Particles before symmetry

Henrique Gomes

March 5, 2026

Abstract

The Standard Model of particle physics is usually cast in *symmetry-first* terms. On this approach, one begins with a symmetry group and postulates matter fields as objects transforming under its representations, without requiring that the group be grounded in, or derived from, independent geometric structures. Recently, a *geometry-first* formulation has been proposed, in which the relevant symmetries are not fundamental. In this paper I extend this approach to two central mechanisms of the Standard Model: spontaneous symmetry breaking and the Yukawa coupling, both essential for particles to acquire mass. These reformulations offer alternative explanations cast in purely geometric terms. In particular, the quantisation of charge arises here as a purely geometric consequence of the tensorial construction of matter fields from the fundamental bundles—a mechanism that is both more general and more transparent than the usual topological account based on the compactness of symmetry groups. More generally, I argue that a symmetry-first account in terms of principal and associated bundles admits a genuine geometry-first counterpart only under certain strict conditions.

Contents

1 Introduction

Mathematically equivalent formulations of a theory can differ profoundly in what they suggest. Feynman made the point in his Nobel Lecture:

> Theories of the known, which are described by different physical ideas may be equivalent in all their predictions and are hence scientifically indistinguishable. However, they are not psychologically identical when trying to move from that base into the unknown.

The subsequent career of Feynman's path integrals—equivalent to Schwinger's operator methods, yet far more fertile—bears this out. So does Minkowski's 1908 recasting of special relativity: it changed no predictions, but supplied the geometric ontology without which general relativity could hardly have been conceived.Theory-reformulation is a live philosophical topic, with positions ranging from deflationary instrumentalism to inflationary fundamentalism ((Hunt, 2025)).[1] What I offer here should interest anyone who is not a thoroughgoing instrumentalist.

[1] According to Hunt (p. 5, ibid), instrumentalism treats reformulation as "merely a different choice of convention, no different in kind than a change in notation"; fundamentalism "proposes a metaphysical picture

The reformulation I propose applies to classical gauge theory as it bears on the Standard Model—less radical than Feynman's or Minkowski's, but nonetheless fruitful, as I will argue.[2] The reformulation does not begin with symmetry so I call it *geometry-first*.

A knowledgeable reader will protest that gauge theory is already geometrical: principal fibre bundles and connections are, after all, the geometer's stock-in-trade. True—but the usual formulation also builds symmetry into its foundations, and with it, an ontology that extends beyond the spaces where matter fields actually live.

Here is what I mean. In particle physics, matter fields are sections of vector bundles. The symmetry-first formulation adds principal fibre bundles on top, positing a structure group that coordinates the various matter sectors. I propose to dispense with this extra layer. If principal bundles appear at all in a geometry-first formulation, they must supervene entirely on the vector bundles—not the other way around.

This is not a minor shift. Representations of Lie groups, Killing forms, Casimir invariants, stabilisers, spontaneous symmetry breaking, gauge-fixing: these are the daily bread of the Standard Model. That the explicit appeal to symmetry at the ground of the explanatory chain might be dispensed with is anything but trivial.

But dispensed with it can be. A geometry-first formulation has recently been proposed in which symmetries are not postulated and principal fibre bundles are unnecessary ((Gomes, 2024, 2025a)). Symmetry groups become implicit: they arise as automorphism groups of fundamental vector bundles. The formulation is available only for gauge groups that are linear, and for representations built from the fundamental representation via tensor products, symmetrisation, and the like. It therefore demands an alignment between symmetry and geometry that the symmetry-first picture does not guarantee.

Even when the two formulations are mathematically equivalent, they come with different ontologies. For the Standard Model, the geometry-first picture posits three fundamental vector bundles over spacetime. Matter fields are sections of tensor products of these bundles. There is no separate space encoding the principal connection.

Change the formulation, and the explanations change with it. Four examples, developed below, show how.

First, in a non-Abelian vacuum Yang–Mills theory with Lie group G, the fundamental dynamical object is no longer a connection ϖ on a G-principal bundle (or its spacetime representative A^I_μ), but the covariant derivative ∇_μ on a vector bundle whose automorphism group is G. This remains true even when no matter fields are present: the affine structure can be dynamical all by itself. Quantum numbers become geometric labels—the internal space a particle inhabits, and its tensor type within that space. (These ideas were explained in (Gomes, 2024); this paper focuses on the next three.)

Second, once symmetry groups drop out of the ground level, 'symmetry-breaking' requires reinterpretation. Vector bosons are replaced by covariant derivatives of fundamental bundles,

similar to David Lewis's posits that some properties belong to an elite set of perfectly natural properties, with physics aiming to provide a partial inventory of these".

[2] It does not apply to all gauge theories: this will matter in Section 5.

which are not on the same footing as matter fields. How could they 'acquire mass' in the usual sense?

Third, Yukawa couplings. In the symmetry-first formulation, Yukawa terms are scalars formed from sections of different associated bundles, requiring explicit bridges between representation spaces—bridges that need not be unique. In the geometric picture, these scalars arise through inner products and contractions that are already determined by the bundle structures.

Fourth, the quantisation of charge. The symmetry-first picture attributes it to the topology of the compact group $U(1)$. The geometric picture derives it from the discrete algebraic structure of tensor powers of the fundamental bundles. This means that even for non-compact groups, where symmetry would permit continuity, geometry enforces discreteness.

Under Hunt (2025)'s lights, such a reformulation is non-trivial: it suggests a different ontology, offers new explanations, and displays an epistemic difference in how it solves problems.[3]

A word of caution about scope. The VB-POV applies to a proper subclass of gauge theories: those whose gauge groups arise as automorphism groups of structured vector spaces, with matter built tensorially from fundamental bundles. As a theoretical framework, this is genuinely narrower than the PFB-POV—it excludes models that the principal-bundle formulation can perfectly well accommodate. But the Standard Model lies within the overlap, and *within* that overlap the two formulations are empirically equivalent. What changes is which features of the physics appear primitive and which derived. Two kinds of claim should therefore be kept apart throughout: claims about *explanatory structure*—the Higgs and Yukawa reformulations of Sections 3–4—and claims about *admissible theory-space*—notably the slack diagnosis and the charge-quantisation discussion of Section 5. The former are reformulations in the strict sense; the latter amount to extra constraints that I argue physicists already tacitly impose.

Section 2 reviews both formulations. Sections 3 and 4 rework the Higgs mechanism and the Yukawa couplings from the geometry-first standpoint. Section 5 argues that the apparent narrowness of the geometry-first picture is its chief virtue. Section 6 draws lessons.

2 Symmetry-first and geometry-first formulations of gauge theory

I begin with the familiar symmetry-first formulation, then introduce the geometry-first alternative. A word on notation: ϖ always denotes the principal connection one-form on P; ∇ denotes a covariant derivative on a vector bundle E; and ω denotes the local connection form on E that appears upon choosing a frame, so that $\nabla = \mathrm{d} + \omega$ with $\omega \in \Gamma(T^*M \otimes \mathsf{End}(E))$. When I speak of the *affine structure* of a vector bundle, I mean the space of covariant derivatives compatible with the fibre structure (e.g. Hermitian metric): this space is an affine space modelled on $\Omega^1(\mathsf{End}(E))$, so that any two compatible covariant derivatives differ by an endomorphism-valued one-form.

[3] "Successful reformulations clarify what we need to know to solve problems, improving our understanding of the world" (ibid. p. 5); "Within [the] shared domain of problems [that they both solve], significant reformulations display an epistemic difference, while trivial reformulations do not" (ibid. p. 9).

2.1 Gauge theory and principal fibre bundles: the symmetry-first formulation

The symmetry-first formulation of the Standard Model is the familiar one. Each fundamental interaction is associated with a symmetry group, taken as the structure group of a principal fibre bundle. Connections on this bundle play the role of vector bosons—the force carriers.

Classical configurations of matter particles charged under a force are sections of vector bundles associated to the principal bundle whose group encodes that force. One may endow these associated bundles with additional structure—a Hermitian inner product on $\mathbb{C}^n$, say—in which case the representations of the structure group need only preserve that structure.

The connection on the principal bundle induces parallel transport on all associated bundles. Crucially, it is the *same* connection in each case: different matter fields charged under the same interaction march in step under parallel transport, probing the same distribution of electroweak or strong forces. Associated vector bundles are distinct entities, but the principal bundle ties them together, acting as their common coordinator (see (Weatherall, 2016) and Figure 1). The primacy of the postulated structure group is what makes this symmetry-first.

Any formulation of gauge theory that introduces symmetries via principal bundles falls under what I call the *principal bundle point of view* (PFB-POV). Technical details are in Appendix A. For now: a principal fibre bundle (P, M, G) is a smooth manifold P equipped with a smooth, free action of a Lie group G, projecting onto a base manifold M (spacetime). Such a bundle codifies the ways G can act on geometric objects over M. The most important such objects for this paper are vector bundles. A vector bundle (E, M, V) assigns to each spacetime point $x \in M$ a copy of a fixed vector space V—the typical fibre. Matter fields are sections: smooth assignments of an element of V to each point of M.

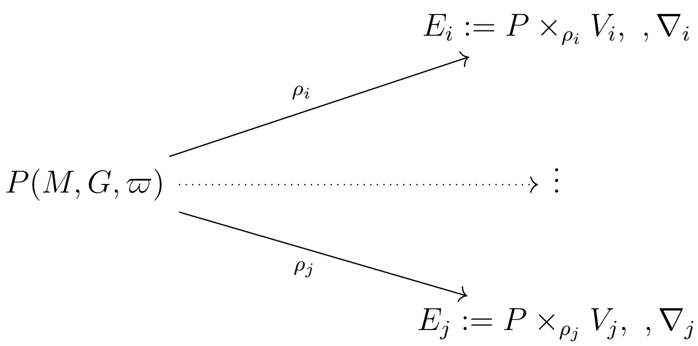

Figure 1: The principal G-bundle , with structure group G, over the manifold M, with a principal connection ϖ (a $\mathfrak{g}$-valued one-form on P), abbreviated by $P(M, G, \omega)$, and its associated vector bundles $E_i := P \times_{\rho_i} V_i$, where $\rho_i : G \to V_i$ is a representation of the Lie group—determined by a particle's quantum numbers—onto the vector space representing the typical fibre, V_i which is linearly isomorphic to $\pi_i^{-1}(x)$, for $x \in M$ and $\pi_i : E \to M$ the projection of the vector bundle onto its base space (spacetime). The covariant derivatives ∇_i are the ones induced by ϖ, as per Equation (A.6). See Appendix A for more details.

Generally, given a principal bundle (P, M, G) and a representation $\rho : G \to GL(V)$ on a typical fibre V, one constructs an associated vector bundle $E = P \times_\rho V$ (with ρ not necessarily faithful or surjective onto $\mathsf{Aut}(V)$[4]). In this framework, each such bundle corresponds to a particle type, its sections $\psi : M \to E$ are admissible field configurations, and the representation labels encode internal quantum numbers such as charge, colour, and weak isospin; for the Standard Model, the relevant gauge group is $SU(3) \times SU(2) \times U(1)$.

The resulting object is the associated vector bundle,

$$E := P \times_\rho V, \qquad (p, v) \sim (g \cdot p,\ \rho(g^{-1})v). \tag{2.1}$$

(See Equation (A.5) for more details.)

The advantage of this construction: fields in bundles associated to the same principal bundle manifestly covary under the symmetry group. (But, as we will see in Section 2.2, it is not the only route to covariance).

The simplest case is $G \simeq GL(V)$, where P is the full frame bundle of a vector bundle E. Enriching the fibre with further structure—a Hermitian metric, a complex volume form, or both—restricts the admissible frames and thereby selects a subgroup of $GL(V)$ as the structure group (see Appendix A for details).

But the notion of principal and associated vector bundles is not limited to these simple cases. G is not obliged to coincide with $\mathsf{Aut}(V)$; it need not be defined through an associated vector bundle. The definition in (2.1) is valid for any group and representation. The structure group of P and the automorphism group of the typical fibre need not be isomorphic. This will be important in Section 5.1.

A further question: do such vector bundles merely covary under the group action, or do they stand in canonical relation to one another? Suppose we are given:

$$E_1 = P \times_{\rho_1} V, \quad E_2 = P \times_{\rho_2} V \tag{2.2}$$

Given a local section of P, i.e. for $U \subset M$ a map $\sigma_U : U \to P$ such that $\pi(\sigma(x)) = x$, for all $x \in U$ (see Appendix A), we can write, for κ_1 a local section of E_1:

$$\kappa_1(x) = [\sigma(x), v(x)]_1, \quad v : U \to V. \tag{2.3}$$

Then the obvious map to consider is:[5]

$$\begin{gathered} T : E_1 \to E_2 \\ \kappa_1 := [\sigma(x), v(x)]_1 \mapsto [\sigma(x), v(x)]_2 =: \kappa_2. \end{gathered} \tag{2.4}$$

So the map acts as the identity on both entries, but nonetheless maps between sections in distinct vector bundles. However, on the right-hand side of (2.4), the representation under which we take equivalence classes is different: it is $\sim_2$ and not $\sim_1$. So is this map well-defined for arbitrary representations ρ_1, ρ_2? The map should be invariant under gauge transformations

[4]Here ρ is an embedding that may fail to be faithful or surjective onto $\mathsf{Aut}(V)$.

[5]I thank Jim Weatherall for suggesting this.

(cf. Eq (2.1)) on both the domain and image. So consider a different representative of the equivalence class on the domain; according to (2.4) we must have:

$$[g(x)\cdot\sigma(x), \rho_1(g^{-1}(x))v(x)]_1 \mapsto [g(x)\cdot\sigma(x), \rho_1(g^{-1}(x))v(x)]_2 \tag{2.5}$$

for any $g : U \to G$. But on E_2, we have the representation ρ_2, and so we must have (omitting dependence on $x \in M$ for clarity):

$$(\sigma, v) \sim_1 (g\cdot\sigma, \rho_1^{-1}(g)v) \sim_2 (\sigma, \rho_2(g)\rho_1^{-1}(g)v) \not\sim_2 (\sigma, v). \tag{2.6}$$

Where the last inequivalence holds iff $\rho_1(g)\rho_2^{-1}(g) \neq \mathbb{1}, \ \forall g$, i.e. the equivalence holds iff $\rho_1 \neq \rho_2$. Thus we find that for the map (2.4) to be well-defined, we must have $\rho_1 = \rho_2$.

Indeed, in physics, we are often faced with situations in which E_1 and E_2 have the same typical fibre, are associated to the same group, and yet have different representations. A simple example is when one of the representations is the trivial, or singleton, one, and the other is the fundamental (or any other).[6] This occurs many times in the Standard Model: for fermions to acquire mass, one must relate sections of bundles that have different representations, since they represent different particles.

In contrast, in the geometry-first picture, all vector bundles that in the symmetry-first formulation would be associated to the same principal bundle are already endowed with natural relations, as we now see.

A second key ingredient in the principal-bundle formalism is the principal connection ϖ. It determines how orbits of the group over neighbouring points of M are related, thereby specifying parallel transport—and hence covariant differentiation—in the associated vector bundles; in the case illustrated by the frame bundle, by determining which frame over one point is mapped to which frame at an adjacent point.

2.2 Gauge theory and vector bundles: the geometry-first formulation

The geometric perspective I develop here dispenses with the principal bundle altogether. This section sets out a formulation of gauge theory that proceeds without gauge potentials, principal bundles, or explicit appeal to gauge symmetries.

The analogy with spacetime clarifies my aim. Consider (M, g, Ξ_i): a smooth Lorentzian manifold (M, g) with various tensor fields Ξ_i on M, living in spaces constructed from the tangent bundle TM. The automorphism group of a typical fibre T_xM is $O(3,1)$ (or $SO(3,1)$ if orientation is background structure). This group becomes explicit once we introduce orthonormal frames. Yet we can say much about g and Ξ_i in a purely geometric, frame-independent manner, without ever mentioning $SO(3,1)$. And if we were to posit a different group acting on TM—$O(2)$, say, rather than $SO(3,1)$—we would need a geometric rationale.

In gauge theory, by contrast, an analogous "frame-free" formulation for the behaviour of matter remains largely undeveloped (cf. (Gomes, 2024; Weatherall, 2016)), and the very

[6] A slightly more sophisticated example is as follows. Let $G = U(1), V = \mathbb{C}^k$, and $\rho_i = n_i$, which acts as $e^{in_i}\mathbb{1}$ on $\mathbb{C}^k$. Then for $n_i \neq n_j$ for $i \neq j$ the map (2.4) is not well-defined, as can easily be verified.

idea of a geometric interpretation of the groups and their representations—for example, the adjoint action of $SU(2)$ on $\mathbb{C}^3$ endowed with an inner product, as opposed to the fundamental representation of $SU(3)$—is seldom raised. We are after a formulation of gauge theories for which these interpretations are transparent.

I call this realisation the *vector bundle point of view* (VB-POV).[7] To motivate it, recall that the main role of ϖ in (P, M, G) is to coordinate covariant derivatives between associated bundles. But what is the physical status of ϖ? Jacobs (2023, p. 41) argues convincingly that it has none:

> Neither the principal bundle nor the [principal] connection on its own represent anything physical. Rather, it is the induced connection on the associated bundle that represents the Yang-Mills field. [But] This approach has difficulties in accounting for distinct matter fields coupled to the same Yang-Mills field.

The issue, as he sees it, is that

> there is no independent Yang-Mills field that the associated bundle connections supervene on. This makes it seem somewhat mysterious that these connections are equivalent. The coordination between associated bundles begs for a 'common cause' in the form of an independently existing Yang-Mills field.[8]

I agree with Jacobs that this is a problem, and in (Gomes, 2024) I showed how to solve it: *Principal bundles are unnecessary if interacting particles are all sections of the same vector bundles or of their tensor products.* Tensor products inherit the same covariant derivative, so parallel transport automatically marches in step. The tensor structure itself provides the 'common cause' for coordination.

In more detail, given two vector bundles, E, E', a covariant derivative on E will induce a covariant derivative on E' whenever E' is equal to a general tensor product involving E and its algebraic dual, E^*. Given E a vector bundle with covariant derivative ∇, and E^* its dual, we define, for sections $\kappa \in \Gamma(E)$ and $\xi \in \Gamma(E^*)$:

$$d(\langle \xi, \kappa \rangle)(X) = \langle \nabla^*_X \xi, \kappa \rangle + \langle \xi, \nabla_X \kappa \rangle, \tag{2.7}$$

where here angle brackets represent contraction. The generalisation to arbitrary tensor products is straightforward due to multilinearity.

The idea, then, is to postulate a family of independent *fundamental* vector bundles, $E_1, \ldots, E_k$, upon which all further structure supervenes. Every field is a section of an appropriate tensor product of these fundamental bundles and their duals, i.e. elements of spaces such as

[7] Other theories could have geometry-first formulations—those based on Cartan geometry, for instance—but my focus is particle physics, where the relevant spaces are vector bundles.

[8] Jacobs instead defends the 'inflationary approach', which: "reifies not the principal bundle but the so-called 'bundle of connections'. The inflationary approach is preferable because it can explain the way in which distinct matter fields couple to the same Yang-Mills field." As I have argued in (Gomes, 2024), I don't believe it is preferable in this sense, but I won't rehash those arguments here.

$\Gamma(E_1 \otimes E_1 \otimes E_j \otimes E_k^*)$. Expressed in abstract-index notation, these fields—together with the corresponding covariant derivatives $\nabla_1, \ldots, \nabla_k$—furnish the entire dynamical content of a gauge theory expressible from the VB-POV. In short, the class of vector bundles reachable from a fundamental bundle is closed under finite direct sums, tensor products, duals, and (anti)symmetrised/exterior powers, with the connection induced functorially in each case.

The resulting picture, for a single fundamental vector bundle, is summarised in Figure 2.

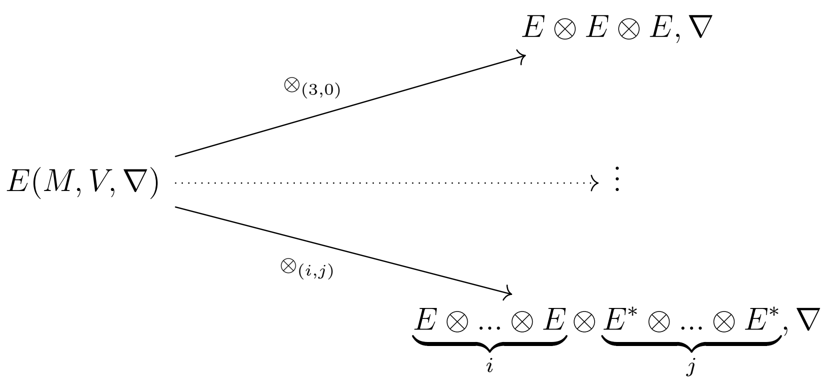

Figure 2: An example of the kinds of bundles that can be formed from a fundamental vector bundle. Each such tensor product—which could include arbitrary symmetrisations—inherits an affine structure from ∇, and each corresponds to a representation of $\mathsf{Aut}(V)$. In the VB-POV, particle types must arise through this kind of structure.

On this view, no gauge groups need be postulated at the ground level. The automorphism groups of the fundamental vector bundles, $\mathsf{Aut}(E_n) \subset \mathsf{End}(E_n)$, are already implied by their internal structure, and the overall gauge group is simply the product $\prod_n \mathsf{Aut}(E_n)$.[9] Should principal fibre bundles be invoked at all, they are entirely *supervenient* on the structure of the vector bundles, which form the subvenience basis. The familiar distinction between Abelian and non-Abelian theories then appears as a distinction between different types of automorphism groups. In particular, one-dimensional vector bundles, whose typical fibres are isomorphic to $\mathbb{C}$, generate Abelian automorphism groups.

Matters of gauge invariance are now also seen under a different lens. Any composite object constructed from the basic dynamical variables—the ∇_n and other tensor fields—will be tensorial, that is, covariant under the corresponding automorphism groups; and any scalar formed from such quantities will be invariant under those groups. In this respect, the description closely parallels that of classical general relativity in modern treatments employing abstract-index notation: such treatments scarcely mention "gauge invariance" or "coordinate invariance"; all they require is that covariance be secured at the ground level.[10]

[9] See (Bleecker, 1981, Ch. 7) for how to 'splice' together the principal bundles with different structure groups.

[10] Upon quantisation, as I argue in (Gomes, 2025b), superpositions of states may require relating objects across distinct classical possibilities. That, I contend, is where "gauge fixing"—or, more broadly, what I call *representational schemes*—enters. Gauge fixings become necessary when a fixed reference across classical states is needed. If such references are physical, they can, incidentally, be used to describe content in a gauge-invariant

This vantage point also reframes the earlier question of whether canonical maps can exist between distinct vector bundles. In the PFB-POV, the natural candidate—Equation (2.1)—is well defined only within a single representation. Matters look different here. Once covariance is secured at the ground level, all vector bundles charged under a given interaction are taken to ascend from a single *fundamental* bundle. In the cases to be explored below, each such fundamental bundle E_n has typical fibre $\mathbb{C}^n$ and is equipped with an inner product $\langle\cdot,\cdot\rangle$ and, where appropriate, a complex orientation (or volume form) ϵ. The various associated bundles then appear not as independently defined objects requiring ad hoc identifications, but as constructions from E_n itself. Their mutual relations are fully accounted for by the standard functorial machinery of geometry. For instance, to contract an element of E^n with one of $E^{n*}\wedge E^{n*}\otimes E^n$, we can use the interior product, which generally is a map:

$$\begin{aligned}\iota : E^n \otimes \bigwedge^m(E^{n*}) &\to \bigwedge^{m-1}(E^{n*})\\ (\xi,\Omega) &\mapsto \iota_\xi\Omega,\end{aligned} \tag{2.8}$$

where $\bigwedge$ is the anti-symmetric product, with $\Omega\in\bigwedge^m(E^{n*})$, and, for any $m-1$-tuple $(\xi_1,\cdots,\xi_{m-1})$ gives

$$\iota_\xi\Omega(\xi_1,\cdots,\xi_{m-1}) = \Omega(\xi,\xi_1,\cdots,\xi_m), \tag{2.9}$$

etc. Similarly, we could use the inner product between the two copies of E^n, and so on.

One might object that a parallel representation-theoretic argument could be mounted. Perhaps there are general ways to relate arbitrary representation spaces that mirror the geometric ones. That may be true, but it is beside the point. The virtue of the geometric route is that it trades purely on geometrical language, speaking directly to those trained in geometry rather than representation theory. The mere availability of a formulation that sidesteps algebraic machinery is already a win.

Still, at first pass the geometry-first picture may seem too narrow to capture the full menagerie of gauge theories. Some theories—those built from exceptional Lie groups—seem to fall outside its reach. And even when a gauge group G is given, it can be nontrivial to reverse-engineer a vector space for which $\mathsf{Aut}(E_x)\simeq G$.[11]

For all that, the Standard Model fits neatly within the geometry-first picture. In the standard PFB-POV every particle field is a section of an associated bundle for a principal bundle with structure group $SU(3)\times SU(2)\times U(1)$, and the representations that appear are tensor products of fundamental representations. The geometry-first alternative is available because the corresponding associated bundles can be constructed geometrically from fundamental vector bundles—via tensor and exterior products, (anti)symmetrisation, determinants, and the like. Namely, we fix once and for all the three fundamental vector bundles (cf. Gomes (2024)):

$$\big(E^3, M, \mathbb{C}^3, \langle\cdot,\cdot\rangle_3, \epsilon_3\big), \qquad \big(E^2, \mathbb{C}^2, \langle\cdot,\cdot\rangle_2, \epsilon_2\big), \qquad \big(E^1, \mathbb{C}, \langle\cdot,\cdot\rangle_1\big). \tag{2.10}$$

(or gauge-fixed) manner, in the traditional PFB-POV sense.

[11]The Peter–Weyl theorem guarantees that $U(n)$ admits nontrivial representations on $\mathbb{C}^m$, but extracting a natural structure on $\mathbb{C}^m$ that renders the action geometrically meaningful may not be straightforward.

I will refer to (2.10) as the *fundamental bundle data*. The gauge group in the VB-POV is then

$$G_{\mathrm{VB}} := \mathsf{Aut}(E^3, \langle\cdot,\cdot\rangle_3, \epsilon_3) \times \mathsf{Aut}(E^2, \langle\cdot,\cdot\rangle_2, \epsilon_2) \times \mathsf{Aut}(E^1, \langle\cdot,\cdot\rangle_1) \cong SU(3) \times SU(2) \times U(1), \tag{2.11}$$

where each factor is the structure-preserving automorphism group of the corresponding fibre. A covariant derivative *on a single* vector bundle suffices to encode *one* fundamental interaction.

With both formulations in hand, the natural next step is to put the VB-POV to work. I begin with the Higgs mechanism; the more familiar PFB-POV treatment can be found in any standard textbook (see e.g. (Bleecker, 1981, Ch. 10.3)).

3 The Higgs mechanism in the geometry-first formulation

In the standard presentation (see, e.g., (Tong, 2025, Ch. 2), (Bleecker, 1981, Ch. 10.3)), the Higgs mechanism is usually described in the language of spontaneous symmetry breaking. One is then compelled to invoke stabilisers, quotients of groups, Killing forms on Lie algebras, Goldstone's theorem, among other tools. In this section I outline an equivalent treatment, differing only in that it is phrased entirely in the language of vector bundles; the essential structure becomes visible without appeal to symmetry. It serves as a proof of principle that, in most applications, rerouting our demonstrations to avoid talk of symmetry does not consign us to intractably long alternatives (a common objection to some well-known reformulations; e.g. the Einstein algebra formulation of general relativity (Geroch, 1972)).

In Section 3.1, I begin with the geometric account of mass acquisition for vector bosons, where the relevant structure is the affine space of covariant derivatives. The Higgs field will here be treated as a fixed background section, without dynamics. In Section 3.2, I then incorporate the Higgs potential and its dynamics, to recover the mass acquisition mechanism for the Higgs itself, now formulated purely in vector-bundle geometry.

3.1 The Higgs field as a background

Let $(E^n, M, \mathbb{C}^n, \langle\cdot,\cdot\rangle_n)$ be a Hermitian vector bundle over a Lorentzian manifold (M, g), with fibres $E^n_x \simeq \mathbb{C}^n$ and $\langle\cdot,\cdot\rangle_n$ an inner product on E^n, induced by a Hermitian inner product in $\mathbb{C}^n$. In this Section we will consider the space of ∇_n, the covariant derivatives on E^n compatible with the Hermitian structure. We will omit the subscript n when it is understood from context, as it will be in this Section, so for now we take $\kappa \in \Gamma(E)$ and $\|\kappa(x)\|^2 := \langle\kappa(x), \kappa(x)\rangle$ (the generalisation to $\kappa \in \Gamma(E^i \otimes \cdots E^j)$ is straightforward). As usual, the norm of a tensor product factorises linearly, i.e. for $\lambda \in \Gamma(TM)$ and $\xi \in \Gamma(E)$, we have:

$$\|\lambda \otimes \xi\|^2 := g(\lambda, \lambda)\langle\xi, \xi\rangle. \tag{3.1}$$

Thus, when taking norms of tensors that are not spacetime scalars, g is used for the spacetime (abstract) indices, but I will leave it implicit, absorbing it into $\langle\, .\, ,\, .\, \rangle$.

As further background structure we will include φ: a vector-valued spacetime scalar field $\varphi \in \Gamma(E)$ with an everywhere non-zero norm, $\|\varphi(x)\| > 0$.[12] Picking out a constant $v > 0$ (later to be the minimum of the potential), we write $\|\varphi(x)\| = (H(x) + v)$, for $H \in C^\infty(M)$, and get

$$\varphi(x) = \|\varphi(x)\| e_0(x) = (H(x) + v) e_0(x), \tag{3.2}$$

where $e_0(x) = \frac{\varphi(x)}{\|\varphi(x)\|}$ is a unit section i.e. $\langle e_0(x), e_0(x)\rangle = 1$, $\forall x \in M$ which is well-defined since $\|\varphi(x)\| > 0$, $\forall x \in M$ (I will omit the dependence on $x \in M$ from now on, as it can be understood from context).

Now all we need to get the qualitative features of the Higgs mechanism for the bosons we can extract from the kinetic term for the Higgs field, i.e.

$$L_{\text{kin},\varphi}(\nabla) = \int_M \mathcal{L}_{\text{kin},\varphi} = \int_M d^4x \, \langle \nabla\varphi, \nabla\varphi\rangle, \tag{3.3}$$

where the caligraphic $\mathcal{L}$ will denote Lagrangian densities.

Here we are considering this as a kinetic term for the Higgs seen as a background structure, φ. In this first step, we only want to assess the mechanism for mass acquisition for the vector bosons, so we add to the Lagrangian the standard kinetic term for the affine structure as well (i.e. the trace of the curvature squared, $\mathcal{L}_\nabla = \mathsf{Tr}(\Omega \wedge *\Omega)$, cf. (A.10)).

Note that:

$$0 = \nabla\langle e_0, e_0\rangle = 2\mathsf{Re}\, \langle e_0, \nabla e_0\rangle, \quad \text{and} \quad \nabla v = 0, \tag{3.4}$$

where Re takes the real component.[13] Using (3.2) and (3.4) we get:

$$\langle \nabla\varphi, \nabla\varphi\rangle = \|\nabla\varphi\|^2 = (\partial H)^2 + (H + v)^2 \langle \nabla e_0, \nabla e_0\rangle, \tag{3.5}$$

where ∂ is the exterior derivative acting on scalars; i.e. it is the gradient.

Thus there are no cross-terms between ∇e_0 and e_0. Moreover, the covariant derivative—and therefore any connection representing it—appears quadratically in the term $(H+v)^2\langle \nabla e_0, \nabla e_0\rangle$. Such terms, when H is assumed small relative to v are to be interpreted as endowing mass to the affine structure (or the connections), as we will see.

First, note that ∇e_0 does not contain all the information in ∇: the components of ∇ that do not enter ∇e_0 are destined to remain 'massless'. From (3.4), only some of the components of ∇ that rotate e_0 into directions orthogonal to e_0 appear in the quadratic term, while those preserving e_0 drop out. In more detail, from (3.4), using $\mathsf{Re}\, \langle\, . \,, \, . \,\rangle$ (see footnote 13), we project ∇e_0 along e_0 and its orthogonal complement to find a geometric description of (3.5) as follows.

Viewing the unit section e_0 as defining a rank-one subbundle $L := \text{span}\{e_0\} \subset E$, we can interpret its orthogonal complement with respect to $\text{Re}\langle \cdot, \cdot\rangle$ $L^\perp$ as a normal bundle to L within

[12] This construction extends smoothly when φ is only almost everywhere non-vanishing: define $e_0 = \varphi/\|\varphi\|$ on the open set where $\varphi \neq 0$, and extend the relevant fields smoothly across the zero set under the usual regularity assumptions.

[13] The induced inner product that appears upon expansion of quadratic terms of the form $\|\xi\|^2$ is $\text{Re}\langle \cdot, \cdot\rangle$, which is effectively what appears in Lagrangians, due to the use of the complex conjugate terms, cf. (Hamilton, 2017, Ch. 8).

E. The covariant derivative ∇e_0 then defines a one-form with values in $L^\perp$; it probes how $L^\perp$ curves within L. Namely, what we can call a mixed *shape operator* for $L^\perp$:[14]

$$K_{e_0}(X) \;:=\; P^\perp_{e_0}\nabla_X e_0 \;=\; (\nabla_X e_0)^\perp = \;\nabla_X e_0 \tag{3.6}$$

where $P^\perp_{e_0} := \mathsf{Id} - \mathsf{Re}\,\langle\, e_0\,,\,.\,\rangle e_0$ is the orthogonal projection onto $L^\perp$ and $X \in TM$. From (3.4) we get

$$\langle \nabla e_0, \nabla e_0\rangle = \;\langle (\nabla e_0)^\perp, (\nabla e_0)^\perp\rangle \;=\; \langle K_{e_0}, K_{e_0}\rangle = \|K_{e_0}\|^2, \tag{3.7}$$

i.e. the norm squared of the second fundamental form, or extrinsic curvature, defined by e_0.

The operator ∇ merely provides the affine structure of the bundle, while the Higgs field singles out an internal direction e_0. From the kinetic term, assuming H small compared to v, we obtain:

$$\langle \nabla\varphi, \nabla\varphi\rangle \approx v^2\|K_{e_0}\|^2. \tag{3.8}$$

That is, the affine structure 'acquires mass' precisely in the directions orthogonal to the Higgs; the mass is proportional to the vacuum expectation value v and to the extrinsic curvature of the corresponding orthogonal sub-bundle. Note that Equation (3.8) is expressed entirely in abstract-index (tensorial) language, with no appeal to symmetries or their breaking. We are, of course, 'breaking the symmetry' by postulating a background field φ. But since any two non-zero Higgs fields define structurally equivalent directions—without further background structures we cannot discern between two such choices—it is the non-zero norm v that is doing the real work.

This concludes the classical, qualitative, account of 'mass acquisition' for the affine structure using the Higgs as a fixed background structure. Two points of contact with the standard formalism should be noted. First, 'mass acquisition' is, at this stage, a statement about the quadratic form (3.8), not yet about gauge bosons in the particle-physics sense. The identification of specific massive and massless components requires writing $\nabla = \mathrm{d} + \omega$ in a frame adapted to e_0; this is carried out in Appendix B. Second, the decomposition above is local: it presupposes that φ is nowhere vanishing. On nontrivial backgrounds the global picture requires patching, but the geometric content of (3.8) is unaffected.

3.2 Dynamical Mass Generation for the Higgs in the Linearised Theory

Above, a background Higgs field sufficed to get the qualitative features of the mechanism; and this we got in (3.8). Now we would like to make the mass generation mechanism more precise in the linearised regime, and include the mechanism required for the Higgs to acquire mass as well.

If in the previous Section the background structure was $(E, M, \mathbb{C}^n, \langle\cdot,\cdot\rangle, \varphi)$, here it will be just $(E, M, \mathbb{C}^n, \langle\cdot,\cdot\rangle)$. And on this background, we want to describe the space of physically

[14] In differential geometry, given a hypersurface $S \subset M$ with unit normal n, the standard shape operator for S is defined as a tensorial operator $K_n : TS \to TS$, namely $K = \nabla n$. Here our operator will mix indices, $K_n : TM \to L^\perp$, but it can still be thought of as an operator describing the shape of $L^\perp$ within L as one goes around M.

distinct, or non-isomorphic configurations of the Higgs field. That is, we want to consider the different *physical* possibilities for a single vector field in $\Gamma(E)$.

We assume E admits at least one unit section and pick one, $e_0(x)$. Again, two different choices e_0, e_0' will be isomorphic: they are structurally indistinguishable within (their copies of) E, since they have the same (unit) norm (and there is no more background structure on E). Thus, in the usual terminology, any two such choices can be mapped to each other by an element of $\mathsf{Aut}(E)$. So, in order to parametrise the space of non-isomorphic unit sections, we can fix the same e_0 across configurations. In the jargon of (Gomes, 2025b; Gomes & Butterfield, 2023; Kabel et al., 2025), we take e_0 to define a 'representational scheme' across physical possibilities, so that each configuration $\varphi \in \Gamma(E)$ is parallel to the same e_0. This is this demonstration's analogue of choosing 'unitary gauge', and it is the only place where symmetry could be mentioned.[15]

Thus the configuration space of non-isomorphic field configurations can be parametrised by scalar functions as follows:

$$\Gamma_P(E) := \{\varphi = (H+v)e_0,\ H \in C^\infty(M)\}. \tag{3.9}$$

(Here v could be absorbed by H, but for what follows it is useful to keep them separate.) So variations of a single section φ would amount to a variation of $H(x)$ as the dynamical variable. That is, according to (3.9), H parametrises the tangent space to $\Gamma_P(E)$: a one-parameter family of configurations $\gamma(t) \in \Gamma_P(E)$ such that $\gamma(0) = w_o \in \Gamma_P(E)$ is written as $\gamma(t) = (Ht+v)e_0$, for some H, so $\frac{d}{dt}_{t=0}\gamma(t) = He_0$. So we write

$$T_{w_o}\Gamma_P(E) = \{He_0,\ H \in C^\infty(M)\}. \tag{3.10}$$

We thus would now like to see H as parametrising the tangent directions to a certain functional space characterising the physically distinct configurations of the Higgs field. But instead of thinking in terms of tangent bundles to configuration spaces, it is more practical to use the standard assumptions of 'low energy', or, equivalently, going into the linearised regime and using expansions in terms of order parameters, i.e. $\mathcal{O}(H) = \mathcal{O}(\nabla) = \mathcal{O}(K_{e_0}) = \varepsilon$.[16] Using $\varphi = (H+v)e_0$ to rewrite the kinetic term (3.5) we once again get

$$\|\nabla\varphi\|^2 = (\partial H)^2 + (H^2 + 2Hv + v^2)\,\|K_{e_0}\|^2, \tag{3.11}$$

and obtain

$$\|\nabla\varphi\|^2 = (\partial H)^2 + v^2\|K_{e_0}\|^2 + \mathcal{O}(\varepsilon^3), \tag{3.12}$$

[15] Indeed, representational schemes can be compared to gauge-fixings (cf. (Gomes, 2025b, Sec. 3.3)), as follows. Consider $\Gamma(E)$, and its sector $\Gamma_0(E)$. Let $\varphi, \varphi' \in \Gamma_0(E)$. The group $\mathsf{Aut}(E)$ acts transitively on the unit normal sections: it can take any internal direction into any other. Therefore, we could, by a suitable gauge transformation on φ, make it collinear with φ'. Once they are collinear, it is a trivial matter to separate out the part that has a given norm, be it e_0, $v'e_0$, or ve_0 (see below, for the definition of v).

[16] In the comparative sense: that $\frac{|H|}{v} \sim \varepsilon << 1$, and *mutatis mutandis* for the appropriate norm on ∇. The pointwise norm on the space of covariant derivatives can be defined in two equivalent ways: either by seeing it as an affine space over the space of connections $\omega \in \Gamma(T^*M \otimes \mathsf{End}(E))$, which requires us to pick a flat covariant derivative ∇_o and then take the norm of $\omega = \nabla - \nabla_o$; or, alternatively, for $v \in T_xM$ of unit g-norm, and $s \in E_x$ of unit $\langle\cdot,\cdot\rangle$-norm, one can define $\|\nabla\|_{\text{op}} := \sup_{v,s}\|\nabla_v s\|$.

as before.

The kinetic term for the Higgs is given by $(\partial H)^2$ whereas v^2 is the mass term for part of the affine structure, as we saw. Now let us turn to the mass of the Higgs, for which we need the Higgs potential.

With the introduction of the potential term, the Lagrangian that endows mass to the affine structure and to the Higgs reads

$$L(\nabla, \varphi) = \int_M \mathrm{d}^4 x \left(\langle \nabla\varphi, \nabla\varphi\rangle + V(\|\varphi\|)\right) + \int_M \mathsf{Tr}(\Omega \wedge *\Omega). \tag{3.13}$$

And we compute the minimum of the potential:

$$\delta V(w_0) = 0, \tag{3.14}$$

where here $V : \Gamma(E) \to \mathbb{R}$ is assumed to be homogeneous on M and depend only on $\|\varphi\|$, and so $w_o(x) = v\, e_0(x) \in \Gamma_0(E)$ with $\|w_o(x)\| = v \neq 0$, $\forall x \in M$. So any potential whose minimum is non-zero suffices for identifying $w_o(x) = v\, e_0(x)$ for some $v > 0$. The familiar 'Mexi [17]

In order to get the mass of the Higgs we must similarly expand $V(\varphi)$ up to $\mathcal{O}(\varepsilon)^2$, thus involving the second variation of V around its minimum, called *the Hessian*. The Hessian is a symmetric linear map that acts fibre-wise; using the notation $\delta\varphi$ for elements of $T_{w_o}\Gamma_0(E)$ we write it as:

$$\mathsf{Hess}(V)_{w_o} := \frac{\delta^2 V}{\delta\varphi^1 \delta\varphi^2}_{|\varphi = w_o} : T_{w_o}\Gamma(E) \to T_{w_o}\Gamma(E), \tag{3.15}$$

where we are taking the minimum to be at w_o, with $\delta\varphi^1, \delta\varphi^2 \in T_{w_o}\Gamma_0(E)$. Thus we have:

$$V(\varphi)_{|\varphi = w_o} = \left(V(\varphi) + \frac{\delta V}{\delta\varphi}\delta\varphi + \langle \delta\varphi_1, \frac{\delta^2 V}{\delta\varphi^1\delta\varphi^2}\delta\varphi_2\rangle + ...\right)_{|\varphi = w_o}. \tag{3.16}$$

The mass terms will be obtained from (3.16) by diagonalising $\mathsf{Hess}(V)_{w_o}$. Thus, as in the standard account, the mass terms for different components of the Higgs along the different directions within $T_{w_0}\Gamma_0(E)$ are given by the eigenvalues of the Hessian with those directions as eigenvectors. But, per (3.10), in the current case we are considering only $\delta\varphi = H(x)e_0 \in T_{w_o}\Gamma_0(E)$: only e_0 can have non-vanishing eigenvectors (the eigenvalues of any completion of e_0 into an orthogonal basis for $\mathsf{Re}\langle\, .\, ,\, .\, \rangle$ (see footnote 13) would vanish for $e_I, I \neq 0$). The value of the surviving eigenvalue along e_0 will depend on the form of $V(\|\varphi\|)$, which is assumed to be translationally invariant (or homogeneous on M). Here we will just call this eigenvalue m_H^2.[18]

[17] Usually, one takes the Higgs vacuum vector w_o to break the symmetry, and then one expands the Higgs fields around this vacuum. But then one must find a way to get rid of the 'Goldstone modes': directions of the Higgs field that are parallel to the gauge orbit of the vacuum. This is where one invokes unitary gauge. But this is awkward: angles are preserved by gauge transformations, so how could a gauge transformation act to eliminate modes that are defined by their inner product with respect to the Higgs vacuum, which is assumed to be a physical object?

[18] In the standard case where $V = -\mu\|\varphi\|^2 + \lambda\|\varphi\|^4$, using (3.9) and (3.10) we can replace $\|\varphi\|^2 \to (H + v)^2$, and $\frac{\delta}{\delta\varphi} \to \frac{\partial}{\partial H}$. So $\mathsf{Hess}(V) \to \frac{\partial^2 V}{\partial H^2}$. Then an easy computation reveals:

$$\frac{\partial V}{\partial H} = -2\mu(H + v) + 4\lambda(H + v)^3 \tag{3.17}$$

Since $\frac{\delta V}{\delta \varphi}(w_o) = 0$, diagonalising the Hessian, so that:

$$\langle \delta\varphi_1, \frac{\delta^2 V}{\delta\varphi^1 \delta\varphi^2} \delta\varphi_2 \rangle = \langle He_0, m_H^2 He_0 \rangle = m_H^2 H^2, \tag{3.19}$$

we obtain from (3.16), up to quadratic order in ε:

$$V(\varphi)_{|\varphi = w_o} = V(w_o) + \frac{1}{2} m_H^2 H^2 + \mathcal{O}(\varepsilon)^3. \tag{3.20}$$

This concludes the geometric gloss on the Higgs mechanism. No *explicit* mention of stabilisers, decomposition of gauge orbits, Killing forms on the Lie algebra, Goldstone's theorem, etc was required—devices indispensable in the standard formulation (see (Hamilton, 2017, Ch. 8.1); (Tong, 2025, Ch. 2.2); (Bleecker, 1981, Ch. 10.3) for comparison).[19] This is just what you would expect from the geometry alone.

In Appendix B I show how from this point it is a trivial matter to reproduce standard results from the familiar or standard approach to gauge theory for the mass acquisition of bosons. At this late stage in the proof, the missing ingredient for the comparison is simply to write the connection ω in terms of preferred representations of the Lie algebras in question.

4 The Yukawa mechanism

The Higgs mechanism endows mass to gauge potentials; the *Yukawa mechanism* endows mass to matter fields.

In the Standard Model, fermion masses cannot be introduced as for scalar fields. A Dirac mass term must couple left- and right-handed chiral fermions, but these transform in inequivalent representations of $G = SU(3) \times SU(2) \times U(1)$. Coupling them directly would violate gauge invariance—the same obstruction we encountered in Section 2.1 concerning canonical maps between associated bundles. The solution is to introduce the Higgs field in such a way that gauge invariance is preserved while fermions acquire effective masses.

I follow (Hamilton, 2017, Ch. 8), whose notation is already geometric. Section 4.1 presents the standard treatment; Section 4.2 argues that the geometry-first picture does better.

4.1 The standard presentation

Fermions are spinors, and for Weyl spinors the inner product is anti-diagonal in chirality: $\overline{\psi}_R \psi_R = 0$. Mass terms must therefore couple left to right: $\overline{\psi}_R \psi_L$. If both chiralities inhabit the same vector bundle and transform in the same representation, one may write

$$\mathcal{L}_{\text{mass}} = -m\overline{\psi}\psi = -m\text{Re}(\overline{\psi}_L \psi_R) \tag{4.1}$$

---

and

$$\frac{\partial^2 V}{\partial H^2} = -2\mu + 12\lambda(H+v)^2, \tag{3.18}$$

Setting $\frac{\partial V}{\partial H} = 0$ and $H = 0$ in (3.17) we obtain $v = \sqrt{\frac{\mu}{2\lambda}}$. Setting $v = \sqrt{\frac{\mu}{2\lambda}}$ and $H = 0$ in (3.18) we obtain $\frac{\partial^2 V}{\partial H^2}(w_o) = 4\mu = m_H^2$.

[19] For instance, the disappearance of Goldstone modes through a choice of gauge is replaced here by having all $\varphi \in \Gamma_0(E)$ parallel to e_0 and the orthogonality relation (3.4).

and this is gauge-invariant. Locally, $\psi_L \in \Gamma(S_L \otimes E)$ and $\psi_R \in \Gamma(S_R \otimes E)$, where (E, M, V) is the vector bundle with representation space V, and S_L is the bundle of left-handed spacetime spinors with typical fibre Δ_L (mutatis mutandis for right-handed).

In the Standard Model, however, fermions are both twisted and chiral: left- and right-handed components transform in inequivalent representations of the gauge group. For instance,

$$e_L \in (\mathbf{1}, \mathbf{2}, -1), \qquad e_R \in (\mathbf{1}, \mathbf{1}, -2).$$

These internal vector bundles are representationally inequivalent; e.g. $\psi_L \in \Gamma(S_L \otimes E_L)$ and $\psi_R \in \Gamma(S_R \otimes E_R)$, have different representation spaces, $V_L \not\cong V_R$. Thus a bilinear such as $\bar{e}_L e_R$ is not gauge-invariant, and a bare mass term as in (4.1) is forbidden. (Table 1, reproduced from (Hamilton, 2017, Table 8.2), shows the representations of $SU(2)_L \times U(1)_Y$ for the fermions and the Higgs in the Standard Model.)

Sector	$SU(2)_L \times U(1)_Y$ **rep.**	**Basis vectors**	**Particle**	T_3	Y	Q
Q_L	$\mathbb{C}^2 \otimes \mathbb{C}_{1/3}$	$\binom{1}{0}$	$\mathbf{u}_L$	$+\frac{1}{2}$	$+\frac{1}{3}$	$+\frac{2}{3}$
		$\binom{0}{1}$	$\mathbf{d}_L$	$-\frac{1}{2}$	$+\frac{1}{3}$	$-\frac{1}{3}$
Q_R	$\mathbb{C} \otimes \mathbb{C}_{4/3}$	1	$\mathbf{u}_R$	0	$+\frac{4}{3}$	$+\frac{2}{3}$
	$\mathbb{C} \otimes \mathbb{C}_{-2/3}$	1	$\mathbf{d}_R$	0	$-\frac{2}{3}$	$-\frac{1}{3}$
L_L	$\mathbb{C}^2 \otimes \mathbb{C}_{-1}$	$\binom{1}{0}$	ν_{eL}	$+\frac{1}{2}$	-1	0
		$\binom{0}{1}$	e_L	$-\frac{1}{2}$	-1	-1
L_R	$\mathbb{C} \otimes \mathbb{C}_{-2}$	1	e_R	0	-2	-1
Higgs φ	$\mathbb{C}^2 \otimes \mathbb{C}_1$	$\binom{1}{0}$	φ^+	$+\frac{1}{2}$	$+1$	$+1$
		$\binom{0}{1}$	φ^0	$-\frac{1}{2}$	$+1$	0
Higgs$_\perp$ φ_c	$\mathbb{C}^2 \otimes \mathbb{C}_{-1}$	$\binom{1}{0}$	$\overline{\varphi}^0$	$+\frac{1}{2}$	-1	0
		$\binom{0}{1}$	-$\overline{\varphi}^+$	$-\frac{1}{2}$	-1	-1

Table 1: First-generation fermion representations under $SU(2)_L \times U(1)_Y$, together with the Higgs doublet and its conjugate. Here boldface on the quarks means each such term is a vector in $\mathbb{C}^3$. (φ^0, φ^+) as well as the left-handed particles are doublets: they can be rotated into each other by an $SU(2)$ transformation. Y is the hypercharge, and T_3 is weak isospin. Here $Q = T_3 + \frac{1}{2}Y$.

Moreover, for V_R, V_L irreducible, unitary, non-isomorphic representations of G, mass pairings, defined as G-invariant maps, $\zeta : V_L \times V_R \to \mathbb{C}$, complex antilinear in the first variable and complex linear in the second (so that they form mass terms), are necessarily trivial (see (Hamilton, 2017, Theorem 7.6.11)).

The remedy is a *Yukawa form*, defined as follows. Let V_L, V_R, W be representation spaces for $G = SU(3) \times SU(2) \times U(1)_Y$. A Yukawa form is a G-invariant trilinear map

$$\tau : V_L \otimes W \otimes V_R \longrightarrow \mathbb{C},$$

antilinear in V_L, real linear in W, linear in V_R.

For the leptons, the relevant $SU(2) \times U(1)$ representations are (from Table 1):

$$V_L = \mathbb{C}^2 \overset{\rho_L}{=} \mathbf{2}_{-1}, \tag{4.2}$$

$$V_R = \mathbb{C} \overset{\rho_R}{=} \mathbf{1}_{-2}, \tag{4.3}$$

$$W = \mathbb{C}^2 \overset{\rho_W}{=} \mathbf{2}_{1}. \tag{4.4}$$

Then, for $l_L : U \to V_L, \phi : U \to W, l_R : U \to V_R$, it is standard to define the Yukawa form as:

$$\tau : V_L \times W \times V_R \longrightarrow \mathbb{C}, \tag{4.5}$$

$$(l_L, \phi, l_R) \longmapsto l_L^\dagger \phi \, l_R, \tag{4.6}$$

which is $SU(2) \times U(1)$ invariant by construction.

The map τ is defined on vector spaces and depends on the choice of trivialisation. To render it invariant, we extend to sections of the associated bundles. Given a section $\sigma(x)$ of an $SU(2) \times U(1)$ principal bundle (cf. (2.3)), the local maps l_L, φ, l_R define global sections $e_L \in \Gamma(S_L \otimes E_L)$, $\varphi \in \Gamma(F)$, and $e_R \in \Gamma(S_R \otimes E_R)$, independent of trivialisation. A left-handed electron is

$$e_L = \psi_L \otimes [\sigma, l_L], \tag{4.7}$$

where $\psi_L \in \Gamma(S_L)$ is a left-handed Weyl spinor, and $\lambda_L := [\sigma, l_L] \in \Gamma(E_L)$, with E_L the vector bundle with typical fibre V_L as in (4.2). Analogous expressions hold, *mutatis mutandis*, for the right-handed field e_R and for the scalar $\varphi = [\sigma, \phi]$.

Since τ is invariant under $SU(2) \times U(1)$, we can define the gauge-invariant map

$$T(e_L, \varphi, e_R) := \tau(l_L, \phi, l_R) = l_L^\dagger \phi \, l_R. \tag{4.8}$$

This construction yields a singlet representation for a spacetime scalar.

There are, however, other invariant maps that also produce a gauge-invariant scalar, and it is not immediately clear how to choose among them. For instance, let H denote the centraliser of G—the subgroup of elements commuting with all $g \in G$. Then for any such map producing a scalar singlet, such as (4.8), one can define a family

$$T'_h(e_L, \varphi, e_R) := \tau'_h(l_L, \phi, l_R) := l_L^\dagger \phi \, h \, l_R, \tag{4.9}$$

with $h \in H$. For $G = SU(2) \times U(1)$, the centraliser is simply $\mathbb{C}$, acting as multiples of the identity. Thus, the ambiguity here amounts to an arbitrary complex factor, which in practice could be absorbed into the Yukawa couplings. For more complicated groups, however, such ambiguities may not be so easily removed. The crude impression is that in the PFB-POV one manufactures invariants in a chosen basis while absorbing ambiguities into coupling constants.

4.2 The geometry-first presentation

In Section 2.1 I argued that there is no canonical map between associated bundles with different representations. Of course we do not *need* such a map to extract scalars—T in (4.8) suffices for comparison with experiment. But the answer is unsatisfying: why this particular map? Could we have found others?

The geometry-first perspective eliminates such ambiguities by treating Yukawa couplings not as equivariant maps on representation spaces, but as natural operations between vector bundles—inner products and contractions of the kind given in Equation (2.8). All we have are structures on the fundamental bundles $(E^n, M, \mathbb{C}^n, \langle\cdot,\cdot\rangle_n)$ for $n = 1, 2, 3$; in particular, E^1 is fixed so that $E^{1\otimes 3}$ has hypercharge $+1$ (so E^1 itself corresponds to $+\frac{1}{3}$). Different particles are sections of different tensor products; quantum numbers become geometric labels. A down-right-handed quark is

$$\mathbf{d}_R \in \Gamma(E^3 \otimes (E^{1*} \otimes E^{1*})), \tag{4.10}$$

whereas vector bosons are replaced by the corresponding affine covariant derivatives, e.g. $\nabla^1, \nabla^2, \nabla^3$ (see (Gomes, 2024, 2025a) for more details).

In this formulation, weak isospin T_3—defined only relative to a chosen basis of the Lie algebra—has no independent geometrical meaning (see Section B.2 and footnote 39). Left-handed fermions are best understood as components of vector fields $\mathbf{Q}_L$ and $\mathbf{L}_L$. The distinction between electron and electron-neutrino, or between up- and down-left quarks, appears only through their couplings to the Higgs. The Higgs field $\boldsymbol{\varphi}$ provides a frame in $\mathbb{C}^2$ that gives T_3 physical significance. The charges in Table 1 presuppose the frame $\varphi = (0, \varphi^0)^T$; only in that frame do left-handed up-quark components appear as $(u_L^I, 0)^T$.[20]

Geometrically, one should define left-handed components as parallel and orthogonal to the Higgs:

$$\mathbf{e}_L := \langle \mathbf{L}_L, e_0\rangle_2\, e_0, \quad e_L := \langle \mathbf{L}_L, e_0\rangle_2\,, \quad \boldsymbol{\nu}_{eL} := \mathbf{L}_L - \boldsymbol{e}_L, \tag{4.11}$$

$$\mathbf{u}_L^I := \langle \mathbf{Q}_L^I, e_0\rangle_2\, e_0, \quad u_L^I := \langle \mathbf{Q}_L^I, e_0\rangle_2\,, \quad \mathbf{d}_L := \mathbf{Q}_L - \mathbf{u}_L, \tag{4.12}$$

where I indexes colour (red, green, blue) and e_0 is the unit Higgs direction from Section 3.1. Thus left-handed up and down quarks are not distinct particles but components of the same field.[21]

Given an orthonormal basis for E^2 aligned with the Higgs, duals are formed by conjugate transpose: for $\xi = (\xi^\perp, \xi^\parallel)^T$, we have $\xi^* = (\overline{\xi}^\perp, \overline{\xi}^\parallel)$. For Hermitian bundles E, F, we write iterated contractions as $\langle\langle\cdot,\cdot\rangle_E, \cdot\rangle_F$.

With an orthonormal frame aligned to the Higgs, the lepton Yukawa term (4.8) becomes (with coupling g_e):

$$T(\mathbf{L}_L, \boldsymbol{\varphi}, e_R) = g_e\langle\langle \mathbf{L}_L, \boldsymbol{\varphi}\rangle_2, e_R\rangle_1 = g_e(v + H)\overline{e}_L e_R, \tag{4.13}$$

[20] This is why Table 1 can mislead: if both Higgs components are retained, the 'up' and 'down' labels have no physical meaning yet.

[21] Some textbooks note this—cf. (Tong, 2025, p. 185)—but usually in the context of unification 'before symmetry breaking'. Here, these just are parallel and orthogonal components of the same field relative to the Higgs; treating them as independent particle fields would be misleading.

where e_R and e_L are Weyl spinors that are internal scalars (e_L is the magnitude of $\mathbf{L}_L$ along the Higgs). The inner product $\langle\cdot,\cdot\rangle_2$ is complex anti-linear in its first entry.[22] (The E^3, E^2, and E^1 factors contract to $\underline{\mathbb{C}}$ via the fibre inner products; see Table 1 for the E^1-exponents.) Mass terms proportional to $g_e v$ emerge for the electron.

The inner products in (4.13) are geometrically natural: they measure internal angles between particle fields on the same spaces. Once the inner product is fixed, there is no geometric justification for scaling by complex numbers or invoking the group centraliser. The inner product does the real work.

Chirality shows up here: only left-handed particles have components in E^2. Neutrinos do not acquire mass—not because they are orthogonal to the Higgs (which they are), but because we have not included right-handed neutrinos. Had we included them, they would couple to what I call the 'symplectic dual' of the Higgs; this is what happens for quarks.[23]

The symplectic dual φ_c (Table 1) recruits a further geometric structure on $\mathbb{C}^2$: an orientation, encoded by the volume form ϵ_{ab}. Whereas the Higgs mechanism used $(E^2, M, \mathbb{C}^2, \langle\cdot,\cdot\rangle_2)$, here we extend to $(E^2, M, \mathbb{C}^2, \langle\cdot,\cdot\rangle_2, \epsilon)$.[24]

Besides the metric, we can use ϵ_{ab} and ϵ^{ab} to raise or lower indices.[25] Writing $J : E^2 \to E^{2*}$ for the conjugate-linear isomorphism $\xi \mapsto \langle\xi,\cdot\rangle$ (conjugate-linear because the inner product is sesquilinear), we define

$$C := \epsilon^\sharp \circ J : E^2 \mapsto E^2 \tag{4.14}$$

$$\xi^a \mapsto \epsilon^{ac} h_{cb} \xi^b \tag{4.15}$$

where h_{ab} is the inner product. The map C is itself conjugate-linear, and so $\boldsymbol{\varphi}_c := C(\boldsymbol{\varphi})$ measures areas orthogonal to $\boldsymbol{\varphi}$—hence 'symplectic dual'.

With $i = 1, 2, 3$ indexing generations, the quark Yukawa term is:[26]

$$T(\mathbf{Q}_L, \boldsymbol{\varphi}, \boldsymbol{d}_R) = Y^d_{ij} \langle\langle\langle \mathbf{Q}^i_L, \mathbf{d}^j_R\rangle_3, \boldsymbol{\varphi}\rangle_2\rangle_1 + Y^u_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{u}^j_R\rangle_3, \boldsymbol{\varphi}_c\rangle_2\rangle_1. \tag{4.16}$$

This is the geometric form of the standard definition (cf. (Hamilton, 2017, Lemma 8.8.4)). In the Higgs-aligned frame with $\varphi^+ = 0$:

$$Y^d_{ij} \langle\langle\langle \mathbf{Q}^i_L, \mathbf{d}^j_R\rangle_3, \boldsymbol{\varphi}\rangle_2\rangle_1 + Y^u_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{u}^j_R\rangle_3, \boldsymbol{\varphi}_c\rangle_2\rangle_1 = (H + v)\left(Y^d_{ij} d^{Ii}_L d^{Ij}_R + Y^u_{ij} u^{Ii}_L u^{Ij}_R\right), \tag{4.17}$$

summing over colour I and generations i, j.

The Yukawa matrices Y mix generations. One can diagonalise Y^u by passing to the *mass basis*, but Y^u and Y^d cannot be diagonalised simultaneously; their mismatch is encoded in the *CKM matrix*.

---

[22] More precisely: $\varphi \in \Gamma(E^2 \otimes^3 E^1)$, $e^*_L \in \Gamma(E^{2*} \otimes^3 E^{1*})$, and $e_R \in \Gamma(\otimes^6 E^1)$; these match to a scalar.

[23] For this reason, lepton Yukawa terms are diagonal in generations: they do not mix electrons with muons or taus.

[24] Under $A \in U(2)$, $\epsilon_{ab} \mapsto \det(A)\epsilon_{ab}$, so $SU(2)$ preserves it. Since $|\det(A)| = 1$ for unitary A, $\det(A) = e^{i\theta}$ encodes orientation. Fixing ϵ_{ab} as geometric data fixes both inner product and orientation on $\mathbb{C}^2$.

[25] Compare the Hodge star in two dimensions: $e_0, e_1 \mapsto -e_1, e_0$, which acts as the matrix $\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$, the same form as ϵ_{ab}.

[26] Unlike leptons and left-handed quarks, right-handed up and down quarks are genuinely distinct particles, having different E^1 components.

One might worry that, from the geometry-first perspective, different generations cannot coherently mix—after all, if they correspond to different masses, they might inhabit distinct bundles. But this worry dissolves: generations are not different bundles but different *sections of the same composite bundle*. All three live in

$$(E^1 \otimes E^2 \otimes E^3) \oplus (E^1 \otimes E^2 \otimes E^3) \oplus (E^1 \otimes E^2 \otimes E^3), \tag{4.18}$$

so each quark field is a section of a tensor product together with a trivial generation factor:

$$\mathbf{Q}_L^i \in \Gamma(S_L \otimes E_3 \otimes E_2 \otimes E_1^{Y_Q} \otimes \mathbb{C}^3_{\text{gen}}), \qquad i = 1, 2, 3.$$

The Yukawa couplings act as endomorphisms on this factor, $Y_u, Y_d \in \Gamma(\mathrm{End}(\mathbb{C}^3_{\text{gen}}))$. Diagonalising either corresponds to choosing a frame in $\mathbb{C}^3_{\text{gen}}$. Since both cannot be diagonalised simultaneously, their relative rotation $V_{\text{CKM}} = U_u^\dagger U_d$ appears as a unitary automorphism of the generation fibre. The CKM matrix is nothing but a geometric rotation within this trivial bundle.

The geometric perspective clarifies one aspect. If up and down left-handed quarks were independent particles rather than components of a single E^2-field, we *could* diagonalise Y_u and Y_d separately. Because they are components of the same field, coupled to φ and φ_c respectively, we cannot. The W bosons represent ∇^2, the covariant derivative on E^2, and so they too naturally mix generations.[27]

5 A defence of the geometry-first formulation

The geometry-first picture binds symmetry to geometry tightly. One might take this for a defect. I shall argue that it is the picture's chief recommendation.

In Section 1 I claimed that the principal-bundle picture tolerates a slack between symmetry and geometry that the geometry-first picture does not. It is time to make good on that claim. The slack is this: in the PFB-POV, the group G, the representation ρ, and the fibre V can be chosen independently, subject only to mutual consistency. The geometry-first picture eliminates this freedom.[28]

I proceed in four steps. First I exhibit the slack through concrete examples (Section 5.1). Then I show that the geometry-first picture has genuine limitations: certain groups—notably the exceptional families—fall outside its natural reach (Section 5.2). Third, I show that the matter content of the Standard Model cannot, by itself, determine the gauge group—but that the VB-POV sidesteps this failure by starting from fundamental bundles (Sections 5.3). Finally, I argue that physicists already tacitly assume the VB-POV's constraints (Section 5.4).

[27] Other aspects remain mysterious. Unlike their left-handed counterparts, right-handed up and down quarks cannot be understood as components of a single field, due to their different E^1 components. If they could, (4.16) would simplify to inner products along and orthogonal to the Higgs. The obstruction is the hypercharge split between u_R and d_R; this assignment, not the geometry-first machinery, prevents a unified coupling.

[28] Throughout this section, $\mathsf{Aut}(V)$ means the group of automorphisms preserving whatever fibre structure is in play (e.g. Hermitian form, and where specified a unit volume form).

5.1 Slack between symmetry and geometry: examples and lessons

Recall from Section 2.1 that an associated bundle $E = P \times_\rho V$ requires only that $\rho(G)$ preserve the structure of V:

$$\rho(G) \subseteq \mathsf{Aut}(V) \subseteq GL(V). \tag{5.1}$$

The geometry-first picture demands more. It requires

$$G \simeq \rho(G) \simeq \mathsf{Aut}(V). \tag{5.2}$$

Both conditions can fail independently, and the PFB-POV tolerates all such failures. Three examples make this concrete.

Example 1 (faithful but not surjective). The condition $G \simeq \rho(G)$ holds iff ρ is *faithful* (injective), so that only $\mathbb{1}$ acts trivially. Consider the gauge group $U(1)$, fibres $\mathbb{C}^3$ with Hermitian inner product, and the representation

$$\rho(\theta) = e^{i\theta}\mathbb{1}.$$

This representation is faithful, so $U(1) \simeq \rho(U(1))$. However, $\rho(U(1))$ is only a one-dimensional subgroup of $GL(\mathbb{C}^3)$. In particular, it is not isomorphic to the full automorphism group of $(\mathbb{C}^3, \langle\, .\, ,\, .\, \rangle)$, which is $U(3)$; equivalently, ρ is not surjective onto $\mathsf{Aut}(V)$.

From the PFB-POV this is perfectly admissible: one may posit a group that preserves the relevant geometric structures without exhausting them. Still, there is a geometric interpretation: the action rotates the complex volume form of $\mathbb{C}^3$.[29] Thus in this case we have

$$G \simeq \rho(G) \subset \mathsf{Aut}(V), \qquad \dim(G) < \dim(\mathsf{Aut}(V)). \tag{5.3}$$

Example 2 (trivial representation). Consider the trivial action of $SU(n)$ on $V = \mathbb{C}^m$. In this case

$$\rho(G) = \mathbb{1} \subset \mathsf{Aut}(V). \tag{5.4}$$

The group G may be larger or smaller than $\mathsf{Aut}(V)$ as a matter of dimension—$\dim SU(n) = n^2 - 1$ while $\dim U(m) = m^2$—but this is irrelevant: one cannot reconstruct G from its representation on V, since the action is trivial. Thus we have

$$G \not\simeq \rho(G), \quad \rho(G) \subset \mathsf{Aut}(V), \tag{5.5}$$

so both conditions in (5.2) fail. Intrinsically, G may be either larger or smaller than $\mathsf{Aut}(V)$, and the group can neither be recovered from, nor recover, the geometry of the fibre.

Example 3 (non-faithful but geometrically admissible). Consider $G = \mathrm{Spin}(4) \cong SU(2) \times SU(2)$ with fibres $V = \mathbb{C}^2$. Spin(4) admits two inequivalent irreducible representations on $\mathbb{C}^2$, corresponding to left- and right-handed chiral spinors. If we pick one of these factors to act, Spin(4) does preserve the structure of $\mathbb{C}^2$, and so the situation is admissible from the PFB-POV. In this case we have

$$G \not\simeq \rho(G) \simeq \mathsf{Aut}(V), \quad \dim(G) > \dim(\mathsf{Aut}(V)). \tag{5.6}$$

[29] That is, it rotates $\Lambda^3\mathbb{C}^3$ (equivalently, the determinant line). Fixing a unit complex volume form reduces $U(3)$ to $SU(3)$; the residual $U(1)$ is the phase on the determinant line. See footnote 24.

The image of the representation matches $\mathsf{Aut}(V) \simeq SU(2)$, yet the full group G is strictly larger.

The upshot is that the two conditions of (5.2) can fail independently. A faithful representation ($G \simeq \rho(G)$) need not be surjective onto $\mathsf{Aut}(V)$; a surjective image ($\rho(G) \simeq \mathsf{Aut}(V)$) need not faithfully represent the full group. The principal-bundle framework tolerates all of these mismatches. So much for the simple cases.

5.1.a A side-by-side comparison: $SU(n)$ gauge theory

It may help to see both formulations at work on a single example. Take a Hermitian vector bundle $(E \to M, \langle \cdot\, , \cdot \rangle, \epsilon)$ with typical fibre $\mathbb{C}^n$, inner product $\langle \cdot\, , \cdot \rangle$, and complex volume form ϵ (i.e. a nowhere-vanishing section of $\Lambda^n E$). The automorphism group preserving both structures is $\mathsf{Aut}(E_x) \simeq SU(n)$.

VB-POV. One begins with E and a compatible covariant derivative ∇. All further structure is derived:

- The *adjoint bundle* is $\mathsf{End}_0(E) := \{A \in \mathsf{End}(E) \mid \mathsf{Tr}(A) = 0,\, A^\dagger = -A\}$, the bundle of traceless anti-Hermitian endomorphisms. Its sections are the "gluon fields" in the sense that $\nabla = \mathrm{d} + \omega$ for $\omega \in \Gamma(T^*M \otimes \mathsf{End}_0(E))$.
- The *determinant line bundle* is $\det(E) := \Lambda^n E$. The volume form ϵ trivialises it; this is what reduces $U(n)$ to $SU(n)$.
- *Matter multiplets* are sections of E (fundamental), E^* (anti-fundamental), $\mathrm{Sym}^k(E)$ (symmetric tensors), $\Lambda^k E$ (antisymmetric), etc. Each inherits a covariant derivative from ∇ by the Leibniz rule.

All of this is determined by $(E, \langle \cdot\, , \cdot \rangle, \epsilon, \nabla)$. No group is postulated; $SU(n)$ appears only when one asks what transformations preserve the data.

PFB-POV. One begins instead with a principal G-bundle (P, M, G) and a principal connection ϖ. To recover the same physics, one *could* take $G = SU(n)$ acting in the fundamental representation $\rho_{\mathbf{n}}$, in which case $E = P \times_{\rho_{\mathbf{n}}} \mathbb{C}^n$ and ∇ is induced by ϖ as in (A.6). In this case the two formulations coincide.

But the PFB-POV does not *require* this choice. One could equally well take:

1. $G = U(n)$ with $\rho = \rho_{\mathbf{n}}$: now $\dim(G) > \dim(\mathsf{Aut}(V))$, and there is an extra $U(1)$ factor that acts trivially on the orientation ϵ—slack in the sense of Example 1.
2. $G = SU(n) \times SU(m)$ with ρ the projection onto the first factor: the second factor acts trivially on E, and cannot be recovered from $\mathsf{Aut}(E_x)$—slack in the sense of Example 2.
3. $G = Spin(2n)$ with ρ a spinor representation onto $\mathbb{C}^{2^{n-1}}$: the map is surjective but not injective, so $G \not\simeq \rho(G)$—slack in the sense of Example 3.

The VB-POV rules out all three by fiat: it insists that the gauge group *be* the automorphism group of the fundamental bundle. But each of these is a perfectly consistent principal-bundle formulation, and each induces a well-defined covariant derivative on the associated bundle. Yet in each case the group G floats free of the geometry of the fibre: it is either too large, too small, or simply different.

5.1.b The slack world

Consider a world with a single matter field: a section of E with typical fibre $V \simeq (\mathbb{C}^n, \langle \cdot, \cdot \rangle)$. The automorphism group preserving only the inner product is $\mathsf{Aut}(V) = U(n)$. But suppose the principal bundle has structure group $G = SU(n)$, acting in the fundamental representation. Then G preserves more structure—an orientation—than the matter fields intrinsically possess.

In the PFB-POV, this is perfectly consistent. The connection ϖ carries physical observables independently of matter. The Yang–Mills equations are valued in $\mathfrak{su}(n)$, not $\mathfrak{u}(n)$; the gauge sector has different degrees of freedom. In vacuum Yang–Mills theory, the connection A_μ is a physical field in its own right, with dynamics governed by G. If $G = U(n)$ rather than $SU(n)$, the vacuum theory has an extra $U(1)$ factor. These are different theories, with different physical content. Moreover, any Lie-algebra directions in $\ker \rho_*$ remain unsourced by matter: the matter current $J(A, \Phi)$ takes values in $\mathrm{im}(\rho_*)$, so the slack directions satisfy vacuum Yang–Mills equations regardless of the matter configuration (cf. Equation (A.11) in Appendix A).

Call this the *slack world*: G has physical significance independently of matter, and $G \neq \mathsf{Aut}(V)$.[30] Evidence for G comes from the gauge sector; evidence for $\mathsf{Aut}(V)$ from the matter sector. The two need not coincide.

But that is not our world—at least, not our Standard Model; we do not observe such a mismatch. To be clear: this is not a complaint that practitioners of the PFB-POV secretly believe the principal connection ϖ is a physical entity—though most attribute physical significance to its spacetime representatives A_μ, or to a section of the affine bundle of connections (Jacobs, 2023; Kobayaschi, 1957); even in vacuum Yang–Mills theory where no matter fields are present. The complaint is that the PFB-POV's representational resources outrun its explanatory needs. This freedom could in principle be observable—as the slack world illustrates—but in the theories physicists actually take seriously, it is not. I return to this point in Section 5.4.

5.2 Fundamental obstacles to equivalence

The preceding examples showed that even within the classical families of linear groups the PFB-POV permits a looseness that the VB-POV forbids. There are, however, other kinds of obstruction—cases in which, irrespective of choices of representation, no choice of geometric

[30] Parallel transport around a closed curve at $x \in M$ defines an element of $\mathsf{Aut}(E_x)$. The subgroup generated by all such curves is the holonomy group $\mathsf{Hol}(\nabla) \subseteq \mathsf{Aut}(V)$. Under natural assumptions, a principal bundle (P', M, G') exists with $G' \simeq \mathsf{Hol}(\nabla)$ (cf. (Michor, 2008, Theo. 17.11)). But if the posited G is larger than $\mathsf{Aut}(V)$, the slack persists.

data on a vector bundle could possibly reproduce the symmetry-first structure. Two such failures are instructive.

(a) Quantisation of charge. One obstacle to equivalence appears even in the simplest setting: the quantisation of charge. In the PFB-POV, the integrality of electric charge—or of hypercharge in the Standard Model—is traced to the topology of the compact structure group $U(1)$. The continuous one-dimensional representations of $U(1)$ are

$$\rho_n(e^{i\theta}) = e^{in\theta}, \qquad n \in \mathbb{Z}, \tag{5.7}$$

with integer labels enforced by periodicity.

In the geometry-first VB-POV the reasoning is inverted. Here one begins not with a Lie group but with a complex line bundle (E, M, V) whose fibre $V \simeq \mathbb{C}$ is equipped with whatever structure one takes to be physically fundamental—a Hermitian inner product if one wishes to preserve probability norms, or nothing beyond the linear structure if not—so that $\mathsf{Aut}(V)$ is $U(1)$ or $\mathbb{C}^{\times}$ accordingly. All further fields are elements of spaces built tensorially from this fundamental bundle, and distinct "charges" correspond to the tensor powers $E^{\otimes n}$ and their duals. The integer label n simply counts how many copies of the fundamental fibre enter the construction; the possible weights therefore form a discrete lattice. Even though, as a group, $\mathsf{Aut}(V)$ may admit continuously many characters, the geometry-first formalism can reproduce only those obtained by finite tensor operations.

The claim is best stated conditionally: *if* the fundamental charged object is a line bundle E (with whatever fibre structure one stipulates) and all charged fields arise as finite tensor powers and duals of E, *then* the charge ratios lie in a lattice generated by that choice. Irrational ratios cannot arise within such a subtheory—indeed the very phrase 'irrational tensor power' is nonsensical.

One caveat qualifies but does not undermine this conclusion. In particle physics one typically restricts to unitary structure so as to preserve Hermitean norm, effectively selecting $U(1) \subset \mathbb{C}^{\times}$. The VB-POV does not eliminate this physical input; it relocates it from the topology of the group to the geometry of the fibre. What changes is the *source* of discreteness: it is the tensorial construction, not compactness, that enforces integer labels. Even if one works with $\mathsf{Aut}(V) \simeq \mathbb{C}^{\times}$—which is non-compact—the lattice structure persists. One might also worry about fractional hypercharges in the Standard Model—E^1 carries hypercharge $+\frac{1}{3}$ in conventional normalization—but this is a matter of which bundle one takes as fundamental. From the VB-POV, E^1 *is* the generator; the charge lattice is $\mathbb{Z}$ by construction, and the conventional fractions are an artefact of matching to the PFB-POV's normalization.

I find this explanation considerably more satisfying than the topological one.

(b) Exceptional Lie groups. A different kind of obstacle arises with the exceptional Lie groups. The issue here is not that the VB-POV's machinery fails to run—one can always build associated vector bundles given any representation of any group—but that the *explanatory direction* I have been advertising becomes strained.

The VB-POV proceeds by fixing a fibre V endowed with invariant geometric or algebraic data—an inner product, symplectic or volume form, or some higher-rank tensor—and defining the gauge group as

$$G = \mathsf{Aut}(V, \text{structure}) \subset GL(V). \tag{5.8}$$

For the classical families this procedure is canonical: the data determine a unique group whose action both preserves and exhausts the geometry of V. For the exceptional families, however, the correspondence between geometry and symmetry becomes less transparent.

Some exceptional groups do admit "matter-first" realisations, but these are formal and not entirely compelling. G_2 arises as the automorphism group of the octonions, and can be realised as the stabiliser of a generic three-form on $\mathbb{R}^7$; F_4 is the automorphism group of the exceptional Jordan algebra $J_3(\mathbb{O})$; E_6 preserves the cubic norm on the same 27-dimensional space. But there are two difficulties. First, the structures involved—octonion multiplication, Jordan algebra products, cubic norms—are not "geometric" in the low-level tensorial sense employed elsewhere in this paper (inner products, volume forms, and the operations of the tensor algebra), even though some of them—such as G_2 as the stabiliser of a three-form—are perfectly geometric in the broader sense of differential geometry. The point is not that these realisations are illegitimate, but that they do not fit the VB-POV's particular explanatory pattern: start with a fibre endowed with familiar linear-algebraic structure, and read off the group. Second, these realisations are not unique: distinct algebraic structures can yield the same abstract group, so the explanatory arrow from geometry to symmetry is not sharp.

The hard case is E_8. Its minimal faithful representation is the *adjoint*, of dimension 248. One can write down a vector bundle whose automorphism group is E_8:

$$E = (\mathfrak{e}_8, \kappa, [\cdot, \cdot]), \tag{5.9}$$

where $\mathfrak{e}_8$, of dimension 248, carries the Killing form κ and the Lie bracket $[\cdot, \cdot]$—but this is just the Lie algebra $\mathfrak{e}_8$ in disguise. There is no small "matter" representation analogous to the fundamental representations of the classical groups. If one tries to run the VB-POV for an E_8 gauge theory, the "fundamental bundle" is essentially the gauge structure itself rather than a space where matter fields naturally reside. The VB machinery still runs, but the claim of explanatory priority—"start with matter bundles, recover gauge group"—loses its force when the fundamental bundle is already the adjoint.

This is a genuine limitation of the geometry-first rhetoric, though not of the VB formalism per se. It suggests that the VB-POV is most naturally suited to theories whose gauge groups are classical (or at least admit low-dimensional faithful representations), and that E_8 grand unified theories—if they are to be taken seriously—may require a different philosophical framing. Whether this counts against the VB-POV or against E_8 unification is a question I leave open.

One might also worry about supersymmetry. A supersymmetric theory requires a symmetry relating bosons to fermions—connecting the gauge-boson sector to matter. Does this not require stepping outside the geometry-first picture? In fact, standard $\mathcal{N} = 1$ supersymmetric gauge theories with classical groups fit comfortably within it. The fundamental data are a Hermitian vector bundle $E \to M$ with compatible covariant derivative, together with the spacetime spin bundle $S \to M$. Gauginos are sections of $S \otimes \mathsf{End}_0(E)$; matter fermions are sections of $S \otimes E$ or $S \otimes E^*$. Supersymmetry transformations mix these fields, but all are sections of bundles constructed tensorially from (E, ∇, S). Supersymmetry does not require treating the connection as a field of the same type as matter; it requires only that the gaugino—a spinor-valued section of the adjoint bundle—be related to matter fermions. And the adjoint bundle $\mathsf{End}_0(E)$ is already part of the tensorial apparatus. Supersymmetry is not a counterexample; it is a stress test that the geometry-first picture passes.

5.3 Recovery of the gauge group from matter bundles

Section 5.1 showed that the gauge group can float free of the geometry of any single matter bundle. But the Standard Model has many particle species, each in its own associated bundle. One might hope that, taken together, they pin down G: that the slack disappears when the full matter content is taken into account. If so, positing G independently would be a convenience, not a substantive addition. But the hope fails. To recover the product group from the totality of matter bundles, one must also specify which (factors of which) bundles are fundamental—precisely what the VB-POV supplies at the outset.

In more detail: in many gauge theories—and routinely in the Standard Model—the fibres of associated bundles take the form of direct sums or tensor products, such as

$$V = V_1 \otimes V_2, \tag{5.10}$$

with different factors of the gauge group acting on different components, and sometimes *trivially* on them. (This is what happens, for instance, when, in the Standard Model, a particle's representation labels include **1** in either of the first two entries or 0 in the last, corresponding to hypercharge). As a result, kernels appear factorwise, and for any given multiplet we have both conditions of (5.2) failing. The Standard Model abounds with such cases: for the right-handed electron bundle, $SU(3)$ and $SU(2)$ act trivially, so no single associated bundle for e_R can recover the full $SU(3) \times SU(2) \times U(1)$; similar things can be said about right-handed quarks, etc. In short, the representation spaces defined by the particles' quantum numbers admit group actions in which entire factors of G play no role—precisely the kind of degeneration illustrated in Eq. (5.5).

From the PFB-POV this is no defect: the standard strategy assigns each particle type to a section of some associated bundle, without attempting to reconstruct the full gauge group from the geometry of any single fibre. This slack poses no internal difficulty, for the two ingredients—the symmetry and the vector bundle structure—are jointly posited and mutually consistent.[31] But if one tries to recover the group from the automorphisms of those associated bundles, the recovery is not merely incomplete; in general, it is inconsistent.

To see this, suppose we are given a collection of associated vector bundles and then want to recover the structure group from the automorphism groups of the associated bundles. Given

[31] If you think that is not how we actually *use* the principal-associated bundle formalism, I agree; see Section 5.4 below.

the collection

$$(P, M, G, \{\rho_i\}_i, \{V_i\}_i), \quad i = 1, 2 \tag{5.11}$$

we can try to reconstruct (P, G, M) from each $L_{\mathrm{adm}_i}(E_i)$ (the bundle of admissible frames for E_i, see Section 2.1). The problem is that we would in general recover *different* groups for different i: constructing the bundle of admissible frames forces a subgroup $G' \subset G \simeq GL(V)$ to act trivially on some subspaces of E_i, so the resulting principal bundle reflects only the subspaces where G' acts non-trivially—and these differ across the associated bundles.

For concreteness, take $V \simeq \mathbb{C}^3 \otimes \mathbb{C}^2$, $G = SU(3) \times SU(2)$ and consider two representations: $\rho_1 = \mathbf{3} \otimes \mathbf{1}$ (a colour triplet, singlet under $SU(2)$), and $\rho_2 = \mathbf{1} \otimes \mathbf{2}$ (a weak doublet, singlet under $SU(3)$). Thus in this case we would recover $G'_1 = SU(3)$ from one bundle, and $G'_2 = SU(2)$ from the other. What we would like, of course, is $SU(3) \times SU(2)$; but the product structure is nowhere to be found at the level of any single associated bundle.

One might hope to do better by considering the *entire* matter content at once. Suppose we form the direct sum of all the associated bundles,

$$E_{\mathrm{tot}} = E_1 \oplus E_2 \oplus \cdots \oplus E_k, \tag{5.12}$$

and try to recover G from the structure of E_{tot}. There are two natural strategies; both fail.

The first strategy examines the full automorphism group $\mathsf{Aut}(E_{\mathrm{tot}})$. This is too large: it includes not only block-diagonal transformations that act independently on each summand, but also off-diagonal maps that mix summands—transformations with no counterpart in gauge theory.

The second strategy restricts to block-diagonal automorphisms: the subgroup of $\mathsf{Aut}(E_{\mathrm{tot}})$ that preserves each summand E_i separately. Call this $\prod_i \mathsf{Aut}(E_i)$. Assuming G acts faithfully on the total matter content (as it does in the Standard Model up to the $\mathbb{Z}_6$ kernel), G embeds injectively into $\prod_i \mathsf{Aut}(E_i)$, so no factor of G is invisible. But knowing that G lives inside this product does not tell you *which* subgroup it is. And crucially, the product decomposition that the block-diagonal strategy hands you is indexed by *particle species* labeled by i—one i per associated bundle—whereas the gauge group is a product indexed by *interactions*: $SU(3)$ for colour, $SU(2)$ for weak isospin, $U(1)$ for hypercharge. These are entirely different decompositions. The species-indexed product $\prod_i \mathsf{Aut}(E_i)$ contains the interaction-indexed product $G = G_1 \times G_2 \times G_3$ as a subgroup, but does not single it out.

To illustrate: a left-handed quark lives in $E_Q = \mathbb{C}^3 \otimes \mathbb{C}^2 \otimes \mathbb{C}$, with $\mathsf{Aut}(E_Q) \simeq U(6)$ (preserving the Hermitian structure on the six-dimensional total fibre). The gauge group $SU(3) \times SU(2) \times U(1)$ embeds in $U(6)$, but so do many other groups. Nothing intrinsic to E_Q singles out the tensor-product decomposition $\mathbb{C}^3 \otimes \mathbb{C}^2 \otimes \mathbb{C}$ from which the product structure of G derives. That decomposition is additional data—it is the specification of fundamental bundles.

Claim. Given a finite collection of Hermitian vector bundles $\{(E_i, \langle\cdot,\cdot\rangle_i)\}_{i=1}^k$, or even the enriched collection $\{(E_i, \langle\cdot,\cdot\rangle_i, \nabla^i)\}_{i=1}^k$ of bundles equipped with compatible covariant derivatives, the product decomposition of the gauge group $G = \prod_\alpha G_\alpha$

into independent interaction sectors is not uniquely determined—even when G acts faithfully on the total matter content.

5.4 Fundamental bundles and the Standard Model

The VB-POV avoids the failure we just witnessed, to reconstruct the full product group G from composite bundles, by working at a different level. One does not try to recover the full gauge group from composite bundles. Instead, each factor of G arises as the automorphism group of a *fundamental* bundle, just as in (2.11). The composite bundles where particles live—tensor products like $E^3 \otimes E^2 \otimes E^1$—are derived from these.[32] But the gauge group is not recovered from the composite fibres; it is built into the fundamental ones from the start.

As it happens, the symmetry group realised in Nature fits the VB-POV's stricter requirements. An instructive wrinkle illustrates how.

The nominal gauge group of the Standard Model is $SU(3) \times SU(2) \times U(1)$. A subgroup $\mathbb{Z}_6 \subset SU(3)\times SU(2)\times U(1)$ acts trivially on the *entire* particle content of the Standard Model—not merely on individual multiplets, but on every representation that appears.[33] The group that acts faithfully on the total matter content is therefore the quotient $SU(3) \times SU(2) \times U(1)/\mathbb{Z}_6$.

From the standpoint of the PFB-POV, this is a genuine ambiguity: one can take either the product or the quotient as the structure group, and only global topological data or further physical input can decide between them. From the standpoint of the VB-POV, by contrast, there is nothing to decide. One begins with three fundamental vector bundles—fibres $\mathbb{C}^3$, $\mathbb{C}^2$, and $\mathbb{C}$, with Hermitian inner products and, for the first two, orientations. Their automorphism groups are $SU(3)$, $SU(2)$, and $U(1)$, and the reconstructed gauge group is their product. The $\mathbb{Z}_6$ kernel is a downstream consequence of the specific tensor products nature happens to employ—a curiosity about which multiplets appear in the Standard Model, not a structural problem for the formalism. The analogy is with $SO(3)$ acting on $\mathbb{R}^3$: certain tensors built from $\mathbb{R}^3$ may be insensitive to certain rotations, but this does not make $SO(3)$ 'too big' for the vector space.

If one instead insists that the reconstructed group act faithfully on the total particle content, one obtains $SU(3) \times SU(2) \times U(1)/\mathbb{Z}_6$. But this quotient is not a product of classical groups acting on standard spaces in any transparent way, and the geometric interpretation that motivates the frame-bundle picture is lost. A faithfulness requirement, far from cleaning up the recovery, makes it less straightforward.

One might wonder whether grand unified theories fare better. In an $SU(5)$ GUT, for instance, there is a single fundamental bundle with fibre $\mathbb{C}^5$, and the entire particle content arises tensorially from it. Here recovery proceeds smoothly—a cleaner fit for the PFB-POV (for the VB-POV the gain is more humble). But $SU(5)$ unification, as far as we know, has failed empirically. The Standard Model, with its product group, its multiple fundamental bundles,

[32] Of course, particles may have no component in a given fundamental direction.

[33] The $\mathbb{Z}_6$ kernel arises from the specific hypercharge assignments of the Standard Model particles. Elements of $SU(3)\times SU(2)\times U(1)$ of the form $(e^{2\pi i/3}\mathbb{1}_3,\, -\mathbb{1}_2,\, e^{i\pi/3})$ and their powers act trivially on all Standard Model fields.

and its $\mathbb{Z}_6$ kernel, is what we have. The geometry-first picture handles it—and a telling fact about practice suggests that physicists already know this.

Recall the slack world of Section 5.1.a. That is not our world. And barring the standard treatment of the vacuum Yang-Mills case, nor is such a world usually included in the menu of physical possibilities. We do not often seriously consider a mismatch between the symmetries of the gauge sector and those of the matter sector. I claim that there is a tacit assumption built into actual practice. Namely, that each factor of G coincides with the automorphism group of a fundamental bundle:

$$G = \prod_i G_i, \quad G_i \simeq \rho(G_i) \simeq \mathsf{Aut}(V_i). \tag{5.13}$$

This is precisely what the VB-POV makes explicit.

Now, I admit physicists are not always consistent about this tacit assumption—being tacit makes it harder to see. After describing the simple case where $G \simeq \mathsf{Aut}(V)$, textbooks typically move to composite cases without noting that the product structure must be *stipulated*. Skinner (2007) is unusually explicit. On p. 141: "Vector bundles are of relevance to physics because a charged matter field is a section of an associated vector bundle." After constructing principal bundles from frame bundles—assuming $G \simeq \mathsf{Aut}(V)$—he adds (p. 142): "The most common Lie groups that arise in physics are indeed matrix Lie groups [of that sort], so the two viewpoints are equivalent." But immediately: "However, in some exotic theories (especially string theory and some grand unified theories) exceptional Lie groups such as E_6 play an important role, so the fundamental picture is really that of principal bundles."

The qualification is telling. Skinner restricts equivalence to theories that "commonly arise in physics"—precisely those satisfying the VB-POV's constraints. What distinguishes those theories? Classical groups acting faithfully in fundamental representations, with matter built tensorially from a small set of fundamental bundles? If those conditions are what make a theory viable, the restriction is not incidental—it is the geometry-first picture in all but name.

6 Conclusions

6.1 Summary

The geometry-first formulation does three things. It posits a different ontology. It offers independent, symmetry-free explanations of familiar mechanisms. And it eliminates the slack between symmetry and geometry that the principal-bundle formulation tolerates. That tighter fit may help explain why certain group–geometry correspondences are realised in our world and others are not.

What changes in the underlying picture is visible even in the simplest case. In vacuum Yang–Mills theory, the fundamental dynamical object is not a connection on a principal bundle but a covariant derivative on a vector bundle—an affine structure whose automorphism group happens to be a Lie group, but whose description need never mention one.

Consider how geometry does the explanatory work of the Higgs mechanism without mentioning symmetry. Since we represent non-isomorphic sections directly from the start, Goldstone modes never appear and never require elimination (see footnote 17). What the symmetry-first formulation calls 'mass acquisition' is simply this: the kinetic term of any section depends on the bundle's affine structure, and the shape operator along the Higgs direction—the extrinsic curvature—is the part that acquires mass.[34] But the covariant derivative along a single section does not exhaust the affine degrees of freedom (when $\dim(E_x) \geq 2$). The degrees absent from the extrinsic curvature remain massless—corresponding, in symmetry-first language, to the unbroken gauge group and its massless photons. Strictly speaking, 'mass acquisition' is a misnomer from the geometry-first standpoint; what could it mean for a covariant derivative to acquire mass? But the phrase acquires clear meaning in the linearised regime.[35]

The Yukawa mechanism, too, gains clarity. The Yukawa form becomes geometrically natural—a fibrewise contraction built from structures already present on the fundamental bundles—rather than a representation-theoretic construction that postulates maps between associated bundles and absorbs ambiguities into coupling constants. And the formulation sharpens physical questions: the quark Yukawa term depends essentially on the orientation of $\mathbb{C}^2$, which explains why $SU(2)$ rather than $U(2)$ appears naturally. For $\mathbb{C}^3$, by contrast, I don't know the analogous mechanism—why does the Standard Model employ $SU(3)$ rather than $U(3)$?[36]

The group is subservient to the geometry, and this changes familiar explanations. Charge quantisation (Section 5.2) is a striking example: discreteness follows from the tensorial construction of matter fields, not from the topology of a compact group. Even if the automorphism group of the fibre were non-compact, the charge lattice would persist—its generator fixed by the choice of fundamental bundle, not by a topological identification in the group.

6.2 Two defences

The geometry-first picture has comparative advantages. The symmetry-free explanations just surveyed—of the Higgs mechanism, the Yukawa couplings, and charge quantisation—share a common source: the geometry-first picture eliminates the slack between symmetry and geometry that the principal-bundle picture tolerates. It makes explicit an assumption—that gauge groups coincide with automorphism groups of fundamental bundles—that the principal-bundle picture leaves tacit, and thereby excludes many symmetry-first models—most of which are never considered in practice anyway.[37]

---

[34] In the linearised regime, no Killing form is needed: the inner product on vector bundles suffices to give dimensions to perturbations of flat covariant derivatives.

[35] Some would hesitate to say gravitons acquire mass merely because a spacetime admits a kinetic term for a constant-norm vector field; yet that they do is the consensus for such theories ((Jacobson, 2008)).

[36] There is a canonical isomorphism $U(3) \simeq (SU(3) \times U(1))/\mathbb{Z}^3$, but the $U(1)$ representations in the Standard Model do not realise it. Benjamin Muntz (p.c.) suggests looking at triality constraints on baryon coupling: colourless three-quark states would not be invariant under $U(3)$.

[37] As discussed in Section 5, there is a widespread tacit assumption that gauge groups simply are automorphism groups of fundamental vector spaces. But this assumption is both unwarranted and inconsistently applied.

This is a defence of the geometry-first picture in absolute terms: I judge it to provide superior explanations—a cleaner ontology—for the class of theories to which it applies. The principal gauge field is not altogether denied reality; in the geometry-first picture, the same physical content is carried by the covariant derivative on a specific vector bundle. But I acknowledge that my preference is staked on ontological commitments that not everyone shares.

Of course, the geometry-first picture also has comparative disadvantages. It cannot accommodate gauge theories built from exceptional Lie groups such as E_8 without losing its explanatory force (Section 5.2). Moreover, the principal-bundle picture's greater flexibility is a genuine virtue when one wants a single framework covering as wide a class of gauge theories as possible, not just the narrower class that includes the Standard Model.

So let me offer a second defence that does not require my aforementioned commitments. What we can say unconditionally is that, for the Standard Model, the predictions stay the same. Why bother, then, with a reformulation that changes no predictions? As Feynman observed, and as Hunt (2025) has argued in detail, equivalent formulations are not always equally fertile. A genuinely different explanation of the same phenomena is already an epistemic gain—it widens the base from which we move into the unknown. The comparative advantages noted above are sufficient to warrant the geometry-first picture's consideration as an alternative, even if it does not supersede the principal-bundle picture for all purposes.

Two morals, then. First, future work in gauge theory might profitably begin with structured vector bundles and the tensors they carry, rather than with groups. Second, a version of Ockham's razor: if the geometry already determines the symmetry, the symmetry is doing no independent work.

Acknowledgements

I would like to especially thank Aldo Riello, Benjamin Muntz, Axel Maas, Silvester Borsboom, Benjamin Feintzeig, and Mark Hamilton for feedback and comments. I would also like to thank David Tong, Oliver Pooley, David Wallace, Caspar Jacobs, Jim Weatherall, Jeremy Butterfield, and Eleanor March for many conversations on this topic.

APPENDIX

A Principal and associated fibre bundles

I will start with the definition of a principal bundle:

Definition 1 (Principal fibre Bundle) (P, M, G) *consists of a smooth manifold P that admits a smooth free action of a (path-connected, semi-simple) Lie group, G: i.e. there is a map $G \times P \to P$ with $(g, p) \mapsto g \cdot p$ for some right (or left, with appropriate changes throughout) action $\cdot$ and such that for each $p \in P$, the isotropy group is the identity (i.e. $G_p := \{g \in G \mid g \cdot p = p\} = \{e\}$). P has a canonical, differentiable, surjective map, called a* *projection, under the equivalence relation $p \sim g \cdot p$, such that $\pi : P \to P/G \simeq M$, where here $\simeq$ stands for a diffeomorphism.*

It follows from the definition that $\pi^{-1}(x) = \{G \cdot p\}$ for $\pi(p) = x$. And so there is a diffeomorphism between G and $\pi^{-1}(x)$, fixed by a choice of $p \in \pi^{-1}(x)$. It also follows (more subtly) from the definition, that local sections of P exist. A local section of P over $U \subset M$ is a map, $\sigma : U \to P$ such that $\pi \circ \sigma = \mathrm{Id}_U$.

Given an element ξ of the Lie-algebra $\mathfrak{g}$, and the action of G on P, we use the exponential to find an action of $\mathfrak{g}$ on P. This defines an embedding of the Lie algebra into the tangent space at each point, given by the *hash* operator: $\sharp_p : \mathfrak{g} \to T_pP$. The image of this embedding we call *the vertical space* V_p at a point $p \in P$: it is tangent to the orbits of the group, and is linearly spanned by vectors of the form

$$\text{for} \quad \xi \in \mathfrak{g}: \quad \xi^\sharp(p) := \frac{d}{dt}_{|t=0}(\exp(t\xi) \cdot p) \in V_p \subset T_pP. \tag{A.1}$$

Vector fields of the form $\xi^\sharp$ for $\xi \in \mathfrak{g}$ are called *fundamental vector fields.*[38]

The vertical spaces are defined canonically from the group action, as in (A.1). But we can define an 'orthogonal' projection operator, $\widehat{V}$ such that:

$$\widehat{V}|_V = \mathsf{Id}|_V, \quad \widehat{V} \circ \widehat{V} = \widehat{V}, \tag{A.2}$$

and defining $H \subset TP$ as $H := \mathsf{ker}(\widehat{V})$. It follows that $\widehat{H} = \mathsf{Id} - \widehat{V}$ and so $\widehat{V} \circ \widehat{H} = \widehat{H} \circ \widehat{V} = 0$. Moreover, since $\pi_* \circ \widehat{V} = 0$ it follows that:

$$\pi_* \circ \widehat{H} = \pi_*. \tag{A.3}$$

The connection-form should be visualized essentially as the projection onto the vertical spaces. The only difference between $\widehat{V}$ and ϖ is that the latter is $\mathfrak{g}$-valued, Thus we get it via the isomorphism between V_p and $\mathfrak{g}$ (ϖ's inverse is $\sharp : \mathfrak{g} \mapsto V \subset TP$). We can define it directly as:

Definition 2 (An principal connection-form) *ϖ is defined as a Lie-algebra valued one form on P, satisfying the following properties:*

$$\varpi(\xi^\sharp) = \xi \qquad \text{and} \qquad {L_g}^*\varpi = \mathrm{Ad}_g\varpi, \tag{A.4}$$

where the adjoint representation of G on $\mathfrak{g}$ is defined as $\mathrm{Ad}_g\xi = g\xi g^{-1}$, for $\xi \in \mathfrak{g}$; ${L_g}^*$ is the pull-back of TP induced by the diffeomorphism $g : P \to P$.

Now, in possession of an principal connection, we can induce a notion of covariant derivative on *associated vector bundles*:

[38] It is important to note that there are vector fields that are vertical and yet are not fundamental, since they may depend on $x \in M$ (or on the orbit).

Definition 3 (Associated Vector Bundle) *A vector bundle over M with typical fibre V, is associated to P with structure group G, is defined as:*

$$E = P \times_\rho V := P \times V/\sim \quad \textit{where} \quad (p, v) \sim (g \cdot p, \rho(g^{-1})v), \tag{A.5}$$

where $\rho : G \to GL(V)$ is a representation of G on V.

One can get a covariant derivative on an associated vector bundle E from ϖ as follows: let $\gamma : I \to M$ be a curve tangent to $\mathbf{v} \in T_xM$, and consider its horizontal lift, γ_h. Let $\sigma : M \to P$ be a local section of P, and define $\kappa(x) = [\sigma(x), \tau(x)]$, for some $\tau : U \to V$. Then

$$\nabla_{\mathbf{v}}\kappa = \frac{d}{dt}[\gamma_h, \tau]. \tag{A.6}$$

Conversely, we can define a horizontal subspace from the covariant derivatives as follows. For $p = e_1, ...e_n \in L(E)$, and for all curves $\gamma \in M$ such that $\mathbf{v} = \dot\gamma(0) \in T_xM$, with $\pi(p) = x$, let $\{e_1(t), ..., e_n(t)\}$ be curves in E such that $\nabla_{\mathbf{v}}(e_i(t)) = 0$. Doing this for each v defines a horizontal subspace.

But we can also obtain the vector bundles more directly as follows:

Definition 4 (Vector Bundle) *A vector bundle (E, M, V) consists of: E a smooth manifold that admits the action of a surjective projection $\pi_E : E \to M$ so that any point of the base space M has a neighborhood, $U \subset M$, such that, for all proper subsets of U, E is locally of the form $\pi^{-1}(U) \simeq U \times V$, where V is a vector space (e.g. $\mathbb{R}^k$, or $\mathbb{C}^k$) which is linearly isomorphic to $\pi^{-1}(x)$, for any $x \in M$.*

Note that the isomorphism between $\pi^{-1}(U)$ and $U \times V$ is not unique, which is why there is no canonical identification of elements of fibres over different points of spacetime. Each choice of isomorphism is called 'a trivialization' of the bundle.

Definition 5 (A section of E) *A section of E is a map $\kappa : M \to E$ such that $\pi_E \circ \kappa = \mathrm{Id}_M$. We denote the space of smooth sections by $\kappa \in \Gamma(E)$.*

Given a vector bundle (E, M, V) a covariant derivative ∇ is an operator:

$$\nabla : \Gamma(E) \to \Gamma(T^*M \otimes E) \tag{A.7}$$

such that the product rule

$$\nabla(f\kappa) = \mathrm{d}f \otimes \kappa + f\nabla\kappa \tag{A.8}$$

is satisfied for all smooth, real (or complex)-valued functions $f \in \Gamma(M)$.

Thus we can define parallel transport as follows:

Definition 6 (Parallel transport in a vector bundle) *Let ∇ be a covariant derivative on (E, M, V), $\mathbf{v} \in E_x$ and $\gamma(t)$ a curve in M such that $\gamma(0) = x$. Then we define the parallel transport along γ as the unique section $\mathbf{v}_h(t)$ of $E|_\gamma$ such that:*

$$\nabla_{\gamma'}\mathbf{v}_h = 0. \tag{A.9}$$

The existence and uniqueness of this map is guaranteed for $\gamma \subset U$ some open subset of M, and it follows from properties of solutions of ordinary differential equations (cf. (Kobayashi & Nomizu, 1963, Ch. II.2)).

Here ∇ is an operator, not a tensor. But by introducing a coordinate frame or basis, we can represent it as such. This is the same as for spacetime covariant derivatives, ∇: it is only upon the introduction of a frame or basis that we find an explicit representation.

To define Ω in terms of ∇, we proceed in the usual way:

Definition 7 (Curvature) *Given a covariant derivative ∇ on a vector bundle E, the* curvature tensor *is the unique multilinear bundle map*

$$\Omega : TM \otimes TM \otimes E \to E \quad : \quad (X, Y, v) \mapsto \Omega(X, Y)\kappa$$

such that for all $X, Y \in TM$ and $\kappa \in \Gamma(E)$,

$$\Omega(X, Y)\kappa = \left(\nabla_X \nabla_Y - \nabla_Y \nabla_X - \nabla_{[X,Y]}\right)\kappa,$$

where $[\cdot, \cdot]$ is the Lie bracket of spacetime vector fields.

We can see the curvature then as an element of $\Omega : TM \otimes TM \otimes \mathsf{End}(E)$, i.e. as a map valued on the endomorphisms of E (the fibre-linear transformations that are not necessarily automorphisms).

The trace operation is defined as $\mathsf{Tr} : \mathsf{End}(E) \to C^\infty(M)$, and so can be included in a Lagrangian specifying the dynamics of ∇. Since $\mathsf{End}(E)$ is closed under composition, we can obtain a Lagrangian 4-form for the action:

$$\mathcal{L} = \mathsf{Tr}(\Omega \wedge *\Omega). \tag{A.10}$$

It will prove useful to know that, given any vector bundle (E, M, V) the bundle of frames for E, called $L(E)$, is itself a principal fibre bundle $(L(E), M, GL(V))$: here elements of $\pi^{-1}(x)$ are linear frames of E_x, and $G \simeq GL(V)$ acts via ρ on the typical fibres. By construction, $E \simeq L(E) \times_\rho V$. Now, for $G' \subset G \simeq GL(V)$ we can partition the points of each orbit in P, $\mathcal{O}_p := Gp$, into orbits of G'. Each such choice gives a principal bundle with group G' and it induces further structure on the associated vector bundle, e.g. an inner product, by selecting which frames are considered orthonormal. This is also a principal fibre bundle, $(L'(E), M, G')$, whose structure group is a proper subgroup of the general linear group, $G' \subset GL(V)$, taken to be the group that preserves the structure of V. This is called the *bundle of admissible frames.*

Sourced Yang–Mills equations and slack gauge directions From the principal-bundle point of view, matter couples to a connection A only through the infinitesimal action of the gauge Lie algebra on the matter fibres. Let $\rho : G \to U(V)$ be the unitary representation used to form the associated bundle $E = P \times_\rho V$, and write $\rho_* : \mathfrak{g} \to \mathfrak{u}(V)$ for its differential. For a Higgs field $\Phi \in \Gamma(E)$, variation of the Yang–Mills–Higgs action with respect to A yields the sourced equation

$$d_A^* F_A \;=\; J_H(A, \Phi), \tag{A.11}$$

where the Higgs current $J_H(A,\Phi) \in \Omega^1(M, \mathrm{Ad}(P))$ is characterised by the Riesz identity

$$\langle \alpha, J_H(A,\Phi)\rangle_{\mathrm{Ad}(P)} \;=\; 2\,\mathrm{Re}\,\langle d_A\Phi, \alpha\cdot\Phi\rangle_E \qquad \text{for all } \alpha \in \Omega^1(M, \mathrm{Ad}(P)). \tag{A.12}$$

Here $\alpha \cdot \Phi$ denotes the pointwise action induced by ρ_*: locally, if $\alpha = \alpha_\mu\, dx^\mu$ with $\alpha_\mu(x) \in \mathfrak{g}$, then $(\alpha\cdot\Phi)_\mu = \rho_*(\alpha_\mu)\Phi$.

This makes explicit what happens in *slack* situations. If $\ker(\rho_*) \neq 0$, then $\xi \cdot \Phi = 0$ for all $\xi \in \ker(\rho_*)$, and hence the right-hand side of (A.12) vanishes whenever α takes values in $\ker(\rho_*)$. Equivalently, the current takes values only in the Lie-subalgebra that actually acts on matter:

$$J_H(A,\Phi) \;\in\; \Omega^1\big(M, \ker(\rho_*)^\perp\big) \;\subset\; \Omega^1(M, \mathrm{Ad}(P)), \tag{A.13}$$

where $\ker(\rho_*)^\perp$ is the orthogonal complement in $\mathfrak{g}$ with respect to the chosen Ad-invariant inner product. Accordingly, (A.11) splits into a sourced equation along $\ker(\rho_*)^\perp$ and a *vacuum* equation along $\ker(\rho_*)$:

$$\pi_{\ker(\rho_*)^\perp}(d_A^* F_A) = J_H(A,\Phi), \qquad \pi_{\ker(\rho_*)}(d_A^* F_A) = 0. \tag{A.14}$$

Any "extra" gauge directions that do not act on the chosen matter content are automatically unsourced. They are empirically idle unless further charged fields are added.

Conversely, if the slack is only *discrete*—i.e. $\ker(\rho)$ is nontrivial but $\ker(\rho_*) = 0$—then the local Euler–Lagrange equation (A.11) is unchanged; the difference concerns only global issues (which principal bundles and large gauge transformations one admits).

B How bosons acquire mass

It is easy to translate Equation (3.12) to the standard formulation's conclusion that the quadratic terms in the (infinitesimal) connection ω would correspond to vector bosons 'acquiring masses'. Introduce a connection $\nabla = \mathrm{d} + \omega$ such that $\mathrm{d}e_0 = 0$ and $\omega \in \Gamma(T^*M \otimes \mathsf{End}(E))$, where $\mathsf{End}(E)$ are the linear endomorphisms of E; so for $\varphi \in \Gamma(E)$, we have $\omega\cdot\varphi \in \Gamma(T^*M \otimes E)$. And thus $\|\nabla e_0\| = \|\omega \cdot e_0\|$. Again, not all components of ω contribute to $\|\omega\cdot e_0\|^2$ in (3.12). In a basis $\{e_I\}_{I=0,\dots,n-1}$ adapted to e_0, we have

$$\nabla e_I = \omega^J{}_I\, e_J, \quad \text{and so} \quad \nabla e_0 = \omega_0^i e_i, \quad \text{with} \quad i \neq 0, \tag{B.1}$$

from the anti-symmetry of the connection. Then

$$\|\nabla\varphi\|^2 = (\partial H)^2 + v^2 \sum_{i\neq 0} (\omega^i{}_0)^2 + \mathcal{O}(\varepsilon^3). \tag{B.2}$$

Hence, only those components of ω that move e_0 (onto the orthogonal directions) 'acquire mass'. The components that preserve e_0, e.g. $\omega^i{}_j, i \neq j$, remain massless. In the group-theoretic language, these would correspond precisely to the stabiliser subgroup of e_0. This suffices to get the mass for some vector bosons, and to ensure the right family remains massless.

Now we can illustrate this in detail with two examples.

B.1 Example: spacetime rotations

Suppose we are dealing with three-dimensional Riemannian manifold, (M, g). Here a general $\mathfrak{so}(3) \simeq \mathfrak{su}(2)$ connection has the form

$$\omega = \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix}. \tag{B.3}$$

If the Higgs unit vector is $e_0 = (1, 0, 0)^T$ (where T here is the transpose, allowing us to write column-vectors in-line), then

$$\omega \cdot e_0 = (0\,,\, \omega_z\,,\, -\omega_y)^T. \tag{B.4}$$

Thus, we would get:

$$\|\nabla\varphi\|^2 = v^2(\omega_y^2 + \omega_z^2). \tag{B.5}$$

So ω_y and ω_z would 'acquire mass', while ω_x would remain 'massless'.

B.2 Electroweak Example

Now we take the example relevant to : $(E^2 \otimes E^1, \langle\, .\,,\, .\,\rangle)$. The covariant derivative on an element $\mathbf{v} \otimes \mathbf{w} \in V \otimes W$ is given by

$$\nabla(\mathbf{v} \otimes \mathbf{w}) = (\nabla^V \mathbf{v}) \otimes \mathbf{w} + \mathbf{v} \otimes \nabla^W \mathbf{w}, \tag{B.6}$$

where ∇^V, ∇^W are covariant derivatives on, in what follows, $V \simeq \mathbb{C}^2, W \simeq \mathbb{C}^1$, respectively.

For the electroweak theory, let $e_0 = e_0^2 \otimes e_0^1 \in \Gamma(E^2 \otimes E^1)$ with $e_0^2 = (0, 1)$, $e_0^1 = 1$. And so we get:

$$\nabla e_0 = \omega \cdot e_0^2 + e_0^2 Z = (\omega + iZ\mathbb{1})e_0^2, \tag{B.7}$$

where ω is the connection for the covariant derivative on $\mathbb{C}^2$ and Z is the connection on $\mathbb{C}$. To complete the comparison with the standard formalism, we choose the weak-isospin eigenbasis, on which the third generator of the $\mathfrak{su}(2)$ algebra, $\mathbb{T}_3$, is diagonal. Omitting the coupling constants for brevity, we can write ω as:[39]

$$\omega = \begin{pmatrix} iW_3 & iW_1 - W_2 \\ iW_1 + W_2 & -iW_3 \end{pmatrix}, \text{and} \quad iZ\mathbb{1} = \begin{pmatrix} iZ & 0 \\ 0 & iZ \end{pmatrix}. \tag{B.8}$$

Applying this to e_0^2 in (B.7) gives

$$\nabla e_0 = \begin{pmatrix} iW_1 - W_2 \\ -iW_3 + iZ \end{pmatrix}. \tag{B.9}$$

[39] Note that this is not the ω written in terms of the spin coefficients, i.e. in terms of an orthonormal frame that includes e_0. That could also be done, and indeed it was done in the previous example $\mathfrak{so}(3) \simeq \mathfrak{su}(2)$, with an orthonormal frame $(0, 1), (0, i), (1, 0), (i, 0)$. Here we are attempting to make contact with the standard notation and formalism and so are using its conventions.

Hence the corresponding quadratic term appearing in (3.12) is

$$\|\nabla e_0\|^2 = W_1^2 + W_2^2 + (Z - W_3)^2. \tag{B.10}$$

Thus W_1, W_2 and the combination $Z - W_3$ acquire mass, while $Z + W_3$ remains massless. The latter is identified with the photon. Of course, had we chosen a different form for e_0^2, we would have obtained different combination of massive and massless bosons. For instance, for $e_0^2 = (1, 0)$ it is easy to see that it would have been $Z + W_3$ that would acquire mass, while $Z - W_3$ would remain massless.

References

Bleecker, D. (1981). *Gauge Theory and Variational Principles.* Dover Publications.

Geroch, R. (1972, dec). Einstein algebras. *Communications in Mathematical Physics*, *26*(4), 271–275. doi: 10.1007/bf01645521

Gomes, H. (2024, October). Gauge Theory Without Principal Fiber Bundles. *Philosophy of Science*, 1–17. doi: 10.1017/psa.2024.49

Gomes, H. (2025a). The Aharonov-Bohm effect: fact and reality. *Philosophy of Physics*.

Gomes, H. (2025b). Representational Schemes for theories with symmetries. *Synthese*.

Gomes, H., & Butterfield, J. (2023, June). The Hole Argument and Beyond: Part II: Treating Non-isomorphic Spacetimes. *Journal of Physics: Conference Series*, *2533*(1), 012003. Retrieved from https://dx.doi.org/10.1088/1742-6596/2533/1/012003 doi: 10.1088/1742-6596/2533/1/012003

Hamilton, M. (2017). *Mathematical Gauge Theory.* Springer International Publishing. doi: 10.1007/978-3-319-68439-0

Hunt, J. (2025). On the Value of Reformulating. *Journal of Philosophy*.

Jacobs, C. (2023). The metaphysics of fibre bundles. *Studies in History and Philosophy of Science*, *97*, 34-43. Retrieved from https://www.sciencedirect.com/science/article/pii/S0039368122001777 doi: https://doi.org/10.1016/j.shpsa.2022.11.010

Jacobson, T. (2008, October). Einstein-æther gravity: a status report. In *Proceedings of from quantum to emergent gravity: Theory and phenomenology — pos(qg-ph).* Sissa Medialab. doi: 10.22323/1.043.0020

Kabel, V., de la Hamette, A.-C., Apadula, L., Cepollaro, C., Gomes, H., Butterfield, J., & Brukner, C. (2025, April). Quantum coordinates, localisation of events, and the quantum hole argument. *Communications Physics*, *8*(1). doi: 10.1038/s42005-025-02084-3

Kobayaschi, S. (1957). Theory of connections. *Annali di Matematica 43, 119–194*.

Kobayashi, S., & Nomizu, K. (1963). *Foundations of differential geometry. Vol I.* Interscience Publishers, a division of John Wiley & Sons, New York-Lond on.

Michor, P. W. (2008). *Topics in differential geometry* (No. volume 93). Providence, Rhode Island: American Mathematical Society. (Includes bibliographical references (pages 479-488) and index. Description based on print version record.)

Skinner, D. (2007). *Quantum Field Theory II.* Cambridge University Press.

Tong, D. (2025). *The Standard Model.* Cambridge University Press.

Weatherall, J. (2016). Fiber bundles, Yang–Mills theory, and general relativity. *Synthese*, *193*(8), 2389–2425. (http://philsci-archive.pitt.edu/11481/)