The Closure Operator Specification and the Realization Spectrum

Specifying the Realization Operator, Locating Its Construction Beneath the D_7 Census, and the First Constraints on Its Spectrum

Keith Taylor *VERSF Theoretical Physics Programme*

Scope note. This paper does ***not*** *construct the realization operator* $\mathfrak{A}$*. It specifies the operator whose construction is required, identifies the closure-architecture result on which that construction depends, and records the first constraints imposed by the* D_7 *reading of the fourteen-generator census. The single most important word in the title is "Specification": the operator is named and constrained here, not built. Its construction is gated on a closure-architecture result (the two-factors-of-two / census question) that is not settled within this paper.*

General Reader Summary

Several strands of the VERSF Standard Model programme point toward one idea: that the three generations of matter correspond to distinct stable "realization depths" inside a closure structure. The electron, muon, and tau would then not be three unrelated species but three rungs of one ladder.

There is a recurring obstacle. The programme describes the ladder but has never exhibited the *machine* whose rungs the ladder is — the operator whose list of allowed values (its spectrum) would be the depths. Without that operator there is no equation to solve, and "realization depth" stays a phrase rather than a calculation.

This paper does not build that machine. It does something more modest and, at this stage, more honest: it writes down the exact specification the machine must meet, and then shows that the machine cannot be built until a *prior* question is answered — what the closure architecture actually is. A separate result, concerning a structure called D_7 (the symmetry group of a seven-sided figure, with fourteen elements), is the thing that fixes that architecture. So the operator's construction is downstream of the census question, not independent of it. The two branches of the programme that had been drifting apart turn out to be one branch.

The paper then takes the first step of reconnaissance. If the closure architecture really is D_7, then the operator must respect that structure, and the representation theory of D_7 already says something striking: D_7 has *exactly three* two-dimensional irreducible representations, and a putative fourth is not new — it collapses onto the third. If generations correspond to those representations, then "exactly three generations" and "no fourth generation" would both follow from the group itself, rather than being assumed. There is one technical proviso the paper is careful about: this clean count holds only if the closure structure's "fourteen" is read as the *size* of the symmetry group rather than as the number of dimensions the operator acts on — the two readings give different mode counts, and the paper shows the clean three is the cheaper of the two.

This is presented as a lead, not a result. The matching of three-to-three is exactly the kind of coincidence that can be a genuine clue or a trap, and the paper says which tests would tell the two apart, rather than banking the match.

Contents

1. **The Demand: Realization Architecture and T2** — the open problem (at most three isolated depths?) the operator must answer to.
2. **The Operator, Specified Not Built** — conditions A1–A4, and why they fix neither cardinality nor termination.
3. **The Construction Dependency** — why 𝔄 cannot be built until the census fixes the closure architecture; the realization and participation branches share one Step 0.
4. **D_7-Leading Reconnaissance** — what the leading census reading imposes on 𝔄.
    - 4.1 — 𝔄 must respect or controllably break D_7.
    - 4.2 — The representation theory is fixed (two 1D, three 2D irreps).
    - 4.2a — The fourteen is an order, not a dimension (the multiplicity argument and GATE-DISCRIM(a′)).
    - 4.3 — What the irrep structure would supply, given the identification.
    - 4.4 — Relation to the other axes; P–D independence supported.
    - 4.5 — Discipline: the matches are leads or traps, not results.
5. **Completion Density and the Decision Protocol** — the pre-registered success/failure criteria and the ordering test.
6. **What This Paper Does and Does Not Claim.**

- **Condition Ledger.**
- **Conclusion.**

Abstract

The Realization Architecture programme proposes that fermion generations correspond to stable realization depths within an underlying closure structure, and the Realization Architecture Conjecture isolates the open problem T2: whether closure admits at most three isolated stable depths. Converting this proposal into a predictive mechanism requires an explicit operator whose spectrum supplies the depths. No such operator has been exhibited.

This paper specifies that operator. We require a realization operator $\mathfrak{A}$ acting on admissible realization states to satisfy closure preservation (A1), finite distinguishability (A2), spectral isolation (A3), and saturation compatibility (A4). We are explicit that **A1–A4 are a specification, not a construction**: they state the properties $\mathfrak{A}$ must possess if it exists and functions, and they do **not** by themselves fix the cardinality of the spectrum. In particular, "exactly three" is not delivered by A1–A4 and remains [Open] at the level of the specification alone.

The principal structural claim is that $\mathfrak{A}$ cannot be read off "closure geometry" until the closure architecture itself is fixed, and that the architecture is fixed by the fourteen-generator census result — the same D_7-versus-formation/preservation question that governs the participation-denominator branch. The construction of $\mathfrak{A}$ is therefore **downstream of the census paper**, which becomes Step 0 for the realization branch as well.

Taking the leading census reading $N_{\text{loop}} = |D_7| = 14$, we record the first constraints the dihedral structure places on $\mathfrak{A}$. The representation theory of D_7 is fixed: its irreducibles are two of dimension 1 and three of dimension 2 (total dimension 8; the order $|D_7| = 14$ is the sum of *squared* dimensions, not the dimension of the irrep list). So D_7 carries **exactly three** two-dimensional irrep types, and a putative fourth satisfies $\rho_4 \cong \rho_3$ rather than being independent. A subtlety governs the cardinality claim: since the irreps total dimension $8 < 14$, any 14-dimensional D_7 representation forces multiplicity ≥ 2, admitting up to six 2D eigenmodes — so "exactly three" holds only if the realization carrier is multiplicity-free, i.e. the census's 14 is the group *order* (D_7 acting on a lower-dimensional carrier) rather than the carrier dimension itself. This sharpens the census's discrimination gate (GATE-DISCRIM(a′)) and hands the census branch a cost asymmetry: the "acts-on" reading yields three clean modes for free, the "elements-are-carrier" reading only by an extra postulate. If realization eigenmodes are identified with the 2D irrep types under the multiplicity-free reading, three-ness and fourth-mode exclusion would follow from the group rather than from A4. This identification is graded [Conjectural]; the irrep counts themselves are [Established]. The paper closes with a pre-registered decision protocol and a condition ledger.

---

1. The Demand: Realization Architecture and T2

The Realization Architecture Conjecture organizes the fermion taxonomy as a product of three independent realization properties,

$F = A \times P \times D$,

with assignment A fixing weak structure, transport participation P fixing colour structure, and realization depth D fixing generation structure. That paper is explicit that the factorization itself — the mutual independence of A, P, D — is a separate conjecture (its C4), and that the depth axis rests on an unresolved problem:

T2. Does closure admit at most three isolated stable realization depths?

T2 is the demand this paper answers to. It is not a demand for a *value* (the masses) but for a *mechanism*: an object whose spectrum *is* the depth structure, so that "three depths," "their ordering," and "no fourth" become statements about that object rather than imported facts.

Grade of the demand: **[Inherited — Realization Architecture Conjecture, T2 Open].**

The missing object is the realization operator. The present paper specifies it and locates its construction; it does not build it.

---

2. The Operator, Specified Not Built

Let $\mathscr{R}$ denote the space of admissible realization states. A realization state is a closure-consistent configuration capable of supporting an observer-invariant realized subject; it carries assignment, transport-participation, refinement, and closure structure. The realization operator acts on the realization coordinate:

$\mathfrak{A} : \mathscr{R} \to \mathscr{R}$.

We require $\mathfrak{A}$ to satisfy four conditions. These are stated as a **specification** — the properties $\mathfrak{A}$ must have if it exists and reproduces the realization-depth picture — and explicitly not as a construction.

A1 — Closure preservation. Admissible realization states map to admissible realization states; closure-complete states do not map to closure-incomplete states. $\mathfrak{A}$ preserves $\mathscr{R}$. **[Specification]**

A2 — Finite distinguishability. The spectrum is discrete; no infinite accumulation of admissible realization classes occurs. **[Specification]**

A3 — Spectral isolation. Distinct realization classes correspond to isolated eigenmodes separated by nonzero gaps. **[Specification]**

A4 — Saturation compatibility. $\mathfrak{A}$ admits the possibility that no further isolated modes exist beyond some finite depth. **[Specification]**

Eigenstates and eigenvalues are written

$\mathfrak{A}\,\psi_n = \lambda_n\,\psi_n$,

with λ_n the realization depth of mode ψ_n, and the central identification is

generation ↔ stable realization eigenmode.

What the specification does not deliver. A2 and A3, imposed as admissibility conditions, make "the spectrum is discrete and isolated" true by construction. They therefore do **not** explain isolation; they assume it. Critically, **A1–A4 place no constraint on the cardinality of the spectrum**: nothing in the four conditions forces three isolated modes rather than two, five, or seventeen. The number three — the entire content of T2 — is untouched at the level of the specification. Likewise A4 *permits* termination but does not *produce* it; reading A4 as fourth-generation exclusion and then re-extracting that exclusion as a result would be circular. The honest status is:

Cardinality of the spectrum (the "three"): **[Open]** under A1–A4 alone. Termination (no fourth mode): **[Open]** under A1–A4 alone; A4 assumes its possibility, it does not establish it.

This is the precise sense in which the operator is specified, not constructed: A1–A4 tell us what $\mathfrak{A}$ must be like, not what $\mathfrak{A}$ *is*, and they leave the load-bearing physical questions open. Those questions can only be answered by an actual action for $\mathfrak{A}$, read off the closure architecture — which is the subject of §3.

3. The Construction Dependency

The phrase "construct $\mathfrak{A}$ from closure geometry" presupposes that the closure geometry is known. It is not yet fixed at the level of detail an operator action requires.

An action for $\mathfrak{A}$ — a definite rule by which $\mathfrak{A}$ acts on a basis of realization states — must be read off the closure architecture: the generators of closure, their composition, their orientation and frame structure. That architecture is exactly what the fourteen-generator census question settles. As recorded across the participation-denominator branch, the count N_loop = 14 is inherited, but its *internal structure* — how the fourteen relate to the seven closure constraints, and where the boundary frame sits — is the unresolved two-factors-of-two question. Until that question is decided, "closure geometry" is a name, not a structure $\mathfrak{A}$ can act on.

The consequence is structural and was not visible in the prior framing of the realization branch:

The construction of $\mathfrak{A}$ is **downstream of the census paper**. The census/D_7 result is Step 0 for the realization branch, not merely for the participation branch.

This links two branches that had been developing in parallel. The participation denominator (the "12") and the realization spectrum (the depths) are not independent problems with a shared vocabulary; they are two consumers of the *same* closure-architecture result. Whatever the census paper returns about the structure of the fourteen generators constrains both. **[Structural claim — Established by dependency analysis; the census result it depends on is itself Open.]**

This also relocates the honest next step. One cannot write the action of $\mathfrak{A}$ before the architecture is pinned. What one *can* do now is reconnaissance: take the leading census reading and ask what it already forces on $\mathfrak{A}$. That is §4.

---

4. D_7-Leading Reconnaissance

The leading reading from the transport analyses is that the fourteen generators are the elements of the dihedral group D_7,

$$N\_loop = |D_7| = 14, \quad D_7 = \langle\, r, s \mid r^7 = s^2 = 1,\ s\, r\, s^{-1} = r^{-1} \,\rangle,$$

rather than a formation/preservation doubling 7_form + 7_preserve. Under this reading the doubling is rotation/reflection structure (seven rotations r_j, seven reflections $s\, r_j$), not a commitment-temporal split. This section records what that structure imposes on $\mathfrak{A}$. Everything here is **doubly conditional**: on the census returning D_7 (not yet established), and on the identifications proposed below (each graded). The group-theoretic facts, by contrast, are fixed.

Two layers of robustness, kept separate. The rotation/reflection element split (seven rotations, seven reflections) is the (a)-capped picture: it reaches the realization carrier $\mathfrak{g}_{14}$ only through the census's GATE-DISCRIM(a), and so inherits that gate's conditionality. The *intrinsic irrep counts* of the abstract group D_7 — two 1D, three 2D — are **not** (a)-capped: they are properties of D_7 as a group and need neither the element split nor any resolution of (a). The reconnaissance results below rest on the intrinsic counts, and are therefore strictly more secure than the element scaffolding that introduces them — *except* at the one point where the irreps meet the realization carrier space, where the cap does bite (§4.2a). That exception is the load-bearing subtlety of this section and is isolated deliberately.

4.1 $\mathfrak{A}$ must respect or controllably break D_7

If the closure architecture carries a D_7 symmetry, an operator $\mathfrak{A}$ read off that architecture must either commute with the D_7 action or break it in a controlled, representation-theoretic way. Its eigenspaces must then organize into D_7-representations. Realization modes must therefore arise from **admissible representations, quotients, or filtrations of D_7** — not from an invented depth ladder imposed by hand. This is the central constraint the D_7 reading supplies: it removes the freedom to posit the spectrum and replaces it with the representation theory of a fixed finite group.

[Constraint — conditional on D_7; Established given D_7.]

4.2 The representation theory is fixed

For odd n, the dihedral group D_n has exactly two one-dimensional irreducible representations and $(n-1)/2$ two-dimensional ones. For $n = 7$ this gives five irreducibles in total. Two statements must be kept apart here, because conflating them is the source of the multiplicity subtlety in §4.2a.

The **irrep list** — five irreducibles, taken once each — has total dimension

$1 + 1 + 2 + 2 + 2 = 8$.

The **order formula** — the sum of *squared* dimensions, i.e. the dimension of the regular representation — is

$\Sigma d_i^2 = 1^2 + 1^2 + 2^2 + 2^2 + 2^2 = 14 = |D_7|$.

These are different numbers: 8 is the dimension of the multiplicity-free direct sum of the irreps; 14 is the group order. They must not be joined by a $\oplus$. The three two-dimensional irreps are indexed by $k \in \{1, 2, 3\}$, with

$\rho_k(r) = \mathrm{diag}(\omega^k, \omega^{-k})$, $\omega = e^{(2\pi i/7)}$, $\chi_k(r^j) = 2\cos(2\pi kj/7)$, $\chi_k(\text{reflections}) = 0$.

These facts are not conjectural; they are the character table of D_7, and they are intrinsic to the abstract group (not (a)-capped). **[Established — representation theory of D_7.]**

4.2a The fourteen is an order, not a dimension

The distinction just drawn is load-bearing, because the census's "14" can be read two ways and the readings diverge on exactly the cardinality claim §4.3 depends on.

If the census's 14 is the group *order* $|D_7|$ (the Σd_i^2 statement), with D_7 **acting on** a realization carrier of lower dimension, then the natural carrier is the multiplicity-free space $2 \oplus 2 \oplus 2$ (dim 6), or dim 8 with the 1D sector attached. The three 2D irreps appear *once each*, and "exactly three modes" is clean.

If instead the census's 14 is read as the *dimension* of the realization carrier (the fourteen elements *are* the carrier $\mathfrak{g}_{14}$), the situation is forced and different. The irreps total only dimension $8 < 14$, so **any** 14-dimensional representation of D_7 must repeat at least one irrep — multiplicity ≥ 2 is unavoidable, not optional. The natural 14-dimensional carrier is the regular representation, in which each 2D irrep appears with multiplicity two:

$\mathrm{reg}(D_7) \cong 1 \oplus 1 \oplus 2\rho_1 \oplus 2\rho_2 \oplus 2\rho_3$, dimension $1 + 1 + 4 + 4 + 4 = 14$.

Now run the identification "eigenmode ↔ 2D irrep" against this carrier. If $\mathfrak{A}$ commutes with the D_7 action and ρ_k appears with multiplicity two, the ρ_k-isotypic component is 4-dimensional, and by Schur's lemma $\mathfrak{A}$ restricted to it acts as a 2×2 matrix **on the multiplicity space** — which carries up to two distinct eigenvalues, i.e. two distinct realization depths both transforming as ρ_k. Multiplicity two therefore admits up to **six** 2D eigenmodes, not three. "Exactly three" is recovered only if one both (i) counts 2D irrep *types* (isotypic components) rather than eigenmodes, **and** (ii) forces $\mathfrak{A}$ to act as a scalar on each multiplicity space. Neither was stated in the original reconnaissance, and (as just shown) no 14-dimensional D_7 representation is multiplicity-free, so the dim-14 reading cannot deliver three clean modes for free.

GATE-DISCRIM(a′). "Exactly three" requires not merely some resolution of the census's GATE-DISCRIM(a), but specifically a realization carrier in which the three 2D irreps are **multiplicity-free** — which is incompatible with $\dim \mathfrak{g}_{14} = 14$. This sharpens the gate into a stated, decidable dependency, and it produces a **cost asymmetry** the realization branch can hand back to the census branch:

- the |G|-**acts-on** horn (14 = group order, lower-dim carrier) yields three clean modes *for free*;
- the **elements-are-carrier** horn (14 = $\dim \mathfrak{g}_{14}$, regular rep) yields three only by *paying an extra postulate* — that $\mathfrak{A}$ acts as a scalar on each multiplicity space — on top of accepting forced multiplicity two.

This is a cost comparison, not a verdict: the elements-are-carrier horn is not refuted (a defender may accept the extra scalar-on-multiplicity postulate), it is shown to be *more expensive*. The realization branch thereby offers the census branch a genuine constraint it did not previously have — *if three clean modes are wanted, the acts-on resolution of GATE-DISCRIM(a) is cheaper* — without legislating how the census branch must resolve. **[Established as a conditional cost asymmetry; the underlying GATE-DISCRIM(a) remains Open.]**

Note the interaction with §4.1: §4.1 leans on the "representations / quotients / filtrations" picture (the acts-on horn), while a naive tie of the irreps to "14" leans on the order/regular-rep picture (the elements horn), and these pull in opposite directions on precisely the cardinality payload. The fix above makes the section internally consistent by committing the cardinality claim to the acts-on, multiplicity-free reading.

4.3 What the irrep structure would supply — if the identification holds

Propose the identification

realization eigenmode ↔ two-dimensional irreducible representation of D_7.

Three consequences follow *from the group*, conditional only on this identification:

(i) Exactly three. There are exactly three two-dimensional irrep *types*, ρ_1, ρ_2, ρ_3. Under the identification — *and* under the multiplicity-free (acts-on) reading required by §4.2a — the number of generations is fixed at three by D_7, not assumed via A4, not fitted. This is precisely the cardinality content that A1–A4 alone could not supply (§2). The caveat from §4.2a is carried explicitly: the count is "three" only when the three 2D irreps appear once each; under the dim-14 regular-rep reading it becomes "up to six" unless $\mathfrak{A}$ is separately forced scalar on each multiplicity space.

(ii) No fourth — by fold-back, not by fiat. A putative fourth two-dimensional mode, $k = 4$, is not independent:

$\rho_4 \cong \rho_3$, since $\cos(2\pi\cdot 4/7) = \cos(2\pi\cdot 3/7)$ (equivalently $\omega^4 = \omega^{-3}$).

The indices fold in pairs $\{k, 7-k\}$: $\{1,6\}$, $\{2,5\}$, $\{3,4\}$, leaving exactly the three distinct irreps $k = 1, 2, 3$. A fourth refinement therefore *folds into an existing realization class* rather than producing a new one. This is the concrete realization of the "fold-back" failure mode hypothesized qualitatively in the Realization Architecture Conjecture (§12) and the saturation papers — there it was one item in a disjunction asserted to be exhaustive; here, under D_7, it is forced by the structure. The distinction in standing must be kept sharp: $\rho_4 \cong \rho_3$ is a *theorem* of representation theory, but "no fourth generation" is that theorem **plus three physics conditionals** — that the census returns D_7, that it resolves to the multiplicity-free acts-on carrier (§4.2a), and that generation ↔ 2D irrep. With those conditionals, fourth-generation exclusion would become a consequence of the group structure rather than an open conjecture; without them, the rep-theory certainty must not be banked for the physics claim.

(iii) Natural ordering label. The three 2D irreps carry a natural label $k = 1, 2, 3$ (the rotation "frequency," via χ_k involving $\cos(2\pi k/7)$). This supplies a *candidate* ordering for the realization depths — but only a candidate (see §4.5).

[Identification graded [Conjectural]; consequences (i)–(iii) Established *given* the identification.]

4.4 Relation to the other axes — independence is supported, not threatened

The D_7 irrep decomposition bears on the Realization Architecture's factorization claim C4 (that A, P, D are independent axes), and on two coincidence-of-numbers hazards.

The 1D sector and assignment (the "2"). D_7 carries exactly two one-dimensional irreps (trivial and sign). If assignment structure A is sourced from the 1D sector and generation structure D from the 2D sector, then the assignment-2 and the generation-3 arise from *one* structure but in **orthogonal irrep sectors** (1-dimensional versus 2-dimensional). This would *support* C4 rather than threaten it: a single architecture generating independent axes is exactly what C4 asserts is possible. **[Conjectural — candidate sourcing of A and D from disjoint irrep sectors.]**

The colour-3 and the generation-3 (are the two 3s the same 3?). This is the §4 analogue of the two-factors-of-two problem one layer up: if the colour triplet and the generation count were the same "3," the independence of P and D would collapse. The D_7 structure answers favourably, but the verdict needs its premise stated. The two 3s are different *kinds* of object: the generation-3 is a **cardinality** (the count of 2D irrep types), while the colour-3 is a **multiplet dimension** (the size of the transport triplet). The narrow established fact is that **D_7 has no three-dimensional irreducible representation** — its irreps are of dimension 1 and 2 only — so the colour triplet cannot live *as a D_7 irrep*. That rules out one specific way the two 3s could coincide (colour sourced as a 3D D_7 irrep), not all conceivable ways. The favourable conclusion therefore carries an unstated premise: that colour is sourced **outside** the D_7 irrep structure (in the transport-channel multiplicity, which is not a D_7 representation). Granting that premise, the colour-3 and the generation-3 have structurally distinct origins, the two 3s are not the same 3, and P–D independence is supported on this axis. **[Supported given D_7, conditional on colour being sourced outside the D_7 irrep structure — not Established; "no 3D irrep" closes one coincidence route, not all.]**

4.5 Discipline: the matches are leads or traps, not results

Three structural matches have appeared: three 2D irrep types for three generations; a fourth that folds into the third; two 1D irreps for binary assignment. Each is exactly the number the physics wants. That is precisely the condition under which a numerical coincidence is most dangerous — the same discipline that, in the participation branch, forbade identifying the surviving channel with the architectural nullity merely because both were "one."

The irrep *counts* are theorems. The *identification* generation ↔ 2D irrep is a conjecture, and it is not corroborated by the count-matching, because the count-matching is what motivated it. The identification earns its grade only from an *independent* test — and the available one is the completion-density ordering. Under the identification, the three modes ρ_1, ρ_2, ρ_3 carry the label $k = 1, 2, 3$; the test is whether the completion densities order *with that label*,

$C(\rho_1) < C(\rho_2) < C(\rho_3)$,

derived from the operator action, independently of the observed masses. If the completion densities come out unordered, or ordered against the k-label, the identification fails even though the counts matched. The count match is the lead; the ordering is the test. The match is not banked. **[Conjectural — identification to be tested by ordering, not confirmed by cardinality.]**

---

5. Completion Density and the Decision Protocol

The Eigenmode Decision programme supplies the ordering object,

$C = p_v / K\_c$,

a completion density associated with a realization mode. The realization hierarchy requires

$C_1 < C_2 < C_3$,

which is **not assumed** but is the principal test of any constructed spectrum. Note that the association of a completion density with each eigenmode, and the coincidence of depth-ordering (by λ_n) with completion-density ordering (by C_n), are themselves conjectural until the operator action makes them computable. **[Conjectural — λ-ordering ≡ C-ordering is an unstated hinge of the hierarchy and is graded here explicitly.]**

The spectrum, once $\mathfrak{A}$ is constructed, is judged by the following pre-registered protocol. Pre-registration is the point: the criteria are fixed *before* the spectrum is computed, so the spectrum cannot fix the criteria.

- **Success.** Exactly three isolated eigenmodes exist and satisfy $C_1 < C_2 < C_3$.
- **Partial success.** Three isolated eigenmodes exist but completion-density ordering is unresolved.
- **Failure.** The spectrum lacks isolation, or the completion-density ordering is violated.
- **Strong failure.** A fourth stable isolated eigenmode exists.

Under the D_7 reading, the protocol acquires teeth before the full construction: §4.3(i)–(ii) would make Success on cardinality and Strong-failure-exclusion structural (theorems of D_7), leaving the completion-density ordering of §4.5 as the live discriminator between Success and Partial success — and as the genuine test of the generation ↔ irrep identification.

This Strong-failure criterion is the empirical face of the two horns of §4.2a, and the two are consistent rather than in tension. Under the multiplicity-free acts-on reading, a fourth *type* is excluded by fold-back ($\rho_4 \cong \rho_3$), so Strong failure cannot arise from a new irrep type. Under the dim-14 elements-are-carrier reading, the multiplicity-two partners — the "extra" modes that the scalar-on-multiplicity postulate would have to suppress — are precisely the candidate fourth isolated eigenmodes. So "a fourth stable isolated eigenmode exists" is exactly the observable signature that would distinguish the two horns: its *presence* would be evidence for the dim-14 reading with multiplicity unsuppressed; its *absence* is what the acts-on reading predicts for free and what the elements reading must pay a postulate to secure. The protocol's Strong-failure test thus prices the same horn §4.2a prices, in observational rather than representation-theoretic terms.

6. What This Paper Does and Does Not Claim

The paper **does** claim:

- a precise specification of $\mathfrak{A}$ via A1–A4, with the explicit caveat that the specification fixes neither cardinality nor termination (§2);
- that the construction of $\mathfrak{A}$ is downstream of the closure-architecture / census result, linking the realization and participation branches to one Step 0 (§3);
- that under the leading D_7 reading, realization modes must arise from representations of D_7, and that — *under the multiplicity-free acts-on carrier* (§4.2a) — the three 2D irreps would supply three-ness and fourth-mode fold-back as theorems given the identification (§4);
- that the dim-14 (regular-rep) reading forces multiplicity two and so does *not* deliver three clean modes without an extra scalar-on-multiplicity postulate, yielding a cost asymmetry (GATE-DISCRIM(a′)) that the realization branch hands back to the census branch (§4.2a);
- that, granting colour is sourced outside the D_7 irrep structure, the colour-3 and generation-3 are structurally distinct, supporting P–D independence (§4.4).

The paper **does not** claim:

- a construction of $\mathfrak{A}$ (no action on a basis is exhibited);
- a computation of the spectrum λ_1, λ_2, λ_3 or the densities C_1, C_2, C_3;
- that the census returns D_7 (this is the leading reading, not a settled result);
- that the generation ↔ 2D-irrep identification holds (it is the conjecture the ordering test must decide);
- a derivation of three generations (this is conditional on both the census result and the identification).

The key sentence: **the present paper does not construct $\mathfrak{A}$; it specifies the operator whose construction is required, identifies the closure-architecture result on which that construction depends, and records the first constraints imposed by the D_7 reading of the fourteen-generator census.**

Condition Ledger

Claim	Grade
Realization depth supplies generation axis (T2 frames the demand)	Inherited (Realization Architecture; T2 Open)
Specification A1–A4 of $\mathfrak{A}$	Specification (properties required, not exhibited)
A1–A4 fix the cardinality of the spectrum ("three")	**Open — they do not**
A1–A4 establish termination (no fourth)	**Open — A4 assumes possibility, does not establish**
$\mathfrak{A}$ constructed (action on a basis)	**Open — not done; the central remaining task**

Claim	Grade
Construction of $\mathfrak{A}$ is downstream of the census/D_7 result	Established (dependency); census result itself Open
Census returns N_loop =	D_7
Irrep list (two 1D, three 2D) has total dim 8;	D_7
D_7 has exactly three 2D irrep *types*	Established
Any dim-14 D_7 representation forces multiplicity ≥ 2 (irreps total dim 8)	Established
"Exactly three modes" requires a multiplicity-free (acts-on) carrier	Established — incompatible with dim $\mathfrak{g}_{14} = 14$ (GATE-DISCRIM(a′))
GATE-DISCRIM(a′): acts-on horn yields three free; elements horn pays a scalar-on-multiplicity postulate	Established as conditional cost asymmetry; GATE-DISCRIM(a) itself Open
Putative fourth 2D mode folds: $\rho_4 \cong \rho_3$	Established (rep theory); "no fourth generation" adds three physics conditionals
D_7 has no 3D irrep (one route to colour-3 = generation-3 closed)	Established (narrow); independence verdict is Supported, not Established
Generation ↔ 2D irrep of D_7	**Conjectural — tested by ordering, not by count-match**
Three-ness and fourth-exclusion as theorems	Conditional (on $D_7 \wedge$ the identification)
Assignment-2 and generation-3 from orthogonal irrep sectors (supports C4)	Conjectural
λ-ordering $\equiv$ C-ordering (hierarchy hinge)	Conjectural
Completion-density ordering $C_1 < C_2 < C_3$	Open — the principal test
Spectrum $\{\lambda_n\}$, densities $\{C_n\}$ computed	Open

Conclusion

The realization branch had reached the point where further progress required an explicit spectral object, and had repeatedly named "realization depth" without exhibiting the operator whose spectrum it is. This paper does not end that wait by constructing the operator. It ends a different and prior confusion: it states exactly what the operator must satisfy, shows that those conditions do **not** by themselves deliver the three-generation count or fourth-mode exclusion, and identifies why — because the operator's action cannot be read off "closure geometry" until the closure architecture is fixed by the census result. The realization spectrum is therefore downstream of the two-factors-of-two question, and the two branches are one.

Taking the leading D_7 reading as reconnaissance, the representation theory of the dihedral group is fixed and suggestive: three two-dimensional irrep types, a fourth that folds into the third, two one-dimensional irreps in an orthogonal sector, and no three-dimensional irrep at all. If generations are those 2D irreps — and if the census's 14 is read as the group *order* with D_7 acting on a multiplicity-free carrier, not as the dimension of the carrier itself — then "exactly three" and "no fourth" would follow from the group rather than from assumption, and the generation-3 would be structurally distinct from the colour-3. That conditional is not incidental: it is the section's sharpest output. The dim-14 regular-rep reading forces each 2D irrep to appear twice, admitting up to six modes unless $\mathfrak{A}$ is separately forced scalar on each multiplicity space — so the realization reconnaissance hands the census branch a constraint it did not have, that three clean modes are cheaper under the "acts-on" resolution of GATE-DISCRIM(a) than under the "elements-are-carrier" one. None of this is banked: the identification is a conjecture whose only honest test is the completion-density ordering, and the count-match — being what motivated the identification — cannot corroborate it.

The chain is now locatable:

census fixes architecture (D_7?) → $\mathfrak{A}$'s action becomes writable → spectrum and densities computed → decision protocol applied.

The next mathematical task is the one this paper deliberately does not attempt: settle the census (Step 0), then write the action of $\mathfrak{A}$ on the D_7-structured realization states, compute its spectrum and completion densities independently of the observed masses, and apply the pre-registered protocol. If the spectrum yields three isolated, correctly ordered modes, realization depth becomes a derived feature of closure architecture. If it does not, §5 says exactly how the proposal fails. Until then, the realization operator has a specification and an address — and the address is the same room the census paper is already standing in.