

Admissibility and the Structure of Physical Symmetry: From Groups to Groupoids

Flyxion
Independent Researcher

2026

Abstract

This essay does not attempt to derive the gauge group of the Standard Model, $SU(3) \times SU(2) \times U(1)$, from any more primitive structure, nor to claim that RSVP or any prior theory of admissibility predicts it. Its aim is narrower: to ask what a symmetry group is doing when it identifies physical states, and what additional structure becomes visible once the assumptions that make a single global group description sufficient are relaxed. Reading the quotient X/G as the removal of distinctions a symmetry structure declares physically inadmissible, we examine three independent ways in which this simple picture proves insufficient: locality, where a varying transformation parameter forces a connection and makes curvature a record of path-dependent transport; topology, where locally consistent transition functions can define a nontrivial cocycle class, so that local admissibility and global trivializability come apart; and vacuum selection, where a dynamically chosen state's stabilizer $H \subsetneq G$ refines admissible equivalence without destroying the underlying symmetry, and where the same broken directions appear as Goldstone modes or are absorbed into the gauge connection depending on whether the symmetry is global or gauged. Each mechanism is shown, via the action groupoid $G \ltimes X \rightrightarrows X$, to retain information a bare group suppresses: source and target, transport and path, and state-relative isotropy. The resulting claim is not that groupoids replace groups but that physical symmetry admits compression to a group description exactly to the extent that this information can be suppressed without losing physically relevant distinctions, which yields a single closing question applicable beyond this essay: which distinctions may be quotiented away without changing what a system can physically distinguish.

1 A Methodological Constraint

Before any of the framework's own vocabulary is introduced, one interpretive move must be ruled out explicitly, because it is the move readers will expect and the move this essay is not making. Group theory did not arise from RSVP, Distinguishability Geometry, or any of the admissibility apparatus developed elsewhere in this program, and no claim to the contrary is intended anywhere below. In particular, nothing here should be read as an assertion that $SU(3) \times SU(2) \times U(1)$ can be recovered from a more fundamental admissibility structure, that the specific representation content of the Standard Model is predicted or explained by that structure, or that RSVP competes with the Standard Model as a physical theory. Group theory is a mature, independently established mathematical subject, and gauge theory is an independently established physical theory built on it. What follows is an interpretive reading of that existing mathematics, not a derivation of it.

The distinction matters because the two moves are easy to conflate and have very different evidential burdens. A derivation claim would require producing the group and its representation content as a theorem of some prior structure; no such theorem is offered here, and none is claimed. An interpretive claim requires something much weaker and much more checkable: that a body of mathematics already in wide use admits a second, consistent reading in a different vocabulary, and that this second reading makes certain features of the mathematics look less like isolated technical facts and more like instances of a single pattern. The test of an interpretive claim of this kind is not whether it predicts anything the original mathematics did not already predict, but whether it organizes what was previously a catalogue of separate results (quotients, connections, curvature, symmetry breaking, obstruction theory) under one argument with a specific, falsifiable-in-form endpoint. That endpoint is stated now so it can be checked against every section that follows rather than assembled retrospectively from a list of loose correspondences:

A symmetry group is the compression of an admissibility groupoid obtained when source, target, transport, and state-relative dependence can all be forgotten without loss.

If this claim is right, it should be possible to specify, for ordinary global symmetry, what its object set, morphism set, composition law, and isotropy structure are, and then show that the transitions to gauge theory, to non-trivial bundle structure, and to dynamically selected vacua are each transi-

tions in which some part of this data, source-and-target dependence, a single global trivialization, or full isotropy, can no longer be forgotten without loss, while an ordinary group persists throughout as exactly the part that remains safely compressible. If the claim is wrong or empty, these transitions will instead look like unrelated changes of subject: locality, topology, and vacuum selection each forcing new mathematics for reasons that have nothing to do with what can and cannot be suppressed. The rest of this essay is organized to make that comparison available at each step rather than only at the end.

2 Quotients as the Removal of Inadmissible Distinctions

Take a group G acting on a space X , $(g, x) \mapsto gx$. The orbit of a point, $[x]_G = \{gx : g \in G\}$, is ordinarily introduced as a bookkeeping device: the set of points related by the symmetry, collapsed for the purpose of counting physically distinct configurations. We propose reading the orbit, and the quotient X/G built from it, as something with independent content rather than as a convenience. A quotient is the space that remains after distinctions the physical theory declares inadmissible have been removed. Two points x and gx are not identical as elements of X ; they are held to be admissibly the same relative to whatever structure G is required to preserve. The quotient does not erase a difference that was never there. It discards a difference that representation introduced and that no admissible operation can use to distinguish the underlying physical configuration.

This gives a precise, restricted role to a notion of admissible equivalence used elsewhere in this program: an equivalence class here is not asserted by fiat but is generated by an explicit, independently specified structure, the group action itself. Where earlier essays in this program treated admissible equivalence as a primitive relation to be characterized case by case, gauge symmetry supplies a case in which the relation is already characterized completely, by the action of G , and can be checked rather than argued for. An observable O on X respects the quotient exactly when $O(gx) = O(x)$ for all $g \in G$; the content of " x and gx are admissibly the same" is nothing more or less than the requirement that every physically meaningful function on X satisfy this identity. Nothing beyond the group action is invoked to license the equivalence, and nothing beyond ordinary invariance theory is required to state it. The interpretive move here is confined to naming what the quotient already does; the mathematics is unchanged.

What this section has not yet done is address locality, and doing so is where the essay's real argument begins. So long as a single $g \in G$ acts on the whole of X at once, the admissibility structure just described is global: one equivalence relation, generated by one group, applies uniformly everywhere. The next section takes up what happens when g is allowed to vary from point to point, because this is exactly where the composition structure that defines an ordinary group stops being available in the form the global case took for granted.

3 From Global Symmetry to Local Admissibility

The distinction between global and local symmetry need not be introduced as a reinterpretation imposed upon gauge theory from outside. It is already present within the historical vocabulary of gauge theory itself, whose local gauge principle originates with Weyl [2] and whose non-Abelian generalization is due to Yang and Mills [3], and which retains a terminology, dating to that earlier literature and preserved in Barnes' treatment [1], of *gauge theories of the first kind* and *gauge theories of the second kind*. In a gauge theory of the first kind, the transformation parameter is constant throughout space-time. In a gauge theory of the second kind, the corresponding parameter is permitted to vary with position. What changes between the two cases is therefore not merely the size of the symmetry group or the choice of representation, but the domain over which a transformation may be specified.

For the familiar $U(1)$ case, consider a matter field transformed by a phase whose parameter is initially constant:

$$\psi(x) \mapsto e^{i\epsilon} \psi(x), \quad \epsilon = \text{const.}$$

Because ϵ is constant, ordinary differentiation respects the transformation:

$$\partial_\mu \psi \mapsto e^{i\epsilon} \partial_\mu \psi.$$

The derivative therefore remains within the same transformation law as the field itself. No additional structure is required to compare the field at neighboring spacetime points.

The situation changes when the transformation becomes one of the second kind:

$$\psi(x) \mapsto e^{i\epsilon(x)} \psi(x).$$

Ordinary differentiation now produces an additional term proportional to the derivative of the transformation parameter,

$$\partial_\mu \left(e^{i\epsilon(x)} \psi(x) \right) = e^{i\epsilon(x)} \left(\partial_\mu \psi + i(\partial_\mu \epsilon) \psi \right).$$

The obstruction is elementary but conceptually decisive. A transformation that is admissible at each point separately does not by itself determine how representatives at neighboring points are to be compared. Local admissibility does not automatically supply admissible transport.

Gauge theory repairs precisely this failure. In the quantum electrodynamics Lagrangian discussed by Barnes [1], one introduces the covariant derivative

$$D_\mu = \partial_\mu + ieA_\mu,$$

with an accompanying transformation of A_μ chosen so that $D_\mu\psi$ transforms covariantly with ψ . The gauge potential is therefore not an arbitrary complication appended to an otherwise complete global symmetry. It is the additional structure required once the admissible transformation is allowed to vary locally.

This observation permits a translation into the language of admissibility without altering the physical construction. A global symmetry identifies a family of transformations that can be applied uniformly across the domain under consideration. A local symmetry assigns admissible transformations point by point. Once that localization is permitted, however, comparison between neighboring representatives requires a rule of transport, and the covariant derivative supplies that rule.

The distinction matters because it prevents "local symmetry" from being treated as merely a larger collection of global transformations. If G is the relevant symmetry group, a global action selects a single element $g \in G$ and applies it throughout the domain. A gauge transformation instead permits

$$g = g(x).$$

The transformation applicable at x need not be identical to the transformation applicable at y . Consequently, the physically relevant structure is no longer exhausted by asking which transformations exist. One must additionally ask where a transformation is defined, between which representatives comparison is licensed, and how such comparisons compose along a trajectory.

This is the point at which the distinction between groups and groupoids becomes substantive rather than terminological. A group supplies a globally composable algebra of transformations around a common identity structure. A groupoid [6, 7] retains identities and inverses while allowing morphisms to have distinct sources and targets, with composition defined only when those boundaries match:

$$x \xrightarrow{f} y, \quad y \xrightarrow{g} z, \quad g \circ f : x \rightarrow z.$$

The restriction is not a defect. It records information a single global group can suppress: which transformation is admissible from which state or local representative, and which transformations may actually be composed.

None of this means that ordinary gauge theory must be renamed “groupoid theory,” nor that the Standard Model gauge groups are somehow mistaken. The claim is more restricted. Once admissibility is indexed by location, state, chart, or trajectory, an ordinary group captures only part of the resulting structure. The group continues to characterize the internal transformations available locally, while a richer object is required to retain the domains, transports, and paths through which those transformations acquire physical comparison.

The gauge potential provides the first indication of this richer structure. Its curvature supplies the second. For electromagnetism,

$$F_{\mu\nu} \equiv \partial_\mu A_\nu - \partial_\nu A_\mu,$$

which in differential-form notation is simply $F = dA$, whereas for a non-Abelian gauge field the corresponding expression becomes

$$F = dA + A \wedge A.$$

The significance of F in the present interpretation is not merely that it is gauge invariant or that it encodes the physical field strength. Curvature measures the nontriviality of transport, in the sense made precise by the general theory of connections on fiber bundles [4, 5]; see also Nakahara [15] for a treatment closer to the present physical setting. If a representative is transported through locally admissible comparisons, the result of composing those transports can depend on the path by which the comparison was made. Infinitesimal transport around a closed loop need not return trivially to its starting representative, and curvature records precisely this failure.

This provides a direct connection to trajectory-first admissibility. Suppose admissibility were completely determined by endpoint membership in a common equivalence class. Then the history by which one representative was related to another could be discarded once the endpoints were known. Curvature demonstrates a setting in which this compression is generally unavailable: the composition of locally admissible transports retains information about the path, and two trajectories sharing endpoints need not therefore be equivalent as transports. Schematically,

local admissibility $\nRightarrow$ path-independent global identification.

Gauge theory consequently supplies a concrete physical instance of a principle that appears more generally in trajectory-first frameworks: admissibility

of each local step does not imply that a composite trajectory can be replaced without loss by a direct relation between its endpoints.

This also clarifies the relation between gauge symmetry and the tangent/normal distinction used elsewhere in this program. At a given representative, infinitesimal gauge transformations specify directions along which the description may change while remaining within the relevant equivalence structure; they behave as admissible tangent directions. A displacement that cannot be generated within that symmetry structure is of a different kind. The distinction is not simply between change and no change, but between change internal to an admissibility class and change that leaves it.

The passage from Barnes' first kind to second kind can therefore be read as a passage from global equivalence to locally indexed equivalence, together with the transport structure necessary to relate its local realizations. The historical terminology is useful precisely because it shows this transition to be internal to gauge theory rather than a distinction manufactured by the admissibility interpretation; it asks only what follows if a distinction gauge theory already recognizes is taken seriously as a statement about the conditions under which transformations may be composed.

From this vantage the guiding question is no longer why physical theories possess symmetry groups, but why, and under what conditions, a locally and trajectory-indexed structure of admissible transformations can be compressed into the action of a group without discarding physically relevant information. An ordinary symmetry group is then not rejected but situated: it is the globally composable limiting case of a more general architecture of admissible transformation.

The next question follows immediately. Local transport can be perfectly well defined on each region while nevertheless failing to admit a single globally trivial realization. Gauge theory therefore leads naturally from local admissibility to transition functions, bundles, holonomy, and obstruction, and it is at this point that the groupoid interpretation intersects most directly with the distinction, developed elsewhere in this program, between local consistency and global gluing.

4 Local Trivialization and the Limits of Gluing

Section 3 established that local admissibility requires a rule of transport, and that curvature measures the failure of such transports to compose trivially around a closed path. That analysis was carried out on a single patch, where one local representative and one connection sufficed throughout. The phys-

ically important case is different: a field configuration is typically described by a family of local representatives, one on each member of a cover $\{U_i\}$ of the space in question, with no guarantee in advance that these local descriptions assemble into a single global one. The question this section takes up is exactly that: given local representatives that are each individually admissible, does the family they form admit a single global representative, or only a coherent pattern of local ones?

On an overlap $U_i \cap U_j$, two local representatives of the same underlying configuration must be related by a transition function

$$g_{ij} : U_i \cap U_j \rightarrow G,$$

and where three patches overlap, consistency of these relations forces the cocycle condition

$$g_{ij} g_{jk} = g_{ik} \quad \text{on } U_i \cap U_j \cap U_k.$$

This condition is precisely what "local consistency" means made checkable: it says that comparing representatives around any triple overlap gives the same answer regardless of which pair of transitions is composed first. It is a purely local requirement, verifiable patch by patch, and it says nothing by itself about whether the patches can be reassembled into one representative valid everywhere.

That gap is the content of this section's central claim:

Local consistency does not imply global triviality.

A family of transition functions satisfying the cocycle condition on every overlap can still fail to be equivalent, by any consistent reassignment of local representatives, to the trivial family $g_{ij} \equiv e$. When no such reassignment exists, the bundle the transition functions define is called nontrivial, and no single global representative can stand in for the whole family without discarding information the local data carries.

Some care is needed here once G is non-Abelian, which is the case of physical interest. The relevant object is the non-Abelian Čech set $H^1(X, G)$ of transition cocycles modulo the equivalence generated by reassigning local trivializations; for non-Abelian G this is in general only a pointed set, not a group, so there is no general zero/nonzero dichotomy to invoke. What survives is a weaker but still exact statement. A cocycle $\{g_{ij}\}$ represents the distinguished, trivial class exactly when it is a coboundary of local trivializations, that is, when there exist local reassignments $\lambda_i : U_i \rightarrow G$ such that the transition functions can be absorbed into them on every overlap. When

no such reassignment exists, the cocycle represents a nontrivial class, and no single global representative can stand in for the whole family without discarding information the local data carries. Nothing in this essay requires developing this apparatus in general; what matters for the present argument is only what a nontrivial class means: not a failure of the local data, which remains perfectly consistent on every overlap, but an obstruction to compressing that data into a single global trivialization.

This is the point at which the essay's admissibility vocabulary and its own prior obstruction-theoretic work meet directly. The TARTAN/CLIO treatment of reconstruction has repeatedly used a notion of a "stall," a point at which local repair or reconstruction data fails to assemble into a coherent global account. The natural first reading of a stall is as a failure: reconstruction was attempted and did not succeed. Bundle theory supplies a case in which that reading is too hasty. A persistent, nontrivial obstruction of this kind is not evidence that the local description was inadequate or that further repair would eventually glue the patches together. It is evidence that no global trivialization exists, full stop, and that the object being described is genuinely, rather than provisionally, non-global. Sometimes the obstruction is not an unfinished reconstruction. It is the structure.

This sharpens rather than weakens the admissibility framework's account of stalls, because it supplies a case with an explicit, checkable invariant distinguishing the two readings. Where the relevant cocycle class can be computed, whether a stall reflects insufficient local data or an authentic global obstruction stops being a matter of interpretation and becomes a matter of calculation. Not every stall arising elsewhere in this program will come with so clean an invariant, but the bundle case shows that the distinction between "not yet reconstructed" and "not reconstructible as one object" is a real distinction with real content, not merely a hedge introduced to excuse an incomplete repair.

The connection to groupoids from Section 3 can now be stated more carefully than a first pass invites. The local admissibility groupoid built from all the local transports and their transition functions is well defined regardless of whether the bundle is trivial, and nothing about its existence presupposes a global trivialization. What a nontrivial bundle shows is not that this groupoid lacks some privileged single object representing the whole space; objects of a groupoid are not in general required to do any such thing, and the obstruction concerns global trivializability rather than the presence or absence of a distinguished object. Nor does nontriviality mean that a group is absent from the picture: a nontrivial principal bundle can perfectly well carry a gauge group, whose transformations are describable as

automorphisms of the bundle. What fails is narrower, and for that reason a stronger claim than "there is no group": the local transport architecture cannot be exhausted by a single globally uniform action on one globally chosen representative. The transition functions g_{ij} are exactly the data that the groupoid retains and that a global trivialization, were one to exist, would render superfluous.

This gives the essay a second, independent route by which global structure can fail to reduce to the action of one symmetry group, distinct from the one considered next. A bundle obstruction is a failure of local representatives to glue into a global one even though every local piece remains individually admissible. Spontaneous symmetry breaking, taken up in the following section, is a different kind of failure: not a failure to glue local data, but a failure of a particular global state to remain invariant under the whole of G . The two phenomena are often discussed as though they belonged to unrelated corners of gauge theory. Read through the present framework, they are two distinct ways in which the naive picture of a single group acting globally and uniformly turns out to be an idealization, one arising from topology and one from the dynamics that select a vacuum.

5 Symmetry Breaking as Constraint Refinement

The preceding sections considered two ways in which the simple picture of a single global symmetry group becomes insufficient. Locality required admissible transformations to be indexed by position and supplemented by a connection specifying their transport. Nontrivial bundle structure then showed that locally consistent representatives need not admit a single global trivialization. There is a third route by which the global picture becomes qualified, and it arises even before topology is considered: the dynamics may select a state that is invariant under less symmetry than the theory that admits it.

Return first to the global setting of Section 2. Let a group G act on a space of states X . Prior to the selection of any particular state, G supplies the relevant symmetry structure: if x and gx differ only by an action of G , the distinction between them is treated according to the equivalence structure generated by that action, and the quotient X/G records the resulting classes. The question now is what becomes of this structure once the dynamics singles out a particular vacuum $v \in X$.

The answer is determined by the stabilizer of v . Define

$$H = \text{Stab}_G(v) \equiv \{g \in G : gv = v\}.$$

The set H is a subgroup of G , containing precisely those transformations that leave the selected vacuum fixed. When $H = G$, the vacuum possesses the full symmetry of the theory. When

$$H \subsetneq G,$$

the vacuum is invariant under only part of that symmetry. This is the familiar setting of spontaneous symmetry breaking [1]: the governing equations or Lagrangian may retain the G symmetry even though a particular vacuum does not. The standard treatment then distinguishes the generators belonging to the unbroken subgroup H from those associated with the broken directions. The present argument changes none of this mathematics. It asks instead what this reduction from G to H means once symmetry is read as a structure of admissible distinctions.

The terminology of "symmetry breaking" can suggest that a symmetry formerly possessed by the theory has been destroyed. In spontaneous breaking this is misleading if taken literally: the symmetry of the underlying theory remains, and what has changed is the symmetry of the selected state. From the perspective developed here, the change is better described as a *refinement of admissible equivalence*. The selection of v supplies an additional constraint, and equivalence must now be evaluated relative not merely to the action of G but also to preservation of the selected vacuum.

Before this additional constraint is imposed, the action of G identifies the symmetry-related possibilities appropriate to the original description. After v has been selected, only $h \in H$ leaves that particular vacuum unchanged; for $g \in G \setminus H$, the transformed state gv is not the same selected vacuum but another point on its orbit. That orbit is naturally identified, under the usual hypotheses, with the homogeneous space G/H . So G/H parametrizes not "transformations that stop being admissible," but the distinct degenerate vacua to which the broken generators move the selected vacuum: it is a manifold of distinct symmetry-related vacuum selections, not a set of forbidden operations.

This formulation permits a useful refinement of the quotient interpretation introduced in Section 2, where quotienting was read as the removal of distinctions a symmetry structure declares physically inadmissible,

quotienting: remove distinctions under a specified equivalence structure;

vacuum selection performs a complementary operation,

constraint refinement: introduce a condition under which previously equivalent directions become

The second operation does not undo the first in any literal set-theoretic sense, nor does vacuum selection invalidate the original G symmetry of the theory. Rather, the additional constraint changes which distinctions are relevant relative to the state that has actually been selected. Equivalence is consequently indexed not only by the transformation structure but by the constraints relative to which that structure is being evaluated.

This is where the tangent language introduced in Section 3 becomes more than an analogy. Let $\mathfrak{g}$ and $\mathfrak{h}$ denote the Lie algebras of G and H . Infinitesimal generators in $\mathfrak{h}$ leave v fixed, whereas the broken generators in the complementary directions of $\mathfrak{g}/\mathfrak{h}$ move v along its orbit. Locally, the tangent space to the vacuum orbit is therefore represented by

$$T_v(G \cdot v) \simeq \mathfrak{g}/\mathfrak{h},$$

subject to the usual qualifications concerning the action and stabilizer.

These are precisely the directions that become important in Goldstone's theorem [8, 9]. For a spontaneously broken continuous global symmetry, fluctuations along the broken directions of the vacuum manifold give rise, under the standard assumptions, to massless Goldstone modes. In the present vocabulary, the tangent directions of the vacuum orbit acquire dynamical realization: what had been a direction generated by symmetry becomes, relative to the selected vacuum, a direction along which the field can physically fluctuate. This gives the tangent/normal distinction a concrete physical instance. Before vacuum selection, the G -orbit describes transformations internal to the symmetry structure. Once v is fixed as the reference vacuum, the stabilizer directions remain invisible at v itself, while the broken directions move the system through distinct points of the vacuum manifold. The relevant distinction is again not between motion and stasis, but between different classes of motion defined relative to a constraint: transformations preserving the selected state, and transformations carrying it elsewhere along its orbit.

The local gauge case modifies this conclusion in a decisive way. When the broken symmetry is global, the broken directions survive as physical Goldstone modes. When the symmetry is gauged, the same counting enters the Higgs mechanism [10, 11, 12, 13, 14] differently: the would-be Goldstone degrees of freedom do not remain as independent physical massless scalar modes but are incorporated into the gauge fields, supplying the additional degrees of freedom associated with massive vector bosons. This returns directly to the transport structure of Section 3. Local symmetry had already required a connection because independently chosen local repre-

representatives could not be compared by ordinary differentiation alone; spontaneous breaking of a gauge symmetry acts within that already localized architecture, so the broken directions become entangled with the gauge connection itself rather than appearing as new propagating fields in the manner of a broken global symmetry. In the conventional shorthand, the gauge field “eats” the Goldstone mode; in the present interpretation, the degrees of freedom associated with motion along the broken directions are transferred into the local transport structure rather than appearing as independent massless excitations.

This difference between the global and gauge cases matters for the admissibility interpretation because it prevents symmetry breaking from being reduced to a purely algebraic replacement $G \rightarrow H$. The same pair (G, H) can have different physical consequences depending on whether the relevant symmetry is global or local. The algebra identifies the broken directions, but the surrounding compositional architecture determines how those directions are physically realized: in a global theory they survive as Goldstone modes; in a gauge theory they participate in the Higgs mechanism and alter the degrees of freedom of the connection. The admissibility structure therefore depends not only on which transformations exist but on whether those transformations act globally or belong to a locally transported gauge structure.

There is a useful groupoid interpretation of this result, though it should not be overstated. Vacuum selection does not convert a group into a groupoid, and G/H is not itself generally a groupoid. Rather, selection changes the isotropy structure of the transformation system. At the chosen vacuum v , the isotropy group is precisely

$$\text{Aut}(v) \equiv \text{Stab}_G(v) = H.$$

Elements of G outside H instead generate arrows from v to other points of its orbit,

$$v \xrightarrow{g} gv.$$

In the action groupoid $G \ltimes X \rightrightarrows X$, whose arrows are pairs $(g, x) : x \rightarrow gx$ composing only when source and target match, $(h, gx) \circ (g, x) = (hg, x)$, spontaneous symmetry breaking can be read as the selection of an object whose isotropy group is a proper subgroup of the full transformation group. The distinction between unbroken and broken generators becomes a distinction between transformations that remain automorphisms of the selected object and transformations that carry that object to another object in its orbit.

This makes the phrase “constraint refinement” exact enough to be tested. The new constraint does not remove arrows from the action groupoid; the transformations generated by G remain present. What changes is their status relative to the selected object: some are endomorphisms of v , $v \xrightarrow{h} v$ for $h \in H$, while others relate v to distinct symmetry-related vacua, $v \xrightarrow{g} gv$ for $g \notin H$. The additional constraint supplied by vacuum selection refines a previously undifferentiated transformation structure according to source, target, and stabilizer, precisely the information that disappears if one records only the abstract group G without its action on the state space.

The result can be summarized by a principle that has appeared in less formal form elsewhere in this program:

Distinction is not absolute; it is induced relative to a constraint structure.

The group G alone does not determine which of its transformations preserve the physically selected state; that determination requires the pair (G, v) , or equivalently the action together with the constraint that singles out v . From this additional information emerges the stabilizer H , the orbit G/H , the distinction between broken and unbroken directions, and, depending on whether the symmetry is global or gauged, the Goldstone or Higgs realization of those directions.

The three preceding sections have therefore exhibited three different mechanisms by which the simplest global symmetry picture ceases to be sufficient. Locality showed that admissible transformations may vary from point to point and consequently require transport. Topology showed that mutually consistent local descriptions may nevertheless fail to admit a single global trivialization. Vacuum selection has now shown that even where a global symmetry G remains a symmetry of the theory, the dynamically selected state can possess only the smaller isotropy group H . These are not three instances of one mathematical failure, and the argument does not require pretending that they are. They are three independent reasons to distinguish the abstract existence of a symmetry group from the richer structure specifying where its transformations act, how they compose, what they transport, and which states they preserve.

Groups have not been displaced by any of this, and nothing in the argument questions the Standard Model’s use of them. What has been challenged is only the assumption that the abstract group by itself exhausts the structure of physical symmetry. Locality adds source and position dependence; topology adds nontrivial gluing data; vacuum selection adds state-relative isotropy. In each case, information suppressed by the bare group

becomes physically consequential. The question a closing section must answer is therefore not whether physical symmetry should be described by groups or by groupoids, but under what conditions the richer, locally and state-indexed structure of admissible transformation can be compressed to a group description without loss.

6 Conclusion: When a Group Is Enough

The argument began with a deliberately restricted question. It did not ask why the Standard Model possesses the particular gauge group $SU(3) \times SU(2) \times U(1)$, nor whether that group can be derived from RSVP or from a prior theory of admissibility. It asked instead what a symmetry group is doing when it identifies physical states, and what additional structure becomes visible when the assumptions that make a single global group description sufficient are progressively relaxed.

The first step was the quotient. For a group G acting on a space X , the passage

$$X \longrightarrow X/G$$

was interpreted as the removal of distinctions that the specified symmetry structure declares inadmissible. This did not introduce a new equivalence relation alongside the group action; it gave an admissibility reading to the equivalence relation the action already generates. Two representatives may remain mathematically distinct in X while belonging to the same physical equivalence class because the relevant observables are invariant under the transformations relating them. In this restricted sense, symmetry determines which representational differences survive as physical distinctions.

The remainder of the essay then examined three independent ways in which this simple global picture acquires additional structure.

The first was locality. Barnes' distinction between gauge transformations of the first and second kinds makes the transition explicit: once a transformation parameter is permitted to vary with position, ordinary differentiation no longer supplies a covariant comparison between neighboring representatives, and a connection becomes necessary. The covariant derivative therefore records something the abstract internal symmetry group alone does not specify: how locally chosen representatives are to be compared across the base space. Curvature adds a further fact — such comparisons need not compose trivially around closed trajectories, so the endpoint relation can fail to exhaust the information contained in transport.

The second was topology. Transition functions may satisfy every local compatibility condition while nevertheless defining a nontrivial cocycle class. Local admissibility and global trivializability are consequently different achievements. A failure to produce one global representative need not indicate unfinished reconstruction; it may be the mathematically correct expression of the object's global structure. The lesson for admissibility is correspondingly precise:

$$\text{local consistency} \not\Rightarrow \text{global trivializability}.$$

The information carried by overlaps and transition functions cannot always be eliminated by choosing a sufficiently clever global representative.

The third was vacuum selection. A theory invariant under G may dynamically select a vacuum v whose stabilizer $H = \text{Stab}_G(v)$ is a proper subgroup of G . The symmetry of the theory is not thereby destroyed; rather, the selected state introduces an additional constraint relative to which transformations acquire different statuses. Elements of H remain automorphisms of v , while elements outside H carry v to other points of its orbit. The homogeneous space G/H records those distinct degenerate vacua, its tangent directions supply the broken directions relevant to Goldstone's theorem, and in the gauged case those directions participate instead in the Higgs mechanism. Distinction is therefore not determined by the abstract group alone but by the group together with its action and the state relative to which invariance is evaluated.

These three mechanisms should not be collapsed into one. Locality is not topology, topology is not spontaneous symmetry breaking, and spontaneous symmetry breaking is not a failure of local transport. Their value for the present argument lies precisely in their independence. Each supplies a different example in which information absent from the bare specification of G becomes physically consequential: locality introduces position and transport; topology introduces gluing and obstruction; vacuum selection introduces state-relative isotropy. Together they show why the abstract symmetry group and the full architecture of physical symmetry should not be identified.

This also permits the thesis stated in Section 1 to be sharpened. A group is, mathematically, a one-object groupoid. A group action $G \curvearrowright X$ already determines the action groupoid $G \ltimes X \rightrightarrows X$ [6], whose arrows retain both the transformation and the object on which it acts. Gauge transport introduces further dependence on localization and path, while nontrivial gluing prevents this local information from always being compressed into a single

global trivialization. The contrast is therefore not properly expressed as a choice between groups and groupoids: groups remain present throughout. The issue is which indices may safely be forgotten.

The central claim of the essay can consequently be stated in its more precise form:

Physical symmetry admits compression to a group description to the extent that the source, target, localization, and trajectory dependence of admissible transformations can be suppressed without losing physically relevant distinctions. Groupoid structure retains this information when it cannot be suppressed.

The word *compression* is important. Passing to a group description need not constitute an approximation or an error. When the suppressed information is irrelevant to the question being asked, the group is exactly the right object. The argument is therefore not that groupoids are more fundamental merely because they retain more information; a representation that refuses every quotient and preserves every distinction is no more illuminating than one that erases distinctions indiscriminately. The relevant question is which distinctions can be removed without changing the physical content under consideration.

This returns the discussion to admissibility. An admissibility structure does not merely divide transformations into those that are permitted and those that are forbidden. It determines which differences matter, which comparisons are licensed, which compositions are defined, and which histories can be identified without loss. In the simplest global case these questions collapse together sufficiently that a group action provides the complete answer. Once locality, transport, topology, or state-relative symmetry enters, they separate: source and target matter, path can matter, choice of local representative can matter even when no individual representative is privileged, and stabilizer structure can matter even while the underlying group remains unchanged.

The resulting picture is therefore not one in which group theory is replaced by an external framework. It is almost the reverse. Gauge theory, bundle theory, and spontaneous symmetry breaking provide independently established mathematical cases in which several distinctions developed abstractly elsewhere in this program already possess exact realizations: equivalence versus identity, local consistency versus global reconstruction, endpoint agreement versus trajectory equivalence, and transformation type versus state-relative isotropy. The value of the comparison lies in those corre-

spondences being constrained by existing mathematics rather than manufactured by analogy.

That observation also establishes the limit of the present essay. Nothing shown here selects the Standard Model gauge group, derives its representations, predicts its coupling constants, or replaces its dynamics. The argument concerns the architecture in which such symmetries are represented. Any stronger claim would require additional mathematics and independent physical evidence. The methodological constraint stated at the beginning therefore survives intact at the end.

What has changed is the status of the initial proposal. The claim that “a symmetry group is the compression of an admissibility groupoid obtained when source, target, transport, and state-relative dependence can all be forgotten without loss” can now be understood not as a claim that physical groups should be discarded in favor of groupoids, but as a claim about the conditions under which information may be forgotten. The action groupoid makes source and target explicit; the connection makes transport explicit; curvature makes path dependence explicit; transition cocycles make gluing explicit; stabilizers make state-relative invariance explicit. A bare group suppresses these indices whenever the problem permits them to be suppressed.

The question with which the analysis closes is therefore also the criterion it proposes for subsequent applications of the admissibility framework:

Which distinctions may be quotiented away without changing what the system can physically distinguish?

When the answer includes source, target, location, and trajectory, a group may be enough. When those distinctions survive, the richer compositional structure is not additional decoration. It is the information that the quotient would have erased.

References

- [1] K. J. Barnes, *Group Theory for the Standard Model of Particle Physics and Beyond*, CRC Press, Taylor & Francis Group, Boca Raton, 2010.
- [2] H. Weyl, “Elektron und Gravitation. I,” *Zeitschrift für Physik* **56** (1929), 330–352.
- [3] C. N. Yang and R. L. Mills, “Conservation of Isotopic Spin and Isotopic Gauge Invariance,” *Physical Review* **96** (1954), 191–195.

- [4] C. Ehresmann, “Les connexions infinitésimales dans un espace fibré différentiable,” in *Colloque de Topologie (Espaces Fibrés)*, Bruxelles, 1950, Georges Thone, Liège; Masson et Cie, Paris, 1951, pp. 29–55.
- [5] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, Vol. I, Interscience Publishers, New York, 1963.
- [6] K. C. H. Mackenzie, *Lie Groupoids and Lie Algebroids in Differential Geometry*, London Mathematical Society Lecture Note Series, Vol. 124, Cambridge University Press, Cambridge, 1987.
- [7] R. Brown, *Topology and Groupoids*, BookSurge Publishing, Charleston, 2006.
- [8] J. Goldstone, “Field Theories with Superconductor Solutions,” *Il Nuovo Cimento* **19** (1961), 154–164.
- [9] J. Goldstone, A. Salam, and S. Weinberg, “Broken Symmetries,” *Physical Review* **127** (1962), 965–970.
- [10] P. W. Higgs, “Broken Symmetries, Massless Particles and Gauge Fields,” *Physics Letters* **12** (1964), 132–133.
- [11] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons,” *Physical Review Letters* **13** (1964), 508–509.
- [12] F. Englert and R. Brout, “Broken Symmetry and the Mass of Gauge Vector Mesons,” *Physical Review Letters* **13** (1964), 321–323.
- [13] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, “Global Conservation Laws and Massless Particles,” *Physical Review Letters* **13** (1964), 585–587.
- [14] T. W. B. Kibble, “Symmetry Breaking in Non-Abelian Gauge Theories,” *Physical Review* **155** (1967), 1554–1561.
- [15] M. Nakahara, *Geometry, Topology and Physics*, 2nd ed., Institute of Physics Publishing, Bristol, 2003.