The Price of Copies

Non-Abelian Structure from Indistinguishable Closure Sectors of the Unique Fold — The First Departure from U(1)

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The central claim, first

The central claim of this paper is simple. If reality is built from a single fundamental structure — the Fold — then all observed multiplicity must arise from different configurations of that one structure. If two such configurations are genuinely identical, then physics cannot depend on which copy is which. The mathematics required to respect that indistinguishability is, wherever the copies can be genuinely exchanged, already non-commutative. The strong and weak forces may therefore exist not because reality contains multiple fundamental objects, but because the one Fold contains multiple indistinguishable closure sectors — the forces as the price reality pays for holding identical copies of the same pattern. Everything else in this paper is the machinery that makes those five sentences precise, tests where they are forced and where they are conditional, and reduces what remains open to a single question with a named slot.

The question

The programme's recent foundational work proves something strict: there is exactly one Fold. Not several kinds of fundamental structure, not a menagerie — one, whose architecture saturates everything observable. At the same time, the world plainly contains many kinds of particle, and the forces that govern them — the strong and weak forces — are built on mathematics richer than the single circle the one Fold's phase provides: their operations *do not commute*, doing A then B differs from doing B then A. So the programme faces a question its own uniqueness theorems make sharp: if there is only one Fold, where does multiplicity come from — and where, in a one-Fold world, could non-commuting forces possibly originate?

This paper's answer begins with the corpus's own ontology. The one Fold admits *closure sectors*: persistent, stable configurations — think of knots tied in a single rope. One rope; many possible knots; and crucially, some knots may come in identical copies — sectors whose intrinsic structure coincides exactly, so that no possible observation can tell copy one from copy two. The paper asks what physics is forced to look like when that happens.

What this paper argues

The programme's oldest principle is that a difference nothing can detect is not a physical difference. If two closure sectors are exact copies, then "which copy is which" is not a fact about the world — and physics is forbidden from depending on it. The paper proves that this prohibition shapes physics in two layers. The first is forced outright: no law, no structure, no bookkeeping may treat copy one differently from copy two — the transport structure must look exactly the same under any relabeling. The second is a genuine physical question the paper names rather than assumes: whether the swap of two identical sectors is something the substrate can actually *do* — whether identical knots can be braided past each other. Where it can, the swap is itself a physical operation, and swapping does not commute with the sectors' individual phases: non-commutativity, the strange algebraic heart of the strong and weak forces, arrives the moment the one Fold contains identical copies it can exchange. The forces richer than electromagnetism would exist not because there are many Folds — there is one — but because the one Fold can tie the same knot more than once, and no fact distinguishes the knots.

The paper then reduces the whole question to its atom. Whatever happens with many identical sectors is generated by what happens with *two*, and for two identical copies the mathematics is completely classified: they can swap and phase (the proven minimum); or blend along one fixed axis; or blend freely — rotate into each other by any amount. Only free blending yields the full structure the Standard Model's forces use, and the gap between the forced minimum and everything is exactly *one element* wide: a single admissible operation that genuinely blends two identical sectors — turns one partway into the other — forces all the rest. The origin of non-abelian forces reduces to one question: *can the substrate blend identical sectors, or only swap them?*

The honest condition

The paper locates that question; it does not answer it. If the substrate can blend, the Standard Model's forces have a home, and the paper derives their characteristic shape in advance: one force-block per class of identical sectors, blocks multiplying together, abelian circles alongside — the product architecture the Standard Model in fact displays, with the classes the corpus already calls species. If the substrate can only swap, the framework cannot reach the strong and weak forces, and the paper says so plainly: that branch is close to fatal for the programme's Standard-Model ambitions, and it is named as the kill-condition it is. Several inputs are stated as imports from the Fold and reconstruction papers, each with a citation slot; one inherited debt — that sector transport preserves the size of amplitudes, which the single-phase case got for free — is surfaced rather than hidden; and one consequence of the one-Fold ontology is faced squarely: with one Fold there is one global phase, so whether identical sectors can carry *independent* phases is itself a slot, with the theorems restated to survive either reading. Nothing in this paper multiplies Folds. The uniqueness theorems are consumed, not contested; the question begins where they end.

Abstract

The corpus's foundational arc establishes Fold Uniqueness and Fold Saturation [Inherited]: exactly one Fold, whose architecture exhausts observable structure. The phase arc establishes that this Fold's transport phase is exactly U(1) [Inherited], and every Standard-Model-facing result to date lives in that abelian sector, with the non-abelian questions reserved to an arc named throughout the corpus and entered by none. The present paper enters it from the only starting point the uniqueness theorems permit: not multiple Folds — the corpus's most defended result forbids them — but multiple admissible **closure sectors** within the unique Fold architecture: the persistent configurations of the Persistent-Fold-Defect ontology, grouped at source into Representation Classes, with species-as-classes [Inherited — the sector ontology, slotted]. The multiplicity throughout is of stable PFD sectors within the one architecture, never of Folds; the paper's classes *are* the source's Representation Classes, with member-indistinguishability *derived* from the identity programme (equal invariant tuples entail indistinguishability, on pain of contradicting the catalogue or violating No-Surplus-Identity — §5's Indistinguishability Lemma) and the existence of $k \geq 2$ classes inherited from the species ontology itself (a singleton-only ontology would collapse species into object — §5's Class Multiplicity Proposition). The entrance question, stated at the corpus's grain: *given one Fold, what transport structure is forced when the Fold admits two or more mutually indistinguishable closure sectors — and where, exactly, would non-abelian structure have to come from?*

Forced. (i) The **compact Lie skeleton** (§4): admissible transport on the n-sector carrier, reversible and norm-preserving, generates a group forced closed by the Boundary Maintenance discipline — a closed subgroup of U(n), hence a compact Lie group [Proven, conditional on the unitarity import U1 and Cartan's theorem [Imported-External]]. The norm-preservation the one-dimensional case obtained for free (the charge paper's flagged gap) is here a genuine import, U1, slotted at the quantum-reconstruction papers. (ii) The **phase substructure** (§4): every sector's amplitude carries the derived global U(1); whether identical sectors carry *independent* relative phases — the diagonal torus T^n — is the sector-addressability import S2, faced as the one-Fold ontology requires: one Fold, one global phase, and sector-resolved phase advance is a slot, not an assumption, with the downstream theorems restated to survive its failure. (iii) The **Indistinguishability Theorem** (§5), split at its true grain: the unconditional half — admissible structure cannot depend on the unwitnessed labeling, so the transport group is relabeling-*covariant*, normalized by the class permutations [Proven, conditional on S3 and no-pre-individuation] — and the membership half: if the exchange of identical sectors is itself an admissible transport (the import **S4**, Exchange Realizability — can the substrate braid identical defects past each other?), then $S_k \subseteq G$, and G is **non-abelian** — for every $k \geq 2$ under the sector-phase import S2, and for $k \geq 3$ under S2's fallback, with the $k = 2$ fallback floor abelian and the weak doublet thereby flagged as exactly where the slots matter most [Proven, conditional as stated]. The covariance/membership split is the symmetry-versus-gauge distinction, derived rather than blurred; non-commutativity remains what a unique Fold containing copies is forced into — at the stated grain, with the gauging imported where the defect ontology can supply it. (iv) The **two-sector reduction** (§6): pairwise mixing subgroups with the torus generate the class group [Imported-External], so the question reduces to the two-sector atom, where the closed subgroups of SU(2) are completely classified [Imported-External]: the forced floor (consuming S4) is the normalizer N(U(1)) under S2 — relative phase plus swap — and any single admissible

element beyond it forces closure to all of SU(2); under S2's fallback the finite exceptional subgroups remain live and the collapse is conditional on the finite-impossibility slot Q2, promoted thereby to load-bearing. **The gap is one element wide, by floor.**

The decision node. Continuous Mixing (CM) (§7) [Open]: whether admissible transport includes one element genuinely blending two identical sectors beyond swap-and-phase. Exhibit it, and the class group closes to U(k) — SU(k) per class, the non-abelian forces arrive (branch M); prove its universal absence, and the forces remain abelian-plus-swap (branch T) — stated plainly as close to fatal for Standard-Model recovery, the largest kill-condition in the corpus, and internal: decided by the substrate's own structure at the slots, not by experiment. Slots at the reconstruction papers (inter-sector superposition; continuity of admissible dynamics) and at the finite-impossibility companion (per-pair generalization); the pressure arguments toward blending carried as [Conjectural], disciplined by the witnessed-wall honesty of the second-integer paper.

The architecture, conditional. The **Class Product Theorem** (§8): mixing forced within classes, unforced across them; under the corpus's surplus-structure economy, $G = \prod_i U(n_i)$ over classes — non-abelian blocks, one per class, abelian circles alongside: the product silhouette SU(3) × SU(2) × U(1) instantiates, with gauge multiplets identified as the corpus's species-classes and the framework explaining product form against simple unification as architecture rather than breaking [Conditional on the economy posture]. The class census — the sizes n_i, why 3 and 2 — is [Open], slotted to the Fold-structure papers; chirality, hypercharge, and breaking are fenced as untouched (§9); the abelian record machinery's non-transfer is inventoried.

Epistemic markers: [Inherited] imported from prior VERSF papers; [Imported-External] imported from standard mathematics outside the programme, carried at the external source's standing; [Proven] established here; [Conditional] holding under stated inputs; [Conjectural] motivated but unproven; [Open] undecided.

Table of Contents

1. Introduction — Opening the Reserved Arc from Uniqueness

A dozen papers in this corpus reserve their hardest questions to "the multi-Fold arc" — and the corpus's own foundational work has since made that name impossible to use innocently. The Fold papers prove uniqueness: there is exactly one Fold, and its architecture saturates observable structure. A paper that opened the reserved arc by supposing several Folds would not be extending the corpus; it would be contradicting its most defended theorem, and a referee would say so in the first paragraph. This paper therefore opens the arc from the only door uniqueness leaves: the multiplicity the Standard Model displays — many particle species, forces richer than one circle — must live *within* the one Fold, and the corpus's own ontology says where: in the **closure sectors**, the persistent configurations of the Persistent-Fold-Defect picture, already grouped at source into Representation Classes, with species identified as classes. The multiplicity this paper trades in is therefore never multiplicity of Folds; it is multiplicity of stable PFD sectors within the one Fold architecture — the sentence that aligns this paper with the ontology papers rather than against the uniqueness ones, and the one the reader should carry through everything below. One rope; many knots; and — the hinge of everything below — possibly identical knots, tied more than once.

The claim, before the machinery, at the altitude it deserves: if reality is built from a single fundamental structure, all observed multiplicity must arise from configurations of that structure; if two configurations are genuinely identical, physics cannot depend on which copy is which; the mathematics that respects this indistinguishability is — wherever the copies can be exchanged — already non-commutative; and the strong and weak forces may therefore exist not because reality contains multiple fundamental objects but because the one Fold contains multiple indistinguishable closure sectors. The forces as the price reality pays for containing identical copies of the same pattern: that sentence is the paper, and the reader should hold it through every normalizer and slot that follows — the derivation exists to determine exactly how much of it is forced, how much is conditional, and on what.

The entrance question, then, is not *what if there are more Folds* but: **does the unique Fold architecture admit two or more mutually indistinguishable persistent closure sectors — and if it does, what transport structure is forced?** The paper's findings divide three ways, in the corpus's standard grain. Forced: a compact Lie group of transports, relabeling-covariant within every class of identical sectors; and — where the substrate can implement the exchange of identical sectors, a named import whose physical reading is whether identical defects can braid past each other — the permutations themselves, which makes the group non-abelian at the stated grain, so that non-commutativity is revealed as not an addition to the framework but a consequence of the one Fold containing copies it can exchange, before any dynamics. Undecidable from principles: whether the mixing of identical sectors is merely discrete — swap — or continuous — blend; the paper reduces this to one element in the two-sector atom, where the complete classification of closed subgroups of SU(2) collapses everything above the full floor, and one named slot governs the thinned floor's residue. Open: that element's existence (the node CM), the class census, and the arc's interior.

Three disclosures govern the paper. First, **the uniqueness compatibility claim, made explicit**: nothing here multiplies Folds; Fold Uniqueness and Fold Saturation are consumed as inherited theorems, and every plural in this paper — sectors, classes, copies — is sectorial, internal to the one architecture. The paper's question is in fact the question uniqueness *creates*: a corpus that proves one Fold owes an account of the world's evident many, and the sector ontology is that account's site. Second, **the unitarity debt**: the charge paper closed the reversibility-versus-norm-preservation gap for free only in one dimension; the n-sector case has no free closure, and norm-preservation is a named import (U1), slotted at the quantum-reconstruction papers. Third, **the one-phase subtlety**: one Fold means one derived global phase circle, so the independence of *relative* phases between identical sectors — taken silently in any many-Fold framing — is here an import of its own (S2), faced with a fallback under which the central theorems are restated and survive. The hardest arc in the programme opens the way the easier ones did: exact about which sentences are theorems, and starting from the corpus's results rather than against them.

2. Inherited Results and the New Debts

Fold Uniqueness and Fold Saturation [Inherited — ⟨cite: Fold papers, the uniqueness and saturation theorems⟩]. Exactly one Fold; its architecture exhausts observable structure. Consumed as the frame of everything: the paper's multiplicity is sectorial by construction, and §11 carries the compatibility refuter.

The sector ontology [Inherited — ⟨cite: Persistent-Fold-Defect papers; Representation Classes; species-as-classes⟩]. The unique Fold admits persistent closure sectors — stable configurations, the corpus's particle ontology — grouped into Representation Classes with species identified as classes. Consumed in §3, where the paper's class is *defined as* the source's Representation Class — the definition imported rather than paralleled, with member-indistinguishability derived in §5 from the identity programme rather than left as a bare slot — and in §5, whose opening grounds the multiplicity claim in this ontology and proves non-degeneracy inherited (the Class Multiplicity Proposition): the multiplicity is never of Folds; it is of stable PFD sectors within the one architecture.

The identity programme [Inherited — ⟨cite: invariant-exhaustion — particle identity exhausted by invariant structure; No-Surplus-Identity⟩]. Identity is exhausted by the invariant catalogue, and unobservable surplus identity is excluded. Consumed in §5's Indistinguishability Lemma, which derives S3's member-indistinguishability from exactly these two principles — the dilemma's two horns.

The single-phase theorem and the abelian arc [Inherited — U(1) paper; charge paper; structure-group paper]: the Fold's transport phase is exactly U(1); carriers respond by integer winding; the gauge identification is discharged at its slot status. Consumed in §4 (the global circle) and as the template for every forcing argument.

No-pre-individuation and the Conservation Lemma [Inherited — at source]: unwitnessed distinctions are not admissible structure. The engine of §5 — its deepest deployment in the corpus, removing not a boundary or a response gap but a *labeling*.

The realized-witness discipline and the surplus-structure economy [Inherited — as deployed across the recent arc]: consumed in §7's motivation discipline and §8's conditional exclusion respectively.

The new debts, as debts: norm-preservation in n dimensions (U1); sector-resolved phase advance (S2); the identity-programme premises behind S3 (invariant-exhaustion and No-Surplus-Identity, from which §5 derives member-indistinguishability); exchange realizability (S4); inter-sector superposition and continuous dynamics (Q1); the per-pair finite-impossibility generalization (Q2); the class census. Each is a slot, none a silent crossing — the widest import base in the corpus, recorded as such.

3. The Carrier and the Imports

The carrier. The amplitudes of the unique Fold's persistent closure sectors: for a class of k identical sectors, the class carrier $\mathbb{C}^k$, with the full carrier the direct sum over classes — internal dimensionality supplied at source by the Fold architecture itself [⟨cite: Fold papers, the carrier structure — the architecture's internal complex dimensionality⟩], so that sector multiplicity inhabits structure the one Fold already possesses rather than structure added to it. Admissible transport acts on this carrier; the object of the paper is the group its actions generate.

U1 (Unitarity of sector transport). *Admissible transport acts on the n-sector carrier by norm-preserving linear maps: elements of U(n).* Reversibility supplies invertibility [Inherited]; norm-preservation is the import, its natural source the conservation of the Born measure under admissible transport, as derived in the quantum-reconstruction arc. [Inherited — ⟨cite: quantum-reconstruction papers, unitary dynamics from admissibility⟩. *Fallback if weaker:* invertibility alone weakens the skeleton from U(n) to a larger linear group, and the downstream classification must be re-derived in that setting; the architecture survives, the specific groups do not.]

S2 (Sector-resolved phase). *Admissible transport can advance one sector's phase with the others idle: the relative phases between sectors are admissible degrees of freedom, and the group contains the diagonal torus T^n.* The one-Fold ontology makes this a genuine question — one Fold supplies one global phase as a theorem; relative phase is more — and the natural source is the same place charge differences live: sectors of different classes already respond to the global phase with different windings [Inherited — charge paper], so *inter-class* relative phase is established; the import's content is *intra-class* relative phase, between identical sectors. [⟨cite: reconstruction papers or PFD papers, the phase structure of multi-sector states⟩. *Fallback if weaker:* if only the global phase and inter-class relative phases are admissible, the forced floor of §6 shrinks from the full normalizer to swap-plus-class-scalars; the Indistinguishability Theorem survives at its true grain (§5c: non-abelian unconditionally for $k \geq 3$; for $k = 2$ the thinned floor is abelian and non-abelianness waits on S2 or CM); the one-element collapse survives conditional on Q2 (§6: the thinned floor leaves the finite exceptional subgroups of SU(2) live, and Q2 is exactly what excludes them); and the floor's description changes.]

S3 (Class structure). *The paper's class is the source's Representation Class — the definition is imported, not paralleled — together with the property the arguments below consume: members of one Representation Class are mutually indistinguishable, with no admissible invariant*

separating members absent realized configuration differences. The definitional import removes the coincidence question entirely; and the property is no longer a bare verification slot: §5's Indistinguishability Lemma derives it from the identity programme — equal invariant tuples entail genuine indistinguishability, on pain of either contradicting the catalogue or violating No-Surplus-Identity — so the slot relocates upstream, to two principles the corpus has already defended together with one definitional premise: that class membership entails equal invariant tuples, which is the source's own meaning of "Representation Class" and is carried by this entry's definitional import. Likewise non-degeneracy — at least one class with $k \geq 2$ — is inherited rather than posited: §5's Class Multiplicity Proposition shows a singleton-only ontology would collapse species into object wholesale, against the source's load-bearing distinction. The **indistinguishability refinement** remains as belt-and-braces: if the lemma's premises were ever weakened at source, the arguments run on the maximal mutually indistinguishable subsets within each Representation Class, with only the census relabeled — a graceful degradation, not a failure mode. [Inherited definition; property derived (§5, Indistinguishability Lemma), conditional on invariant-exhaustion and No-Surplus-Identity at source — ⟨cite: identity programme papers; Species-as-Classes paper⟩. The census — class sizes n_i, in particular 3 and 2 — is [Open] at the species inventory.]

S4 (Exchange Realizability). *The physical exchange of two identical closure sectors is an admissible transport: the permutation of a class pair is implemented in G.* Label-covariance alone does not supply this (§5a, with the diagonal torus as the standing counterexample), and the import is the difference between a symmetry and a gauged symmetry: that physics is blind to which identical sector is which is forced by no-pre-individuation; that the blindness is *implemented by transports* — that the substrate can actually carry one identical defect past its twin — is a further physical fact. Natural sources: the PFD papers (whether the substrate braids identical defects) and Q1's continuity (a continuous admissible path implementing the exchange would supply S4 and bear on CM at one stroke). [⟨cite: PFD papers, defect exchange; or reconstruction papers via Q1⟩. *Fallback if absent:* the membership half of §5 lapses and only covariance survives — G normalized by S_k, the swap absent rather than forced — a graceful degradation to the theorem's unconditional half, recorded in §11.]

The external imports, recorded once: Cartan's closed-subgroup theorem; the classification of closed subgroups of SU(2) — finite; conjugates of U(1); the normalizer N(U(1)); SU(2) entire; and the generation theorem — the diagonal torus with the SU(2) blocks on all coordinate pairs generates U(k). [All Imported-External, at their standing.]

4. The Compact Lie Skeleton and the Phase Substructure

The skeleton. Admissible transports compose and invert [Inherited]; under U1 each is an element of U(n); the generated subgroup is forced closed by the Boundary Maintenance discipline, exactly as at the single-phase source — a boundary between the generated group and its closure is a boundary no admissible invariant maintains [Inherited mechanism, applied]. The admissible transport group is therefore

$G \subseteq U(n)$, closed,

and by Cartan's theorem a compact Lie group. What survives every generalization of the single-phase theorem is *compact Lie*: the dimensional and abelian clauses fall away, the compactness and the closure do not. [Proven, conditional on U1 and the inherited closure mechanism.]

The phase substructure. The global circle is inherited outright: every sector's amplitude carries the Fold's derived U(1), and the scalars $e^{i\theta}\cdot\mathbb{1}$ lie in G [Inherited]. Inter-class relative phase is inherited through the charge lattice: classes respond with different windings, so transport already separates class phases [Inherited — charge paper]. Intra-class relative phase — the diagonal torus within a class of identical sectors — is S2, with its fallback as stated. Under S2, G contains the full diagonal torus T^n and is of maximal rank in U(n); under the fallback, G contains the global circle, the inter-class phases, and (from §5, under S4) the permutations — a thinner floor on which §6's collapse runs conditional on Q2. The paper carries both readings explicitly so that the source can select between them rather than the paper presuming. [Proven at the stated conditionality; the bifurcation is the honesty, not a weakness.]

5. The Indistinguishability Theorem — Non-Abelian Structure, Forced

Why identical sectors are expected — multiplicity inherited, not introduced. Before the theorem, the paper discharges the question a referee should ask first: why should classes with $k \geq 2$ identical sectors exist at all, and why should their members be *genuinely* indistinguishable rather than merely similar? Neither is this paper's posit; both are the corpus's existing ontology, taken seriously. The chain **Fold → Persistent Fold Defect → Representation Class → Species** is already in place at source: the unique Fold admits persistent defect sectors; the Representation Classes group multiple realizations of identical invariant structure — that is what makes them classes rather than labels for singletons; and the species programme treats particle identity as class identity, so that "two particles of the same species" is, at source, already "two realizations within one class." The multiplicity is never of Folds: **it is multiplicity of stable PFD sectors within the one Fold architecture.**

Two formal observations convert the appeal to ontology into inherited results.

Proposition (Class Multiplicity Necessity). *The species ontology presupposes classes with multiple realizations.* A Representation Class containing a single admissible realization is observationally indistinguishable from an individual object: for it, the distinction between species and object collapses. The Species-as-Classes programme maintains that distinction as load-bearing — a species is an equivalence class of realizations, not a primitive object — and an ontology in which every class were a singleton would collapse species into object wholesale, leaving the class apparatus with no work to do. Non-degeneracy — at least one class with $k \geq 2$ — is therefore not an additional hypothesis but the natural content of "class" in the source ontology. The census of *which* classes carry multiplicity remains open, and nothing here claims every class does: singleton classes are admissible, the ontology's analogue of singlet representations; what the ontology cannot be is singleton throughout. One escape is closed in advance: a defender of a singleton-only ontology might reply that the species/object distinction survives on counterfactual work — a species *could* have had multiple realizations even if no

actual class does. The corpus's own tools answer: an apparatus doing only modal work is surplus structure under the economy posture, and the witnessed-distinction discipline disfavors distinctions that no realized configuration ever instantiates — and in any case the empirical anchor below carries the actual load. [Proven at the strength of the source's species/object distinction — ⟨cite: Species-as-Classes paper, the distinction and its role⟩.]

Indistinguishability Lemma (equal invariants entail indistinguishability). *Members of one class with equal invariant tuples are genuinely mutually indistinguishable.* Let two realizations D_1, D_2 satisfy $\mathcal{I}(D_1) = \mathcal{I}(D_2)$, and suppose some further fact distinguished them. Either the fact is observable — in which case it is admissible invariant content and belongs in the catalogue, contradicting the equality of the tuples, since identity is exhausted by invariant structure [Inherited — the identity programme] — or it is unobservable, in which case it is surplus identity, excluded by the No-Surplus-Identity discipline [Inherited]. There is no third horn: equal tuples entail genuine indistinguishability, and "merely similar" is not a status the ontology permits members of one class to occupy. [Proven, conditional on invariant-exhaustion and No-Surplus-Identity at source, and on the antecedent that class membership entails equal invariant tuples — the source's own meaning of "Representation Class," carried by the definitional import of §3 — ⟨cite: the identity programme papers⟩.]

The lemma upgrades S3's standing (§3): member-indistinguishability is no longer a bare verification slot but a derived consequence of the identity programme, with the slot relocated upstream to two principles the corpus has already defended at length. And with both results in hand, the burden reverses. The empirical world already presents repeated realizations of apparently identical structure — three colour states in one representation class, two weak-doublet components in another — so the question confronting the paper is not *why assume copies?* but the one the remainder answers: the world contains them, and the ontology explains why it must; what transport structure do they force? What this paper adds is not the sectors but the question their identity forces on transport. [The chain: Inherited; the proposition and lemma: Proven at stated conditionality; the anchor: calibration, no marker.]

Here is the paper's central derivation, and the corpus's deepest application of its oldest principle — now stated in the ontology the Fold papers force.

Theorem (Indistinguishability forces relabeling covariance; exchange realizability forces membership). *Let a class contain $k \geq 2$ mutually indistinguishable closure sectors of the unique Fold (S3). Then: (a) — the unconditional half — the admissible transport group is relabeling-covariant: $\sigma G \sigma^{-1} = G$ for every permutation σ of the class, so G is normalized by S_k and its form is constrained accordingly; (b) — the membership half — if the exchange of identical sectors is an admissible transport (S4), then $S_k \subseteq G$; and (c) G is non-abelian: for every $k \geq 2$ under S2, and for every $k \geq 3$ under S2's fallback, with the $k = 2$ fallback floor abelian and its non-abelianness waiting on S2 or CM.*

The argument for (a). By S3 the labeling is unwitnessed; by no-pre-individuation it is not a fact, and no admissible structure may depend on it. The transport structure, described in relabeled coordinates, must therefore be the same structure: $\sigma G \sigma^{-1} = G$. This much is forced — and this much is *all* that label-independence forces. The stronger conclusion $\sigma \in G$ does not follow, and

the paper's own materials show why twice over. The diagonal torus is relabeling-covariant, abelian, and contains no swap: a standing counterexample to any claimed entailment from covariance to membership. And the §7 discipline applies in advance: a boundary in the group at "swap excluded" is witnessed — the swap, were it admissible, would act distinguishably — so the Conservation Lemma permits exactly that boundary, and §5 may not consume what §7 declines. [Proven — (a) — conditional on S3 and no-pre-individuation at source.]

The argument for (b). What converts covariance into membership is a physical fact, not a logical one: that the substrate can implement the exchange — carry one identical defect past its twin and return the configuration to itself up to the swap of amplitudes. That is S4, an import with its slot (§3), and under it the relabelings are transports: $S_k \subseteq G$ by permutation action on the class carrier. The distinction being drawn is one the conventional literature lives with and rarely derives: identical-particle permutations are *symmetries* — with superselection and statistics consequences — and whether the symmetry is *gauged*, implemented by the transport structure itself, is a further fact about the world. The paper derives the distinction rather than blurring it: covariance is forced, gauging is imported, and the import has a physical reading — can identical defects braid? — that the source papers can answer. [Proven — (b) — conditional additionally on S4.]

The argument for (c). Under S2 the floor contains the intra-class torus, and the swap conjugates one sector's circle to the other's: G is non-abelian for every $k \geq 2$. Under the fallback, S_k is non-abelian by itself for $k \geq 3$, which suffices; for $k = 2$ the thinned floor — the swap together with class-scalar phases — is abelian, since the swap commutes with everything scalar on the class, and the paper states the true grain rather than papering it: the $k = 2$ fallback floor is abelian, and non-abelianness there waits on S2 or on CM's blend. The irony is physically pointed and is owned: the weak doublet is the $k = 2$ case, so the Standard Model's SU(2) is precisely where the fallback bites hardest — one more reason S2's slot matters. [Proven — (c) — at the stated grain, conditional on S4 throughout.]

Three remarks fix the result's meaning. First, its *direction*: conventionally, non-abelian gauge symmetry is postulated; here it is forced at the stated grain — the framework is not permitted an abelian theory of identical sectors once exchange is realizable and the floor is rich enough, and the conditions under which it would escape are themselves named slots, not silences. A unique Fold that ties the same knot twice, under the discipline that nothing may depend on unwitnessable facts, generates relabeling covariance outright and non-abelian transport wherever the substrate can braid its copies. Second, its *ancestry and its fit*: this is the identical-particle principle located rather than imposed — the covariance half derived from no-pre-individuation (physics blind to which copy is which), the gauging of that blindness imported as S4's physical content — so the corpus derives the symmetry-versus-gauge distinction instead of blurring it. The Standard Model's gauge multiplets, on this reading, are the classes, and the symmetry is what indistinguishability forces, gauged where exchange is realizable. Third, its *limit*: permutations are discrete. The theorem forces at most the swap, never the blend; the forced floor is far below U(k); and whether the substrate blends is §§6–7's business. The floor is non-abelian where stated; the Standard Model needs the ceiling.

6. The Two-Sector Atom — Classification and the One-Element Gap

The reduction. A class is generated by its pairs: U(k) is generated by its diagonal torus together with the SU(2) blocks on coordinate pairs [Imported-External], so the class group is determined by the torus structure (§4) and the pairwise mixing subgroups $M_{ij} = G \cap SU(2)_{ij}$. Within a class, indistinguishability makes all pairs conjugate; one pair speaks for all, and the entire question reduces to the **two-sector atom**: the admissible mixing group of two identical closure sectors.

The classification and the floor. The closed subgroups of SU(2) are completely classified [Imported-External]: finite; circles; the normalizer N(U(1)); SU(2). Both floors consume S4 — the swap sits in the floor only where exchange is realizable; absent S4 the forced content is the phase structure alone, with covariance still constraining G's form (§5a). Under S2, the floor is the relative-phase circle plus the swap — exactly N(U(1)) — and the floor's circle excludes the finite and bare-circle options outright, leaving a two-point decision:

M confined to the normalizer (swap-and-phase) — or M = SU(2) (full blending).

Under S2's fallback, the floor is the swap plus class scalars — thinner, and for k = 2 abelian (§5c) — and the finite closed subgroups of SU(2) remain live, which changes the gap statement below.

The one-element gap, by floor. Under S2's full floor the collapse is unconditional: there is no closed group strictly between N(U(1)) and SU(2), so any single admissible blend generates SU(2) entire. Under the fallback floor the honest statement is a trifurcation, not a collapse: the exceptional finite subgroups — the binary tetrahedral, octahedral, and icosahedral groups — contain swap-type elements together with genuine off-normalizer blends, so a blend at a suitable rational angle generates, with the swap, a finite group: closed, unconfined to any normalizer, and not SU(2). What excludes the finite branch is exactly Q2 — the per-pair generalization of finite-mixing impossibility — thereby promoted from narrowing to load-bearing on the fallback branch: under the fallback, *one blend plus Q2* forces SU(2). The headline survives at its stated grain: **one admissible blend decides the atom — unconditionally under S2's floor, conditional on Q2 under the fallback — the atom decides the class, and the classes decide the forces.** [Proven — the reduction, the floors, and the gap-by-floor — conditional as stated; classification [Imported-External].]

The single-phase paper's structure recurs at the new level, and the recurrence should be recorded: there, infinitude forced density in the circle and closure did the rest; here, one off-normalizer element forces density in SU(2) and closure does the rest. Floor, classification, collapse — the corpus's forcing pattern is apparently scale-invariant, and the present instance reduces the largest reserved question in the programme to the existence of a single group element.

7. The Decision Node — Continuous Mixing

Continuous Mixing (CM) — *the decision node of this paper* [Open]

Admissible transport includes at least one element that genuinely blends two identical closure sectors: an element of the pair's SU(2) outside the swap-and-phase floor — a partial rotation of one sector's amplitude into its indistinguishable twin's. (Negation: every admissible mixing of identical sectors is swap-and-phase.)

Slot Q1 (inter-sector superposition and continuous dynamics). The reconstruction papers' state space and dynamics: whether superpositions across identical sectors are admissible states, and whether admissible transport connects them continuously — a continuous admissible path from the identity to the swap passes through blends and decides CM positively at a stroke. This is the node's most probable decider, and the paper says so: a Hilbert-space reconstruction with continuous unitary dynamics on the class carrier contains blends unless something at source confines the dynamics to the normalizer, which would itself be a remarkable finding demanding its own mechanism. [⟨cite: quantum-reconstruction papers, the multi-sector state space and the continuity of admissible dynamics⟩]

Slot Q2 (finite-mixing impossibility, per pair). The companion's finite-holonomy-impossibility, generalized to the pair's mixing subgroup. Under S2's floor, Q2 narrows without deciding — an infinite closed M is a circle, the normalizer, or SU(2). Under S2's fallback, Q2 is load-bearing (§6): the thinned floor leaves the finite exceptional subgroups live, Q2 is what excludes them, and the fallback's one-element collapse is conditional on exactly this slot. [⟨cite: companion paper; hypotheses checked against the pair sector⟩]

Motivation toward blending, carried as motivation. Two pressures point at branch M; neither is a proof, and the discipline of the second-integer paper governs both. The wall pressure: a hard admissibility boundary at the normalizer — swap admissible, every partial blend forbidden — sits inside a connected ambient group and invites the question of what maintains it; but distinct group elements act distinguishably, the wall is therefore *witnessed*, and the Conservation Lemma does not remove witnessed boundaries — the pressure is economic, not derivational [Conjectural]. The interference pressure: if intra-class superpositions are physical and their interference is part of the derived Born phenomenology, transport rotating the interference basis is hard to exclude — and "hard to exclude" is a mood, with the corpus's recorded standard for moods. The bias is at most mechanism-level pending Q1; the marker is the marker.

The branches, prepared. *Branch M (blending):* one element exists; the atom closes to SU(2) — outright under S2's floor, with Q2 under the fallback (§6); a class of size k closes to U(k) under S2, to SU(k) with the shared circles under the fallback; the non-abelian forces arrive with their conventional local structure — the class's transport connection is an SU(k)-type gauge field, and the reserved questions (occupancy, embedding, weights-for-windings) acquire their arena. *Branch T (swap-only):* the multi-sector forces are abelian-plus-swap — a gauge theory of the normalizer, nothing like SU(2) or SU(3) — and the kill-condition is stated at full strength: **branch T is close to fatal for the programme's Standard-Model ambitions**, the largest named kill-condition in the corpus and an internal one, decided by the substrate's own structure at Q1–Q2 rather than by any experiment. A programme becomes scientific in proportion to what can

kill it; this paper arms the largest weapon yet pointed at its own programme, and does so on purpose.

The verification posture. Branch M by existence — one blend, plausibly already implicit in Q1's continuous dynamics; branch T by universality — confinement to the normalizer, a strong constraint that would itself demand a mechanism. The asymmetry mirrors SA's; the paper declines to lean beyond its sources, and the sources are unread at the slots. [Open; the node, the one-element criterion, and both branches are the deliverables.]

8. The Class Product Theorem — The Shape of the Forces

Within classes, mixing is forced (§5) at least to the swap and conditionally (CM) to the full unitary group. *Across* classes, the situation inverts, and the inversion generates the Standard Model's most distinctive architectural feature.

Across classes, mixing is unforced. The theorem's engine — the unwitnessed labeling — has no purchase across classes: sectors of different classes are distinguishable by definition (S3); their labels are witnessed facts; no relabeling symmetry is forced. Cross-class mixing is not forbidden by §§4–6; it is *unforced*, and the corpus's posture for unforced structure applies: surplus structure carries the burden of exhibiting its work [Inherited — the economy posture]. The exclusion is carried at exactly the posture's strength:

Class Product Theorem. *Under the economy posture, the admissible transport group is the product of the per-class groups:*

$G = \prod_i G\_class_i$, with $G\_class_i = U(n_i)$ under branch M and S2.

Under S2 the class factors carry their own circles and the product reads $\prod_i U(n_i)$; under S2's fallback the intra-class circles are absent and the product reads $\prod_i SU(n_i)$ together with the shared global and inter-class circles — the bifurcation discipline of §4, carried through.

[Conditional — on the posture, marked as posture; a cross-class mixing exhibited at source with independent admissible work amends the product to the exhibited extent, and the refuter is exactly that exhibition. One known candidate for such an exhibition is flagged now rather than discovered later: the flavour-mixing phenomenology (CKM/PMNS) involves transitions across what may be classes, and whether those transitions are transport-structural or dynamical-interactional is a question the posture hands to the flavour arc — the product theorem governs the *gauge* group, and the flag prevents its overreading.]

The shape, read against the target. Under branch M: $\prod_i U(n_i)$ — non-abelian blocks, one per class, abelian circles alongside. That is the Standard Model's silhouette: $SU(3) \times SU(2) \times U(1)$ is a product of class blocks with abelian factors, the gauge multiplets are the classes — the corpus's species — and the framework derives the *form* while leaving the *instance* to the census (why classes of sizes 3 and 2 — [Open] at the species inventory) and to the abelian consolidation (how the per-class circles and the derived electromagnetic U(1) relate — the hypercharge question, named, not attempted). The calibration against convention: standard model-building selects its

group from the infinite catalogue by fit; here the product-of-blocks form is forced by the distinguishability partition of the sector inventory, and the simple unification groups — SU(5), SO(10) — are exactly what the architecture does *not* produce. The framework, if right, explains why the forces come as a product and not as one unified block: the observed product form, which unification programmes treat as the debris of breaking, is here the architecture itself. [The shape: Proven under branch M and the posture; the calibration: no marker.]

9. What Does Not Transfer and What Is Not Touched

The abelian arc's machinery does not generalize naively. The **Cap Lemma is abelian-only by its own clause** [second-integer paper §3]: non-abelian composition is path-ordered, interior-edge cancellation needs commutativity it lacks, and the record/winding apparatus — the second integer, the Bounding Lemma, the sector analysis — awaits non-abelian reconstruction, with the conventional analogues (path-ordered exponentials, higher characteristic classes) marking territory the substrate derivation has not entered. The **charge lattice generalizes to weight lattices**: the integer q is the one-dimensional shadow of representation weights, and the multi-sector response classification is the representation theory of G — downstream, where the occupancy question lives. Three Standard-Model structures are **untouched, by name**: **chirality** — the weak SU(2) acts on left-handed doublets, a parity structure nothing in sector-counting addresses, and the single largest gap between the derived shape and the physical instance; **hypercharge** — the abelian consolidation and the embedding $Q = T_3 + Y/2$; **symmetry breaking** — the Higgs sector and the masses. And one structure is *deliberately not claimed even under branch M*: that the class carrier's SU(k) acts as the Standard Model's corresponding factor acts — the identification of derived class groups with physical force factors is the arc's interior, requiring the census, the chirality account, and the coupling structure, none present here. The paper opens the arc; it does not traverse it.

10. Position in the Programme

Inherited exposure. The widest import base in the corpus: the Fold papers (uniqueness, saturation, sector ontology, carrier dimensionality) at U-slots; the reconstruction papers at U1 and Q1; the companion at Q2; the abelian arc entire, at its accumulated slot statuses.

Compatibility with uniqueness, stated as a result. Every plural in this paper is sectorial: the carrier is the one Fold's internal structure; the classes are its species; the groups are transport on its amplitudes. No step multiplies Folds, and the paper's question — whence multiplicity, given uniqueness — is the question the uniqueness theorems create rather than a hedge against them. [Proven by inspection of the construction; the refuter in §11.]

Assignment-neutrality. The inputs are catalogue-level and definitional throughout — the skeleton quantifies over transports, the theorem over labelings, the product over the partition — and none consumes which path carries which holonomy. The arc opens independent of Transport-Completeness, extending the independence chain unbroken. [Proven, given the stated grain.]

The diagram, extended:

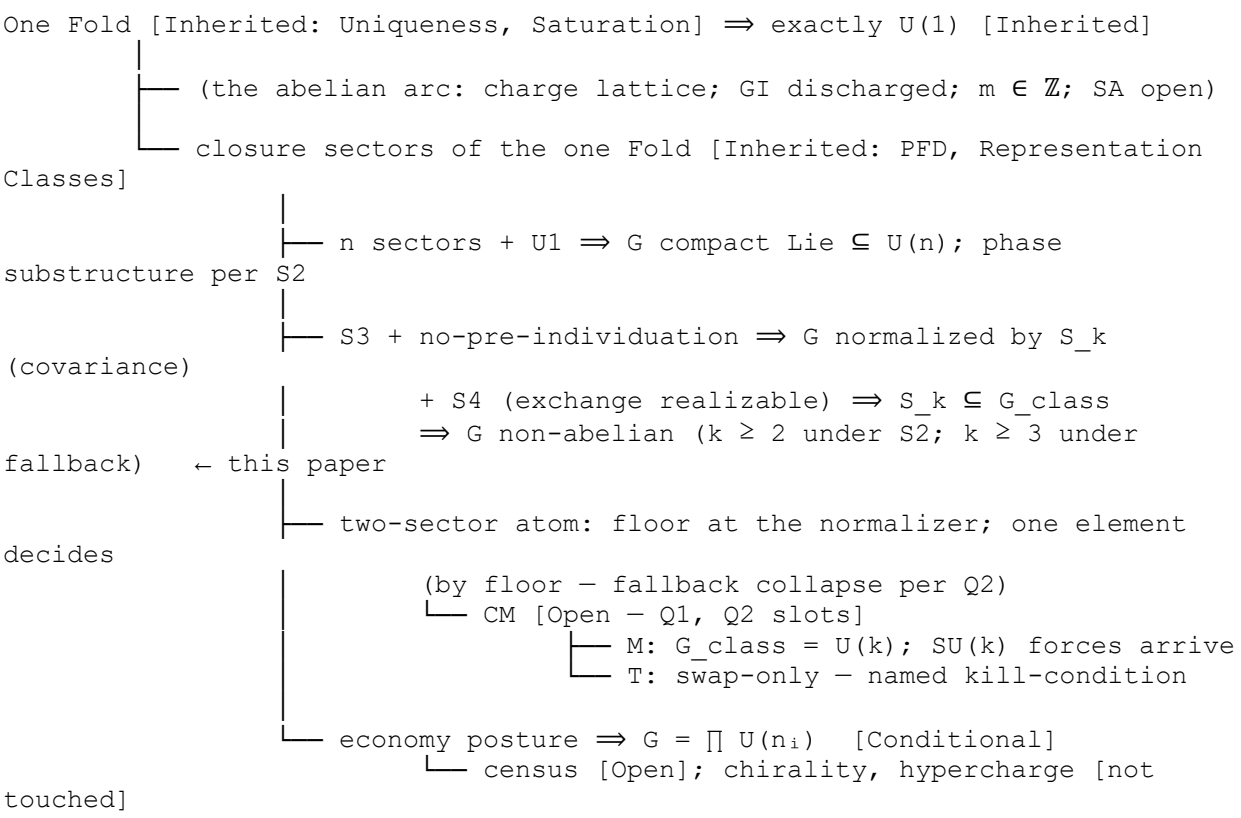

11. What Would Refute or Decide This

Against the uniqueness compatibility (§10). Exhibit a step at which the construction's plurality is not sectorial — a second Fold smuggled in as a "sector." Standing: the carrier is the one Fold's internal dimensionality at source; the refuter's burden is to show the source's sector ontology itself violates saturation, which would strike the Fold papers, not this one.

Against the Indistinguishability Theorem (§5). Exhibit an admissible invariant distinguishing sectors within a class absent realized configuration differences — which contradicts the equality of invariant tuples at the catalogue (the Indistinguishability Lemma's first horn) and, under S3's refinement clause, relocates the sector to its own refined class: a degradation of the census, not of the theorem — or exhibit a sense in which admissible structure may depend on an unwitnessed labeling, which strikes no-pre-individuation at source. The covariance half is exactly as secure as the corpus's deepest principle. A third, honest entry: exhibit that exchange is not an admissible transport at source — S4's failure — which removes the membership half and leaves covariance standing: the swap moves from forced to absent, a graceful degradation to the theorem's unconditional half rather than a strike on no-pre-individuation. And the grounding pair (§5, opening) fall only with their premises: strike invariant-exhaustion, No-Surplus-Identity, or the species/object distinction at source — each a strike on the identity programme and the species ontology, not on this paper.

Against the skeleton and floor (§4, §6). U1's fallback (invertible-not-unitary) reopens the group theory in a larger setting; S2's fallback thins the floor with the collapse surviving, as stated — both are graceful degradations, not refutations, and are recorded as such.

Deciders of CM (§7). Q1's reading — continuous admissible dynamics on the class carrier decides M nearly immediately, or confinement-to-the-normalizer is discovered and demands its own mechanism; Q2 narrows; a universality proof of confinement decides T and activates the kill-condition.

Against the Class Product Theorem (§8). Exhibit cross-class mixing at source with independent admissible work; the product amends to the exhibition, exactly as the posture provides — with the flavour-mixing flag of §8 marking the place to look first.

Empirically, at the horizon. Branch M's product form, against simple unification: proton decay at unification rates would evidence the single-block architecture this framework does not produce; its continued absence is the framework's quiet friend. Horizon, not lever.

12. What the Paper Establishes

Established (conditionally, with conditions named):

- The reframing as a result: the reserved arc opened *from* Fold Uniqueness rather than against it — multiplicity located in the closure sectors of the one Fold and grounded as the PFD ontology's existing content (multiplicity of stable defect sectors, never of Folds); the paper's class *defined as* the source's Representation Class, with member-indistinguishability derived and non-degeneracy inherited (§5, opening) and the refinement fallback retained as belt-and-braces (§3). [The compatibility: Proven by construction; the property: derived, conditional at the identity programme's source.]
- The grounding pair: the Class Multiplicity Proposition — the species ontology presupposes classes with multiple realizations, since a singleton-only ontology collapses species into object against the source's load-bearing distinction, with singleton classes themselves remaining admissible — and the Indistinguishability Lemma — equal invariant tuples entail genuine indistinguishability, the distinguishing-fact dilemma closing both horns via the catalogue and No-Surplus-Identity; together they convert S3 from assumption to inheritance and reverse the burden: the world already contains copies, the ontology explains why it must, and the question is what transport they force (§5, opening). [Proven, conditional on the species/object distinction, invariant-exhaustion, No-Surplus-Identity, and the definitional entailment that class membership implies equal invariant tuples, at source.]
- The compact Lie skeleton: G closed in U(n), compact Lie (§4). [Proven, conditional on U1 and the inherited closure mechanism.]
- The phase substructure, bifurcated honestly: global circle inherited; inter-class relative phase inherited through the charge lattice; intra-class relative phase as the import S2, with both readings carried and the downstream theorems surviving either (§4). [Proven at stated conditionality.]

- The unitarity debt, surfaced: the one-dimensional free closure does not generalize; U1 stated as the import it always was (§3). [Disclosure; the slot is the obligation.]
- The Indistinguishability Theorem, split at its true grain: covariance — G normalized by S_k — forced unconditionally by no-pre-individuation, with the diagonal torus as the standing counterexample to anything stronger; membership — S_k ⊆ G — conditional on the new import S4 (Exchange Realizability: identical defects can be braided); non-abelianness for $k \geq 2$ under S2 and $k \geq 3$ under the fallback, with the $k = 2$ fallback floor abelian and the weak-doublet flag planted; the symmetry-versus-gauge distinction derived rather than blurred (§5). [Proven at the stated grain, conditional on S3, S4, and inherited principles.]
- The two-sector reduction and the one-element gap, by floor: pairs generate; the floors under each phase reading, both consuming S4; under S2 the collapse is unconditional — no closed group between the normalizer and SU(2); under the fallback it is a trifurcation with the finite exceptional subgroups live and Q2 promoted to load-bearing — one blend plus Q2 forces SU(2) (§6). [Proven, conditional as stated; classification [Imported-External].]
- The decision node CM with slots Q1–Q2, the one-element criterion, motivation at [Conjectural] under the witnessed-wall discipline, both branches prepared, and branch T's kill-condition named at full strength — the corpus's largest, and internal (§7). [Open; the node is the deliverable.]
- The Class Product Theorem: mixing forced within classes, unforced across; under the posture $G = \prod U(n_i)$ given S2, and $\prod SU(n_i)$ with the shared circles under the fallback — the Standard Model's product silhouette with multiplets-as-classes, product form explained as architecture against unification's breaking story, and the flavour-mixing flag planted against overreading (§8). [Conditional on the posture, marked as posture.]
- The non-transfer inventory and the fence: path-ordering against the Cap Lemma; weights for windings; chirality, hypercharge, breaking, and the force-factor identification — named, fenced (§9).
- Uniqueness compatibility and assignment-neutrality (§10). [Proven as stated.]

Not established (open, out of scope, or pending):

- CM — blend or swap; slots Q1–Q2 (§7);
- the class census — the sizes n_i, hence the specific groups (§3, §8) [Open at the species inventory]; the identity-programme premises (invariant-exhaustion; No-Surplus-Identity) and the species/object distinction, at their citations;
- chirality, hypercharge, the embedding, breaking, occupancy, and the identification of class groups with physical force factors (§9) — the arc's interior;
- the non-abelian record machinery (§9);
- all imports at their slots: U1, S2, S3, S4, Q1, Q2 — the widest base in the corpus, every slot named.

The honest summary: starting from the corpus's own uniqueness theorems rather than against them, the paper proves that the one Fold cannot have an abelian theory of its identical closure sectors wherever it can exchange them — covariance forced outright, the swap forced under exchange realizability, non-commutativity arriving at the stated grain before any dynamics —

reduces the question of how much more is forced to one blending element in the two-sector atom, by floor, derives the product-of-blocks architecture of the forces conditionally on a named posture, and arms the largest kill-condition in the corpus against its own programme. The reserved arc is open, from the right door; its first theorem is on the board; its first decision is one element wide, by floor.

13. Conclusion

The forces beyond electromagnetism have always carried an air of arbitrariness: the Standard Model's SU(3) and SU(2) are chosen from an infinite catalogue because they fit, and the question *why these, why non-commutative at all* is conventionally answered by nothing. This paper relocates the question — and it does so from inside the corpus's strictest result. There is one Fold; the uniqueness theorems are not loosened here by a hair. But a unique Fold that admits persistent closure sectors can tie the same knot more than once, and the moment it does, the programme's oldest discipline takes hold: nothing physical may depend on which identical knot is which — and where the knots can be braided past each other, that prohibition, run through the substrate's transport structure, *is* a non-abelian symmetry. The strong and weak forces, on this account, are not evidence of hidden multiplicity at the foundations. They are the bookkeeping of indistinguishability within a foundation that is, and remains, one thing — what it costs a unique Fold to contain copies.

How much it costs is the open question, and the paper has reduced it to its smallest possible form. Two identical sectors support, at the forced minimum, phase — and, where exchange is realizable, swap, each minimum with its named slot; the complete classification of what two copies can carry leaves exactly nothing between the full floor and everything, and exactly one slot between the thinned floor and the same conclusion; and a single admissible blend — one transport that turns identical sectors partway into each other — collapses the question to its maximal answer, class by class, force by force. One element, most plausibly waiting in the reconstruction arc's continuous dynamics, and the Standard Model's silhouette — non-abelian blocks in product, abelian circles alongside, multiplets that are simply the species — stands as the forced shape of sector transport. The proven absence of that element would be the programme's largest internal defeat, and the paper names it as such, because a framework that cannot lose is not playing the game this corpus was built to play.

One Fold forced a circle, and the circle forced the first integer. The same one Fold, containing copies, forces at least the right to be confused about which copy is which — and that right has the algebra of forces. Whether the substrate exercises it fully is now a question with slots, a procedure, and a one-element answer — asked, this time, from the uniqueness the corpus fought to prove, instead of around it.