NOTE-01 / CONCEPTUAL RESEARCH NOTE

Why Yang–Mills Theory Has Fibers

J. Scott Little

Source: v3, July 2026 Publication treatment: 8 September 2026

Research notes — not peer reviewed. Trust, but verify. Always.

How do gauge fibers over spacetime differ from reconstruction fibers over thickness?

An introduction to gauge connections, Wilson observables, and the second base that appears when thickness becomes part of a reconstruction.

Scope and status

The historical one-loop calculation is not recovered here. This conceptual note does not construct a thin theory or establish its mass spectrum.

This edition adds a reading guide and publication notices. The source's mathematical status and historical provenance remain attached to its claims.

A way into the note

The central distinction

A gauge fiber holds internal frames at a spacetime point. A fixed-thickness fiber collects reconstruction objects at one thickness. Keeping their bases explicit prevents two different uses of “fiber” from being mistaken for one object.

A concrete example

Compare two local frames related by a gauge transformation, then compare two realizations lying at the same thickness. The first comparison changes a presentation; the second belongs to the reconstruction diagram.

What this approach gains

The note gains a way to discuss local data, global choices, and scale compatibility together. It explains why a well-defined physical space at each stage need not occupy one canonically fixed ambient representation.

Reading route

Read the abstract and historical opening, then follow the principal-bundle discussion into the fixed-thickness reconstruction. Keep a two-column list: spacetime base / thickness base.

Citation: J. Scott Little, Why Yang–Mills Theory Has Fibers, v3, July 2026; website production edition, 8 September 2026. Cite theorem and printed body-page numbers from the note itself.

Why Yang–Mills Theory Has Fibers

Where Are the Gluons? From Poincaré Particles to Fixed-Thickness Reconstruction

J. Scott Little

Research Notes on Yang–Mills Reconstruction, Note 0

July 2026

Abstract

Yang–Mills theory is often introduced through a Lie-algebra-valued potential A_μ and its curvature $F_{\mu\nu}$. That notation is efficient, but it can hide the geometric reason the theory is a gauge theory at all. The field is not fundamentally a function taking values in one globally preferred copy of a Lie algebra. It is a connection on a principal bundle, and its purpose is to compare internal frames attached to different spacetime points. Curvature records the failure of these comparisons to close around loops, while Wilson observables record the corresponding holonomy in a gauge-invariant manner.

The positive-thickness reconstruction program introduces another fibered structure. A cofinal family of Wilson-tube systems is organized by an intrinsic thickness map, and the objects lying over one value of the thickness form a fixed-thickness fiber. This reconstruction fiber is not the same object as a principal gauge fiber. Nevertheless, both arise from the same structural problem: local data are available over a base, but a coherent global choice need not exist. This note explains that relationship before any detailed discussion of scale twisting or the three reconstruction sectors. Its purpose is conceptual and architectural; it does not construct the zero-thickness theory, prove a spectrum condition, or establish a Yang–Mills mass gap.

The historical motivation was not originally expressed in the language of transport, fibers, or higher descent. It began with a more direct question: where are the gluons in the physical one-particle irreducible representations of the Poincaré group, if Haag's theorem prevents the interacting theory from being identified with a free-gluon interaction picture? A later one-loop projection-operator calculation in a difficult noncovariant algebraic gauge suggested that even the embedded physical subspace could receive quantum corrections. The original calculation is no longer available. This note therefore records the lesson rather than claiming to reproduce the historical result: a physical Hilbert space may be well defined at each stage without being canonically fixed inside one kinematical representation across scale.

Section −1: Lessons in the Fog Before a High Mountain

This section is not a reconstruction of lost calculations. It is a record of the questions that survived them. The dates are approximate, the terminology is partly retrospective, and no modern language of transport or higher descent is attributed to the earlier work. In 2006 the question was not "what is the correct connection on a family of physical Hilbert spaces?" It was simply:

Where are the gluons in the physical one-particle spectrum?

2006–2007: where are the gluons?

The first question arose from the tension between two familiar descriptions. On one side, the perturbative gauge potential is decomposed into modes that look like massless helicity-one quanta. On the other, a relativistic particle is identified through a one-particle irreducible unitary representation of the Poincaré group carried by the physical Hilbert space. The natural question was therefore whether the gluon modes used in perturbation theory were actually present in the physical one-particle representation.

Haag's theorem sharpened the difficulty. If an interacting relativistic field cannot be obtained by a global unitary interaction-picture identification with the corresponding free field, then one should not assume from the outset that the physical Hilbert space is a dressed free-gluon Fock space. This did not prove that gluons were absent, and it did not yet suggest a fibered answer. It made the ontology unclear:

$$\text{gauge-potential mode} \quad \overset{?}{=} \quad \text{physical Poincaré particle.}$$

The eventual reconstruction program reverses the usual order. It first seeks a gauge-invariant local observable system and its vacuum representation, and only then asks what one-particle or scattering interpretation is supported by the resulting Poincaré spectrum.

2007: the projection operator moved

A second question emerged from a one-loop Yang–Mills calculation using Klauder's projection-operator method in a noncovariant algebraic gauge. The gauge was technically severe. The calculation is no longer available in its original form and was never communicated to Klauder or prepared for publication. What survived was the structural observation: after the loop correction, the operator used to project from the kinematical space to the physical subspace no longer looked like a fixed projector multiplied by a harmless scalar renormalization.

The lesson should be stated conservatively. Suppose the physical space is selected by a family of constraints $G_a(\mu)$ and a corresponding positive constraint operator

$$Q(\mu) = \sum_a G_a(\mu)^* G_a(\mu), \qquad \mathcal{H}_{\text{phys}}(\mu) = \ker Q(\mu).$$

When zero is spectrally isolated, the physical projection can be written

$$P(\mu) = \mathbf{1}_{\{0\}}(Q(\mu)).$$

A purely invertible recombination of the same constraints leaves the common kernel fixed. A correction that changes the kernel, however, changes the embedded physical subspace:

$$P(\mu) \neq P(\mu_0) \qquad \text{inside} \qquad \mathcal{H}_{\text{kin}}.$$

This does not imply that the physical theory is inconsistent or that no abstract isomorphism can exist between the ranges. It means that literal identity inside one fixed kinematical representation is no longer automatic. Comparison requires additional structure.

In modern terms, a sufficiently regular family of projections carries a Kato parallel transport. Along a chosen scale path, a unitary $U(\mu, \mu_0)$ may be defined by

$$\frac{dU}{d\mu} = [\dot{P}(\mu), P(\mu)]U, \qquad U(\mu_0, \mu_0) = I,$$

so that

$$P(\mu) = U(\mu, \mu_0)P(\mu_0)U(\mu, \mu_0)^*.$$

Nothing like this language was being used in the original calculation. It is a retrospective explanation of the lesson: the invariant object may be a family of physical subspaces with a rule of comparison, rather than one subspace held fixed by convention.

2008: a path integral could be defined before the theory was global

After graduation, a path integration over a Yang–Mills jet manifold was constructed using the continuous-time regularization associated with Klauder and Daubechies. The metric entering that regularization—later remembered in connection with Klauder's language of a "shadow" or metrical structure—made the local integration procedure mathematically meaningful. But the result carried substantial qualifications and relied on global assertions that were not themselves produced by the local regularization.

That experience separated two problems that are often conflated:

$$\boxed{\text{a well-defined regularized integral} \not\Rightarrow \text{a globally reconstructed quantum field theory.}}$$

A measure or limiting prescription may exist in each selected presentation while compatibility under chart change, gauge comparison, regional restriction, and refinement remains unresolved. The later descent language was not present, but the missing theorem was already a gluing theorem.

2009–2011: entropy outside a black hole was too large

A later calculation of Yang–Mills entropy outside a black hole, using Bogoliubov transformations together with jet formation, produced more entropy than expected. The historical calculation should not be retroactively turned into a theorem about edge modes or algebraic entropy. Its durable lesson was more elementary: the physical gauge system did not behave as though the global Hilbert space admitted an innocent factorization into independent inside and outside tensor factors.

The problem can be stated without claiming its final diagnosis:

$$\mathcal{H}_{\text{phys}} \overset{?}{=} \mathcal{H}_{\text{inside}} \widehat{\otimes} \mathcal{H}_{\text{outside}}.$$

Gauss constraints, boundary flux, and the distinction between kinematical and physical modes all make this equality suspect. This was one route toward the later conviction that entropy should be attached to local normal algebras and states, not inferred only by counting modes in a presumed global tensor factorization.

2010–2012: an infinite string was still too rigid

String-localized field constructions offered another possible home for the missing gauge excitation. But the semi-infinite string appeared too rigid for the scattering structure being sought, and the resulting picture did not resemble the known organization of amplitudes. This was not a criticism of the mathematical validity of string localization. It was a failure of that object to answer this particular question.

Moving from point localization to string localization still begins with a field or particle carrier and then changes its localization. The later observable-first approach changes the logical direction:

$$\text{local gauge-invariant observables} \longrightarrow \text{vacuum representation} \longrightarrow \text{spectrum and particles.}$$

2012–2014: amplitudes, positive geometry, and the vacuum shadow

The path through the amplituhedron, Grassmannians, Yangian symmetry, and finally operator algebras suggested that a scattering amplitude might not be the primitive object. The visible scattering surface appeared instead as though it carried a shadow cast by a deeper vacuum organization. By 2014 the working intuition was roughly

$$\boxed{\text{scattering data are a visible trace of the vacuum structure,}}$$

not a complete specification of that structure.

The statement that the vacuum was "not globally separable" should also be read historically rather than as a claim that the Hilbert space lacked a countable dense set. The intended difficulty was that the vacuum did not look like one globally preferred product state, tensor factorization, or trivially identified representation across all local descriptions.

The question without a date: how do local W^*-algebras cohere?

At some later stage the problem took an operator-algebraic form. If local physics is represented by von Neumann algebras

$$O \longmapsto \mathcal{M}(O),$$

how do these algebras, their normal states, and their preferred Hilbert-space realizations arrange themselves into one theory? The word "coherently" was not initially a technical term. It named the unease that inclusion alone did not settle the compatibility of representations, projections, modular data, and vacuum choices.

This was the decisive conceptual shift. The question was no longer only whether one had found the correct global Hilbert space. It was whether the local physical objects and their comparison maps formed a globally consistent system, and what information remained when they failed to be strictly identical.

June 2025: scale became part of the base

The generalized renormalization-group idea added spatial thickness or resolution as a genuine direction of comparison. One no longer had merely an operator acting on a fixed state space. One had a family

$$\varepsilon \longmapsto (\mathcal{M}_\varepsilon(O), \mathcal{H}_\varepsilon, P_\varepsilon, q_\varepsilon)$$

and had to ask what related the objects at distinct values of ε. The old projection drift had returned in a more mature form. The problem was not simply that an operator ran with scale. The selected physical realization itself could vary, and its comparisons could carry nontrivial coherence data.

The persistent question

The chronology can be compressed into one sequence:

$$\begin{aligned}
\text{Where are the gluons?} &\longrightarrow \text{Is the projected physical subspace fixed?}\\
&\longrightarrow \text{Does a local path integral globalize?}\\
&\longrightarrow \text{Why does regional entropy overcount?}\\
&\longrightarrow \text{Why does a new localization carrier not recover scattering?}\\
&\longrightarrow \text{Is scattering a shadow of the vacuum?}\\
&\longrightarrow \text{How do local } W^*\text{-realizations cohere?}\\
&\longrightarrow \text{How are they compared across scale?}
\end{aligned}$$

The fiber language is therefore not an ornamental reformulation of the classical principal bundle. It is the eventual answer to a repeated failure of strict global identification. Each attempted single carrier—one-particle space, projected Hilbert subspace, path-integral chart, horizon tensor product, string-localized field, scattering geometry, or global vacuum representation—captured part of the theory while leaving a comparison problem unresolved.

Central thesis. Yang–Mills theory is fibered because it describes the comparison of internal frames across spacetime, not the evolution of a field in one globally fixed internal frame. The reconstruction becomes fibered again when thickness is made part of the base data.

1 Why begin with fibers?

The word *fiber* can sound like technical packaging added after the physical content is known. In Yang–Mills theory the opposite is closer to the truth. The fiber description explains what the gauge potential is, why gauge transformations have their familiar form, why curvature is the natural field strength, and why holonomy gives the basic nonlocal observable.

The usual local formula

$$A = A^a_\mu T_a \, dx^\mu$$

looks like an ordinary one-form on spacetime with values in a Lie algebra $\mathfrak{g}$. That expression is valid inside a chosen gauge trivialization. It does not by itself say that the same basis $\{T_a\}$ has been selected canonically and globally at every point. A gauge choice is precisely a local choice of internal frame. Changing that frame changes the coefficients of A, although it does not change the underlying connection.

This distinction matters even when the principal bundle happens to be trivializable. A global trivialization may exist mathematically without being physically preferred. Gauge covariance expresses the fact that the physics cannot depend on which internal frame was chosen to write the local formulas.

Principle 1 (Gauge theory is a theory of comparison)**.** *A Yang–Mills connection is not primarily an internal vector attached to each spacetime point. It is a rule for comparing internal data attached to different points. Curvature measures the path dependence of that comparison.*

The later fixed-thickness fibers should be read against this background. They do not add a new hidden dimension to spacetime. They organize families of admissible quantum reconstruction data over a new base variable: the intrinsic thickness of the Wilson-tube observables.

2 The classical gauge fiber

Let M be spacetime and let G be a compact Lie group with Lie algebra $\mathfrak{g}$. A classical Yang–Mills configuration is naturally formulated on a smooth principal bundle

$$G \longrightarrow P \xrightarrow{\pi} M.$$

For each $x \in M$, the fiber

$$P_x = \pi^{-1}(x)$$

is a G-torsor: it has a free and transitive right action of G, but no preferred identity element. Choosing a point of P_x is choosing an internal frame at x.

That absence of a preferred identity is physically important. One may compare two choices in the same fiber by a unique element of G, but no choice is intrinsically distinguished. A local section

$$s_\alpha : U_\alpha \longrightarrow P$$

selects one frame at every point of an open set $U_\alpha \subset M$. On an overlap $U_\alpha \cap U_\beta$, two sections differ by a transition function

$$s_\beta(x) = s_\alpha(x) g_{\alpha\beta}(x), \qquad g_{\alpha\beta} : U_\alpha \cap U_\beta \longrightarrow G.$$

The cocycle relation

$$g_{\alpha\beta}g_{\beta\gamma} = g_{\alpha\gamma}$$

on triple overlaps expresses the consistency required to assemble the local trivializations into one principal bundle.

The gauge group is the group of vertical automorphisms of P. In a local trivialization, a gauge transformation is represented by a function

$$u_\alpha : U_\alpha \longrightarrow G,$$

but the invariant object is the bundle automorphism, not the chosen coordinate expression.

2.1 Associated fibers and physical fields

Given a representation $\rho : G \to \mathrm{GL}(V)$, the associated bundle

$$E = P \times_\rho V \longrightarrow M$$

has fiber E_x isomorphic to V, but not canonically identified with one fixed copy of V until a local gauge is chosen. Matter fields are sections of such associated bundles. Infinitesimal gauge parameters and adjoint-valued fields are naturally sections of the adjoint bundle

$$\operatorname{ad} P = P \times_{\mathrm{Ad}} \mathfrak{g}.$$

Thus the statement that a field is "V-valued" or "$\mathfrak{g}$-valued" is local shorthand for a section of a bundle whose fibers are locally modeled on V or $\mathfrak{g}$. The local vector space is fixed up to the action of the gauge group; the global identifications are part of the geometry.

3 The connection: comparison between fibers

A principal connection on P selects a horizontal subspace

$$H_pP \subset T_pP$$

at every $p \in P$, equivariantly under the right action of G. Equivalently, it may be represented by a connection one-form

$$\omega \in \Omega^1(P; \mathfrak{g}).$$

The connection determines parallel transport. Given a path $\gamma : [0, 1] \to M$, it provides an equivariant map

$$\mathsf{PT}_\gamma : P_{\gamma(0)} \longrightarrow P_{\gamma(1)}.$$

This is the operational meaning of the gauge field: it tells us how to compare internal frames over separated spacetime points.

After choosing a local section s_α, the connection becomes the familiar local gauge potential

$$A_\alpha = s_\alpha^*\omega \in \Omega^1(U_\alpha; \mathfrak{g}).$$

On an overlap, the local potentials satisfy

$$A_\beta = g_{\alpha\beta}^{-1} A_\alpha g_{\alpha\beta} + g_{\alpha\beta}^{-1} dg_{\alpha\beta}.$$

The inhomogeneous term is not an inconvenience appended to the theory. It is the coordinate expression of changing the local frame in which the same connection is written.

Remark 1. *A connection is not a canonical identification of all fibers with one another. Parallel transport depends on a path. The failure of path independence is curvature.*

4 Curvature: the obstruction detected by loops

The curvature of the connection is

$$F_A = dA + \frac{1}{2}[A \wedge A]$$

in a local trivialization. Under a gauge transformation u, it transforms covariantly:

$$F_A \longmapsto u^{-1} F_A u.$$

The Yang–Mills action is built from an $\mathrm{Ad}(G)$-invariant inner product on $\mathfrak{g}$:

$$S_{\mathrm{YM}}[A] = \frac{1}{2g_{\mathrm{YM}}^2} \int_M \langle F_A \wedge * F_A \rangle.$$

Its Euler–Lagrange equation is

$$D_A * F_A = 0.$$

The local formula for F_A is important, but its geometric meaning is more revealing. Transport an internal frame around a small loop. If the frame returns unchanged, the connection is flat along that loop. If it returns rotated, the mismatch is controlled infinitesimally by curvature.

For a closed loop C, parallel transport produces a holonomy element

$$\mathrm{Hol}_A(C) \in G$$

defined up to conjugation. In a representation ρ, the Wilson loop

$$W_\rho(C) = \mathrm{Tr}_\rho \, \mathrm{Hol}_A(C)$$

is gauge invariant. Wilson observables therefore arise directly from the fibered nature of the theory: one compares a fiber with itself after transport around a closed path and removes the residual frame dependence by taking a conjugation-invariant function.

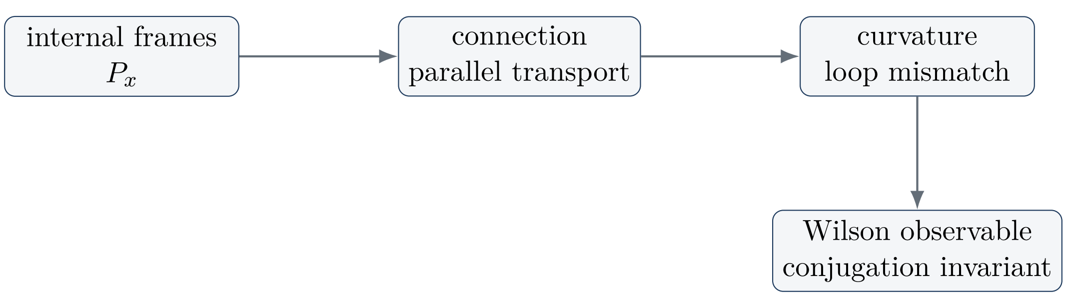

The sequence is structural:

$$\boxed{\text{fibers} \longrightarrow \text{connection} \longrightarrow \text{curvature and holonomy} \longrightarrow \text{gauge-invariant observables.}}$$

5 Why both C^*- and W^*-algebras are needed

The historical progression toward von Neumann algebras should not be described as a rejection of C^*-algebras. The two completions answer different questions.

For a region O, let $\mathfrak{A}(O)$ denote the norm-closed algebra generated by bounded gauge-invariant observables. The quasi-local C^*-algebra

$$\mathfrak{A} = \overline{\bigcup_{O} \mathfrak{A}(O)}^{\|\cdot\|}$$

records the representation-independent bounded observable content. It is the natural level for locality, isotony, norm control, and comparison of different states and representations.

After a state ω is selected, its GNS representation $(\pi_\omega, \mathcal{H}_\omega, \Omega_\omega)$ produces the local von Neumann algebra

$$\mathcal{M}_\omega(O) = \pi_\omega(\mathfrak{A}(O))''.$$

The weak closure is required for the structures that drove the historical questions: normal states, support and central projections, spectral projections, conditional expectations, modular operators, relative entropy, closed forms, and operators affiliated with the represented algebra. In particular, a scale-dependent physical projection $P(\mu)$ is naturally a projection in a represented operator environment, and its corner

$$P(\mu)\mathcal{M}(O)P(\mu)$$

is a W^*-algebra acting on the selected physical range.

Thus

$$\boxed{\begin{aligned} C^*\text{-net} &= \text{representation-independent bounded observables,} \\ W^*\text{-net} &= \text{the local-normal realization seen by a selected state.} \end{aligned}}$$

The C^*-layer should not be erased, because different states may yield inequivalent weak closures. The W^*-layer cannot be omitted, because the projection, entropy, modular, and spectral questions are invisible if one refuses to pass to a represented weakly closed algebra.

6 How the gauge connection survives reconstruction

The passage from a classical gauge field to an algebraic quantum theory does not eliminate the gauge connection. It changes the form in which its comparison structure remains visible. The relevant progression is schematically

$$\boxed{A \longrightarrow F_A \longrightarrow \mathrm{Hol}_A \longrightarrow W_\rho(\gamma) \longrightarrow \text{representation transport} \longrightarrow \text{coherence class.}}$$

This chain is not a sequence of literal identifications. Each term preserves a different invariant aspect of the preceding comparison problem.

6.1 From the gauge connection to invariant transport data

Let

$$G \longrightarrow P \xrightarrow{\pi} M$$

be a principal G-bundle. A connection A supplies the rule for comparing internal frames over distinct spacetime points. In a local trivialization over $U_\alpha \subset M$, it is represented by

$$A_\alpha \in \Omega^1(U_\alpha, \mathfrak{g}),$$

and on an overlap one has

$$A_\beta = g_{\alpha\beta}^{-1} A_\alpha g_{\alpha\beta} + g_{\alpha\beta}^{-1} dg_{\alpha\beta}.$$

The curvature

$$F_A = dA + A \wedge A$$

measures the infinitesimal failure of path-independent comparison. Parallel transport along a path $\gamma : x \to y$ gives a map between the corresponding principal fibers, while a closed path produces holonomy

$$\mathrm{Hol}_A(\gamma) \in G$$

up to conjugation at the base point. Applying a class function gives the Wilson observable

$$W_\rho(\gamma) = \mathrm{Tr}_\rho\big(\mathrm{Hol}_A(\gamma)\big).$$

Thus the first survival mechanism is direct:

> The gauge-dependent local potential is replaced by gauge-invariant transport data. Wilson observables retain the observable effect of the connection without retaining a preferred local gauge frame.

This also clarifies the original gluon question without deciding it by terminology. A perturbative gluon is a local linearized quantum of the connection in a chosen gauge and representation. Whether an isolated gluon occurs in the physical one-particle Poincaré spectrum is a later spectral question about the reconstructed gauge-invariant theory. The connection can remain fully present through curvature and holonomy even if no isolated gluon vector appears in that spectrum.

6.2 Why thick Wilson observables still remember the connection

The positive-thickness reconstruction does not begin with idealized zero-width loops. It begins with gauge-invariant observables localized in controlled tubular neighborhoods. Positive thickness is an analytic and localization parameter, not a replacement for the gauge geometry.

A Wilson tube can still retain information generated by the connection:

- parallel transport along the longitudinal direction;
- curvature detected by controlled loop and surface deformations;
- linking and center-flux information around an incision or excluded locus;

- compatibility of local gauge presentations;
- stability under refinement of the representative at fixed intrinsic support thickness.

The connection is therefore not reintroduced as a globally chosen operator-valued potential. Its gauge-invariant transport content is carried by the Wilson-generated observable system.

6.3 From observables to represented local systems

Let b denote an admissible local presentation and let

$$\mathfrak{A}_b(O)$$

be its bounded local observable algebra. After a selected state or complete-ground functional is represented, one obtains

$$(\pi_b, \mathcal{H}_b, \Omega_b), \qquad \mathcal{M}_b(O) = \pi_b(\mathfrak{A}_b(O))''.$$

The represented package may also contain a complete-ground projection, closed graph and quadratic-form data, Gauss–Ward operators or forms, and normal-state information:

$$\mathcal{R}_b = (\mathfrak{A}_b, \mathcal{M}_b, \mathcal{H}_b, \pi_b, P_{\Omega,b}, q_b, \ldots).$$

Different admissible charts, regions, Cauchy presentations, refinement representatives, or thicknesses can yield equivalent observable systems without supplying a preferred literal identification of their represented Hilbert realizations. Their comparisons may therefore be implemented by normal maps, unitaries, partial isometries, correspondences, graph maps, or coefficient-line-enriched intertwiners.

A second connection-like structure can appear when a selected projection $P(b)$ varies regularly over a parameter base. Under stable rank and spectral isolation hypotheses, the ranges form a Hilbert subbundle with Grassmann–Kato connection

$$\nabla\psi = P\, d\psi,$$

and path transport generated by $[dP, P]$. This is not the Yang–Mills gauge potential A:

> The Yang–Mills connection compares internal gauge frames. The induced projection connection compares selected physical subspaces. The more general reconstruction transport compares complete local physical realizations.

The first structure motivates the latter two, but they must not be conflated.

6.4 Why a PDE treatment is not sufficient

One might ask whether the reconstruction can be reduced to a partial differential equation. If the local Yang–Mills equations are well posed and solutions propagate between suitable Cauchy surfaces, why is a separate fiber, transport, and descent structure required?

A PDE theorem and a reconstruction theorem address different levels. A PDE evolves data inside an already selected geometric and analytic realization. A reconstruction theorem must determine whether the different local realizations used to formulate those equations belong to one coherent physical system.

For a presentation b, suppose

$$\mathcal{P}_b u_b = 0$$

is well posed. Existence, uniqueness, regularity, finite propagation, and continuous dependence for u_b do not by themselves identify it with a solution constructed in another presentation b'. That requires a comparison map

$$\mathsf{T}_a : \mathcal{X}_b \longrightarrow \mathcal{X}_{b'}, \qquad a : b \longrightarrow b',$$

together with a proof that the comparison preserves the structures used by the physical reconstruction.

6.4.1 Admissible transport

Let Γ_{cof} denote schematically the selected admissible presentation category. Its objects contain only those combinations of region, adapted chart, admissible Cauchy section or development, refinement representative, and positive thickness allowed by the declared control structure. Its arrows are only the permitted regional, chart, Cauchy, refinement, and scale comparisons.

An *admissible transport* along an arrow $a : b \to b'$ is a comparison map, unitary, correspondence, or graph-space map

$$\mathsf{T}_a : \mathcal{X}_b \longrightarrow \mathcal{X}_{b'}$$

that preserves the relevant reconstruction package. According to the object being transported, admissibility includes:

- localization and support control;
- gauge invariance and compatibility with the Gauss–Ward package;
- preservation of the relevant graph, form, and operator domains;
- covariance of the normal observable action;
- compatibility with regional restriction and admissible Cauchy propagation;
- the selected uniform refinement and scale estimates;
- compatibility with coefficient lines, multiplication maps, and associators already present in representation descent.

At the differential level, exact covariance may take the form

$$\mathcal{P}_{b'} \mathsf{T}_a = \mathsf{T}_a \mathcal{P}_b,$$

while a controlled branch may satisfy

$$\mathcal{P}_{b'} \mathsf{T}_a - \mathsf{T}_a \mathcal{P}_b = \mathcal{E}_a,$$

with $\mathcal{E}_a$ controlled on the declared packet and graph domains. This proves compatibility with an individual comparison arrow. It does not prove coherence of all comparison paths.

6.4.2 Propagation does not imply descent

For composable arrows

$$b_0 \xrightarrow{a} b_1 \xrightarrow{c} b_2,$$

one must compare

$$\mathsf{T}_c\mathsf{T}_a \qquad \text{with} \qquad \mathsf{T}_{ca}.$$

PDE uniqueness can identify two solutions only after they have been placed in a common realization and assigned compatible initial data. It does not construct that realization, determine the comparison maps, or prove their composition laws.

In a strict system,

$$\mathsf{T}_c\mathsf{T}_a = \mathsf{T}_{ca}.$$

In a coefficient-enriched system, one may instead have

$$\mathsf{T}_c\mathsf{T}_a \cong \mu_{c,a}\mathsf{T}_{ca},$$

with higher compatibility controlled by an associator. The local equations can be exactly covariant under every individual arrow while this associator remains nontrivial.

> A PDE propagates local data along an already selected comparison. Descent determines whether all comparisons agree, satisfy the cocycle laws, and define a strict, twisted, or obstructed global object.

The represented W^*-system also contains information not exhausted by a field equation: normal states and preduals, complete-ground projections, modular and entropy data, closed forms, affiliated operators, and selected domains. Two realizations may induce the same adjoint action on bounded observables while their Hilbert implementers differ by a commutant unitary or coefficient line. Hence

$$\text{equivalent local PDE evolution} \not\Longrightarrow \text{strictly identical represented realization.}$$

6.5 Why reconstruction has a second base

The classical gauge bundle is organized over spacetime. The reconstructed quantum system depends additionally on the presentation through which the localized observable is resolved and on the positive thickness at which it is defined. Schematically, an admissible base object has the form

$$b = (O, \alpha, \mathsf{X}, r, \varepsilon),$$

where O is a local region, α an adapted chart, X an admissible Cauchy presentation, r a selected refinement representative, and $\varepsilon > 0$ the intrinsic Wilson-tube thickness.

Thus there are two bases with different roles:

spacetime base	reconstruction base
$x \in M$	$b = (O, \alpha, \mathsf{X}, r, \varepsilon)$
internal gauge frame	local physical realization
spacetime parallel transport	admissible comparison maps
ordinary holonomy	descent and coherence data.

The second base does not replace spacetime and thickness is not a fifth spacetime coordinate. It records the controlled choices required to construct and compare the quantum realizations.

6.6 The thickness map and fixed-thickness fibers

The intrinsic thickness must be distinguished from the mesh used to resolve an observable. A representative can be refined while its physical support thickness remains fixed. The zero-thickness problem begins only when distinct positive thicknesses are compared and one asks for controlled behavior as $\varepsilon \downarrow 0$.

Let

$$e : \mathrm{Ob}(\Gamma_{\mathrm{cof}}) \longrightarrow E_{\mathrm{cof}} \subset (0, \infty)$$

be the thickness map. For $\varepsilon \in E_{\mathrm{cof}}$, define

$$\Gamma_{\mathrm{cof}}(\varepsilon) = e^{-1}(\varepsilon).$$

This fixed-thickness fiber contains the admissible presentations having the same intrinsic support thickness. It may contain several regions, charts, Cauchy presentations, and refinement representatives.

The total reconstruction consequently has two directions of comparison:

$$\begin{aligned} &\text{within-fiber transport:} \quad b, b' \in \Gamma_{\mathrm{cof}}(\varepsilon), \\ &\text{across-scale transport:} \quad b \in \Gamma_{\mathrm{cof}}(\varepsilon), \quad b' \in \Gamma_{\mathrm{cof}}(\varepsilon'), \quad \varepsilon \neq \varepsilon'. \end{aligned}$$

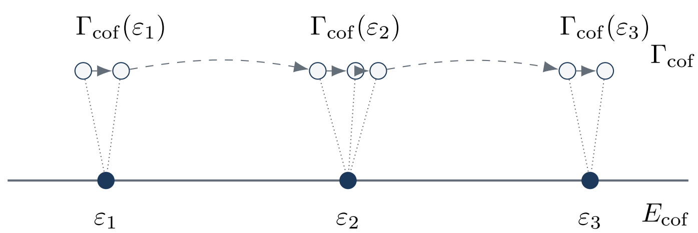

The fixed-thickness fiber is not the principal gauge fiber P_x. It is an inverse image in the reconstruction base. Its naturality is inherited from the same general problem: local data are available, comparison is required, and a preferred global identification need not exist.

6.7 Why cohomology begins to appear

Suppose admissible arrows

$$a : b_0 \longrightarrow b_1, \qquad c : b_1 \longrightarrow b_2$$

carry implementation maps T_a and T_c. The bounded observable maps may compose strictly, while the represented implementations satisfy only

$$T_c T_a \cong \mu_{c,a} T_{ca}.$$

For three composable arrows, the two bracketings of the multiplication data may differ by an associator

$$\Omega_{d,c,a}.$$

In a scalar realization $\Omega_{d,c,a} \in U(1)$; more generally it belongs to the relevant banded coefficient system.

The associativity law imposes a cocycle equation. Changing local implementers changes the representative Ω by a coboundary. The invariant content is therefore a cohomology class, written schematically as

$$[\Omega]^b_{\mathrm{cs}}.$$

Cohomology records the part of the composition defect that cannot be removed by changing the permitted local choices.

This class does not necessarily measure failure of the bounded observable net. Central factors disappear under the adjoint action:

$$\mathrm{Ad}(\lambda U)(A) = UAU^*, \qquad \lambda \in U(1).$$

The observable maps can therefore remain strict while their selected Hilbert implementers retain nontrivial coefficient-line or associator data.

The total nerve of Γ_{cof} contains simplices lying inside one fixed-thickness fiber, simplices formed from scale arrows, and mixed chart–scale simplices. Restriction along

$$j_\varepsilon : \Gamma_{\mathrm{cof}}(\varepsilon) \hookrightarrow \Gamma_{\mathrm{cof}}$$

gives the fixed-thickness class

$$j_\varepsilon^*[\Omega]^b_{\mathrm{cs}} = [\Omega_\varepsilon]_{\mathrm{band}}.$$

The total class can therefore contain information absent from every individual fiber.

6.8 Two coherence questions, not one

The second base separates the reconstruction into two logically different questions:

Can the represented data be strictified at each fixed thickness?
Can those strictifications be chosen coherently across thickness?

A fixed-thickness strictification at ε removes the restricted class

$$[\Omega_\varepsilon]_{\mathrm{band}}.$$

A total chart–scale strictification removes

$$[\Omega]^b_{\mathrm{cs}}.$$

Restriction immediately gives the one-way implication

$$[\Omega]^b_{\mathrm{cs}} = 0 \quad \Longrightarrow \quad [\Omega_\varepsilon]_{\mathrm{band}} = 0 \quad \text{for every } \varepsilon.$$

The converse is not automatic:

$$([\Omega_\varepsilon]_{\mathrm{band}} = 0 \text{ for every } \varepsilon) \not\Longrightarrow [\Omega]^b_{\mathrm{cs}} = 0.$$

Indeed, choose a trivialization β_ε in every fiber. Along a scale arrow $s : \varepsilon \to \varepsilon'$, one must compare the transported choice $s_*\beta_\varepsilon$ with $\beta_{\varepsilon'}$. Their discrepancy, together with the coefficient datum of the scale arrow, defines a relative scale cochain. Its invariant remainder is a relative class, written schematically as

$$[\tau]_{\mathrm{rel}}.$$

It measures the failure of the separate fiber trivializations to fit together along scale and mixed chart–scale diagrams.

The mixed comparison is visible in a square of admissible presentations:

$$\begin{tikzcd}
b_{\alpha,\varepsilon} \arrow[r, "c^{\varepsilon}_{\alpha\beta}"] \arrow[d, "s_\alpha"'] & b_{\beta,\varepsilon} \arrow[d, "s_\beta"] \\
b_{\alpha,\varepsilon'} \arrow[r, "c^{\varepsilon'}_{\alpha\beta}"'] & b_{\beta,\varepsilon'}
\end{tikzcd}$$

At the observable level the two routes can have the same adjoint action. At the represented level their implementers can differ by a coefficient factor. That factor is invisible after restriction to either horizontal fiber; it belongs to the geometry relating the fibers.

This produces three structural regimes:

	$[\Omega_\varepsilon]_{\mathrm{band}}$	$[\Omega]^b_{\mathrm{cs}}$
globally strict	0 for all ε	0
strict fiberwise, twisted across scale	0 for all ε	$\neq 0$
twisted within a fiber	$\neq 0$ for some ε	$\neq 0$.

This is only the structural preview. The full three-sector formulation and its analytic consequences belong to the next note.

6.9 What has and has not been identified

The preceding construction relates several levels of Yang–Mills geometry, but it does not identify them literally.

- The fixed-thickness fiber $\Gamma_{\mathrm{cof}}(\varepsilon)$ is not the principal-bundle fiber P_x.
- The admissible comparison maps between local represented systems are not, in general, parallel transport by the original Yang–Mills connection A.
- The chart–scale class $[\Omega]^b_{\mathrm{cs}}$ is not asserted to be the curvature or a characteristic class of the classical principal bundle.
- Covariance of every local PDE under every admissible arrow does not imply strict representation descent.
- Cohomological strictifiability does not by itself prove analytic effectiveness of graph spaces, forms, domains, or scale limits.
- A coherent positive-thickness reconstruction does not by itself establish a zero-thickness Haag–Kastler net, a global Poincaré representation, a spectrum condition, an isolated one-particle sector, or a Yang–Mills mass gap.

The fixed-thickness fibers are not the fibers of the original principal gauge bundle. They arise at a later categorical level of the quantum reconstruction. Nevertheless, their appearance is not accidental. Yang–Mills theory begins with locally defined internal data whose comparison along spacetime paths is governed by a connection and may carry nontrivial holonomy. After passing to gauge-invariant Wilson observables and their local-normal representations, the reconstruction encounters the same structural problem one level higher: local physical realizations must be compared across charts, admissible Cauchy presentations, refinements, and scales, and those comparisons may fail to compose strictly. The reconstruction therefore replaces internal frames by local representations, ordinary parallel transport by admissible representation comparison, and ordinary holonomy by chart–scale coherence classes.

The connection has not disappeared. Its gauge-dependent local description has been removed, while its fundamental comparison problem survives first in Wilson holonomy and then, at a higher level, in the coherence structure of the reconstructed quantum theory.

7 The conceptual chain

The relationship developed in this note can be summarized as follows.

Yang–Mills assigns internal gauge frames over spacetime.

A connection compares those fibers; curvature and holonomy measure the failure of path-independent comparison.

Gauge-invariant observables remove dependence on the chosen local frame, but physical implementations can retain nontrivial gluing data.

Positive-thickness reconstruction organizes local quantum data over a thickness base E_{cof}.

Fixed-thickness fibers record what can be assembled at one ε; scale arrows record compatibility between different ε.

The distinction between coherence within a fiber and coherence across fibers leads to the later three-sector trichotomy.

The deepest continuity is therefore not an identity of objects but an identity of questions:

> At every level, which local choices are merely coordinates, which comparisons carry physical information, and what obstructs replacing the full comparison system by one globally strict object?

8 Conclusion and transition to the next note

Yang–Mills theory has a fiber description because gauge symmetry is the freedom to choose local internal frames, and the gauge field is the connection that compares those frames. Curvature, holonomy, and Wilson observables are not incidental constructions placed on top of this geometry; they are the natural quantities produced by it.

The reconstruction program inherits that local-to-global logic and applies it to quantum objects. It then introduces thickness as an additional organizing base. The fixed-thickness fibers collect the admissible local-normal and represented data available at one intrinsic support thickness, while scale arrows ask whether those fibers can be related coherently as the thickness tends toward zero.

The next note begins at that point. It will explain why

strict at every fixed thickness

is weaker than

$$\text{strict coherently across thickness,}$$

and why this distinction produces three reconstruction sectors rather than two.

A A modern reconstruction of the projection-drift mechanism

This appendix does not reproduce the lost one-loop calculation. It gives a clean operator-theoretic model of the lesson that the calculation appeared to contain.

Let $\mathcal{H}_{\rm kin}$ be a Hilbert space and let $G_a(\mu)$ be closed constraint operators with a common form domain. Assume the nonnegative form sum

$$Q(\mu) = \sum_a G_a(\mu)^* G_a(\mu)$$

defines a self-adjoint operator. Suppose zero is isolated from the remainder of the spectrum uniformly on a parameter interval. Then

$$P(\mu) = \mathbf{1}_{\{0\}}(Q(\mu))$$

is the spectral projection onto

$$\mathcal{H}_{\rm phys}(\mu) = \ker Q(\mu) = \bigcap_a \ker G_a(\mu).$$

If $Q(\mu)$ varies differentiably in a norm-resolvent sense, the Riesz formula

$$P(\mu) = \frac{1}{2\pi i}\oint_\Gamma (z - Q(\mu))^{-1}\, dz$$

gives

$$\dot{P}(\mu) = \frac{1}{2\pi i}\oint_\Gamma (z - Q(\mu))^{-1}\dot{Q}(\mu)(z - Q(\mu))^{-1}\, dz,$$

where Γ encloses only the zero spectral island.

Proposition 1 (Conditional projection drft)**.** *If*

$$(1 - P(\mu))\dot{Q}(\mu)P(\mu) \neq 0,$$

then $\dot{P}(\mu) \neq 0$. Consequently the physical subspace is not literally constant as an embedded subspace of $\mathcal{H}_{\rm kin}$.

Proof. If $\dot{P} = 0$, differentiating $QP = 0$ gives $\dot{Q}P + Q\dot{P} = \dot{Q}P = 0$. Applying $1 - P$ on the left yields $(1 - P)\dot{Q}P = 0$, contrary to the hypothesis. □

The converse need not be expressed solely through the displayed off-diagonal block without additional spectral hypotheses, but the criterion is sufficient for the intended diagnostic. It separates two situations.

First, if the renormalized constraints are related by an invertible recombination

$$G_a(\mu) = M_a{}^b(\mu) G_b(\mu_0)$$

with compatible domains, then the common kernel is unchanged and the physical projection may remain fixed. Second, if the correction acts nontrivially on the old kernel, then the physical range drifts.

When the gap persists and $P(\mu)$ is differentiable, Kato transport solves

$$\dot{U}(\mu) = [\dot{P}(\mu), P(\mu)]U(\mu), \qquad U(\mu_0) = I,$$

and satisfies

$$U(\mu)P(\mu_0)U(\mu)^* = P(\mu).$$

Thus projection drift does not imply absence of comparison. It implies that comparison is additional data rather than literal equality. On a multi-parameter base, the induced connection may have nontrivial holonomy. If the spectral gap closes or the rank changes, even the smooth Hilbert-bundle picture can fail and a weaker correspondence-based formulation may be needed.

Remark 2 (Relation to the reproduced manuscript)**.** *The language of an "operator ideal generated by unbounded constraints" is heuristic unless one first passes to bounded transforms or specifies a common operator-algebraic domain. Likewise, the claim that mixing outside such an ideal forbids every canonical unitary identification is too strong: a regular projection path itself supplies Kato transport. The defensible conclusion is more precise and more useful for the later theory: renormalization can obstruct* strict constancy *of the embedded physical subspace, thereby making transport and its coherence part of the structure.*

Claim boundary

This research note is explanatory. It records the geometric and conceptual motivation for the fiber language used in the positive-thickness reconstruction program. It does not claim that every quantum descent obstruction is determined by the topology of the classical gauge bundle. It does not construct a final zero-thickness Haag–Kastler net, a global Poincaré representation, a spectrum condition, vacuum non-leakage, or a Yang–Mills mass gap.

References

[1] C. N. Yang and R. L. Mills, "Conservation of isotopic spin and isotopic gauge invariance," *Physical Review* **96** (1954), 191–195.

[2] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, Vol. I, Wiley, New York, 1963.

[3] M. F. Atiyah and R. Bott, "The Yang–Mills equations over Riemann surfaces," *Philosophical Transactions of the Royal Society of London A* **308** (1983), 523–615.

[4] J. C. Baez and J. P. Muniain, *Gauge Fields, Knots and Gravity*, World Scientific, Singapore, 1994.

[5] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed., Springer, Berlin, 1996.

[6] J. R. Klauder, "Universal procedure for enforcing quantum constraints," `arXiv:hep-th/9901010` (1999).

[7] J. R. Klauder, "Quantization of constrained systems," `arXiv:hep-th/0003297` (2000).

[8] J. R. Klauder, "Metrical quantization," `arXiv:quant-ph/9804009` (1998).

[9] J. R. Klauder and S. V. Shabanov, "An introduction to coordinate-free quantization and its application to constrained systems," `arXiv:quant-ph/9804049` (1998).

[10] T. Kato, *Perturbation Theory for Linear Operators*, 2nd ed., Springer, Berlin, 1976.

[11] J. S. Little, "Thick spacetime reconstruction for Yang–Mills Wilson-tube systems: Local von Neumann nets and center-banded realizations," research manuscript, 2026.