The Origin of the Fourteen-Generator Census

Closure Architecture Beneath SU(3) × SU(2) × U(1), Dihedral Transport, and a Decidable Test of the Doubling

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The VERSF Standard Model programme derives a substantial portion of the gauge architecture of particle physics from information-theoretic principles. The Minimal Internal Symmetry programme, the Distinguishability Conservation programme, and the associated closure studies obtain the appearance of the Standard Model gauge structure

SU(3) × SU(2) × U(1)

from admissibility constraints, distinguishability conservation, Fisher geometry, anomaly exclusion, and entropy minimization, rather than treating that structure as an empirical input.

A separate body of work addresses generations, realization depth, hierarchy, and saturation. A recurring lesson of that work is that several explanatory tasks once treated together are in fact distinct. Generation explains why nature exhibits a finite set of stable species classes. Refinement explains hierarchy and localization. Saturation explains why the admissible sequence terminates. Closure architecture explains the structural machinery through which admissibility is implemented.

This paper concerns the fourth task only.

Many closure analyses inherit a *fourteen-generator response structure* attached to the $K = 7$ closure architecture. The count appears throughout the programme, yet its origin has not been examined on its own terms. The arithmetic is obvious — $14 = 2 \times 7$ — but the *meaning* of the factor of two is not. The purpose of this paper is therefore not to explain fermion generations, mass hierarchy, or fourth-generation exclusion. It asks a more basic question:

Why does a seven-element closure architecture produce fourteen generators rather than seven?

The contribution is threefold. First, the paper isolates the fourteen-generator census as a problem of *closure architecture*, distinct from generation, refinement, and saturation. Second, it converts the competing interpretations of the doubling into a decidable test on the inherited structure.

Third, it runs that test against the transport-group result, which identifies the response structure as a dihedral group of order fourteen — and the verdict that returns is not the binary the test was first posed to settle. The transport result separates two questions that were being run as one: the *origin* of the doubling, which it answers geometrically, and the *symmetry* of the doubling, which it answers asymmetrically. The honest reading vindicates the geometric interpretation on origin and the dynamical interpretation on structure — that is, its prediction that the two halves are unlike, not its claim about where the doubling comes from — and shows the original either/or framing was too coarse to let both be true at once.

Contents

Abstract

- Introduction
- Four Explanatory Layers
- 2.1 Closure architecture
- 2.2 Generation
- 2.3 Refinement
- 2.4 Saturation
- Closure Architecture Beneath the Gauge Programme
- The Inherited K = 7 Architecture
- The Origin of the Doubling: Two Routes
- A Decidable Discriminator and the Transport-Group Evidence
- 6.1 The Transport-Group Evidence: a Split Verdict
- The Two Factors of Two, and the Status of the Double-Counting Protection
- Consequences and Scope
- Condition Ledger and Epistemic Status

Conclusion

Abstract

The VERSF Standard Model programme contains two partially independent explanatory tracks. The first derives admissible structural architecture and yields the gauge-sector structure SU(3) × SU(2) × U(1) from distinguishability conservation, Fisher geometry, anomaly exclusion, and entropy minimization. The second addresses the existence of generations, hierarchy, and saturation. These tracks should not be conflated: generation concerns isolated realization classes, refinement concerns hierarchy and localization, saturation concerns termination, and closure architecture concerns admissibility machinery.

This paper isolates the fourteen-generator census as a problem of closure architecture. It examines the inherited relation N_loop = 2K associated with the K = 7 architecture and investigates the origin of the doubling. Two interpretations are placed in competition. **Route A** interprets the doubling geometrically, as oriented traversals of seven closure constraints. **Route B** interprets it dynamically, as a separation between closure *formation* and closure *preservation*.

A decidable discriminator is established on the symmetry of the inherited fourteen-generator structure, and the transport-group programme supplies the first inspection. It identifies the response structure as the dihedral group

$G \cong D_7$, $|D_7| = 14 = 7$ rotations + 7 reflections

(D_n denotes the symmetry group of the regular n-gon, of order 2n, so $|D_7| = 14$). This result does not settle the original binary; it dissolves it, by separating two questions that Routes A and B were each answering jointly:

- **Origin.** Is the factor of two a geometric symmetry-group structure or a dynamical formation/preservation process? The transport identification answers *geometric*: the factor of two is the quotient $\mathbb{Z}_2 \cong D_7 / C_7$. This is the thesis of Route A. [Supported, capped by the open identification $|D_7| = \dim \mathfrak{g}_{14}$]
- **Symmetry.** Is the realized doubling a clean exchange involution of matched septets, or an asymmetric relation between inequivalent ones? The transport identification answers *asymmetric*: the seven rotations form the cyclic subgroup C_7 and the seven reflections its coset, inequivalent in element type, with no matched-commutator exchange automorphism. This asymmetry is the structural inequivalence Route B predicted. [Supported, capped by the same identification; free of the commutator gate (b)]

The honest summary is therefore not "Route A wins." It is that the origin question goes to A and the symmetry question goes to B, and the either/or framing of §5–6 was too coarse to admit both. The paper grades its claims accordingly, flags the load-bearing identification $|D_7| = \dim \mathfrak{g}_{14}$ as Open and capping everything downstream, and shows the participation double-counting protection to be conditional on a single open check (GATE-FRAME) carried by the now-leading geometric reading.

1. Introduction

The derivation of the Standard Model gauge structure within VERSF proceeds through a sequence of architectural results. The Minimal Internal Symmetry programme establishes the emergence of admissible internal sectors. The Distinguishability Conservation programme establishes the necessity of gauge structure. The Closing the Interfaces programme reduces the residual empirical content required to connect these structures to observation. Together these establish the appearance of

SU(3) × SU(2) × U(1)

as the admissible gauge architecture of the low-energy world. [Inherited]

These results operate at the level of *structural admissibility*. They explain why particular classes of internal structure exist. They do not explain fermion generations, they do not explain mass hierarchy, and they do not explain the termination of the realization sequence. These questions belong to distinct explanatory layers, and the paper begins by separating those layers explicitly so that the fourteen-generator census can be placed correctly within them.

2. Four Explanatory Layers

The programme separates four explanatory tasks. The separation is organizational, but it is load-bearing: a count assigned to the wrong layer imports the wrong obligations.

2.1 Closure architecture. Concerns admissibility machinery — the structural framework from which admissible sectors emerge. The fourteen-generator census belongs to this layer.

2.2 Generation. Concerns isolated realization classes — why nature exhibits a finite collection of stable species classes. The Generation Theorem, the Eigenmode Decision, and the Individuation–Isolation programme belong primarily here.

2.3 Refinement. Concerns localization, hierarchy, amplification, and scale — why distinct realization classes carry distinct characteristic scales. Mass hierarchy belongs primarily here.

2.4 Saturation. Concerns terminality — why the admissible realization sequence terminates and why additional stable classes fail to appear.

The boundary between (2.1) and the remaining three layers is the operative one for this paper. The fourteen-generator census is frequently invoked in contexts that belong to generation or refinement, and the conflation has obscured its origin. [Proven, organizational]

3. Closure Architecture Beneath the Gauge Programme

The closure-architecture layer sits *beneath* the gauge-architecture programme. The gauge derivations establish admissible internal sectors, distinguishability-preserving transport, gauge connections, uniqueness of the Abelian sector, and the emergence of SU(3), SU(2), and U(1). [Inherited]

Each of these constructions depends on admissibility structure. Any inherited closure count therefore belongs conceptually to the same architectural layer as the admissibility machinery, not to the generation or hierarchy programmes that sit above it. The fourteen-generator census is accordingly interpreted as part of the machinery underlying admissible structure. [Conditional — depends on the layer assignment of §2 being exhaustive for the constructions that invoke the count.]

4. The Inherited K = 7 Architecture

The architecture is reviewed only to fix notation; its derivation is inherited and not re-argued.

- six boundary constraints,
- one hub constraint,
- nullity-one mode,
- fourteen-generator response structure,
- $N_{loop} = 2K$ with $K = 7$, giving $N_{loop} = 14$.

The *existence* of the fourteen-generator census is [Inherited, Proven within the programme]. Its *interpretation* is the subject of the remainder of the paper.

5. The Origin of the Doubling: Two Routes

The central question is: why does K = 7 produce fourteen? Two interpretations are placed in competition. As §6 will show, each route bundles together two logically separate questions — the *origin* of the factor of two and the *symmetry* of its realization — and the transport result answers those two questions differently. The two routes are stated here in their original, bundled form, because that is the form the discriminator was built to test.

Route A — Oriented traversals (geometric origin, symmetric realization). The fourteen arise as

$$14 = 7 \times 2$$

through directional realization of seven constraints: each closure constraint contributes an inbound and an outbound traversal. The doubling is *geometric in origin* and a *clean exchange* in realization. [Origin: Leading — see §6.1; Realization: not supported — see §6.1]

Route B — Formation and preservation (dynamical origin, asymmetric realization). The fourteen arise as

$$14 = 7_{form} + 7_{preserve}$$

through the distinction between closure formation and closure preservation: seven generators establish closure and seven preserve it against reversal. The doubling is *dynamical in origin* and *asymmetric* in realization. [Origin: not supported — see §6.1; Realization: Supported — see §6.1]

A prior-probability remark is owed to Route A's origin claim, and is stated rather than buried. The inherited relation is phrased as a *loop* count, $N_{loop} = 2K$, and loops are natively oriented objects: a loop carries a traversal sense. **The default reading of the origin of $N_{loop} = 2K$ is therefore geometric.** Orientation is the parsimonious account of a doubled loop census, and it requires no structure beyond what the loop already carries. A dynamical origin must earn its existence by exhibiting structure that a geometric symmetry cannot supply. The next section states the test, and then reports the transport-group evidence that runs it — with a verdict that splits these two routes along their two bundled questions rather than electing one whole.

---

6. A Decidable Discriminator and the Transport-Group Evidence

The two routes are often discussed as a contest between "geometry" and "dynamics." That framing is not decidable, and — as the transport result will show — it is also too coarse, because it bundles the *origin* of the doubling together with the *symmetry* of its realization. The discriminator is built on the symmetry of the inherited fourteen-generator structure, and the transport result then forces the two questions apart.

Write $\mathfrak{g}_{14}$ for the inherited fourteen-generator structure of the $K = 7$ architecture, and let C denote a distinguished generator associated with closure completion (the hub constraint together with the nullity-one mode). Suppose $\mathfrak{g}_{14}$ splits into two seven-dimensional subspaces,

$$\mathfrak{g}_{14} = S_a \oplus S_b , \quad \dim S_a = \dim S_b = 7 .$$

A clean exchange (the symmetric realization). If inbound and outbound traversals are mirror images, there exists an order-two automorphism

$\sigma : \mathfrak{g}_{14} \to \mathfrak{g}_{14}$, $\sigma^2 = \mathrm{id}$, $\sigma(S\_a) = S\_b$,

under which the two septets carry *matched* commutation relations with C. The doubling is a reflection: S_a and S_b are equivalent up to orientation reversal.

An asymmetric realization. If the two septets are inequivalent in type and no matched-commutator exchange automorphism exists, the relation between them is not a reflection. Preservation presupposes commitment — a closure is not preserved unless it is committed — so one realization of asymmetry is a directed commitment map $\delta : \Gamma\_form \to \Gamma\_preserve$ with no symmetric inverse (a filtration $F^0 \supseteq F^1$, or a nilpotent relation expressing that committed closure does not un-form). But a directed map is not the only way the septets can be inequivalent, as the transport result will show.

These possibilities yield a decision procedure. A naïve trichotomy conflates origin with realization and omits the case the transport result realizes; the correct procedure separates the two axes — origin ∈ {geometric, dynamical}, realization ∈ {symmetric exchange, asymmetric} — and reasons over the cross-terms.

Theorem 6.1 (Doubling Discriminator). Let $\mathfrak{g}_{14}$ denote the inherited fourteen-generator structure of the K = 7 architecture, let C denote the distinguished closure-completion generator (the hub constraint together with the nullity-one mode), and suppose $\mathfrak{g}_{14} = S\_a \oplus S\_b$ with dim S_a = dim S_b = 7.

(A — geometric origin, symmetric realization.) If $\mathfrak{g}_{14}$ admits an order-two automorphism σ with $\sigma^2 = \mathrm{id}$, $\sigma(S\_a) = S\_b$, and matched commutators $[S\_a, C] \cong [S\_b, C]$, the doubling is geometric in origin and a clean exchange in realization: orientation reversal of matched septets. This is pure Route A.

(A′ — geometric origin, asymmetric realization.) If the factor of two arises from a symmetry-group structure (geometric origin) but no matched-commutator exchange automorphism exists — the two septets being inequivalent in type, related by a coset/subgroup or other non-matched involutive structure — then the origin is geometric while the realization is asymmetric. This is a positive, substantive outcome, not an obstruction. It awards the origin question to Route A and the symmetry question to Route B.

(B — dynamical origin, directed realization.) If no symmetry-group origin exists, and the surviving relation between S_a and S_b is a directed graded or nilpotent map $\delta : \Gamma\_form \to \Gamma\_preserve$ with no symmetric inverse, the doubling is dynamical: pure Route B.

(O — obstructed.) If the structure is none of the above — neither symmetric exchange, nor geometric-asymmetric, nor directed — the question is genuinely [Open] and the obstruction is identified.

The two axes give four cross-terms, and the enumeration is exhaustive because the fourth is provably empty. (A) is geometric-symmetric; (A′) is geometric-asymmetric; (B) is dynamical-asymmetric. The remaining cross-term — dynamical-symmetric — is uninstantiable by the meaning of "dynamical" here: a formation/preservation origin makes the relation directed, because preservation presupposes commitment, and a directed relation is asymmetric. A dynamical origin therefore forces an asymmetric realization, so a matched-exchange dynamical doubling cannot arise. (O) catches the genuine remainder: structure with neither a cleanly geometric nor a cleanly dynamical origin. (A′) is accordingly not a wastebasket but a positive structural finding — a geometric origin with an asymmetric realization — and, as §6.1 shows, the one the transport result returns. [Conditional — the procedure is well-posed given $\mathfrak{g}_{14}$; the outcome is the subject of §6.1.]

The value of the test is that it removes the contest from the realm of preference. "The doubling is dynamical, not geometric" is, on its own, narrative. "The inherited structure has a geometric origin and an asymmetric realization" is a fact about a fixed inherited object, and it is decidable from that object.

The test also disciplines two algebraic assumptions the formation–preservation proposal carries and that should not be assumed for free:

1. **The direct sum.** $\Gamma\_total = \Gamma\_form \oplus \Gamma\_preserve$ presupposes trivial intersection and independent span. If preservation is "resisting the reversal of formation," the natural structure may instead be graded or paired. The $\oplus$ is therefore [Conjectural], not notational.
2. **The equal split.** $\dim \Gamma\_form = \dim \Gamma\_preserve = 7$ is presently *required* (it must reproduce 2×7), not *derived*. The equal split is [Conjectural], motivated only by reproducing the inherited total.

6.1 The Transport-Group Evidence: a Split Verdict

The transport-group programme supplies the missing structure. It identifies the response structure as a dihedral group:

$$G \cong D_7 \,,\ |D_7| = 14 \,,\ \text{with } 14 = 7 \text{ rotations} + 7 \text{ reflections} \,.$$

Read against Theorem 6.1, this is outcome **(A′)**, not (A) — and it returns a split verdict.

The origin question goes to Route A. D_7 is the symmetry group of the regular seven-gon. Identifying the fourteen with $|D_7|$ makes the doubling geometric in origin — Route A's thesis — with the factor of two an honest group-theoretic quotient

$$\mathbb{Z}_2 \cong D_7 / C_7 \,,$$

where C_7 is the cyclic subgroup of seven rotations and the seven reflections form its single nontrivial coset. A dynamical formation/preservation process is not needed to account for the count; a symmetry group does it. [Supported — but see the cap below.]

The symmetry question goes to Route B. The clean exchange of outcome (A) is *not* realized, and this can be seen without the explicit algebra. The coset exchange is carried by left translation, a group action, not an automorphism; the natural automorphism — conjugation by a reflection — fixes each septet setwise, inverting the rotations and permuting the reflections. So no automorphism exchanges S_a and S_b, and clause (i) of (A) fails. And the two septets are inequivalent in *type* — seven elements of order seven forming a subgroup, versus seven involutions forming a coset — so the matched-commutator clause (ii) fails as well, visibly, from element order alone. That type-inequivalence is exactly the structural distinction Route B predicted between the two septets. This inequivalence is cap-free only for D_7's own rotation-set and reflection-set; transporting it onto the septets of $\mathfrak{g}_{14}$ requires those septets to be the images of D_7's rotation and reflection sets — which is again the identification noted below. [Supported, capped by GATE-DISCRIM(a); free of GATE-DISCRIM(b).]

The load-bearing caveat. Both halves of the verdict — origin and symmetry alike — rest on identifying the fourteen of $|D_7|$ with the dimension of the generator census $\mathfrak{g}_{14}$: that the fourteen group elements *are* the fourteen generators, rather than D_7 merely *acting on* a possibly different-dimensional $\mathfrak{g}_{14}$. The distinction is concrete: the regular representation of D_7 is itself fourteen-dimensional and carries no canonical 7 + 7 partition of its own, so the subgroup/coset split that delivers both the quotient $\mathbb{Z}_2 \cong D_7 / C_7$ and the septet asymmetry reaches $\mathfrak{g}_{14}$ only through this identification. That identification is **[Open · GATE-DISCRIM(a)]**, and it is the link on which everything in this section hangs. Until it is at least conditionally argued, "the factor of two is $\mathbb{Z}_2 \cong D_7 / C_7$" is a numerical match promoted to a structural identity. By the programme's inheritance-cap discipline, both the origin and the symmetry claims inherit this open ceiling. The legitimate distinction the verdict can draw is against GATE-DISCRIM(b), the commutators: the symmetry half is free of (b) — the order-7-versus-order-2 inequivalence needs no commutator computation — whereas full algebraic confirmation of (A′) does wait on (b). Independence of (b) is not independence of (a).

The honest summary. The transport result does not elect Route A whole and demote Route B to a gloss. It separates the bundle: the *origin* of the doubling is geometric (Route A), and the *symmetry* of the doubling is asymmetric (the inequivalence Route B predicted). The binary of §5 was too coarse to let both be true at once. What is genuinely retired is not Route B but the *dynamical-origin* claim it carried — formation/preservation is unnecessary as the *source* of the count. What survives, and is now positively supported, is Route B's *structural* prediction that the two septets are inequivalent. A formation/preservation reading may still be offered as one interpretation of that geometric asymmetry (rotations as the formed sector, reflections as the preserving coset), but it now lives *inside* the dihedral structure as an optional gloss, carrying no independent content unless the explicit commutators relative to C show otherwise.

---

7. The Two Factors of Two, and the Status of the Double-Counting Protection

The participation analyses identify two distinct factors of two, and the relationship between them governs whether later constructions double-count.

Factor A: the loop doubling, $N\_loop = 2K$, identified by §6.1 as the quotient $\mathbb{Z}_2 \cong D_7 / C_7$.
Factor B: the boundary-frame dimension, $\dim \Gamma\partial = 2$.

The double-counting risk is specific: a participation construction that uses *both* a loop-derived factor of two and the boundary-frame factor of two counts the same structure twice **iff Factor A and Factor B are the same factor of two**. The protection of the participation programme therefore reduces to whether the two factors are distinct.

The operative reading is the geometric one, and on that reading the distinctness of Factor A from Factor B is not automatic; it carries an open check. The loop factor $\mathbb{Z}_2 \cong D_7 / C_7$ is distinct from $\dim \Gamma\partial = 2$ *provided the dihedral reflection is not itself sourced from the boundary frame* — provided the reflection coset is not the two directions of $\Gamma\partial$. This is **[Open · GATE-FRAME]**. The branch that would have made the distinctness automatic — a dynamical origin, with formation/preservation carrying no a priori tie to $\Gamma\partial$ — is precisely the origin claim §6.1 found unnecessary, and is not the operative reading.

The honest status of the protection is therefore: **Conditional, with one open check (GATE-FRAME) carried by the leading geometric reading.** The transport identification sharpens that check from a vague worry into a precise question — is the reflection generator of D_7 independent of the boundary frame, or induced by it? — but it does not close it. What can be said firmly is that the threat is not *immediate*: it reduces to a single, stated, decidable question rather than an open-ended concern.

---

8. Consequences and Scope

The paper does **not** derive generation count, mass hierarchy, fourth-generation exclusion, or participation denominators. None of these is touched. The denominator-twelve question, in particular, survives untouched and on its own merits.

What the paper does is organizational and structural:

- It assigns the fourteen-generator census to the closure-architecture layer — the same layer that underlies the emergence of SU(3), SU(2), and U(1) — and away from the generation, refinement, and saturation layers. [Proven, organizational]

- It supplies a decidable test for the doubling, separating the origin question from the symmetry question, and corrects an earlier exhaustiveness error by admitting the geometric-asymmetric outcome (A′). [Conditional]
- It runs the test against the transport identification $G \cong D_7$ and returns a split verdict: geometric origin (Route A), asymmetric realization (Route B). Both halves inherit the open ceiling of the identification $|D_7| = \dim \mathfrak{g}_{14}$. [Supported, capped]
- It places the participation double-counting protection at Conditional, with one open check (GATE-FRAME) on the leading geometric reading. [Conditional]

The fourteen-generator census should accordingly be read as admissibility machinery with a dihedral geometric origin and an asymmetric internal structure — not as a statement about generation count, hierarchy depth, or saturation.

---

9. Condition Ledger and Epistemic Status

Inherited (not re-argued here).

- The K = 7 architecture: six boundary constraints, one hub constraint, nullity-one mode. [Proven within the programme]
- The fourteen-generator census and N_loop = 2K. [Proven within the programme]
- The gauge derivations of SU(3) × SU(2) × U(1). [Inherited]
- The transport-group identification $G \cong D_7$ with $|D_7| = 14 = 7$ rotations + 7 reflections. [Inherited]

Established here.

- Closure architecture is a distinct explanatory layer; the census belongs to it. [Proven, organizational]
- The discriminator separates origin from realization, with four substantive outcomes (A), (A′), (B), (O). [Conditional — test well-posed]
- The transport result realizes outcome (A′): geometric origin, asymmetric realization. [Supported, capped by GATE-DISCRIM(a)]
- Origin verdict: the doubling's source is geometric, the factor of two being $\mathbb{Z}_2 \cong D_7 / C_7$. [Supported, capped by GATE-DISCRIM(a)]
- Symmetry verdict: the realized doubling is asymmetric — the two septets inequivalent in type (order-7 subgroup vs order-2 coset), with no matched-commutator exchange automorphism. This vindicates Route B's structural prediction. [Supported, capped by GATE-DISCRIM(a); free of GATE-DISCRIM(b)]
- The original §5 binary (geometric-symmetric vs dynamical-asymmetric) is too coarse; both true answers cannot be expressed within it. [Established by the split verdict]

Conjectural / unnecessary.

- Route B's *dynamical-origin* claim — formation/preservation as the source of the count — is unnecessary given the geometric origin. [Not supported as origin]
- Formation/preservation as a *gloss on the geometric asymmetry* (rotations formed, reflections preserving). [Conjectural — carries no independent content pending the explicit commutators]
- The direct-sum structure $\Gamma\_total = \Gamma\_form \oplus \Gamma\_preserve$ (vs graded/paired). [Conjectural]
- The equal split dim Γ_form = dim $\Gamma\_preserve = 7$. [Conjectural — motivated only by reproducing 2×7]

Open / gate-tagged dependencies.

- **[Open · GATE-DISCRIM(a)]** Identification of the group order $|D_7| = 14$ with the dimension of the explicit fourteen-generator algebra $\mathfrak{g}_{14}$ — that the fourteen group elements are the fourteen generators, not that D_7 merely acts on a possibly different-dimensional $\mathfrak{g}_{14}$. This is the load-bearing link; every claim in §6.1 inherits its ceiling. Until conditionally argued, the dihedral identification is a numerical match, not a structural identity.
- **[Open · GATE-DISCRIM(b)]** The commutator structure of $\mathfrak{g}_{14}$ relative to C, which would confirm outcome (A′) at the algebraic level and either close or reopen a formation/preservation gloss. Note that the *symmetry verdict* (asymmetry) does not wait on this: the order-7-vs-order-2 inequivalence is visible from element type alone.
- **[Open · GATE-FRAME]** Is the dihedral reflection (the $\mathbb{Z}_2 \cong D_7 / C_7$) intrinsic to the transport group, or induced by the boundary frame (reflection coset sourced from the two directions of $\Gamma\partial$)? This is the single residual exception to the double-counting protection of §7, carried by the leading geometric reading.

Scorecard.

Claim	Status
Closure architecture is distinct from generation / refinement / saturation	Proven (organizational)
The doubling has a geometric (symmetry-group) origin, not a dynamical one	Supported, capped by GATE-DISCRIM(a)
The realized doubling is a clean exchange involution (matched septets)	Disfavored, capped by GATE-DISCRIM(a) — (A) excluded only under the identification
The realized doubling is asymmetric (inequivalent septets) — Route B's prediction	Supported, capped by GATE-DISCRIM(a) (order-7 vs order-2; free of (b))
The §5 binary (geometric-symmetric vs dynamical-asymmetric) is adequate	Refuted — outcome (A′) realized
Identification	D_7
Participation double-counting protection	Conditional — one open check (GATE-FRAME) on leading branch

Programme-scorecard contribution (M / C format). The row this paper contributes to the programme status scorecard, in its two-column convention — M = mechanism maturity, C = completeness of the chain to a measured fact, under the inheritance cap (C cannot exceed its weakest necessary link). The result is closure-architecture infrastructure: it feeds the participation constructions and does not itself reach an observable, so it sits with the architecture entries rather than the gauge-sector chains.

Result	M	C	Depends on
Origin of the closure doubling — the 14-generator census (N_loop = 2K)	★★☆☆☆	★☆☆☆☆	K = 7 census (inherited); D_7 transport (inherited); GATE-DISCRIM(a) identification (open, load-bearing); GATE-DISCRIM(b) commutators (open); GATE-FRAME (open)

M sits at two stars, not three, under the inheritance cap: the mechanism's load-bearing link, the identification $|D_7| = \dim \mathfrak{g}_{14}$, is open, and stripped of that unproven identification the content reduces to "a finite group of the right order exists." M rises to three once the identification is conditionally argued, and toward four once the commutators confirm (A′). C inherits the low ceiling of the open participation / O1 chain the result ultimately feeds.

One amendment propagates upward: the participation / denominator-twelve row's double-counting protection is conditional on GATE-FRAME — the loop factor $\mathbb{Z}_2 \cong D_7 / C_7$ must be shown distinct from the boundary-frame factor $\dim \Gamma\partial = 2$. This is a conditional clearance, not a completion, so the participation row's M and C are unchanged.

Conclusion

Closure architecture, generation, refinement, and saturation perform distinct explanatory roles, and the fourteen-generator census belongs to the closure-architecture layer — the same layer that underlies the derivation of SU(3) × SU(2) × U(1). The census is a statement about admissibility machinery, not about generation count, hierarchy, or saturation.

The origin of its doubling is no longer evenly balanced between two readings — but neither does one reading win whole. The paper supplies a decidable discriminator that, once corrected to separate the *origin* of the doubling from the *symmetry* of its realization, admits a geometric-asymmetric outcome (A′) that the earlier trichotomy had wrongly filed as indeterminacy. The transport-group identification $G \cong D_7$ realizes exactly that outcome. Its origin verdict goes to the geometric reading: the factor of two is the quotient $\mathbb{Z}_2 \cong D_7 / C_7$, and no dynamical process is needed to produce the count. Its symmetry verdict goes to the dynamical reading's structural prediction: the two septets are inequivalent in type — an order-seven subgroup and an order-two coset — with no matched-commutator exchange automorphism. The either/or framing of §5 was too coarse to hold both answers at once; the corrected reading holds them together.

The verdict is supported, not closed, and the cap is explicit. Both halves rest on identifying the order fourteen of D_7 with the dimension of the generator census $\mathfrak{g}_{14}$, which remains open and load-bearing; until it is conditionally argued, the dihedral picture is a numerical match raised to a structural identity, and every downstream claim inherits that ceiling. The participation double-counting protection is correspondingly conditional, reduced to a single open check — whether the dihedral reflection is sourced from the boundary frame — rather than cleared outright.

Whether the doubling is geometric or dynamical was, from the outset, an algebraic question with a definite answer in the inherited structure. The transport result returns the larger part of that answer, and returns it in two parts: the doubling's origin is geometric, and its realized structure is asymmetric.