states the layer ontology closes both grains to odd weight (class grain symmetric by nature; member grain forbidden standing asymmetry by the Generation Theorem), which would force $\chi = 0$ — **but only if the contrast is required to be standing**. The corpus supplies a named non-standing candidate, the transient refinement-driven access operator, which lies outside the standing no-go entirely. So the programme turns on **two** gates in order: **(B0-prior)** is the χ contrast required to be a standing quantity, or may it be transient? — and, only if standing is required, **(B0)** does the standing carrier permit reflection-odd weight at all? If standing is required and B0 returns No, $\chi = 0$ is a theorem and the mechanism is **falsified** (degenerate first-generation masses). If transient is allowed, the standing no-go audits the wrong sector, no falsification follows, and the question becomes the dynamical one of whether the transient operator carries odd weight (B2), after which the representation question (B1) for E_1 remains. The standing-sector indicators lean toward No; whether that No bites depends entirely on B0-prior, which the corpus has not settled. **[Reframing — carrier contradiction dissolved at the carrier level, conditional on the adopted layer ontology; E_1 demoted to candidate readout; the no-go located as the unshown odd-sector weight of the standing state, and gated behind the standing-vs-transient question B0-prior that decides whether B0's likely No is fatal or misaimed.]**

Table of Contents

- Summary for the General Reader
 - Abstract
1. What Went Wrong: "Carrier" Used at Three Levels
 2. The Original Conflict, Stated Exactly
 3. The Three Layers
 4. Why E_1 Is Excluded from the Transport-Current Sector
 5. The Demotion of E_1 to a Readout Representation
 6. What a_+ and a_- Can and Cannot Be
 7. The Contrast No-Go and Why the Layer Split Sharpens It
 8. The Bridge Ladder — and the Gate Before It
 9. Status Table
 10. Conclusion: The Gate the Programme Turns On
- Appendix A — The Transport-Current Exclusion (C vs E_1)
 - Appendix B — The Symmetric-State Contrast Collapse
 - Appendix C — Referee Objections and Replies

1. What Went Wrong: "Carrier" Used at Three Levels

The earlier construction said: the physical fold carrier is $E_1 \subset H^1(W_7; \mathbb{C})$, the two-dimensional sector of the six-dimensional wheel cycle space on which the up/down access contrast χ lives. A corpus audit shows this conflates objects that the later papers force apart. The word *carrier* was doing three jobs:

- naming the **transport-current** bookkeeping (how structure is carried step to step);
- naming the **standing** structure that owns species and generation;
- naming the **readout basis** in which an access contrast is expressed.

These are not the same level, and E_1 was placed at the wrong one. The fix is not to argue E_1 into or out of the homology — that debate was a category error — but to separate the layers and ask, of each, what actually lives there. When that is done, the apparent contradiction between the two strands mostly dissolves, and E_1 lands where it belongs: a candidate readout representation, not a carrier.

2. The Original Conflict, Stated Exactly

Strand 1 — large cycle/access workspace. The χ -construction treats W_7 as a graph with no filled 2-cells, so its cycle space is six-dimensional, $\dim H^1(W_7; \mathbb{C}) = 12 - 7 + 1 = 6$, and an E_1 doublet lives inside it (the $k = \pm 1$ standing pair on the rim hexagon).

Strand 2 — K=7 transport 2-complex. The sequential-transport strand treats the K=7 architecture as a 2-complex whose physical transport homology collapses to the outer cycle,

$$H_1(W_7) \cong \mathbb{Z}\langle[C]\rangle,$$

and restricts admissible spoke corrections to the outer-cycle direction,

$$\Delta_i = \lambda_i C,$$

so the admissible transport-current sector is one-dimensional at the homology-shadow level.

Stated as carrier claims, these conflict: Strand 1's "the fold carrier is the full six-dimensional cycle workspace" against Strand 2's "the physical transport homology is one-dimensional." The earlier paper sided with Strand 1 and was therefore in genuine trouble — Strand 2 explicitly reduces H^1 to the rim generator via the 2-cell structure, while the bare graph has rank 6 only because it has no 2-cells. The two cannot both be the physical carrier.

The resolution is not to pick a winner. It is to notice that neither strand was naming the fold/species carrier. Strand 2's $[C]$ is the transport-current *shadow*; Strand 1's six-dimensional space is an *access workspace*. The thing that owns the fold — species, generation, standing response — is a third object the later corpus supplies.

3. The Three Layers

The Carrier Theorem states that quantized standing response attaches to the **minimal transport-stable subject**. Under bath transport, individual members are not standing transport-stable subjects — their identity is realized, not standing — but the **class** is. Member grain lies *below* quantization grain. This forces a three-layer ontology:

Layer 1 — Transport-current workspace (shadow). The σ /spoke-current sector. Collapses to the outer-cycle class $[C]$. One-dimensional. **This is not the fold/species carrier** — it is the bookkeeping shadow of the transport dynamics.

Layer 2 — Class-grain standing carrier. The PFD representation class. **This is where standing response, generation, and species identity live.** The minimal transport-stable subject; the genuine carrier.

Layer 3 — Member-grain registration workspace. Seats/members can be probed, but they *read* class-owned structure. By the Realization Theorem, under its registration premises a member-grain record reads q/k — class winding over class capacity — permanently and identically across seats. Members own no independent standing carrier content.

The fold/species carrier is Layer 2. The transport-current $[C]$ is Layer 1, its shadow. And the readout — where E_1 was placed — is at most a *basis* over Layer 2, expressed through Layer 3 registrations. The old paper put the fold carrier in a sector of the bare-graph cycle space, which is

neither Layer 1 nor Layer 2: it was a workspace floating free of the standing structure that actually owns the fold.

4. Why E_1 Is Excluded from the Transport-Current Sector

For completeness, the exclusion that started the audit is exact and worth keeping, now correctly scoped to Layer 1. Let T_0, \dots, T_5 be the six triangular cycles. The outer cycle is their uniform sum,

$$C = T_0 + T_1 + \dots + T_5,$$

the rotationally invariant $k = 0$ mode. The E_1 modes are the wave-like combinations

$$T_0 + \omega T_1 + \omega^2 T_2 + \dots \text{ and its conjugate, } \omega = e^{i\pi/3},$$

which are *not* multiples of C (distinct C_6 Fourier modes; orthogonal). So once admissible transport is restricted to $\mathbb{Z}\langle C \rangle$, the E_1 pair is absent from the transport-current sector (Appendix A). This is real, but it is now a statement about **Layer 1 only**: E_1 is not a transport current. It says nothing for or against E_1 as a *readout representation* over Layer 2 — which is the role the paper now assigns it, and a different question entirely.

(Qualification, kept honest: the admissible-cycle rule is stated as a restriction on spoke-correction currents, and the specific 2-cell boundary relations producing $H_1 \cong \mathbb{Z}\langle [C] \rangle$ are not laid out clearly enough in the corpus to verify independently. They sit beside the χ paper's no-2-cells treatment. The layer split makes this no longer decisive: whether Layer 1 is exactly one-dimensional, it is the *shadow*, not the fold carrier, so the fold question does not hinge on the unverified quotient.)

5. The Demotion of E_1 to a Readout Representation

E_1 cannot continue to be claimed as the physical carrier. It is the wrong level: the carrier is the class-grain standing structure (Layer 2), and E_1 is a sector of a cycle/access workspace, not a representation of class-owned standing structure — not yet, anyway. The honest status:

E_1 is a candidate readout representation. The $k = \pm 1$ (equivalently the "1, 5") boundary-harmonic doublet is a candidate basis in which a binary fold-access registration over the class carrier might be expressed. It is **not** the underlying transport-current carrier (Layer 1, which is $[C]$) and **not** the class carrier itself (Layer 2). **[Candidate — representation/readout basis; not established as carrier; the representation bridge B1 is open.]**

This is the safe, true statement. It does not assert E_1 floats as an independent physical workspace, which the layer ontology forbids. To upgrade E_1 from candidate readout to genuine content, one would have to prove

$$E_1 \rightarrow \text{class-grain fold/access split,}$$

not merely $E_1 \subset H^1(W_7)$. That is bridge B1 (§8). Until it is proved, E_1 is retained only as a representation hypothesis, and the paper does not lean on it for any physical conclusion.

6. What a_+ and a_- Can and Cannot Be

The access readings a_+ , a_- were the operational heart of the χ mechanism: $\chi = \ln a_+ - \ln a_-$, and χ is what becomes the up/down susceptibility. The layer split forces a reinterpretation of what they are — this is not notation, it is ontology.

What they cannot be. The Generation Theorem forbids a standing generation invariant at member grain: such an invariant must live at class grain or above. So a_+ , a_- **cannot be two member-owned fold amplitudes** — two independent standing weights belonging to individual seats. Under bath transport member identity is realized, not standing, and a member-level standing asymmetry is exactly the forbidden object. The old reading of a_+ , a_- as primitive member-owned fold weights is dead.

What they must be instead. One of:

- **class-grain response components** of the class-owned standing carrier; or
- **member-grain registration outputs** that read class-owned structure (Realization Theorem); or
- **representation-basis projections** of the class-owned carrier onto a readout basis (the candidate role of E_1).

In every admissible reading, a_+ and a_- are *derived from class facts*, not primitive seat variables. The safe statement: **a_+ , a_- are candidate registration/projection weights of class-owned standing fold structure, not primitive fold-carrier weights.** This is a substantial correction to the χ mechanism's foundation, and it has a consequence the next section shows is severe.

7. The Contrast No-Go and Why the Layer Split Sharpens It

Here is the cost, and it is the paper's central result. Whatever a_+ , a_- are, they are read by the trace form $a_{\pm} = \text{Tr}(P_{\pm} R)$ with P_+ , P_- the two fold projectors and R the standing access state. The corpus fixes the reflection S exchanging the fold lines:

$$S P_+ S = P_-.$$

For **any** reflection-invariant standing state, $S R S = R$, the trace readings are forced equal:

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

hence

$$\chi = \ln a_+ - \ln a_- = 0, a_+/a_- = 1.$$

This is exact, not approximate, and holds for every symmetric R (verified over random symmetric positive states, Appendix B). Against the physical first-generation ratio $m_u/m_d \approx 0.4596$ — i.e. $\chi \approx -0.78$ — the symmetric state predicts zero contrast and unit ratio. No mass splitting at all.

A caveat on the target, which in fact strengthens the no-go. Quark masses are not directly measured the way the electron or proton mass is; they are renormalization-scheme- and scale-dependent inferred parameters, and any single quoted ratio (here the standard low-energy/current-quark value) carries that dependence. The no-go does **not** rest on the precise number 0.4596. It requires only the scheme-independent fact that $m_u \neq m_d$ — the up and down quark are not degenerate in any scheme. The mechanism on a symmetric state predicts ratio exactly 1 (degeneracy); nature requires not-1. So the argument is robust to all scheme- and scale-dependence: it is about the qualitative gap between predicted degeneracy and observed non-degeneracy, not about hitting a specific value.

The gap is older than this paper — but it is not what "a theorem all along" would suggest. It would be an overstatement to say the three-gate construction failed to notice that the symmetry kills the contrast. It noticed precisely, and answered by design: Prop 2.1's observable split placed χ *outside* the D_6 -invariant algebra, as an explicitly odd functional, $\ell_\chi \circ J = -\ell_\chi$. So the readout was put in the odd sector deliberately; the earlier paper did not miss the symmetry. The honest location of the gap is therefore tighter, and the paper must not slide past it: **placing an odd functional on the table is not the same as showing there is anything in the odd sector for it to read.** The reflection S here is the same S used as transport B and as the realiser of the exchange involution J (bridge (c), $SP+S = P-$) — the operator that made the fold sector "genuine." What the three-gate paper asserted was odd-sector *access*; what it never showed is that the **committed standing state carries odd-sector weight** for that access to see. An odd readout applied to a state with zero odd weight returns zero. That — not "the construction secretly fails" — is the precise, defensible form of the no-go: the readout's odd-sector placement was legitimate; the unverified step is that the standing state has odd weight at all. The layer split's contribution is to make that unverified step impossible to keep parking outside an algebra.

And the layer split appears to leave the odd weight nowhere to live — among *standing* states, and even there only under an assumption that is itself the open question. Here is the squeeze, stated plainly because if it closes it is fatal — but stated with its two jaws graded differently, because they are not equally firm. The contrast needs reflection-**odd weight** in whatever state the readout sees — without it, the odd functional reads zero. Among **standing** states ($SRS = R$), the three-layer ontology has two jaws:

- **Member jaw (firm).** The member grain is forbidden, by the Generation Theorem, from carrying *standing* asymmetry at all. This is a theorem of the corpus, not a naturality claim; the member jaw is shut.
- **Class jaw (not established as forced).** The class-grain standing state is, *under the current corpus reading*, naturally reflection-symmetric — argued (§3, Carrier Theorem) to be the standing structure of the *whole class*. But whether this naturality is **forced** or merely **typical** is not settled here; it is precisely the B0 question below. The class jaw is shut only under an unproven symmetry-naturality assumption.

So the standing-sector squeeze is **not airtight**: one jaw is a theorem, the other is an assumption. If the class-grain symmetry is forced, both jaws shut, no standing layer supplies odd weight, and $\chi = 0$ on every standing state. If it is merely typical, the class jaw was never shut — some class standing states may be legitimately odd, and B2 has somewhere to live even within the standing sector. The paper must not assume the answer: it poses B0 as exactly the question of whether the class jaw is forced or typical, and does not assert it is forced.

But there is a second, independent reason the squeeze may not close, prior even to the class-jaw question: the contrast need not be *required to be standing* at all. There is a named candidate for odd weight that is **not** a standing object: the **transient access operator** (Candidate F of the Pushforward Verdict), depth-dependent and refinement-driven, neither a symmetric class standing state nor a forbidden member-grain standing asymmetry. The entire no-go is a statement about *standing* states; a transient amplitude lives outside its scope. So the either/or has the horns the paper states rather than buries:

Either (i) the contrast must be a standing quantity *and* the class-grain symmetry is forced — then both standing jaws shut, $\chi = 0$ is forced, and the mechanism is falsified; **or** (ii) the class-grain symmetry is merely typical — then odd class standing states exist and B2 is well-posed in the standing sector; **or** (iii) the contrast may be **transient**, carried by the refinement-driven access operator — then the standing-state squeeze is the wrong audit entirely, and the question is the *dynamical* one of whether the transient operator carries odd weight. The paper asserts none of these. It flags that (i), (ii), (iii) are not yet distinguished, that the corpus supplies the (iii) candidate, and that (i) — the only fatal horn — requires *both* a standing-mandate (B0-prior) *and* a forced class symmetry (B0), neither established.

8. The Bridge Ladder — and the Gate Before It

The paper's honest deliverable is a carrier contradiction **dissolved at the carrier level, conditional on the adopted layer ontology** (the unverified $H_1 \cong \mathbb{Z}\langle[C]\rangle$ quotient downgraded to Layer-1 bookkeeping, not formally settled), a load-bearing gap located precisely (the standing state's odd weight was never shown, only the readout's odd-sector placement), and a ladder whose bottom is **two** questions that must be settled in order — and the first determines whether the second can falsify anything.

Because the bottom gate turns on the words "standing" and "transient," they need operational criteria, not just the conceptual contrast:

Definitions for B0-prior.

- **Standing quantity**: a quantity belonging to the transport-stable class carrier and invariant under the relevant standing-state conditions — part of class identity, preserved across refinement.
- **Transient quantity**: a refinement-driven, access-dependent, or dynamical registration quantity *not* required to be a standing invariant of the class carrier — present during access/refinement but not part of standing class identity.

- **Test:** if χ must be preserved as part of class identity, it is standing (the no-go applies). If χ appears only during access/refinement/registration, it may be transient (the no-go does not reach it). If χ is *both* required for class identity *and* generated during access — an ambiguous case the corpus may license — then the standing/transient split is not clean and B0-prior has a genuine third branch to resolve, not a binary to pick.

Gate B0-prior — Sector (must be settled first). Is the physical χ contrast *required to be a standing quantity*, or may it be **transient** (carried by the refinement-driven access operator, Candidate F)?

- **Standing required:** then odd weight must live in a standing state, and B0 below becomes load-bearing and potentially fatal.
- **Transient allowed:** then the standing-state no-go audits the wrong sector; the squeeze does not close; the question shifts from the ontological B0 to the dynamical one of whether the transient access operator carries odd weight. **No falsification follows from the standing-sector audit in this branch.**

[Open — prior to B0; decides whether B0's No is fatal or irrelevant. The corpus supplies the transient candidate, so this branch is live, not hypothetical.]

Gate B0 — Standing permission (load-bearing only if B0-prior = standing required). Determine, from the Carrier, Generation, and Realization theorems *as stated*, whether the class-grain standing carrier admits reflection-odd weight at all.

- **No (forbidden) and standing required:** $SRS = R$ always, $\chi = 0$ is a **theorem**, degenerate up/down masses, **clean falsification**.
- **No (forbidden) but transient allowed:** the standing sector is odd-free, but the contrast lives elsewhere; not a falsification — the audit was aimed at the wrong sector.
- **Yes (permitted):** "naturally symmetric" (§7) weakens to "generically symmetric," and B2 becomes a well-posed standing-sector dynamical question.

[Open — binary, answerable from the three theorems; fatal *only* in conjunction with B0-prior = standing required. Standing-sector indicators lean No; whether that No bites depends entirely on B0-prior.]

Bridge B2 — Symmetry breaking (conditional on a sector that permits odd weight). In whichever sector B0-prior selects, derive a state carrying odd weight — a standing class state with $SRS \neq R$ (if standing is required and permitted), or a transient access operator with odd weight (if transient is allowed) — from independently motivated dynamics, so that $a_+ \neq a_-$ and $\chi \neq 0$. The derived object must not import the quark ratio it aims to explain — no fitted offset, no normalization drawn from the target — or the bridge is circular. [Open — conditional on B0-prior/B0; the no-go stands on symmetric *standing* states regardless.]

Bridge B1 — Representation (conditional on B2). Prove that the E_1 boundary-harmonic doublet genuinely *represents* the observable binary fold-access split over the carrier — i.e. $E_1 \rightarrow$

fold/access split, not merely $E_1 \subset H^1(W_7)$. Until then E_1 is a candidate readout basis, load-bearing for nothing physical. [**Open — representation bridge.**]

The ladder is climbed bottom-up, and the paper does **not** assert the upper rungs are reachable — nor, crucially, does it assert the falsification. B0-prior is the true first question: it decides whether the standing-sector no-go is a falsification (standing required) or a misaimed audit (transient allowed). Only if standing is required does B0's likely No become fatal. So the next substantive work is **not** B2, not the descent-machine reconciliation, and not even B0 alone — it is **B0-prior**: settle whether the physical χ contrast must be standing or may be transient. That single distinction decides whether the programme is one audit away from falsification or one audit away from relocating the contrast into a sector the no-go never covered.

9. Status Table

Item	Status	Grade
$\dim H^1(W_7; \mathbb{C}) = 6$, D_6 -decomp $A_2 \oplus B_1 \oplus E_1 \oplus E_2$	Exact (prior audit)	[Proven]
Outer cycle $[C] = A_2$ (uniform $k=0$)	$r(C)=C$, $m(C)=-C$	[Proven]
$E_1 = k=\pm 1$ doublet, \neq multiple of C	Distinct Fourier modes	[Proven]
E_1 excluded from transport-current sector $\mathbb{Z}\langle C \rangle$	Appendix A	[Proven — Layer 1 only]
$H_1(W_7) \cong \mathbb{Z}\langle [C] \rangle$ via 2-cell quotient	2-cell relations not laid out to verify; beside no-2-cells treatment	[Open — corpus inconsistent; not decisive after layer split]
Three-layer ontology (shadow / class / registration)	Carrier, Realization, Generation Theorems	[Adopted — from later corpus]
Fold/species carrier = class-grain standing structure	Carrier Theorem (class is minimal transport-stable subject)	[Adopted]
E_1 as physical carrier	Withdrawn — wrong level	[Rejected]
E_1 as candidate readout representation	Retained as hypothesis	[Candidate — bridge B1 open]
a_+ , a_- as member-owned fold amplitudes	Forbidden by Generation Theorem at member grain	[Rejected]
a_+ , a_- as class registration/projection weights	The admissible reading	[Adopted — reinterpretation]
Symmetric state $\Rightarrow a_+ = a_- \Rightarrow \chi = 0$	Exact ($SP_+S=P_-$), Appendix B	[Proven — the no-go]
No-go's origin: readout placed odd by design (Prop 2.1, $\ell \chi \circ J = -\ell \chi$)	Three-gate paper <i>anticipated</i> the symmetry; placed χ odd deliberately	[Proven — not a missed theorem]

Item	Status	Grade
Unshown step: standing state carries odd-sector weight	Odd readout on zero-odd-weight state returns 0	[Open — the actual gap]
Class-grain carrier is naturally symmetric	Standing structure of whole class	[Argued — and possibly too strong; see B0]
Both grains closed to odd standing amplitude	Class symmetric by nature; member forbidden by Generation Theorem	[Argued — the ontology squeeze]
B0-prior — is the χ contrast required to be standing, or may it be transient?	Decides whether B0's No is fatal or misaimed; corpus supplies transient candidate (access operator)	[Open — the true first question; prior to B0]
B0 — does the standing carrier permit reflection-odd weight?	Binary; from Carrier/Generation/Realization theorems; fatal <i>only if</i> B0-prior = standing required	[Open — indicators lean No; bites only under B0-prior=standing]
Symmetry-breaking standing state R, $\chi \neq 0$ (B2)	Conditional on B0 = Yes; not derived	[Open — conditional on B0]
$E_1 \rightarrow$ class-grain fold/access split (B1)	Conditional on B2	[Open — representation bridge]
χ = physical quark susceptibility	Untouched upstream bridge	[Open]

10. Conclusion: The Gate the Programme Turns On

The carrier contradiction is mostly dissolved — at the carrier level and conditional on the adopted layer ontology, not formally settled at the homology-detail level — and dissolved honestly: not by inflating a readout space to rescue E_1 , but by recognizing that "carrier" named three different layers. The transport-current sector ([C], Layer 1) is the dynamics' shadow. The fold/species carrier is the class-grain standing structure (Layer 2). Member seats register class-owned content (Layer 3). E_1 , the old paper's physical carrier, is demoted to a **candidate readout representation** over Layer 2, and the access weights a_+ , a_- are reinterpreted from primitive member amplitudes — forbidden at member grain — to registration/projection weights of class-owned structure. The two strands stop contradicting each other; the programme no longer has to force E_1 to survive as physical homology. That is the resolved part, and it is real.

But the resolution exposes something heavier than it fixes, and honesty requires stating it precisely — neither overstated nor hidden. The contrast no-go is not a problem the carrier audit introduced, but neither is it a theorem the three-gate construction missed: that paper anticipated the symmetry and placed χ outside the invariant algebra by design, as an odd functional (Prop 2.1, $\ell_{\chi} \circ J = -\ell_{\chi}$). The unshown step is tighter. Placing an odd readout is not the same as showing the committed standing state has odd-sector weight for it to read; an odd functional on a zero-odd-weight state returns zero. Among **standing** states the layer ontology closes both grains to odd weight (class grain symmetric by nature, member grain forbidden standing asymmetry by the Generation Theorem) — which forces $\chi = 0$ *if the contrast must be standing*. But the corpus

names a non-standing candidate, the transient refinement-driven access operator, outside the standing no-go's scope. So the squeeze closes only under an unsettled premise.

So the programme does not have a comfortable open bridge. It has a gate that *may* be a wall — but whether it is a wall depends on a prior question the paper must not skip. Six papers ago the open question was a vague "carrier identification." It has now been driven to two sharp, finite questions that must be answered in order:

(B0-prior) Is the physical χ contrast required to be a standing quantity, or may it be transient? (B0) — only if standing is required — does the standing carrier permit reflection-odd weight, yes or no?

If standing is required and B0 returns No — if the Carrier, Generation, and Realization theorems force symmetric standing states — then $\chi = 0$ is a theorem not a gap, the mechanism predicts degenerate first-generation masses, and the programme is **falsified**, cleanly and publishably. But if the contrast may be **transient**, the standing-state no-go audits the wrong sector: the corpus's own transient access operator lies outside it, the squeeze does not close, and the question is no longer the ontological "can a standing state be odd" but the dynamical "does the transient operator carry odd weight" (B2). The indicators that lean toward No lean only *within the standing sector*; they say nothing about the transient one.

The next paper is therefore neither the descent-machine reconciliation nor B2 nor even B0 alone — all three presuppose the sector is settled. It is **B0-prior**: determine whether the physical χ contrast must be standing or may be transient. That single distinction decides everything downstream — whether the likely standing-sector No is a falsification or a misaimed audit, and whether the programme is one step from being falsified or one step from relocating its contrast into a sector the no-go never reached. The audit must be run from the corpus forward, prepared to publish a falsification if the contrast is forced standing and the standing sector is odd-free — and equally prepared to find that the question was aimed one sector off. The series' credibility rests on not prejudging which.

E_1 is not the carrier; the class is. The three-gate paper put the readout in the odd sector by design — what it never showed is that the standing state has odd weight to read. Among standing states there may be none; whether that is fatal depends on whether the contrast must be standing at all. That — not a flat falsification — is the one question left.

Appendix A — The Transport-Current Exclusion (C vs E_1)

In the f-basis (triangular cycles), $C = \sum T_i$ is the constant $k = 0$ Fourier mode (rotation eigenvalue +1, reflection-odd $\Rightarrow A_2$). The E_1 modes are $\sum \omega^i T_i$ and its conjugate, $\omega = e^{i\pi/3}$ (rotation eigenvalues $e^{\pm i\pi/3} \Rightarrow E_1$). Distinct C_6 Fourier modes are orthogonal: $\langle E_1, C \rangle = 0$ (machine precision). Hence restriction to $\mathbb{Z}\langle C \rangle$ retains only A_2 and excludes E_1 . This is a Layer-1 (transport-current) statement; it does not bear on E_1 as a readout representation over the class carrier.

Appendix B — The Symmetric-State Contrast Collapse

With P_+ , P_- the fold projectors and S the reflection exchanging them, $SP_+S = P_-$ (verified). For any reflection-invariant state $SRS = R$:

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

so $\chi = \ln a_+ - \ln a_- = 0$ exactly. Tested over 10^4 random symmetric positive-definite R on the E_1 doublet: $a_+ = a_-$ in every case, zero exceptions. Physical target $\chi = \ln(0.4596) \approx -0.78$; symmetric-state prediction $\chi = 0$.

Appendix C — Referee Objections and Replies

Objection 1 — You inflated the readout carrier to save E_1 . The opposite: E_1 is *demoted*, not saved. The paper withdraws the carrier claim entirely and relocates the fold to the class-grain standing structure. E_1 is retained only as a candidate readout basis, load-bearing for no physical conclusion, with its upgrade (B1) explicitly open.

Objection 2 — The layer split is ad hoc, introduced to dodge the contradiction. The three layers are not introduced here; they are the Carrier, Realization, and Generation Theorems of the later corpus. The split is *applied*, not invented, and it is applied symmetrically — it costs the paper its old E_1 result and sharpens the contrast no-go against the programme. An ad hoc rescue would not damage its own mechanism; this one does.

Objection 3 — The 2-cell quotient $H_1 \cong \mathbb{Z}\langle[C]\rangle$ is still unverified, so the contradiction isn't really resolved. Correct that it is unverified, and the paper says so. But after the layer split the fold question no longer hinges on it: Layer 1 (transport-current shadow) is not the fold carrier whether or not it is exactly one-dimensional. The unverified quotient is downgraded from decisive to a Layer-1 housekeeping question.

Objection 4 — If a_+ , a_- aren't member amplitudes, the whole χ readout is undefined. It is *reinterpreted*, not undefined: a_+ , a_- become registration/projection weights of class-owned structure (Realization Theorem). The trace form $a_{\pm} = \text{Tr}(P_{\pm}R)$ survives; what changes is that R is a class-grain standing state, not a member configuration. The cost is real — and it is exactly the contrast no-go of §7, which the paper promotes rather than hides.

Objection 5 — The contrast no-go was already known; you've added nothing. Two things are sharpened. First, the gap is located precisely: the three-gate paper legitimately placed χ in the odd sector (Prop 2.1), so the unverified step is not the readout's placement but the *unshown odd-sector weight of the committed standing state* — an odd functional on a zero-odd-weight state returns zero, and that the state has odd weight was asserted, never shown. Second, the layer split closes both *standing* grains to odd weight, which would force $\chi = 0$ — but only if the contrast must be standing. The same corpus that poses the squeeze supplies the escape (the transient access operator), so the paper does not claim falsification; it claims the no-go is real on standing states and gated behind the standing-vs-transient question (B0-prior).

Objection 6 — Does this kill the programme? No — it relocates the hardest problem into view. The carrier contradiction is dissolved at the carrier level (conditional on the layer ontology); the fold sits at class grain; E_1 is retained only as a candidate readout representation. What remains *first* is not B2 but **B0-prior**: determine whether the physical χ contrast must be a standing quantity or may be transient. If standing is required, B0 decides whether the standing carrier can carry reflection-odd weight; if it cannot, the mechanism is cleanly falsified. If transient contrast is allowed, the no-go has audited the wrong sector, and the next dynamical problem becomes whether the transient refinement-driven access operator carries odd weight without importing the observed quark ratio. B2 is reached only after that sector question is settled.