Particles before symmetry

Henrique Gomes

Oriel College, University of Oxford

October 21, 2025

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 $U(1)$. 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 conditions.

Contents

1 Introduction

Should we value mathematically equivalent formulations of a theory? Feynman (1994, p. 127), in his characteristic style, has a persuasive answer to this question:

> Every theoretical physicist who is any good knows six or seven different theoretical representations for exactly the same physics. He knows that they are all equivalent, and that nobody is ever going to be able to decide which one is right at that level, but he keeps them in his head, hoping that they will give him different ideas for guessing.

In his Nobel Prize Lecture ("The Development of the Space-Time View of QED", 1965), he further reflects on the value of alternative ways of thinking about a theory:

> 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. [...] If every individual student follows the same current fashion then the variety of hypotheses being generated [...] is limited.

The conduit of Feynman's reflections was his path integral formalism, which was mathematically equivalent to an earlier approach (Schwinger's) to quantum field theory. The subsequent career of path integrals is evidence of the validity of his arguments.

Another classic case further illustrates the point: Minkowski's 1908 recasting of Einstein's special relativity into the language of four-dimensional spacetime geometry. Again, the underlying physics was unchanged, but the shift in formulation was decisive for future developments. Einstein initially dismissed Minkowski's treatment as "überflüssige Gelehrsamkeit" (superfluous erudition), yet by 1912 he had conceded that only the spacetime formulation revealed the true essence of the theory.[1] What Minkowski introduced was not new predictions but a new ontology: space and time no longer standing apart, but merged into a single structure. And it was precisely this geometrical vantage point that enabled the later generalisation to general relativity (cf. (Stachel, 2002, p. 226)).

The value of theory-reformulation is a live philosophical topic, admitting a spectrum of positions, from the maximally deflationary instrumentalism to the maximally inflationary fundamentalism (see (Hunt, 2025) for a recent appraisal).[2] In this paper, I aim to provide a reformulation that I judge to be valuable under the lights of all but the most deflationary instrumentalist.

The reformulation will be of classical gauge theory as it applies to the Standard Model of particle physics, so not quantum electrodynamics or special relativity, and of a less radical nature than either Feynman's or Minkowski's reformulations. Namely, I aim to provide a reformulation of the Standard Model that does not rely on an initial postulate about symmetry; I will call it *geometry-first*.

The knowledgeable reader will be quick to point out that gauge theory is already highly geometrical: principal fibre bundles and connections are, after all, the stock-in-trade of the geometer. However, the usual formulation also incorporates symmetry into its foundations, and with it, an ontology that extends beyond the spaces where matter fields actually reside.

For theories of particle physics, I will assume the spaces where matter fields reside are vector bundles and that the spaces that introduce symmetry at a ground level are principal fibre bundles. Principal bundles can still appear in a geometry-first formulation, but if they do, since they are not part of the fundamental ontology, they must be entirely supervenient on the structure of the vector bundles, which is the subvenience basis for the geometry-first formulation. I do *not* demand that the geometry-first formulation should be superior in every respect, or practically advantageous. I will start with a more modest aim: to offer an alternative perspective that clarifies some features of particle theory while omitting symmetry at the ground of the explanatory chain. And, lest I sound too irenic, in Section 5 I will defend my preference for the geometry-first formulation more openly.

[1] In 1912 Einstein wrote to Sommerfeld: "I have come to value greatly the four-dimensional formalism of Minkowski, which I had previously considered unnecessary erudition. In the meantime, I have also become convinced that only this formalism brings out the true essence of the theory." (quoted in (Holton, 1974, p. 263))

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

Needless to say, symmetry is the cornerstone of particle physics. Representations of Lie groups, Casimir invariants, spontaneous symmetry breaking, gauge-fixing: these are the daily bread of the Standard Model. (This much will be obvious to anyone familiar with the field, so I need not belabor the point.) That the associated principal bundles—and with them the explicit appeal to symmetry at the ground of the explanatory chain—might be dispensed with is therefore anything but trivial.

However, a new geometry-first formulation has recently been proposed, in which symmetries are not postulated and principal fibre bundles are unnecessary (Gomes, 2024, 2025a). In the alternative formulation, the symmetry groups are only implicit: they arise as the automorphism groups of vector bundles. The geometry-first formulation is generally available as an alternative only for gauge groups that are linear, and for representations obtained from the fundamental representation (when it is unique) via tensor and direct products, symmetrisation, and other similar operations. Thus the geometry-first formulation demands an alignment between symmetry groups and structure of the vector spaces where matter resides that is not always guaranteed from the symmetry-first perspective, with its largely independent principal fibre bundles and representation spaces. (In Section 5 I will have more to say about the restricted equivalence between geometry-first and symmetry-first formulations).

Even in the cases that admit the two formulations as mathematically equivalent, the geometry-first formulation comes with a significantly different ontology: for the Standard Model of particle physics, it consists of three fundamental vector bundles over spacetime where the various matter fields reside (as sections of tensor products). There is no need for a separate space to encode the principal connections.

Change the formulation, and the explanations change with it. Three examples show how familiar features of particle physics acquire alternative interpretations. 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 fibre bundle (or its spacetime representative A^I_μ), but the covariant derivative D_μ on a vector bundle whose automorphism group corresponds to G—and this remains true even if no vector fields are present to be differentiated: the affine structure of the vector bundle can be dynamical all by itself. In this setting, a particle's quantum numbers become geometric labels: the internal space it inhabits, and the tensor type it has within that space. (These ideas were explained in (Gomes, 2024); so this paper will focus on the next two examples.)

Second, once symmetry groups drop out of the ground level of the explanatory chain, the very notion of 'symmetry-breaking' requires reinterpretation. Vector bosons A^I_μ are replaced by covariant derivatives of the fundamental bundles, which are not on the same footing as matter fields, and it is no longer clear how they could 'acquire mass' in the usual sense. Third, consider the Yukawa couplings. In the symmetry-first formulation, Yukawa terms are scalars formed from sections of different associated bundles, requiring explicit 'bridges' between them. In the geometry-first picture, by contrast, the fundamental objects are vector bundles and particle fields are represented by sections of suitable tensor bundles constructed from the fundamental vector bundles. Scalars then arise naturally through inner products and contractions between vectors and their duals.

Another striking example of the explanatory reach of the geometry-first framework is the quantisation of charge. Where the symmetry-first picture attributes it to the topology of the compact group $U(1)$, the geometry-first picture derives it from the discrete algebraic structure of the tensor powers of the fundamental vector bundles. The explanation is at once more general and more direct: geometry enforces discreteness even in settings where symmetry would permit continuity (i.e. for non-compact structure groups).

These examples show how a geometry-first perspective reshapes explanations. Under Hunt (2025)'s lights then, this reformulation is non-trivial because it suggests a different ontology, offers new explanations, and displays an epistemic difference in how it solves some problems in common with the usual formulation.[3]

But the real strength of the approach lies in the methodological discipline it suggests. Where the principal–bundle picture allows a loose fit, or slack, between symmetry and geometry, the *vector–bundle point of view* (POV) ties the two tightly together. Section 5 argues that this apparent restriction is in fact a virtue: it narrows the space of admissible theories in a way that clarifies the ontology of gauge theory and still encompasses our best physics.

Here is how I will proceed. Section 2 introduces both the familiar principal–bundle formulation, which gives a point-of-view on gauge theories (PFB-POV), and the alternative VB-POV. Section 3 gives the alternative account of the Higgs mechanism. Section 4 does not attempt a full reformulation of the Yukawa mechanism, but argues that its interpretation is more transparent in the geometry-first approach. Section 5 develops the methodological defense of the VB-POV. Finally, Section 6 draws the broader morals.

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

Here I will give brief overviews of both the familiar, symmetry-first, and of the less familiar, geometry-first formulations of gauge theory. I will start with the more familiar and then introduce the novel.

In the PFB-POV I write the principal connection one-form as ϖ; in the VB-POV I write the affine covariant derivative operator as ∇ and the associated-bundle connection as ω.

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

In short, the symmetry-first formulation of the Standard Model is the familiar one: each fundamental interaction is associated with a symmetry group, which is taken as the structure group of a principal fibre bundle. Connections on this bundle then play the role of the vector bosons—the "force carriers."

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

Classical configurations of matter particles charged under a force are described by sections of vector bundles associated to the principal bundle whose group encodes that force. One may endow these associated bundles with additional structure (for instance, a Hermitian inner product on $\mathbb{C}^n$); in such cases, the representations of the structure group are only required to preserve that structure.

The connection on the principal fibre bundle induces parallel transport on all associated bundles. Crucially, it is the *same* connection that governs transport in each case, ensuring that different matter fields charged under the same interaction remain coordinated: they all "march in step" under parallel transport, probing the same distributions of electroweak or strong forces. Thus, while associated vector bundles are distinct entities, they are tied together by the principal bundle, which acts as their common coordinator (see (Weatherall, 2016) and Figure 1). The primacy of the postulated structure group, and the central role it plays, is what makes this a "symmetry-first" formulation.

This is the familiar exposition, but, as we will see, it can be applied more generally. I will say more general symmetry-first formulations of gauge theories in which the symmetries are introduced as part of the principal fibre bundles fall under the *principal bundle point of view* (PFB-POV).

More technical details are provided in Appendix A, for now, it is sufficient to say that, more generally and in mathematical language, 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. Intuitively, such a bundle codifies the ways in which the symmetry group G can act on geometric objects defined over M. In this paper, I will focus on one especially important class of such objects: *vector bundles*. A vector bundle (E, M, V) assigns to each spacetime point $x \in M$ a copy of a fixed vector space V, called the *typical fibre*. Sections of vector bundles are smooth assignments of an element of V to each point of M, and matter fields are precisely such sections.

Generally, given a typical fibre V and a principal bundle (P, M, G), we can define *associated vector bundles* using a representation, i.e. $\rho : G \hookrightarrow GL(V)$ which is determined by G, V, and the particle lables, or quantum numbers.[4] Matter fields are represented as sections of vector bundles associated to the principal bundle whose structure group is the gauge group of the theory (e.g. $SU(3) \times SU(2) \times U(1)$).

The broad idea can be illustrated with the case of a faithful and transitive representation of $G \simeq GL(V)$. We take the points of P to be frames for V, understood as ordered tuples—say $V \simeq \mathbb{C}^n$—over each point of M, so that the group action corresponds to a change of frame on P and a corresponding change of components of a vector in V. Since vector fields are frame-independent, one quotients pairs (p, v), with $v \in V$ and $p \in P$, by the simultaneous action of the group on P and on V (via ρ). The resulting object is the associated vector bundle,

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

[4]Here ρ is understood as an embedding that is, in general, only a partial homomorphism, since it may fail to be faithful or to act transitively. This will be important in Section 5.

$$\begin{tikzcd}
 & E_i := P \times_{\rho_i} V_i, \ , \nabla_i \\
P(M, G, \varpi) \arrow[ur, "\rho_i"] \arrow[r, dotted] \arrow[dr, "\rho_j"'] & \vdots \\
 & E_j := P \times_{\rho_j} V_j, \ , \nabla_j
\end{tikzcd}$$

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.

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

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 (by determining which frame over one point is mapped to which frame at an adjacent point).

For general gauge theories, each associated vector bundle then corresponds to a possible particle type, classified by the representation labels that serve as the particle's quantum numbers. To accommodate more constructions allowed by the PFB-POV, we must refine the broad construction above by introducing additional geometric structure on E.

When introducing the broad idea, we used the fact that, given E, the frame bundle $L(E)$ is a principal fibre bundle (P, M, G) with $G \simeq GL(V)$. More general principal bundles—call them $L_\rho(E)$—arise when we consider *admissible frames* for E, i.e. frames that respect additional structure on E. For instance, if we endow E with a Hermitian metric, we can restrict $L(E)$ to the sub-bundle $L_\rho(E)$ of orthonormal frames and recover (P, M, G) with $G \simeq U(V)$, acting in the fundamental representation. We could also endow each fibre with a complex volume element ε, and if P encodes only the symmetry transformations that preserve this coarser geometric structure, we obtain (P, M, G) with $G \simeq SL(V)$—transformations of unit determinant. Alternatively, we could consider transformations preserving both the Hermitian inner product and the volume form, giving $G = SU(V)$, or those rotating only the volume form, giving $G = U(1)$, and so on.

So we can form (P, M, G) by picking out subgroups of the group of automorphisms of the typical fiber, $\mathsf{Aut}(V)$. These (P, M, G) are still obtained from the bundle of frames $L(E)$ by restriction to admissible frames, but G is not obliged to coincide with $\mathsf{Aut}(V)$. More generally, the structure group of P and the automorphism group of the typical fibre of its associated vector bundles need not be isomorphic—a good thing, as I will argue in Section 5.

Thus, generally, each representation $\rho : G \to GL(V)$ defines an associated vector bundle $E = P \times_\rho V$, whose sections $\psi : M \to E$ represent admissible field configurations for a particle type transforming according to ρ. The fibre V thus represents the space of internal degrees of freedom (e.g. charge, colour, weak isospin, chirality, etc.), and the representation labels play the role of the particle's internal "quantum numbers." Let me repeat the advantage of defining vector bundles in this way: it becomes transparent that fields inhabiting bundles associated to the same principal bundle covary under the action of the symmetry group. But it is not the only transparent way of ensuring this covariance.

A further, related question is whether such vector bundles merely covary under the group action, or whether they in fact stand in a canonical relation to one another. Answering this will be important to establish a relative advantage of the VB-POV, in later Sections.

Thus, 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 ξ_1 a local section of E_1:

$$\xi_1(x) = [\sigma(x), v(x)]_1, \quad v : U \to V. \tag{2.3}$$

Then the obvious map to consider is:[5]

$$\begin{aligned} T : E_1 &\to E_2 \\ \xi_1 := [\sigma(x), v(x)]_1 &\mapsto [\sigma(x), v(x)]_2 =: \xi_2. \end{aligned} \tag{2.4}$$

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 (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 $\rho_1 = \rho_2$.

[5]I thank Jim Weatherall for suggesting this.

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 vector-bundle point of view, all the vector bundles that, in the symmetry-first formulation, would be associated to the same principal bundle, are already endowed with a natural relation, as we will now see. Later, in Section 4, we will see how the issue of relating different vector bundles which in the PFB-POV are associated to the same principal bundle can arise in practice.

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

The geometric perspective I want to develop aims to dispense with the principal fibre bundle altogether. In this Section I set out a formulation of gauge theory that proceeds without gauge potentials, principal bundles, or explicit appeal to gauge symmetries.

The analogy with spacetime helps clarify my aim. Consider (M, g, Ξ_i), where (M, g) is a smooth Lorentzian manifold and the Ξ_i are various tensor fields on M, i.e. objects 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 treated as background structure). This group becomes explicit once we introduce orthonormal frames. Yet much can be said about g and Ξ_i in a purely geometric, frame-independent manner, without any reference to $SO(3,1)$. And if instead we were to posit a different group acting on TM—say $O(2)$ and not $SO(3,1)$—a geometrical rationale would be required to justify that action.

In gauge theory, by contrast, an analogous "frame-free" formulation for the behavior of matter remains largely undeveloped (cf. (Gomes, 2024; Weatherall, 2016)), and the very 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 will introduce a realisation of the geometry-first formulation of a certain class of gauge theories, which I will call *the vector bundle point of view* (VB-POV).[7] To motivate the VB-POV from interpretational issues with the PFB-POV, let me recall that the main role of ϖ in (P, M, G) is to coordinate covariant derivatives between different associated vector bundles. But what is the physical status of ϖ? Jacobs (2023, p. 41) convincingly argues they don't have one; he concludes:

[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.

[7]Other kinds of theories could also have a geometry-first formulation, e.g. those based on Cartan geometry, but here my focus is on particle physics, for which the relevant value spaces are vector bundles.

> 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 an issue and in (Gomes, 2024) I showed that it can be overcome. The introduction of (P, M, G) is unnecessary if particles that interact are all sections of the same vector bundles or of tensor products of the same vector bundles. Tensor products over a vector bundle inherit the same covariant derivatives by construction. In this case, parallel transport of the vector bundles in question automatically march in step. In this case we have at a hand a natural 'common cause' for the coordination of covariant derivatives, without the introduction of principal bundles. This is a realisation of the geometry-first aim, from a VB-POV.

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^*. In more detail, given E a vector bundle with covariant derivative D, 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 $\Gamma(E_1 \otimes E_1 \otimes E_j \otimes E^*_k)$. 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.

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

[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.

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]

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 systematic constructions from E_n itself. Their mutual relations are fully accounted for by the standard functorial machinery of geometry: tensor products, (anti)symmetrisation, dualisation, projections into tensor factors, contractions, interior and inner products, and so forth. 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:

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

where Λ is the anti-symmetric product, with $\Omega \in \Lambda^m(E^{n*})$, 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 for associated vector bundles could be mounted, mirroring the geometric one I have just given. Perhaps given

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

arbitrary different representations of the same group, for arbitrary vector representation spaces, there are systematic ways to relate these representation spaces that mirror the ones I displayed above. That may well be true, but it is beside the point. Even if such arguments exist—and I have not found or worked one out, and am skeptical that one exists (for representations that are not faithful and transitive)—the virtue of the geometric route is that it trades purely on geometrical language, and so it speaks directly to a community trained in geometry rather than in group and representation theory. The mere availability of a geometric formulation that sidesteps representation theory or other more technical algebraic machinery is already a win. My aim, after all, is to broaden the borders of the subject, making it accessible to different habits of thought.

Still, for all its merit, at first pass the VB-POV may seem too narrow to capture the full menagerie of gauge theories employed in physics. Some theories—those built from the exceptional Lie groups, for example—fall outside its reach. And even when a gauge group G is given, it is often a nontrivial matter to "reverse-engineer" a vector space structure for which $\mathsf{Aut}(E_x) \simeq G$. How, for instance, does one coax $SO(4)$ out of a space whose typical fibre is $\mathbb{C}^3$; or $U(1)$ out of a space whose typical fibre is $\mathbb{C}^n$ with $n \neq 1$? I will have much more to say about all this in Section 5.[11]

For all that, the Standard Model of particle physics fits neatly within the VB-POV. In the PFB-POV in which the Standard Model is usually described, every particle field is a section of an associated bundle for a principal fibre bundle whose structure group is $SU(3)\times SU(2)\times U(1)$, and the fundamental representation of each one of these component subgroups ($SU(n)$ or $U(n)$ for appropriate n) appears for some such section or other. The VB-POV alternative is available because under any representation of $U(n)$, the corresponding associated bundles can just as well be constructed by geometric means from the fundamental vector bundle—via tensor and exterior products, (anti)symmetrization, determinants, and the like. The VB-POV alternative is compelling (as I will expand on in Section 5), because the particular combination of representation, groups, and vector spaces is particularly suited for a geometrical interpretation. In such cases, a covariant derivative on a single vector bundle suffices to encode one fundamental interaction, while the various particle fields appear as sections of the appropriate derived bundles (e.g. tensor products).

Having surveyed both approaches to gauge theory—the symmetry-first PFB-POV and the geometry-first VB-POV—I now turn to the Higgs mechanism. My aim is to present it from within the VB-POV, while relegating to Appendix B a sketch of the more familiar PFB-POV treatment, which can be found in any standard textbook.

[11] The Peter–Weyl theorem guarantees that $U(n)$ admits nontrivial representations on $\mathbb{C}^m$, but extracting from this a natural structure on $\mathbb{C}^m$ that renders the action geometrically meaningful is anything but straightforward.

3 The Higgs mechanism in the geometry-first formulation

The proof, they say, is in the eating of the pudding. So here, to prove that the geometry-first perspective embodied by the VB-POV is sufficiently different to the PFB-POV to merit attention, I will provide a stand-alone derivation of the (classical) Higgs mechanism.

In the standard presentation (cf. e.g. (Tong, 2025, Ch. 2), the Higgs mechanism is often described in terms of spontaneous symmetry breaking, and one must employ Goldstone's theorem, gauge fixing (e.g. unitary gauge), etc. I give a brief overview of that presentation in Appendix B. Here I will outline an alternative approach, phrased purely in the geometric language of vector bundles, which makes the essential structure transparent without appeal to symmetry-breaking jargon.

3.1 The non-linearised Higgs field

Let $(E^n, M, \mathbb{C}^n, \langle\cdot,\cdot\rangle_n, \nabla_n)$ be a Hermitian vector bundle over a manifold M, with fibres $E^n_x \simeq \mathbb{C}^n$ and $\langle\cdot,\cdot\rangle_n$ an inner product on E^n, which is compatible with ∇_n, the covariant derivative on E^n. We will omit the subscript when it is understood from context, as it will be in this Section, so for now we take $\varphi \in \Gamma(E)$ (the generalisation to $\varphi \in \Gamma(E^i \otimes \cdots E^j)$ is straightforward, as we will see). So φ is a vector-valued spacetime scalar field, satisfying

$$\min_{x\in M} \|\varphi(x)\| = v', \tag{3.1}$$

for some constant $v' > 0$. We write $\|\varphi(x)\| = (\Delta + v')$, for $\Delta \in C^\infty_+(M)$ (the positive real-valued smooth scalar functions on M), and get

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

where $e_0 = \frac{\varphi}{\|\varphi\|}$ is a unit section, well-defined since $\|\varphi\| > v' > 0$, and $\langle e_0, e_0\rangle = 1$.

All we need to get the qualitative features of the Higgs mechanism, we can extract from the existence of such a norm around which to expand a kinetic term of the Higgs field in the Lagrangian of the theory, i.e.

$$L_{\text{kin}}(\varphi) = \int \langle \nabla\varphi, \nabla\varphi\rangle. \tag{3.3}$$

Note that:

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

where Re takes the real component. Using (3.2) and (3.4) the kinetic term reads

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

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

Thus the covariant derivative—and therefore any connection representing it—appears quadratically in the term $v'^2\langle\nabla e_0, \nabla e_0\rangle$. However, ∇e_0 does not contain all the information in ∇: the components of ∇ that do not enter this term remain 'massless'. Only 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 and remain 'massless'. From (3.4), using $\mathsf{Re}\langle\,.\,,\,.\,\rangle$ we project ∇ along e_0 and its orthogonal complement: $\nabla^{\parallel}$ and $\nabla^{\perp}$, where, as per (3.4), only $\nabla^{\parallel}$ is absent from the kinetic term, and so will remain 'massless'.

Conceptually, this geometric presentation of the Higgs mechanism makes the essential structure transparent. The operator ∇ merely provides the affine structure of the bundle, while the Higgs field singles out an internal direction e_0. The dynamics of that affine structure—once the curvature term in Equation (A.10) is included in the Lagrangian—couple to v' precisely for the components that rotate e_0. With everything expressed entirely in abstract-index (tensorial) language, no symmetry can be 'broken'.

From a more pragmatic, particle-physics standpoint, however, there is no need for scare quotes around 'mass acquisition': the coupling of the linearised connection to a constant in the Lagrangian is indistinguishable from a mass term.

In sum, the geometric formulation already captures the qualitative behaviour of the Higgs mechanism in full generality, without any linearisation. No mention of stabilisers, gauge orbits, or gauge fixing was required—devices indispensable in the standard formulation (see (Hamilton, 2017, Ch. 8.1); (Tong, 2025, Ch. 2.2) for comparison). For instance, the disappearance of Goldstone modes through a choice of gauge is replaced here by the orthogonality relation (3.4). This concludes the classical, non-linearised, qualitative account of the 'mass acquisition' mechanism.[12]

3.2 Mass Generation in the Linearised Theory

Introduce a connection $\nabla = d+\omega$ such that $de_0 = 0$ and $\omega \in \Gamma(T^*M \otimes \mathsf{End}(E))$, where $\mathsf{End}(E)$ are the linear endomorphisms of E; so for $\xi \in \Gamma(E)$, we have $\omega \cdot \xi \in \Gamma(T^*M \otimes E)$.

Now, above I said that having a constant norm around which to expand the Higgs mechanism sufficed to get the qualitative features of the mechanism. But we don't have access to such a minimum via any dynamical feature of the theory, and so instead we involve the Higgs potential, and this will give us quantitative accuracy for the various masses, etc. Now the Higgs Lagrangian reads

$$L_\varphi = \int \langle \nabla\varphi, \nabla\varphi\rangle + V(\|\varphi\|) \tag{3.6}$$

[12]One may reasonably argue that these symmetry concepts—such as gauge-fixing—may be required when we introduce quantum mechanics. Moving to the quantum domain, one must consider the whole configuration space of the Higgs field. Then the analysis applies to a sector $\Gamma_0(E)$ for which one of the configurations has an absolute minimum over all the others (i.e. $\min_{\varphi\in\Gamma_0(E), x\in M} \|\varphi(x)\| = v'$). Here is how far my concession would go (see footnote 10): in a sum over configurations, we use e_0 as the anchor, or 'representational scheme' across physical possibilities; cf; (Gomes, 2025b; Kabel et al., 2025). And indeed, representational schemes can be compared to gauge-fixings (cf. (Gomes, 2025b, Sec. 3.3)). A translation of this idea to the gauge terminology would go as follows: consider $\Gamma(E^2)$, and its sector $\Gamma_0(E^2)$. Let $\varphi, \varphi' \in \Gamma_0(E^2)$. The group $\mathsf{Aut}(E^2)$ 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 from the rest.

We call $v \neq 0$ the minimum of the potential, but it need not coincide with v'. Defining $v' - v =: c$, for v a spacetime-independent (i.e. 'translation-invariant') minimum of the Higgs potential, we rewrite (3.2) as

$$\varphi(x) = (H(x) + v)e_0, \tag{3.7}$$

where $H(x) = \Delta(x) + c$. If we assume that c and Δ are of the same order—which is implied by the norm of the Higgs field not deviating too much from v—since $c = (v' - v) < 0$ and $\Delta(x) > 0$, then $H(x)$ can be both positive or negative, i.e. $H \in C^\infty(M)$.[13] Then from (3.5)

$$\|\nabla\varphi\|^2 = (\partial H)^2 + (H^2 + 2Hv + v^2)\,\|\omega \cdot e_0\|^2, \tag{3.8}$$

where, as usual, the norm of a tensor product factorises linearly, i.e. for each basis element $\lambda \otimes \xi \in \Gamma(T^*M \otimes E)$, we have:

$$\|\lambda \otimes \xi\| := \|\lambda\|_M \|\xi\|_E. \tag{3.9}$$

But to unclutter notation I will omit the subscripts when understood from context.

Further assuming that $\mathcal{O}(H) = \mathcal{O}(\omega) = \varepsilon$,[14] yields

$$\|\nabla\varphi\|^2 = (\partial H)^2 + v^2\|\omega \cdot e_0\|^2 + \mathcal{O}(\varepsilon^3). \tag{3.10}$$

Here we see clearly how the quadratic terms in the connection ω would correspond to vector bosons 'acquiring masses'; again, without invoking unitary gauge or Goldstone's theorem.

But as I said, not all components of ω contribute to $\|\omega{\cdot}e_0\|^2$ in (3.10). In a basis $\{e_I\}_{I=0,\dots,n-1}$ adapted to e_0, we have

$$\nabla e_I = \omega^J{}_I\, e_J, \quad \text{and so} \quad \nabla e_0 = \omega^i_0 e_i, \quad \text{with} \quad i \neq 0, \tag{3.11}$$

from the anti-symmetry of the connection. Then

$$\|\nabla\varphi\|^2 = (\partial H)^2 + v^2 \sum_{i\neq 0} (\omega^i{}_0)^2 + \mathcal{O}(\varepsilon^3). \tag{3.12}$$

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 concludes the geometric derivation of the Higgs mechanism. Let us now see how it reproduces standard results from the familiar or standard approach to gauge theory. The missing ingredient for the comparison is to write the connection ω in terms of preferred representations of the Lie algebras in question. I will start by providing an example (that is indeed isomorphic to $su(2)$) before showing how the usual endpoint of the Higgs mechanism for gauge bosons is recovered.

[13] We could of course have started directly from (3.7), by again assuming that: (i) the potential depended only on the norm of the Higgs field; (ii) that the minimum of the potential was non-zero and spacetime independent; and (iii) that the norm of the Higgs field did not deviate too much from this minimum, in particular, that it was also non-zero everywhere. I find the order of assumptions made in my presentation clearer, because they can be easily stated outside the linearised regime and extract what I think are the essential qualitative features.

[14] In the comparative sense: that $\frac{|H|}{v} \sim \varepsilon << 1$, and *mutatis mutandis* for the appropriate norm on ω.

3.2.a Example: (M^3, g)

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

$$\omega = \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix}. \tag{3.13}$$

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

$$\omega \cdot e_0 = \begin{pmatrix} 0 \\ \omega_z \\ -\omega_y \end{pmatrix}. \tag{3.14}$$

Thus, we would get:

$$\|\nabla\varphi\|^2 = v^2(\omega_y^2 + \omega_z^2). \tag{3.15}$$

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

3.2.b Electroweak Example: $\mathbb{C}^2 \otimes \mathbb{C}^1$

The covariant derivative on an element $\mathbf{v} \otimes \mathbf{w} \in V \otimes W$ is given by

$$\nabla(\mathbf{v} \otimes \mathbf{w}) = (\nabla^V \mathbf{v}) \otimes \mathbf{w} + \mathbf{v} \otimes \nabla^W \mathbf{w}, \tag{3.16}$$

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:

$$\nabla e_0 = \omega \cdot e_0^2 + e_0^2 Z = (\omega + iZ\mathbb{1})e_0^2, \tag{3.17}$$

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:[15]

$$\omega = \begin{pmatrix} iW_3 & iW_1 - W_2 \\ iW_1 + W_2 & -iW_3 \end{pmatrix}, \text{and} \quad iZ\mathbb{1} = \begin{pmatrix} iZ & 0 \\ 0 & iZ \end{pmatrix}. \tag{3.18}$$

[15] 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 $(0,1), (0,i), (1,0), (i,0)$, for the inner product $\mathrm{Re}\langle\cdot,\cdot\rangle$, which is effectively what appears in Lagrangians, due to the use of the complex conjugate terms, cf. (Hamilton, 2017, Ch. 8). Here we are attempting to make contact with the standard notation and formalism and so are using its conventions.

Applying this to e_0^2 in (3.17) gives

$$\nabla e_0 = \begin{pmatrix} iW_1 - W_2 \\ -iW_3 + iZ \end{pmatrix}. \tag{3.19}$$

Hence the corresponding quadratic term appearing in (3.10) is

$$\|\nabla e_0\|^2 = W_1^2 + W_2^2 + (Z - W_3)^2. \tag{3.20}$$

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.

4 The Yukawa mechanism

Whereas the Higgs mechanism is used to 'endow mass' to the gauge potentials, the *Yukawa mechanism* is used to endow mass to the matter fields—here we needn't use scare-quotes!

In the Standard Model fermion masses cannot be introduced as they can for real or complex-valued scalar fields. First of all, a Dirac mass term must couple left- and right-handed chiral fermions; moreover, the two chiralities are mapped into internal spaces that transform differently under the gauge group $G = SU(3) \times SU(2) \times U(1)$, so coupling them would violate gauge invariance: this is related to the issue we saw in Section 2.1 about canonical isomorphisms between associated vector bundles with different representations. The solution is to introduce the Higgs field ϕ, in such a way that gauge invariance is preserved, while the fermions acquire effective masses. This is the *Yukawa mechanism*.

Here I will essentially follow the treatment given in (Hamilton, 2017, Ch. 8), whose notation and general approach is already much closer to the geometric approach that I'm pursuing here (as compared to the treatment of more familiar textbooks, for instance, the one given in (Skinner, 2007; Tong, 2025), which use representation theory more heavily). So I will call the treatment to be followed here 'the standard' treatment of the Yukawa mechanism. In Section 4.1 I will describe the obstruction to the formulation of mass terms for fermions, and its resolution in this, geometric-friendly but still 'standard', exposition. Then in Section 4.2 I will discuss what I think is explanatorily unsatisfactory about this resolution, and say why I take the VB-POV to provide a more transparent explanation.

4.1 The 'standard' presentation of the Yukawa mechanism

In more detail, here is the obstruction to the formulation of mass terms for fermions. Fermions are spinors, but for Weyl spinors, the inner product is anti-diagonal in the left and right basis: $\overline{\psi}_R\psi_R = 0$, and so, in order to extract mass terms we must couple left to right-handed spinors:

$\overline{\psi}_R\psi_L$. Thus, if both ψ_L and ψ_R are valued in the same internal space, i.e. in the same vector bundle, and are in the same representation, one may add mass terms of the form:

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

and this will be gauge invariant since ψ_L and ψ_R transform in the same representation of the gauge group. I.e. locally, $\psi_L \in \Gamma(S_L \otimes E)$, $\psi_R \in \Gamma(S_R \otimes E)$, where (E, M, V) is the vector bundle with the representation space V of the gauge group in question, and S_L is the bundle of left-handed spinors over spacetime, whose typical fibre space is called Δ_L (*mutatis mutandis* for right-handed spinors).

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, $V_L \not\simeq V_R$. Thus a bilinear such as $\overline{e}_L e_R$ 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.)

Moreover, for V_R, V_L irreducible, unitary, non-isomorphic representations of G, mass pairings, defined as G-invariant maps, $\kappa : V_L \times V_R \to \mathbb{C}$, 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. What do these maps look like, more precisely? Let us look at an example. Consider the $SU(2) \times U(1)$ representations for the leptons (taken 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.

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 $Q = T_3 + \frac{1}{2}Y$.

The map τ is defined on vector spaces, and therefore depends on the choice of trivialisation of the vector bundles, i.e. on local decompositions of the form $E_{|U} \simeq U \times V$. To render this construction invariant, we must extend τ to sections of the associated vector bundles. This is straightforward. Given a section $\sigma(x)$ of an $SU(2) \times U(1)$ principal bundle (cf. (2.3)), we can use the local maps $l_L : U \to V_L$, $\varphi : U \to W$, and $l_R : U \to V_R$ introduced above to form the corresponding global sections $e_L \in \Gamma(S_L \otimes E_L)$, $\varphi \in \Gamma(F)$, and $e_R \in \Gamma(S_R \otimes E_R)$, which are independent of any choice of trivialisation. For instance, a left-handed electron (I will discuss its 'up' and 'down' components shortly) is represented as

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

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 possible invariant maps that achieve the same result, 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.[16] The crude impression is that in the PFB-POV one manufactures invariants in a chosen basis while absorbing ambiguities into coupling constants.

Eliminating such extraneous factors is a matter of conceptual housekeeping: the geometry-first perspective accomplishes this by treating the Yukawa couplings not as equivariant maps on the underlying representation spaces, but as natural operations between the relevant vector bundles, e.g. as given in Equation (2.2) in Section 2.8.

4.2 The VB-POV presentation of the Yukawa mechanism

In Section 2.1 I argued that there was no canonical relation between associated vector bundles corresponding to different representations of the principal bundle; but of course we don't *need* such a map to get a scalar out of sections of these bundles. All we need is that T, given in (4.8), is *a* gauge-invariant map between associated vector bundles, with τ a map between the representation spaces; and presenting one such map is sufficient for comparison with experiments. Nonetheless, lacking the canonical relation, I find this answer unsatisfactory, because opaque: why this particular map? Couldn't we have found others? And how should we interpret them?

I take the geometric, VB-POV, to provide a more transparent interpretation of what the map T represents, and what other choices would represent. Again, in the geometry-first formulation, all we have are structures in the fundamental vector bundle spaces. The fundamental vector spaces are given by $(E^n, M, \mathbb{C}^n, \langle\cdot,\cdot\rangle_n)$, for $n = 1, 2, 3$ (we will include orientation as further structure below, when we look at the Yukawa form for quarks). Different particles are merely different sections of different tensor products for these fundamental vector spaces. We replace 'quantum numbers' by a geometric characterisation of a given particle. Thus, for instance, a down-right-handed quark (of any of the three generations, but here we assume the first) is given by:

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

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).

[16] For a single Higgs doublet with SM hypercharges, the residual $C^\times$ freedom in the centraliser can be absorbed by field rephasings and a Yukawa reparametrisation, hence is physically inessential. In multi-Higgs or extended-gauge settings, the centraliser may act non-trivially on distinct invariant maps, leaving physically inequivalent couplings.

In this formulation, weak isospin T_3—defined only with respect to a chosen basis of the Lie algebra—has no independent geometrical meaning (see 3.2.b and footnote 15). Accordingly, left-handed fermions, together with φ^+ and φ^0 (also $SU(2)$-doublets), are best understood simply as components of the vector fields $\mathbf{Q}_L$ and $\mathbf{L}_L$ (where I have introduced boldface to emphasise their vectorial character). The familiar distinction between, say, the electron and the electron-neutrino, or between the up- and down-left quarks, does not arise at this level: it appears only through their couplings with the Higgs. The Higgs field $\boldsymbol{\varphi}$ provides a frame within $\mathbb{C}^2$ that endows T_3—and hence these component fields—with physical significance. The charges listed in Table 1 are already adapted to this frame, since they presuppose the choice $\varphi = (0, \varphi^0)^T$ (i.e. $\varphi^+ = 0$); for example, only in that frame do the left-handed up-quark components appear as $(u_L^I, 0)^T$.[17]

Thus, geometrically, it makes more sense to define the left-handed components of both leptons and quarks as parallel and orthogonal to the Higgs according to the inner product on E^2, i.e.:

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

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

where capital I indicates color components (i.e. red, green and blue; but note that $\mathbf{Q}_L^I$ is still a vector in E^2) in an orthonormal frame of $\mathbb{C}^3$ and I used the notation e_0 (without boldface, to avoid confusion with the electron) for the unit-direction of the Higgs, introduced in Section 3.1 (not to be confused with the left-handed electron, $\mathbf{e}_L$).[18]

Before we give the geometric interpretation of (4.8), and of the corresponding form for quarks, note that, given an orthonormal basis for E^2, we can form duals: for $\xi = \xi^\perp e_\perp + \xi^\parallel e_0 = (\xi^\perp, \xi^\parallel)^T$ (e.g. $\mathbf{e}_L = \mathbf{L}_L^\parallel$, $\boldsymbol{\nu}_{eL} = \mathbf{L}_L^\perp$) the dual takes the conjugate of the transpose, so:

$$((\xi^\perp, \xi^\parallel)^T)^* = (\bar{\xi}^\perp, \bar{\xi}^\parallel). \tag{4.13}$$

Moreover, given Hermitian bundles E, F, we use $\langle \cdot, \cdot \rangle_E$, $\langle \cdot, \cdot \rangle_F$ for fibrewise pairings, and write iterated contractions as $\langle \langle \, . \, , \, . \, \rangle_E, \, . \, \rangle_F$, with obvious extensions to tensor products.

Now using (4.13) and an orthonormal frame aligned with the Higgs (3.7), the Yukawa term for the leptons in Equation (4.8) now can be stated directly using (trivialisation-independent) sections of the vector bundles, without the need to involve the sections σ of the principal bundles, and reads (including a coupling constant, g_e):

$$T(\mathbf{L}_L, \boldsymbol{\varphi}, e_R) = g_e \langle \langle \mathbf{L}_L, \boldsymbol{\varphi} \rangle_2, e_R \rangle_1 = g_e (v + H) \bar{e}_L e_R, \tag{4.14}$$

where the first equality gives the 'standard' definition; e_R and e_L are, right and left-handed (resp.) Weyl spinor but internal scalars (e_L is the magnitude of $\mathbf{L}_L$ along the Higgs, as in

[17] This explains why Table 1, reproduced from (Hamilton, 2017, Table 8.2), can be misleading: if both components of the Higgs are retained, the up and down components of the left-handed quarks and leptons do not yet have any physical meaning.

[18] Not many textbooks that I have encountered emphasise this point—(Tong, 2025) is an exception. And none describe it geometrically as I did here.

(4.11)); $\langle .\,,\,.\rangle_2$ is complex anti-linear in the first entry acting on the relevant spaces as:

$$\langle .\,,\,.\rangle_2 : E^2 \otimes (\otimes^3 E^{1*}) \times E^2 \otimes (\otimes^3 E^1) \to (\otimes^6 E^1). \tag{4.15}$$

And $\langle .\,,\,.\rangle_1$ is just the scalar inner product in $\mathbb{C}$.[19] From (4.14) we can see how mass terms, proportional to $g_e v$ (as well as interactions with the Higgs field) emerge for the electron.

Geometrically, the inner products in (4.14) are a very natural way to obtain scalars: we are measuring 'internal angles' between the different particles seen as vector fields on the same spaces. Thus I take this form of (4.14), namely $\langle\langle \mathbf{L}_L, \boldsymbol{\varphi}\rangle_2, e_R\rangle_1$, to be a more transparent interpretation of the Yukawa term for leptons. Once one determines the inner product, there is no geometric justification to scale each such term by a complex number (or, more generally, to find alternative maps involving the group centraliser).

Chirality shows up in the fact that only left-handed particles have components in E^2; right-handed particles do not. e_R couples to the $\mathbb{C}^1$ components of $\mathbf{L}_L$ and $\boldsymbol{\varphi}$. Neutrinos don't acquire mass, but not because they are orthogonal to the Higgs—which they are—but because we have not included right-handed neutrinos in our particle content. Had we included right-handed neutrinos, they would acquire mass by coupling to the 'symplectic dual' of the neutrino; this is what happens in the case of quarks.[20]

This 'symplectic dual', called φ_c on Table 1, is obtained by recruiting another geometric structure that we can equip $\mathbb{C}^2$ with (besides the Hermitean inner product): an orientation. This implies we can use the totally anti-symmetric form, or the volume form, ϵ_{ab}, as part of the geometrical structure. In other words, whereas the Higgs mechanism, described in Section 3, used the structure $(E^2, M, \mathbb{C}^2, \langle\cdot,\cdot\rangle_2)$, here we extend that to the structure $(E^2, M, \mathbb{C}^2, \langle\cdot,\cdot\rangle_2, \epsilon)$.[21]

Now, besides the metric, we can use ϵ_{ab} and its inverse ϵ^{ab} to raise or lower indices.[22] Thus if we call the isomorphism $J : E^2 \to E^{2*}$ which acts as $\xi \mapsto \langle\xi,\cdot\rangle$ we have:

$$C := \epsilon^\sharp \circ J : E^2 \mapsto E^2 \tag{4.16}$$

$$\xi^a \mapsto \epsilon^{ac} h_{cb} \xi^b \tag{4.17}$$

where we used, in abstract index notation, h_{ab} as the inner product on E^2. Thus we call

$$\boldsymbol{\varphi}_c := C(\boldsymbol{\varphi}); \tag{4.18}$$

[19]This is slightly misleading: what we have here is that $\varphi \in \Gamma(E^2 \otimes^3 E^1)$, i.e. the third tensor product of E^1, which is still one-dimensional, $e_L^* \in \Gamma(E^{2*} \otimes^3 E^{1*})$, and $e_R \in \Gamma(\otimes^6 E^1)$. This is why they match to a scalar.

[20]Because of this feature, the Yukawa terms for leptons are diagonal in generations: these mass terms don't mix, say electrons with muons and taus.

[21]Under $A \in U(2)$, ϵ_{ab} is taken to transform as $\epsilon_{ab} \mapsto \det(A)\epsilon_{ab}$. So $SU(2)$ preserves it. Moreover, since $AA^\dagger = \mathbb{1}$ for any $A \in U(n)$, we know that $\det(A)\det(A^\dagger) = |\det(A)| = 1$, so $\det(A) = e^{i\theta}$ denotes an orientation change the $\mathbb{C}^n$. Using ϵ_{ab} as a geometric datum then implies we have a fixed orientation, as well as an inner product, on $\mathbb{C}^2$.

[22]Indeed, in standard differential geometry, we can find a similar sort of operator acting on two dimensions: the Hodge star: which would take a basis $e_0, e_1 \mapsto -e_1, e_0$, respectively, so *its action on vectors* can be written in this frame as a matrix operator: $* = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$, which is of the same form as ϵ_{ab}.

it can be seen as a measure on the 'areas' orthogonal to $\boldsymbol{\varphi}$.which is why I called it a 'symplectic dual'.

Denoting the generation by an index $i = 1, 2, 3$, we then have, for the total Yukawa coupling term for quarks:[23]

$$T(\mathbf{Q}_L, \boldsymbol{\varphi}, \boldsymbol{d}_R) = Y^d_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{d}^j_R\rangle_3, \boldsymbol{\varphi}\rangle_2\rangle_1 + Y^u_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{u}^j_R\rangle_3, \boldsymbol{\varphi}_c\rangle_2\rangle_1. \tag{4.19}$$

The boldface on lowercase letters is used to indicate that these are vector fields. This is the geometric form of the 'standard' definition (cf. (Hamilton, 2017, Lemma 8.8.4)); it is, in the VB-POV, what really counts.

Nonetheless, as often is the case in physics, we can extract more information by introducing a frame: here, once more it is convenient, in order to compare with standard presentations, to choose the orthonormal frame (3.7) for the Higgs, which gives the components for the quarks along and orthogonal to the Higgs (given in Equation (4.11)) as in Table 1, as well as $\varphi^+ = 0$. Then:

$$Y^d_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{d}^j_R\rangle_3, \boldsymbol{\varphi}\rangle_2\rangle_1 + Y^u_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{u}^j_R\rangle_3, \boldsymbol{\varphi}_c\rangle_2\rangle_1 = (H + v)\left(Y^d_{ij} d^{Ii}_L d^{Ij}_R + Y^u_{ij} u^{Ii}_L u^{Ij}_R\right), \tag{4.20}$$

where now all variables are scalar (and we are summing over the color indices, I, as well as over the generations i, j).

The Yukawa matrices Y are generically non-diagonal, i.e. they mix generations of quarks. One can always find linear combinations of quarks such that, say, Y^u is diagonal; this defines what is called the *mass basis*. But Y^u and Y^d cannot be diagonalised simultaneously, and the residual mixing is encoded in the *Cabibbo–Kobayashi–Maskawa* (CKM) matrix. Most textbooks (cf. (Hamilton, 2017, p. 515)) then explain that the CKM matrix describes the physical effects of left-handed quark mixing across generations, from the 'mass eigenstate basis' to the 'weak eigenstate basis' (the latter being the one we have used here). It then "follows that the interactions with the W-bosons can connect quarks from different generations if the CKM matrix is not diagonal" (ibid).

One might have thought that, from the geometry-first perspective, the three generations of quarks could not coherently mix. After all, if different generations correspond to sections associated with different masses, they might seem to inhabit distinct vector bundles, and sections of distinct bundles cannot be linearly combined. Yet this worry dissolves once we note that the generations do not correspond to different bundles at all, but to different *sections of the same composite bundle*. That is, all the three generations live in a direct sum, e.g.

$$(E^1 \otimes E^2 \otimes E^3) \oplus (E^1 \otimes E^2 \otimes E^3) \oplus (E^1 \otimes E^2 \otimes E^3). \tag{4.21}$$

So each quark field is a section of a tensor product of the fundamental bundles, together with a trivial "generation" factor,

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

[23]Note that here, unlike for the leptons and the left-handed quarks, the up and down right-handed quarks are genuinely different particles, since they have different components in E^1.

and likewise for u_R^i and d_R^i. We can use the linear structure of the fundamental bundles to form linear combinations of the three generations of quark. The Yukawa couplings then act as endomorphisms on this generation factor,

$$Y_u,\, Y_d \in \Gamma(\mathrm{End}(\mathbb{C}^3_{\text{gen}})).$$

Diagonalising Y_u or Y_d corresponds simply to choosing an orthonormal frame in $\mathbb{C}^3_{\text{gen}}$, i.e. passing to the *mass basis*. But since Y_u and Y_d cannot be diagonalised simultaneously, their relative rotation

$$V_{\text{CKM}} = U_u^\dagger U_d$$

appears as a unitary automorphism of the generation fibre. Thus, from the VB-POV, the Cabibbo–Kobayashi–Maskawa matrix is nothing but a geometric rotation within the trivial generation bundle, fully compatible with the underlying fibre structure.

In the geometric perspective one aspect of the situation is more transparent. If the up and down left-handed quarks were truly independent particles—i.e. distinct fields rather than components of the same field (usually called a doublet) in E^2—we *could* diagonalise Y_u and Y_d separately. But because they are components of the same E^2-field, and because these components are coupled to different fields (e.g. φ, φ_c), we can't. Correspondingly, the W bosons represent ∇^2, the covariant derivative on E^2, and so they, too, naturally mix generations when they couple to the relevant currents.[24]

5 A defense of the geometry-first formulation

To motivate the methodological defense of the VB-POV, let me recall an illustrative example from the previous Section. The Yukawa coupling for quarks depends essentially on the orientation of $\mathbb{C}^2$: it requires the introduction of $\boldsymbol{\phi}_c$, which encodes the oriented area orthogonal to the Higgs. This shows how, in the VB-POV, the reduction from $U(2)$ to $SU(2)$ arises directly from subvening geometric structures. The broader moral is already clear: the VB-POV ties the existence of symmetry groups much more tightly to the underlying geometry than the symmetry-first formulation requires. And while this may look like a liability of the VB-POV, it in fact strengthens its case over the PFB-POV.

This Section develops that defense in four stages. First, in Section 5.1, I show how the symmetry-first, or PFB, formulation permits a looseness between geometry and symmetry—

[24] But other aspects of the Yukawa mechanism remain, to my eyes, mysterious from the VB-POV. For instance, it is a little disappointing that, unlike their left-handed counterparts, up and down right-handed quarks can't be straightforwardly understood as components of a single vector field, due to their different components in $\mathbb{C}^1$. If they could be so understood, in place of (4.19), we would have the simpler:

$$T(\mathbf{Q}_L, \boldsymbol{\varphi}, \mathbf{Q}_R) = Y^d_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{Q}^j_R\rangle_3, \boldsymbol{\varphi}\rangle_2\rangle_1 + Y^u_{ij}\langle\langle\langle \mathbf{Q}^i_L, \mathbf{Q}^j_R\rangle_3, \boldsymbol{\varphi}_c\rangle_2\rangle_1, \tag{4.22}$$

which only takes the components of the same inner product along and orthogonal to the Higgs. But the obstruction is the hypercharge split between u_R and d_R; it is this assignment, not the VB-POV machinery, that prevents a single right-handed multiplet with a unified coupling.

a slack that can manifest even in perfectly ordinary gauge-theoretic settings. I analyse the conditions under which symmetry groups fail to capture the full geometry of the fibres and illustrate this failure through three representative examples that typify the range of possibilities allowed by the PFB-POV. Second, in Section 5.2, I exhibit two cases of groups for which the VB-POV could not possibly recover the full flexibility of the PFB-POV. One of these cases shows that quantisation of charge is more fundamental in the VB-POV. Third, in Section 5.3, I turn to composite fibres such as those appearing in the Standard Model, where different factors of the gauge group act on different tensor components—sometimes trivially. Here the slack of the PFB-POV becomes manifest: the symmetry group can no longer be recovered from the geometry of any single associated bundle. The VB-POV restores coherence by grounding each gauge factor in the automorphism group of its own fundamental vector bundle. Finally, in Section 5.4, I turn to the methodological upshot: the VB-POV's refusal to countenance such slack is not a constraint to be lamented but a virtue to be embraced, since it aligns the theory's ontology with its genuine explanatory structure.

5.1 Slack between symmetry and geometry: examples and lessons

The first step is to recall how the two points of view differ in their basic ingredients. In the VB-POV, one can *only* define groups G as isomorphic to $\mathsf{Aut}(V)$ (of course, one is under no obligation to invoke this group; it is not part of the ground level of explanation). By contrast, in the PFB-POV there are in principle three separate ingredients: not only G (or P) and V, but also the representation ρ. P is just a manifold admitting the free and proper action of a Lie group, and an associated vector bundle can be built by using any representation of that group on any vector space V. Thought in this way, why should the geometry of V reflect the group that acts on P?

Here again the spacetime analogy is instructive. The automorphism group of the tangent bundle equipped with a Lorentzian metric is $SO(3,1)$ (or $O(3,1)$), a fact that becomes explicit once orthonormal frames are introduced. Yet most of spacetime geometry can be developed without ever invoking $SO(3,1)$; this is precisely why it serves as an analogue for the geometry-first VB-POV of gauge theory. And in the spacetime case, if one were to posit some other group acting on TM—say $O(2)$ or $SU(n)$—a clear geometric rationale would be required: which feature of TM could prompt one to consider such an action?

The situation is different in the symmetry-first approach to gauge theory. In general, one posits a vector space V, a group G, and an action ρ, without requiring that G transparently reflect the structure of V. In the equivalence relation that defines associated bundles, (2.1), g is an element of the structure group of the PFB, but $\rho(G)$ need not coincide with the automorphism group of the typical fibre V. The only requirement is that the gauge group preserve the structure of the vector spaces:

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

To sum up the contrast: in the spacetime case we have an identification,

$$SO(3,1) = G \simeq \rho(G) \simeq \mathsf{Aut}(T_xM), \tag{5.2}$$

whereas in general gauge theories the corresponding condition,

$$G \simeq \rho(G) \simeq \mathsf{Aut}(V), \tag{5.3}$$

generally fails. But it is upheld when the symmetry group is reconstructible from the geometry, as it would have to be in the VB-POV. Consistently, from the PFB-POV, (5.3) holds when G has a unique action on V via the fundamental representation.

Now, $G \simeq \rho(G)$ means the action of the group is faithful. Faithfulness alone, however, is not sufficient for the group to capture the full geometry of the fibre. To encode that geometry, the action must also be *transitive* on each fibre: it should be able to carry any element of a given fibre to any other.[25]

When $\mathsf{Aut}(V)$ acts transitively—as it does in the cases of interest here—the condition that $\rho(G)$ encode the geometry can equivalently be expressed by demanding that $\rho(G)$ be *surjective* onto $\mathsf{Aut}(V)$. Only when both conditions are met—faithfulness and transitivity—does the group G fully reflect the fibre's geometry.[26]

The theme of this Section is that the PFB-POV only loosely links symmetry and geometry: neither one determines the other; both conditions of (5.3) can independently fail. Let us now make this looseness concrete with three perfectly ordinary PFB-POV examples: $U(1)$ acting on $\mathbb{C}^3$ by scalar multiples of the identity; $SO(4)$ acting on $\mathbb{C}^2$ via a spinor map; and the trivial action of $SU(n)$ on $\mathbb{C}^m$. Each case packs a lesson.

Example 1 (faithful but not transitive). The condition $G \simeq \rho(G)$ holds iff ρ is *faithful* (injective), so that only $\mathbf{1}$ acts trivially. This is satisfied when the gauge group is $U(1)$, the fibres are $\mathbb{C}^3$ with Hermitian inner product, and the representation is $\rho_y(\theta) = e^{iy\theta}\mathbf{1}$ with $y \in \mathbb{N}$. But $U(1)$ is clearly not isomorphic to the full automorphism group of $(\mathbb{C}^3, \langle\, .\, ,\, .\, \rangle)$, which is $U(3)$: 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 $\mathbb{C}^3$.[27] Thus in this case we have

$$G \simeq \rho(G) \subset \mathsf{Aut}(V), \qquad \dim(G) < \dim(\mathsf{Aut}(V)). \tag{5.4}$$

A similar situation arises in ordinary differential geometry, with the so-called *surfaces of revolution*. These surfaces possess rotational symmetry, yet that symmetry does not determine

[25] On the broader question of recovering the geometry of a vector space from the action of subgroups of $GL(V)$, or, more broadly, the geometry of a space from the action of pseudo-groups, see (Barrett & Manchak, 2024; Wallace, 2019) .

[26] So, I will assume here that when (5.3) holds—e.g. for $G \simeq O(4)$ acting on $V \simeq \mathbb{R}^4$ with the Euclidean inner product via the fundamental representation—one can reconstruct (P, G, M) from $L_\rho(E)$ (the bundle of admissible frames for E, see Section 2.1).

[27] 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 21.

their full geometry—think of the diversity of shapes (or vases) that share the same axis of rotation. Lesson: even with a faithful representation, the full geometry of the fibre may not be recovered from $\rho(G)$ alone.

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

$$\rho(G) = \mathbf{1} \subset \mathsf{Aut}(V), \quad \dim(G) \leq \dim(\mathsf{Aut}(V)) \text{ if } m \leq n, \quad \dim(G) \geq \dim(\mathsf{Aut}(V)) \text{ if } m \geq n. \tag{5.5}$$

Even if $n = m$, one cannot reconstruct G from its representation on V, since the action is trivial. Thus we have

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

so both conditions in (5.3) 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. Lesson: with a trivial representation, the fibre carries no information about G, and G imposes no structure on the fibre.

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

$$G \not\simeq \rho(G) \simeq \mathsf{Aut}(V), \quad \dim(G) > \dim(\mathsf{Aut}(V)). \tag{5.7}$$

The image of the representation matches $\mathsf{Aut}(V)$, yet the full group G is strictly larger. Lesson: even when $\rho(G) \simeq \mathsf{Aut}(V)$, the embedding of G may lack a clear geometric rationale.

Let us sum up the examples. In the $U(1)$ case on $\mathbb{C}^3$ (Eq. (5.4)), the representation was faithful, so $G \simeq \rho(G)$. Though the full geometric data of the fibre could not be recovered, it was at least plausible that the group could be grounded in appropriate geometric structures on the fibre, as we saw. By contrast, in the cases of $SO(4)$ acting on $\mathbb{C}^2$ and $SU(n)$ acting trivially on $\mathbb{C}^m$, the representation is not faithful: $G \not\simeq \rho(G)$, so no unique reconstruction of the group from the fibre is possible, even in principle.

What we can abstract from the examples is that, even when $\mathsf{Aut}(V)$ acts transitively, the two necessary conditions for the group to reflect the geometry of the fibre, as expressed in Equation (5.3) can fail independently. We may have $G \not\simeq \rho(G)$, or $\rho(G) \subset \mathsf{Aut}(V)$ (or, more generally, we could have $\rho(G)$ fail to act transitively). And even when $G \simeq \mathsf{Aut}(V)$ with $\mathsf{Aut}(V)$ acting transitively, the link between $\mathsf{Aut}(V)$ and G—namely $\rho(G)$—may still fail. These cases illustrate the PFB-POV's tolerance for symmetry groups that exceed what can be informed by geometry, or fall short of fixing the geometry.

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. In the cases considered there, equivalence between the two formulations fails because the symmetry-first framework tolerates degenerate or unfaithful actions. There are, however, other kinds of obstruction—cases in which, irrespective of choices of representation, no choice of geometric 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

$$\rho_n(e^{i\theta}) = e^{in\theta}, \qquad n \in \mathbb{Z}, \tag{5.8}$$

with integer labels enforced by periodicity: since $e^{i(\theta+2\pi)} = e^{i\theta}$, any continuous homomorphism $\rho : U(1) \to \mathbb{C}^\times$ must satisfy $e^{i2\pi n} = 1$. If, however, one replaces $U(1)$ by its universal cover $\mathbb{R}$, the general continuous characters into $\mathbb{C}^\times$ are

$$\rho_\lambda(t) = e^{\lambda t}, \qquad \lambda \in \mathbb{C}, \tag{5.9}$$

so there is a continuum of representations (including varying modulus when $\Re\lambda \neq 0$). Restricting to unitary characters forces $\lambda = iq$ with $q \in \mathbb{R}$, but without that extra structure there is no quantisation at all. Thus, in the symmetry-first framework, discreteness arises only from a topological identification built into the choice of group, not from the bare formalism.

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 $V \simeq \mathbb{C}$ carries no additional metric or orientation, so that $\mathsf{Aut}(V) \simeq \mathbb{C}^\times$. All further fields are 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 $\mathsf{Aut}(V)$ admits continuously many characters, the geometry-first formalism can reproduce only those obtained by finite tensor operations. Quantisation of charge thus follows not from topology or compactness but from the internal combinatorics of the tensor algebra.[28]

(b) Exceptional Lie groups. A different obstacle to equivalence arises with the exceptional Lie groups. 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

$$G = \mathsf{Aut}(V, \text{data}) \subset GL(V). \tag{5.10}$$

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 tenuous.

[28] See (Derdzinski, 1992, Ch. 4.13) for a similar take on the quantization of charge.

Groups such as F_4 and E_6 can be represented as automorphism or determinant-preserving groups of exceptional algebras—F_4 as $\mathrm{Aut}(J_3(\mathbb{O}))$, the automorphism group of the 27-dimensional real vector space of Hermitian 3×3 octonionic matrices, and E_6 as the group preserving the cubic norm on that same space. But these realisations are not unique: distinct and equally natural geometric structures, sometimes on spaces of different dimensions, can yield the same abstract group. Many of the exceptional groups admit no unique "fundamental" representation in the sense required by the VB-POV: E_6 possesses two inequivalent complex 27-dimensional representations, E_7 has several minimal ones (of dimensions 56 and 133), and E_8 has none smaller *than the adjoint* 248. The choice of fibre V is therefore underdetermined even when the group is fixed.

5.3 Composite fibres and the Standard Model

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

$$V = V_1 \otimes V_2, \tag{5.11}$$

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 (5.3) 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. (5.6).

From the PFB-POV this is no defect: it is essential to the standard strategy of assigning 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. But if one nonetheless tries to recover the group from the automorphisms (of the associated bundles corresponding to each particle's labels) themselves, not only would one fail to get the right verdict; in general, the recovery would not even be consistent.

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.

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 $SU(3)$). Given the collection

$$(P, M, G, \{\rho_i\}_i, \{V_i\}_i), \quad i = 1, 2 \tag{5.12}$$

we can try to reconstruct (P, G, M) from each $L_{\rho_i}(E_i)$ (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. 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. Naively, taking the product of automorphisms groups of the associated vector bundles won't work either: by merely adding another particle whose representation labels included, say, **2** for $SU(2)$, we would get repetition, e.g. $SU(2) \times SU(2)$.

The VB-POV avoids this difficulty by shifting the level at which geometry fixes symmetry. Each gauge-group *factor* must arise as the automorphism group of a *fundamental* (or atomic) vector bundle, not of the composite associated bundles corresponding to particle multiplets with varying representation labels, or quantum numbers. These fundamental bundles serve as the basic building blocks: from them, one can reconstruct each factor ($SU(3)$, $SU(2)$, etc.) and then deduce their actions on the tensor products and direct sums that, in the VB-POV, characterise the various particles. At that composite stage, however, the group is no longer straightforwardly recoverable from geometry.

By contrast, the symmetry-first PFB-POV begins by postulating a group G and constructing associated bundles for each particle, often combining several of the VB-POV's geometric building blocks at once. Since it has no mandate to recover symmetry from geometry, the PFB-POV naturally allows a loose fit, or slack, between the structure group G of P and the geometry of the associated vector bundles V that represent the particles. This slack poses no internal difficulty, for the PFB-POV treats the two ingredients—the symmetry and the geometry—as jointly posited rather than mutually constraining.[29] Fortuitously, however, in our world the symmetry group realised in Nature happens to fit the VB-POV's stricter account of symmetry quite snugly. It need not have done so.

5.4 Methodological closure: why tight identification is a virtue

To close this Section, let me underscore this last point with a simple case. Above I gave several examples where the VB-POV would fail. It is easy to imagine such a world, admissible under the PFB-POV, in which there is a single kind of particle: a section of E (with typical fibre V), together with a principal bundle (P, M, G) where $\dim(G) < \dim(\mathsf{Aut}(V))$. In the PFB-POV, G has physical significance and may differ from the automorphism group relevant to the matter fields. For a concrete instance, take any example from the previous Section. If you want to remain closer to current physics, suppose that $V \simeq (\mathbb{C}^n, \langle\, . \,, \,.\, \rangle)$, with $\mathsf{Aut}(V) = U(n)$ because V is not endowed with an orientation (so vectors differing by an orientation are indistinguishable), and let $G = SU(n)$ acting in the fundamental representation.

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

Suppose it were physically possible in this world to obtain independent evidence for a non-trivial connection ϖ. That should be possible, since all hands agree that Yang–Mills theory in vacuum is a genuine physical theory, with measurable observables. Then one could, in principle, have independent evidence for a symmetry group inferred from ϖ, alongside a geometrical structure for the fermions invariant under a larger automorphism group, also in principle empirically accessible. In such a world there would be more structure preserved under parallel transport—for instance an orientation—than the particles intrinsically possess (see footnote 30.)

But that's not our world; at least it is not our Standard Model.

In the VB-POV, connections are derived from the covariant derivatives ∇_i on the *fundamental vector bundles*. No physical effect is *directly* tied to the principal connection ϖ, since ϖ does not figure in the ontology. In this view, what experiments would probe in the slack-world above is ∇ (or at most ω, the $\mathsf{End}(E)$-valued representative of ∇), not ϖ. Therefore, from the geometry-first perspective, no such slack is possible: all we have is E, and, implicitly, if you want, $\mathsf{Aut}(V)$.[30]

The slack-world scenario is implausible; but its very implausibility points to a tacit assumption already built into our use of principal–associated formalisms, which in practice never exploit the full flexibility of the PFB-POV. The assumption is precisely that such formalisms ultimately rest on underlying *fundamental* (or atomic) vector bundles for which

$$G \simeq \rho(G) \simeq \mathsf{Aut}(V). \tag{5.13}$$

In other words, the standard use of principal and associated fibre bundles often tacitly presupposes the commitments of the geometry-first formulation of gauge theory.

This tacit assumption likely stems from an implicit epistemic argument: that we cannot, even within the PFB-POV, obtain independent evidence for the connection and for the matter fields. But I resist such an argument, along with the tacit assumption. In the PFB-POV, $\mathsf{Aut}(V)$ could, in principle, be revealed by methods that do not rely on the principal connection or on parallel transport—for example, through the structure of bound states or potential terms. But I need not get bogged down in operational or epistemological questions about what kind of experiment could in principle do this, or what would count as independent empirical evidence for a section of E and a connection on P. The point is simply that the gauge-invariant content of the principal connection is part of the physical apparatus of a gauge theory formulated in

[30] In neither case are we required to recover the entire $\mathsf{Aut}(V)$—or G, in the PFB-POV—from parallel transport. In more detail, we can see parallel transport around a closed curve starting (and ending) at $x \in M$ as an element $g \in \mathsf{Aut}(E_x)$. If we take all the closed curves, this generates a subgroup of $\mathsf{Aut}(E_x)$ called $\mathsf{Hol}_{(x)}(\mathrm{D})$. It can be shown that, on a simply-connected region, the holonomy depends on x only up to conjugation by a group element. Thus it is customary to refer to the path-independent $\mathsf{Hol}(\mathrm{D})$ as the *the holonomy group* $\mathsf{Hol}(\mathrm{D})$. It follows that, for two linearly isomorphic bundles, $E, \widetilde{E}$, $\mathsf{Hol}(\mathrm{D}) = \mathsf{Hol}(\widetilde{\mathrm{D}})$. Under some further natural assumptions, it can also be shown that, given a connection D, one can find a new principal bundle (P', M, G'), with a connection ϖ', such that the holonomy group is isomorphic (as a G-torsor) to the structure group $G' \subseteq \mathsf{Aut}(V)$, and E is an associated bundle to P' with D being the induced connection from ϖ' (cf. (Michor, 2008, Theo. 17.11)).

the PFB-POV. In that framework, the relevant group has physical significance and can differ from the automorphism group of the matter fields.

Moreover, physicists are not always consistent about this tacit assumption—perhaps the main reason it is so hard to tease it out from the literature. Textbooks often begin with simplified cases where there is only one compact, semi-simple group and one vector space admitting the fundamental representation, as in the VB-POV. Yet I have never encountered an explicit mandate for always constructing principal and associated bundles "piecewise" in the manner the VB-POV advocates. If such a mandate exists, it remains tacit, and the literature is inconsistent about whether it is even intended. For after describing the relationship between P and E in the simple case where $P = L_\rho(E)$, where $G \simeq \rho(G) \simeq \mathsf{Aut}(V)$, expositions usually move straight to composite cases, where particles are assigned representation labels defining associated vector bundles under a product group such as $SU(3)\times SU(2)\times U(1)$ in the Standard Model.

For instance, (Skinner, 2007) is among the textbooks that most clearly illustrate this line of exposition. Here we find perhaps the closest thing to an articulation of the tacit assumption—but he then goes on to deny it. On page 141 he writes: "Vector bundles are of relevance to physics because a charged matter field is a section of an associated vector bundle." So far, his associated vector bundles can carry different quantum numbers, and so are not fundamental, in my sense; for that they would only be allowed to carry the fundamental representation, when it exists and is unique. But after constructing principal fibre bundles with matrix Lie groups from frame bundles of a vector space—i.e. assuming the fundamental representation and the validity of (5.3)—he says (ibid., p. 142): "The most common Lie groups that arise in physics are indeed matrix Lie groups [of that sort], so the two viewpoints [PFB and VB] are equivalent." Here we see a restriction of the comparison to the cases that commonly arise in physics. I would resist saying they are 'equivalent': I take him to mean that the VB-POV can account for the cases of interest in physics. However, regarding which picture is more fundamental, he adds (ibid., p. 142): "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." Here he refers to the fact that theories whose structure groups are certain exceptional Lie groups, such as E_6 (see Section 5.2), explicitly sever the link to the automorphism groups of V.

Thus, methodologically, once you consider the space of more general theories, as Skinner admits, the link between the VB-POV and the PFB-POV is severed; and once it is severed, the door reopens to all the examples of this section—including the slack world itself.

Summing up: upon reconstructing symmetry groups, the VB-POV insists that each gauge subgroup factor be the automorphism group of its corresponding fundamental vector bundle. I do not see this tight identification as a limitation of the VB-POV; on the contrary, it is its chief virtue. By grounding group actions in geometry, the VB-POV rules out many possible theories. That restriction is methodologically valuable, so long as it still encompasses our best physics—and in this case, it does.

6 Conclusions

Feynman's Nobel Prize lecture, with which I began, reflected on his alternative formulation of quantum electrodynamics via path integrals. That formulation, like Minkowski's introduction of spacetime—and indeed many other mathematically equivalent yet conceptually transformative innovations scattered through the history of physics—proved invaluable. They provided new explanations and opened new directions for development.[31] I do not expect the geometry-first formulation of gauge theory developed here will ascend to comparable heights, nor do I expect it will become orthodoxy, as Feynman's and Minkowski's eventually did.

But I do not want to understate what has been gained. The geometry-first formulation postulates a different ontology, offers independent 'symmetry-free' explanations of familiar mechanisms in gauge theory, and, above all, *eliminates the slack* between symmetry and geometry. That tighter fit might be taken to explain why certain group–geometry correspondences—correspondences between properties of vector bosons and matter fermions—are realised in our world and others are not.

Let me summarise the virtues of the independent explanations.

In brief, the Higgs field is a nowhere-vanishing section of a vector bundle with approximately constant norm. The component of the Higgs field carrying this constant nonzero norm plays the role of the *Higgs vacuum*. If the field couples dynamically to the affine structure of the vector bundle, and if the affine structure is itself dynamical, the existence of such a section is enough: the geometry alone performs the explanatory work of the Higgs mechanism that symmetry was thought indispensable for:

(H.i) Goldstone modes never appear here, and so never require elimination. The reason is simple: the constant magnitude of the Higgs vacuum section ensures that it is orthogonal to its covariant derivative.

(H.ii) What in the symmetry-first formulation is described as the 'acquisition of mass' by vector bosons is, in this geometry-first account, nothing more than the non-vanishing of the (covariant) kinetic energy of the Higgs vacuum. In other words, it is nothing more than the run-of-the-mill idea that the kinetic term of the Higgs depends on the affine structure of the vector bundle. In this formulation, then, talk of 'mass acquisition' may strike a geometry-first militant—say, a relativist—as misplaced. From the VB-POV, what would it even mean for the affine structure, or the covariant derivative ∇, to 'acquire mass'?[32]

(H.iii) The covariant derivative along a single section of a vector bundle does not depend on all the affine degrees of freedom of the bundle (for $\dim(E_x) \geq 2$). The absent degrees of freedom correspond, in the symmetry-first idiom, to the unbroken gauge group, giving rise at the perturbative level to the massless photons.

Turning to the Yukawa mechanism, discussed in Section 4: I argued that standard presen-

[31] See Hunt (2025) for more examples of virtuous reformulations of physical and mathematical theories.

[32] To be sure, some would hesitate to say that gravitons acquire mass merely because a spacetime, or a collection thereof, admits a kinetic term for a vector field of constant norm; yet that is precisely the consensus for such theories (Jacobson, 2008).

tations are explanatorily 'opaque,' and offered instead a more transparent geometric version of the Yukawa form itself. I readily admit that my judgements about what is opaque may stem from a general preference for geometric explanations, *simpliciter*. But the point remains: as emphasised in Section 2.2, the mere availability of a geometric argument that bypasses representation theory is grist to my mill. The aim, after all, is to open the subject to a different community, with different, more geometric ways of thinking.

In this spirit, the geometric formalism brings out certain features of the Yukawa mechanisms that, to my knowledge, have not been emphasised in the literature. (That does not mean they are controversial; perhaps they are simply too minor to warrant mention in standard presentations.) It also raises new questions.

(Y.i) In the geometric picture, the left-handed up and down quarks, and likewise the electron and electron-neutrino, are not distinct particles at all. They are simply the parallel and orthogonal components, with respect to the Higgs direction, of the (first-generation) left-handed quark fields and leptons.[33]

(Y.ii) In the VB-POV the Yukawa form is *geometrically natural*. It is the fibrewise contraction to a scalar density built from the intrinsic structures already present on the fundamental bundles: the Hermitian pairings on colour and generation bundles, and the oriented area form on the weak bundle. By contrast, in the PFB-POV one manufactures invariants in a chosen basis while absorbing ambiguities into coupling constants.

(Y.iii) The geometric formulation also sharpens certain physical questions. In particular, the quark Yukawa term depends essentially on the orientation of $\mathbb{C}^2$: up-type quarks couple to φ_c, which encodes the oriented area orthogonal to the Higgs field. This explains why the geometry-first picture naturally singles out $SU(2)$ rather than $U(2)$. By contrast, for $\mathbb{C}^3$ there appears to be no analogous mechanism: why does the Standard Model employ $SU(3)$ rather than $U(3)$?[34] Such open questions are emblematic of the broader methodological moral: when symmetry is reconstructed from geometry, its explanatory role becomes both sharper and more constrained than in the symmetry-first account.

Section 5 develops this methodological point.Where the PFB-POV tolerates a loose fit between the structure group and the geometry of its associated bundles, the VB-POV requires each gauge-group factor to coincide with the automorphism group of a fundamental vector bundle.[35] In the PFB-POV this looseness gives the gauge group an independent physical role: it can, in principle if not in practice, differ from the automorphism group relevant to the matter

[33] That they cannot represent physically distinct particles before symmetry breaking is noted in some textbooks—(Tong, 2025, p. 185) is an exception more than the rule—but I have not seen this parallel/orthogonal decomposition relative to the Higgs section made explicit anywhere.

[34] There is a canonical isomorphism $U(3) \simeq (SU(3) \times U(1))/\mathbb{Z}^3$, but the representations of $U(1)$ in the Standard Model do not seem to realise this isomorphism. Benjamin Muntz (p.c.) has suggested that the place to look may be in the triality constraints on baryon coupling: colourless states built from three quarks would not be invariant under the full $U(3)$. Indeed, naively, from the structure endowed to $\mathbb{C}^3$ we would expect scalars built of quark–anti-quark pairs, via $\langle\, .\, ,\, .\, \rangle$, and others built of three quarks, via ϵ, and nothing more.

[35] This automorphism group might still not fully encode the geometry of the corresponding vector space; it plausibly does so when the group acts transitively. See footnote 25.

fields, allowing a certain independence between vector bosons and matter fermions. The VB-POV removes this independence by binding the properties of bosons and fermions to a common geometric source.

But it is important to note that, although the VB-POV tightens the slack between symmetry and geometry, there remains a clear explanatory direction: the group is subservient to the geometry. A striking illustration is provided by the quantisation of charge. In the standard, symmetry-first picture, the fact that charge comes in discrete values is traced to the topology of a compact group such as $U(1)$. In the geometry-first picture, by contrast, it follows from the discrete tensorial structure through which matter fields are assembled from the fundamental bundles. Even if the relevant group were non-compact, the tensorial construction would still yield charges labelled by integers. The mechanism is therefore both more general and more geometric than the topological account: the quantisation of charge arises from geometry itself, rather than from the global structure of the symmetry group.

The VB-POV thereby excludes many symmetry-first models.[36] Fortunately, that narrowing of theoretical scope is (currently) a virtue rather than a liability, for our best (current) physics—the Standard Model—falls squarely within its bounds.

To close, the methodological lesson of eliminating the slack between symmetry and geometry yields two broader morals. First, that future developments of gauge theory might do well to begin with structured, fundamental vector bundles and the tensors they carry. Second, that Occam's razor—if it has an edge here—trims our ontology down to the VB-POV.

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.*

[36] Most of which are never considered in practice anyway. Indeed, as discussed in Section 5, there is a widespread tacit assumption that the gauge groups of principal bundles simply are the automorphism groups of the fundamental vector spaces. But this assumption is both unwarranted and inconsistently applied.

$G_p := \{g \in G \mid g \cdot p = p\} = \{e\}$). 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.*[37]

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:*

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

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*:

[37] 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. Suppose $\kappa(x) = [p, v]$. Then

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

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 D is an operator:

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

such that the product rule

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

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 D 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:*

$$\mathrm{D}_{\gamma'}\mathbf{v}_h = 0. \tag{A.9}$$

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 D 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 D, we proceed in the usual way:

Definition 7 (Curvature) *Given a covariant derivative* D *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* $\kappa \in \Gamma(E)$*,*

$$\Omega(X, Y)\kappa = \left(\mathrm{D}_X\mathrm{D}_Y - \mathrm{D}_Y\mathrm{D}_X - \mathrm{D}_{[X,Y]}\right)\kappa,$$

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 fiber-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 composisiton, 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*, e.g. of orthonormal frames. Conversely, if V has more than just the structure of a linear vector space, e.g. if it is endowed with an inner product, it will induce a subgroup $G' \subset GL(V)$ on P that respects that structure.

B A sketch of the standard exposition of the Higgs Mechanism

Before turning to our geometric reformulation, we briefly review the conventional mathematical account of the Higgs mechanism, as in Hamilton (Hamilton, 2017, Ch. 8). This will allow us to highlight the points at which symmetry groups, stabilisers, and coset spaces enter essentially.

We begin with a compact Lie group G acting unitarily on a complex vector space W (the Higgs vector space). A Higgs potential of the form

$$V(w) = -\mu\|w\|^2 + \lambda\|w\|^4, \quad \mu, \lambda > 0,$$

is G-invariant, and has minima along a sphere

$$\mathcal{M}_{\text{vac}} = \{w \in W : \|w\| = v\}, \quad v = \sqrt{\mu/2\lambda}.$$

Thus the set of vacua is itself a homogeneous G-space:

$$\mathcal{M}_{\text{vac}} \cong G/H,$$

where $H = G_{w_0}$ is the stabiliser (isotropy subgroup) of a chosen vacuum vector $w_0 \in W$. Already here the reasoning is group-theoretic: the possible vacua are classified by subgroup data (G, H).

A vacuum configuration is given by a constant section Φ_0 of the Higgs bundle, with $\Phi_0(x) = w_0$ for all $x \in M$. The unbroken subgroup H is compact (as a closed subgroup of G). If $H \subsetneq G$, the gauge theory is said to be *spontaneously broken* (Hamilton, 2017, Def. 8.1.6). The Higgs condensate Φ_0 is the non-zero background field in which other particles propagate, and is invariant only under $H \subset G$. Again, the classification of broken versus unbroken symmetries is a stabiliser argument.

Perturbations of the Higgs field $\Phi = \Phi_0 + \tilde{\phi}$ decompose relative to the tangent space at w_0:

$$T_{w_0}W \cong T_{w_0}(G \cdot w_0) \oplus (T_{w_0}(G \cdot w_0))^{\perp}.$$

Group theory guarantees this orthogonal splitting (Hamilton, 2017, Lem. 8.1.12). One then expands $\tilde{\phi}$ in an eigenbasis of the Hessian:

$$\tilde{\phi} = \frac{1}{\sqrt{2}} \sum_{i=1}^{d} \pi_i e_i \; + \; \frac{1}{\sqrt{2}} \sum_{j=1}^{2n-d} \sigma_j f_j,$$

with $\{e_i\}$ tangent to the orbit $G \cdot w_0$ and $\{f_j\}$ orthogonal. The π_i are massless scalar fields: the Nambu–Goldstone bosons. The σ_j are massive scalars: the Higgs bosons (Hamilton, 2017, Def. 8.1.14). This is precisely Goldstone's theorem: $\dim(G/H)$ massless scalars, deduced from the group structure of the vacuum manifold.

Physically the Goldstone bosons are unobservable, since they can be gauged away. Mathematically this is formalised by the *unitary gauge* (Hamilton, 2017, Def. 8.1.18, Thm. 8.1.20). One uses a physical gauge transformation $\gamma : M \to G$ to rotate the Higgs field entirely into the fixed direction w_0:

$$\Phi(x) \mapsto \gamma(x) \cdot \Phi(x) = (0, \ldots, 0, v + h(x)).$$

By definition, in unitary gauge the shifted Higgs field is orthogonal to the orbit $G \cdot w_0$, and the Nambu–Goldstone bosons vanish. This step is again an essentially group-theoretic argument, exploiting the transitivity of the G-action on the vacuum manifold.

Let $\mathfrak{g}$ be the Lie algebra of G, with $\mathfrak{h} = \text{Lie}(H)$ the subalgebra of unbroken generators, i.e. the stabilisers of the Higgs. With respect to an invariant scalar product, decompose

$$\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{h}^\perp.$$

The $\mathfrak{h}^\perp$ directions correspond to broken generators. It is exactly these components of the gauge field A_μ that acquire mass terms from the kinetic energy of the Higgs. Conversely, the $\mathfrak{h}$-components remain massless.

References

Barrett, T. W., & Manchak, J. B. (2024, December). On Privileged Coordinates and Kleinian Methods. *Erkenntnis*. doi: 10.1007/s10670-024-00914-4

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

Derdzinski, A. (1992). *Geometry of the Standard Model of Elementary Particles.* Springer Berlin Heidelberg. doi: 10.1007/978-3-642-50310-8

Feynman, R. P. (1994). *The character of physical law* (Modern Library ed ed.). New York: Modern Library. (Originally published in hardcover by the British Broadcasting Corporation in 1965 and in paperback by M.I.T. Press in 1967"–T.p. verso)

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*.

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

Holton, G. (1974). Thematic origins of scientific thought: Kepler to Einstein. *Philosophy of Science*, *41*(4), 415–418. doi: 10.1086/288604

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

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.

Stachel, J. J. (2002). *Einstein from 'B' to 'Z'* (No. 9). Boston: Birkhäuser. (Includes bibliographical references)

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

Wallace, D. (2019). Who's Afraid of Coordinate Systems? An Essay on Representation of Spacetime Structure. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *67*, 125–136. doi: 10.1016/j.shpsb.2017.07.002

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