Thick Wilson Loops in 2+1 Yang–Mills

Projection, admissible thin transport, and the silence of the gerbe

Can the admissible positive-thickness construction recover the known $2+1$ Yang–Mills mechanisms while correctly refusing to manufacture the $3+1$ associator obstruction?

OBJECT CLASS	GAT Research Notebook
CODE / VERSION	RN-YM-01 / v1.1
PRIMARY STATUS	**CONDITIONAL** — the finite projection, Wilson operator, relative class, and comparison calculation are exact; the full thick–thin analytic bridge remains conditional on stated nonleakage and Mosco hypotheses.
SELECTED BRANCH	Source-free, positive-thickness Hamiltonian carriers with compact lattice gauge group, Klauder–Shabanov projection, and a selected KKN infrared recovery branch.
AUTHOR / DATE	J. Scott Little — Independent Researcher, Columbus, Ohio / September 2026
CANONICAL OUTPUT	PDF generated from the GAT notebook source

CLAIM BOUNDARY

This notebook does not claim an unrestricted continuum construction, a complete $2+1$ mass-gap proof, or an exact full-theory string tension. It also does not exclude every ambient or enriched gerbe.

OPEN EDGE

Test whether any noncanonical enrichment can inhabit the surviving relative incision class without changing the dynamics of pure Yang–Mills.

HOW TO USE THIS NOTEBOOK

PREREQUISITES	Compact-group representation theory, elementary lattice gauge theory, closed quadratic forms, and basic relative cohomology.
CENTRAL TECHNIQUE	Follow one Wilson observable through projection, refinement, coercive nonleakage, Mosco convergence, heat-kernel reduction, and comparison-cube evaluation.
TARGET INSIGHT	Learn to distinguish existence from collapse, binary charge from ternary obstruction, and observable silence from topological disappearance.

Notebook orientation

QUESTION UNDER TEST

Can the same admissible machinery that detects a nontrivial associator in the $3+1$ construction recover the correct $2+1$ physics while calculating a trivial selected associator?

We construct bounded positive-thickness Wilson-loop operators for Hamiltonian Yang–Mills theory in $2+1$ dimensions by combining normalized character multiplication, heat regularization, and the gauge-invariant projection operator of Klauder and Shabanov. The projectors are organized over an admissible refinement system, and their naturality yields transported Ward identities. An exact telescoping formula separates failure of Wilson-operator naturality into projection, heat-transport, and holonomy-comparison defects, making the relation among the established methods explicit. We isolate a reduced Casimir coercivity element and state a transverse coercivity hypothesis sufficient for nonleakage. Mosco convergence, together with a uniform-energy estimate required when the collar time also tends to zero, then gives a conditional thick–thin reconstruction theorem. On the Karabali–Kim–Nair infrared branch, the limiting carrier recovers the familiar two-dimensional heat-kernel area mechanism.

The spacetime incision behaves differently from its $3+1$-dimensional counterpart. Its absolute local third cohomology vanishes, while a relative orientation class survives through Thom descent. A linked Wilson–'t Hooft pair retains a nontrivial finite center phase, but radial refinement transports that phase flatly; the selected three-direction associator is therefore trivial. The result is a relative bordism lift with surviving binary charge data, not an automatically generated gerbe. This provides a controlled benchmark showing that admissible-fibration methods distinguish the geometry appropriate to the dimension rather than importing the $3+1$ answer by force.

The gain is stronger than recovery alone. The admissible record is strictly finer than the endpoint observable: it retains the provenance of the Wilson operator, the reason particular data disappear, the relative class hidden by the projected boundary value, and the binary comparison maps needed to test the associator. In this precise relative sense the construction is more primitive than the single-carrier formulations it compares, although no claim is made that it is the most primitive possible layer.

CLAIM BOUNDARY

The fixed-thickness projection, boundedness, Ward compatibility, reduced Casimir coercivity, relative incision class, and electric–magnetic comparison are finite or operator-theoretic statements. Full positive-thickness coercivity, Mosco–Riesz reconstruction, and the identification with the KKN heat-kernel branch are conditional. No claim is made here of a complete continuum construction, an unrestricted $2+1$ mass-gap proof, or an exact full-theory string tension.

1 Introduction

Yang–Mills theory in $2+1$ dimensions is not solved in the unrestricted continuum sense. It is, however, a sharply controlled benchmark: the coupling is dimensionful, confinement and a mass gap are strongly supported by lattice calculations, and the Hamiltonian program of Karabali, Kim, and Nair (KKN) produces a successful infrared vacuum functional and string-tension formula [5, 6, 4]. A reconstruction mechanism intended for $3+1$ dimensions should therefore recover the appropriate $2+1$ mechanisms while also explaining which higher comparison structures disappear when the dimension is lowered.

The point of using $2+1$ is not merely to rehearse an easier theory. It is to ask whether the machinery knows that the answer should change. A dimensional test is meaningful only if the construction can recover what is present without manufacturing what is not.

The organizing principle used here is that the finite, positive-thickness, thin, represented, and physical objects need not be identified at the outset. They are placed in fibers over an admissible refinement category. The main questions are then typed questions: which operators exist in each fiber, which maps transport them, which identities descend, which classes survive collapse, and which possible obstructions remain uninhabited.

The gauge projection is not new. We use the group-averaging projector reviewed by Shabanov and developed with Klauder, including its application to Kogut–Susskind lattice gauge theory [15, 16]. The new role assigned to it is functorial: we require a family of projectors to be natural under admissible refinement of a two-dimensional time-zero lattice.

> **TECHNIQUE: READING ACROSS CARRIERS**
>
> A useful way to read this notebook is to resist asking where "the" Wilson loop lives. At finite lattice spacing it is a bounded multiplication operator. After projection it is a physical operator. Along refinement it is a lineage of compatible operators. In the thin limit it acts on a different carrier. On the infrared branch its expectation is evaluated by a two-dimensional heat kernel. These are related realizations, not literally one unchanging object. The purpose of admissible transport is to state which identifications are earned.

2 What the established methods see

This note uses several established constructions, but it asks a question that none of them was designed to ask by itself. The issue is not which formalism is the correct one. The issue is what remains visible when one passes from one formalism, scale, or carrier to another.

GAT is therefore not another quantization prescription placed beside lattice Hamiltonian theory, gauge fixing, the KKN variables, or the two-dimensional heat kernel. It is a comparison architecture. Its additional objects are the admissible maps between already meaningful descriptions and the defects of their composition.

2.1 The fixed-lattice Hamiltonian picture

For a finite spatial lattice Λ, the unreduced carrier is

$$\mathcal{H}_\Lambda = L^2\left(G^{E(\Lambda)}, \prod_{e\in E(\Lambda)} \mathrm{d}U_e\right),$$

and the lattice gauge group $G^{V(\Lambda)}$ acts at the ends of each edge by

$$U_e \longmapsto g_{s(e)} U_e g_{t(e)}^{-1}.$$

The Kogut–Susskind Hamiltonian separates into electric and magnetic pieces,

$$H_\Lambda = H_{E,\Lambda} + H_{B,\Lambda},$$

where the electric term is built from group Laplacians on edge variables and the magnetic term from plaquette holonomies [7]. At a fixed lattice this supplies a concrete Hilbert space, a concrete constraint, and a concrete dynamics.

What it does not automatically supply is a canonical comparison among all choices of lattice, collar thickness, regional decomposition, and continuum identification. One can construct refinement maps, but their compatibility with projection, spectral selection, and incision data becomes an additional problem. GAT begins at precisely that additional problem. It regards

$$\Lambda, \epsilon \longmapsto (\mathcal{H}_{\Lambda,\epsilon}, Q_{\Lambda,\epsilon})$$

as fibered data and asks which arrows between these fibers are admissible.

> **LEARNING CHECKPOINT: THE FINITE CARRIER**
>
> Take a lattice with E oriented edges and V vertices. Verify that a gauge transformation needs one element of G at each vertex, while a kinematic state depends on one element of G at each edge. This is why the gauge group is G^V but the unreduced configuration carrier is $L^2(G^E)$.

2.2 Gauge fixing, BRST language, and projection

The Faddeev–Popov and BRST traditions encode the gauge quotient through gauge conditions, determinant or ghost data, and a cohomological selection of physical states [3]. They are indispensable, especially in perturbative field theory. Their global difficulty is also familiar: a single gauge condition need not define a global section of the orbit map, and Gribov copies make the geometry of the quotient impossible to ignore.

Klauder and Shabanov take a different route. They impose the quantum constraint by averaging the unitary gauge action. The physical subspace is the range of an operator, not the image of a chosen global gauge slice. Shabanov's analysis also emphasizes that the geometry and operator ordering on the orbit space survive into the physical dynamics [15, 16].

That is why the projection method is the natural input here. It permits the question

$$S_{\epsilon'\epsilon}\mathbb{E}_\epsilon \stackrel{?}{=} \mathbb{E}_{\epsilon'} S_{\epsilon'\epsilon}$$

to be asked before a global coordinate choice has identified the two physical spaces. This is not an assertion that gauge fixing or BRST is incorrect. It is a decision to retain the quotient as an operatorial object while the refinement geometry is still visible.

> **COMMON MISTAKE: CONSTRAINT IMPOSITION IS NOT GAUGE SLICING**
>
> The equality $\mathcal{H}^{\text{phys}} = \operatorname{Ran}\mathbb{E}$ does not choose one representative from every classical gauge orbit. It selects the invariant quantum vectors. Confusing those operations hides the reason group averaging remains meaningful when no global gauge slice exists.

2.3 The KKN gauge-invariant Hamiltonian

The KKN construction reorganizes the continuum Hamiltonian theory in terms of gauge-invariant matrix variables [5, 6, 13]. In complex coordinates on the spatial surface, a connection is parametrized locally by

$$A_z = -\partial M\, M^{-1}, \qquad A_{\bar{z}} = M^{\dagger -1}\bar{\partial} M^{\dagger}.$$

Under an ordinary gauge transformation $M \mapsto gM$, the Hermitian field

$$H = M^{\dagger} M$$

is gauge invariant. The physical inner product is not a flat integral in H; its measure contains the Wess–Zumino–Witten weight

$$\mathrm{d}\mu(\mathcal{C}) = \mathrm{d}\mu(H)\, e^{2c_A S_{\mathrm{WZW}}(H)}.$$

The regularized kinetic operator then produces the natural mass scale

$$m = \frac{e^2 c_A}{2\pi}.$$

At long wavelength, the leading vacuum measure reduces the relevant expectation values to a two-dimensional Yang–Mills theory with an effective coupling. This is the source of the KKN area-law calculation and its successful string-tension estimate.

The KKN formalism sees far more dynamics than the reduced character sector used in our coercivity lemma. Conversely, it normally presents the final gauge-invariant carrier after the relevant variables and measure have been chosen. The present note does not rederive KKN. It uses the KKN infrared branch as a recovery target: if the admissible thick–thin machinery is sound, its reduced limit should reproduce the same heat-kernel mechanism without claiming that the finite-thickness construction is globally identical to the KKN carrier.

> **TECHNIQUE: CHANGING VARIABLES WITHOUT FLATTENING THE MEASURE**
>
> The field $H = M^{\dagger} M$ is gauge invariant, but its physical measure is not ordinary flat integration over H. The WZW factor records the Jacobian and the orbit-space geometry created by the change of variables. The mass scale does not arise merely because A was rewritten as M; it arises through the regularized kinetic operator acting with that nontrivial measure. This is why the KKN result is a recovery target rather than a symbolic change of notation.

2.4 Two-dimensional Yang–Mills as an area-enriched theory

Two-dimensional Yang–Mills is exactly organized by heat kernels and gluing. On a boundary circle its natural carrier is the class-function space

$$L^2(G)^{\mathrm{Ad}\, G},$$

and a cylinder of area A acts by

$$Z(\mathrm{Cyl}_A) = e^{-AQ_C}.$$

Gluing cylinders adds their areas and composes their heat operators. Disk and pair-of-pants amplitudes are governed by the same character expansion [18, 1].

This is close to a bordism field theory, but it is not an ordinary metric-free TQFT: the area is retained as structure. That distinction matters here. The small collar area a_ϵ is driven to zero so its cylinder operator tends to the identity. The area $A(D)$ of the physical spanning disk is held fixed and continues to carry the area law. The same heat kernel appears in both places, but the two areas have different mathematical jobs.

LEARNING CHECKPOINT: SEMIGROUP GLUING

Apply the cylinder rule twice:

$$e^{-A_2 Q_C} e^{-A_1 Q_C} = e^{-(A_1+A_2) Q_C}.$$

The equality is the analytic form of gluing two cylinders. Now notice the asymmetry in the present problem: $a_\epsilon \to 0$ for the collar, while $A(D)$ is fixed for the spanning disk. Sending both areas to zero would erase the area-law calculation rather than derive it.

2.5 Loop algebra and Chern–Simons framing

The Wilson–'t Hooft algebra detects linked electric and magnetic operators by means of a center character. This powerful topological statement survives in Hamiltonian and lattice descriptions. It is, however, binary: it compares two ordered insertions.

Chern–Simons theory goes further. Wilson lines become framed ribbon objects, and braiding, fusion, and associators are organized by a nontrivial tensor category [17]. But Chern–Simons theory has its own level and its own topological dynamics. Importing that structure into the shrinking collar would alter pure Yang–Mills rather than explain its limit.

GAT asks the intermediate question: does the binary center phase vary when it is transported through the independent regional and refinement directions? If it varies coherently, it may inhabit a three-cocycle. If it is transported unchanged, the comparison cube is flat. In $2+1$ the latter occurs.

TECHNIQUE: BINARY VERSUS TERNARY INFORMATION

A commutation phase compares two ordered operations. An associator compares two ways of composing three operations. A nontrivial binary phase therefore does not by itself imply a nontrivial gerbe witness. One must transport the binary comparison through an independent direction and calculate the resulting cube holonomy. That extra comparison is the decisive dimensional test in this notebook.

2.6 The initial amplification supplied by GAT

The relationship among the methods can be summarized as follows.

Method	Primary carrier	What it makes visible	GAT use
Kogut–Susskind	$L^2(G^{E(\Lambda)})$	Finite Hamiltonian, Gauss law, plaquette dynamics	Organize changing lattices and thicknesses
Klauder–Shabanov	$\operatorname{Ran} \mathbb{E}$	Gauge-invariant physical subspace without a global gauge slice	Test projector naturality and Ward transport
KKN	Gauge-invariant $H = M^\dagger M$ with WZW measure	Continuum orbit-space dynamics and infrared mass scale	Supply the controlled infrared recovery target
Two-dimensional YM	$L^2(G)^{\operatorname{Ad} G}$	Exact heat-kernel gluing and area dependence	Identify the reduced cylinder and disk amplitudes
Wilson–'t Hooft algebra	Electric and magnetic loop sectors	Center phase from linking	Supply the surviving binary comparison
Chern–Simons	Framed ribbon category	Braiding, framing, and categorical associators	Clarify which extra topological data pure YM does not acquire
Admissible fibrations	Distributed carriers with selected transports	Naturality, nonleakage, comparison defects, and survival under collapse	Keep the passages themselves as mathematical data

The amplification is therefore specific. Earlier methods determine the states, operators, measures, and amplitudes appropriate to their chosen carriers. GAT retains the passages among those carriers long enough to distinguish four possibilities that a final strictified presentation may merge: data that were never present, data removed by projection, data canceled by a Ward identity, and data preserved through collapse but invisible to the final observable.

3 What more is gained: retained mathematical information

The strongest justified claim is not merely that GAT gives another derivation of a Wilson-loop limit. On the selected branch it carries a *strictly finer record* than the final projected operator or its expectation value. The ordinary endpoint record may be written schematically as

$$\mathfrak{O}_0 = (\mathcal{H}_0^{\mathrm{phys}}, Q_0, W_{R,0}, \langle W_{R,0} \rangle).$$

The admissible record contains this endpoint together with the family from which it descended:

$$\mathfrak{A} = \big(\{\mathcal{H}_i, \mathbb{E}_i, Q_i, W_{R,i}\}_i, \{S_{ji}\}_{i\to j}, \{\mathfrak{D}_{ji}\}_{i\to j}, [B_i, \partial B_i], \Xi_2, \Xi_3\big).$$

Here the index i can carry lattice, thickness, regional, and selected spectral information. Passing from $\mathfrak{A}$ to $\mathfrak{O}_0$ forgets the intermediate carriers, the comparison maps, their defects, and any class annihilated by the chosen observable.

Proposition 3.1 (Relative information dominance)**.** *For the selected $2+1$ construction, the endpoint Wilson operator does not determine the admissible record from which it descends. In particular, the projected endpoint can be insensitive to the relative incision generator even though that generator is preserved by every admissible radial refinement. Consequently, the admissible record is strictly more informative than its single-carrier Wilson observable.*

Proof. The endpoint observable is obtained by projection, thin transport, and restriction to the selected carrier. The two boundary-face contributions are canceled there by the transported Ward identity. Nevertheless, Theorem 12.1 identifies a nonzero class $[B_\epsilon, \partial B_\epsilon]$ preserved under radial refinement. Since the endpoint observable records the cancellation but does not determine this relative generator, information has been lost in the passage $\mathfrak{A} \mapsto \mathfrak{O}_0$. No canonical inverse can reconstruct the full admissible record from that endpoint alone. □

The proposition makes "more primitive" precise in a relative, rather than absolute, sense. Say that a description X precedes a description Y in the information order when Y is obtained from X by specified projection, descent, limiting, or forgetting operations, while X cannot be canonically reconstructed from Y. Then the present admissible system precedes the final single-carrier observable. It is pre-projection in its kinematic entries, pre-limit in its positive-thickness entries, and pre-strictification in its comparison entries. This is a mathematical statement about the direction of the maps; it is not a claim that GAT is the ultimate ontology of Yang–Mills.

What is gained is therefore explicit:

1. **Provenance.** The limiting Wilson operator is accompanied by the finite operators, projections, and transports that produced it.
2. **Mechanism of disappearance.** The construction distinguishes NEVER PRESENT, PROJECTED OUT, WARD-CANCELED, DESCENDED, and OBSERVABLE-SILENT. Equality of final observables does not merge these five histories.
3. **Latent regional topology.** A relative class can survive even when its boundary representative vanishes after projection. The topology is not inferred from the final Wilson value; it is transported independently.
4. **Order of operations.** Projection, refinement, transverse selection, and the thin limit are separately typed. Their commutation is a claim to prove, not a notational convenience.
5. **Negative higher-structure information.** Because the binary comparison maps have been retained, $\Xi_3 = 1$ is a calculated flatness result. A single operator algebra can only make that question disappear in advance.
6. **Controlled extensibility.** A different theory or a richer branch can alter one comparison map without forcing a rewrite of every endpoint formula. The architecture identifies exactly where genuinely new structure would have to enter.

> **WHAT BECAME VISIBLE**
>
> This is why the construction can show something that an observable-first presentation does not show. The additional information is not a new numerical Wilson-loop value. It is the mathematical history of that value: which finite object existed, which quotient was taken, which identity caused a boundary term to vanish, which class survived that vanishing, and which higher comparison was tested and found flat.

The claim also has a deliberate ceiling. There may be still finer layers: microlocal data, jets of the comparison maps, measure-level provenance, higher comparison cells, or other enrichments not retained here. GAT should therefore be called *more primitive than the compared endpoint formulations on this branch*, not the most primitive possible description. Its demonstrated advantage is that it delays irreversible identifications long enough to measure what those identifications erase.

4 The explicit limit architecture

Several limits occur in the argument, and they are not interchangeable. The notation $\epsilon \downarrow 0$ is only one part of the architecture. The table below states the variable, topology, mathematical job, and status of every limiting passage used or deliberately withheld.

Passage	Mode of convergence	What it does	Status here
Admissible refinement	Directed arrows $i \to j$	Compares changing lattices, regions, and thicknesses without first identifying their carriers	Exact on the selected system; no unrestricted continuum net is claimed
Transverse collapse	$\epsilon \downarrow 0$ with $\|(1 - Z_\epsilon)\Psi_\epsilon\| = O(\epsilon)$	Removes positive-thickness excitations at bounded energy	Conditional on the uniform coercive estimate
Mosco form limit	Weak liminf plus strong recovery sequence	Identifies q_0 and gives strong-resolvent convergence of Q_ϵ	Conditional on (A3)–(A4) and the chosen varying-carrier identifications
Collar-time limit	$a_\epsilon \downarrow 0$ and strong convergence on a uniform-energy core	Sends $e^{-a_\epsilon Q_\epsilon/2}$ to the identity even while Q_ϵ varies	Conditional; controlled by (6)
Wilson-operator limit	Strong operator convergence	Produces $\mathbb{E}_0 M_{w_R,0} \mathbb{E}_0$ from $W_{R,\epsilon}$	Conditional Theorem 10.1
Riesz projection limit	Strong convergence of spectral projections	Transports a selected isolated spectral sector	Requires an additional uniform spectral window
KKN infrared passage	Long-wavelength selected branch	Identifies the two-dimensional heat-kernel mechanism and effective coupling	Recovery target, not derived here from the unrestricted theory
Radial comparison	$\epsilon' \le \epsilon$, then cofinal collapse	Preserves $[B_\epsilon, \partial B_\epsilon]$ and the scalar Ξ_2	Exact for the stated orientation-preserving radial maps
Physical area	$A(D)$ held fixed	Retains the heat-kernel propagation and area exponent	Explicitly not sent to zero
Infinite volume / full continuum	Not taken	Would remove the remaining lattice, volume, and branch restrictions	Outside the claim boundary

The order is part of the theorem. First construct $W_{R,\epsilon}$ and impose gauge projection at fixed $\epsilon > 0$. Next prove naturality, Ward transport, and the coercive nonleakage estimate along admissible arrows. Only then take the Mosco and collar-time limits together. Finally identify the resulting selected carrier with the KKN infrared heat-kernel branch. Symbolically,

$$\boxed{\begin{array}{c} \text{finite existence} \longrightarrow \text{projected lineage} \longrightarrow \text{nonleakage} \\ \operatorname*{s\text{-}lim}_{\substack{\epsilon\downarrow 0\\ a_\epsilon \downarrow 0}} W_{R,\epsilon} \longrightarrow \text{infrared identification} \end{array}}$$

where the displayed simultaneous strong limit is shorthand for Theorem 10.1, not a free interchange of two limits.

COMMON MISTAKE: DO NOT REPLACE THE ARCHITECTURE BY "TAKE EPSILON TO ZERO"

Setting $\epsilon = 0$ at the beginning destroys the positive-thickness carrier on which boundedness, projection naturality, the relative regional class, and the nonleakage estimate are established. Sending $A(D)$ to zero as well would also erase the physical area law. The result depends on retaining the finite objects long enough to prove which of their data survive collapse.

5 Admissible positive-thickness data

Let $\mathsf{I}^{\mathrm{adm}}$ be a directed category of positive thicknesses $\epsilon > 0$. An object carries the data

$$(\mathcal{H}_\epsilon, \mathcal{G}_\epsilon, U_\epsilon, q_\epsilon, Q_\epsilon, Z_\epsilon),$$

where $\mathcal{H}_\epsilon$ is a Hilbert carrier, $\mathcal{G}_\epsilon$ is a compact lattice gauge group, U_ϵ is a unitary representation, q_ϵ is a densely defined closed nonnegative form with generator Q_ϵ, and Z_ϵ selects the transverse zero mode. An admissible arrow $\epsilon \to \epsilon'$ carries a bounded transport map

$$S_{\epsilon'\epsilon} : \mathcal{H}_\epsilon \longrightarrow \mathcal{H}_{\epsilon'}$$

and a continuous gauge homomorphism

$$\rho_{\epsilon'\epsilon} : \mathcal{G}_\epsilon \longrightarrow \mathcal{G}_{\epsilon'}.$$

We impose the following working hypotheses.

(A1) Gauge equivariance:
$$S_{\epsilon'\epsilon} U_\epsilon(g) = U_{\epsilon'}(\rho_{\epsilon'\epsilon} g) S_{\epsilon'\epsilon}.$$

(A2) Haar compatibility: $(\rho_{\epsilon'\epsilon})_* \mu_\epsilon = \mu_{\epsilon'}$ on the selected gauge image.

(A3) Form compatibility sufficient for Mosco convergence after the chosen identifications of varying carriers.

(A4) Uniform transverse coercivity as stated in Section 9.

(A5) Compatibility of the normalized Wilson multipliers on the selected lineages.

The assumptions separate what is finite and exact from what is required to reach the continuum thin carrier.

WHAT BECAME VISIBLE

The tuple $(\mathcal{H}_\epsilon, \mathcal{G}_\epsilon, U_\epsilon, q_\epsilon, Q_\epsilon, Z_\epsilon)$ is a compact record of six different jobs: states, gauge transformations, their action, energy, the energy generator, and the chosen transverse zero mode. Listing all six prevents an argument about one of them from being silently reused as an argument about another.

This separation matters. Existence is earned at positive thickness. Collapse is a second theorem, with its own hypotheses, and no conclusion below is allowed to pass silently from one side of that distinction to the other.

6 The Klauder–Shabanov projection and Ward identities

For normalized Haar measure μ_ϵ, define

$$\mathbb{E}_\epsilon = \int_{\mathcal{G}_\epsilon} U_\epsilon(g)\, \mathrm{d}\mu_\epsilon(g). \tag{1}$$

This is the lattice specialization of the group-averaging projector

$$\widehat{P} = \int_{\mathcal{G}} \mathrm{d}\mu_{\mathcal{G}}(\omega) e^{i\omega^a \widehat{\sigma}_a}$$

used by Klauder and Shabanov [15, 16]. The projector is therefore asked to do precisely what it was designed to do: construct the physical carrier before a gauge choice or a thin-limit identification has erased the geometry of the passage.

Proposition 6.1 (Projector naturality)**.** *Under* (A1)–(A2),

$$S_{\epsilon'\epsilon}\mathbb{E}_\epsilon = \mathbb{E}_{\epsilon'} S_{\epsilon'\epsilon}.$$

Hence the physical spaces $\mathcal{H}_\epsilon^{\mathrm{phys}} = \operatorname{Ran} \mathbb{E}_\epsilon$ *form an admissible subfibration.*

Proof. Move $S_{\epsilon'\epsilon}$ through the Bochner integral by boundedness, use equivariance, and then push normalized Haar measure forward through $\rho_{\epsilon'\epsilon}$. □

For a student, it is worth writing the three lines that the short proof compresses:

$$\begin{aligned} S_{\epsilon'\epsilon}\mathbb{E}_\epsilon &= \int_{\mathcal{G}_\epsilon} S_{\epsilon'\epsilon} U_\epsilon(g)\ \mathrm{d}\mu_\epsilon(g) \\ &= \int_{\mathcal{G}_\epsilon} U_{\epsilon'}(\rho_{\epsilon'\epsilon} g) S_{\epsilon'\epsilon}\ \mathrm{d}\mu_\epsilon(g) \\ &= \mathbb{E}_{\epsilon'} S_{\epsilon'\epsilon}. \end{aligned}$$

The first equality uses boundedness to exchange the map and the Bochner integral, the second uses equivariance, and the third is exactly where Haar compatibility enters. Each hypothesis has one visible job.

Differentiating the equivariance relation at the identity yields the transported Ward relation

$$S_{\epsilon'\epsilon} G_\epsilon(\xi) = G_{\epsilon'}(\,\mathrm{d}\rho_{\epsilon'\epsilon}\xi) S_{\epsilon'\epsilon}, \tag{2}$$

on a common invariant core. Thus Gauss-law identities are not separately reimposed after refinement; they are transported with the physical subspaces.

If Haar compatibility is approximate rather than exact, the defect is controlled by total variation:

$$\big\| S\mathbb{E} - \mathbb{E}' S \big\| \le \|S\|\, \|\rho_* \mu - \mu'\|_{\mathrm{TV}}.$$

LEARNING CHECKPOINT: WARD TRANSPORT

Differentiate the equivariance identity along $g(t) = \exp(t\xi)$ at $t = 0$. Track the domain carefully: the generators may be unbounded, so the identity is first asserted on a common invariant core. This is the infinitesimal origin of (2); it is not obtained by differentiating an arbitrary operator identity on all of Hilbert space.

7 Existence of thick Wilson loops

Let R be an irreducible unitary representation of the compact group G and normalize its character by

$$w_R(U) = \frac{\chi_R(U)}{d_R}.$$

Since $|\chi_R(U)| \le d_R$, multiplication by w_R satisfies

$$\|M_{w_R}\| \le 1.$$

TECHNIQUE: NORMALIZING THE CHARACTER

For a unitary $d_R \times d_R$ matrix, the trace has absolute value at most d_R. Dividing by d_R turns the character into a uniformly bounded class function. This elementary normalization is what makes multiplication a contraction before any continuum argument begins. Without it, the operator would still be bounded at fixed R, but the clean norm-one estimate would be lost.

For collar area $a_\epsilon > 0$, define

$$\mathcal{W}_{R,\epsilon} = \mathbb{E}_\epsilon e^{-a_\epsilon Q_\epsilon/2} M_{w_R} e^{-a_\epsilon Q_\epsilon/2} \mathbb{E}_\epsilon. \tag{3}$$

Theorem 7.1 (Fixed-thickness existence). *For every admissible $\epsilon > 0$, $\mathcal{W}_{R,\epsilon}$ is a bounded physical operator and*

$$\|\mathcal{W}_{R,\epsilon}\| \le 1.$$

It is compatible with gauge projection and the transported Ward identities.

Proof. Haar averaging is an orthogonal projection, the heat semigroup is a contraction, and M_{w_R} is a contraction. The norm estimate follows by composition. Gauge compatibility follows from Proposition 6.1, the gauge invariance of the character, and (2). □

WHAT BECAME VISIBLE

Existence is obtained by composing three familiar contractions: projection, heat evolution, and normalized-character multiplication. The new issue is not whether any one factor exists. It is whether the factors continue to match as the lattice, collar, and physical carrier change together.

8 One loop through the comparison architecture

The preceding survey can now be read as an actual calculation. Fix a spatial loop $C \subset \Sigma^2$ bounding a disk D, choose a positive collar thickness ϵ, and let C_ϵ be a lattice representative. The same observable appears differently at successive stages, but the stages should not be identified prematurely.

8.1 Kinematic holonomy and physical projection

On the unreduced lattice carrier, the Wilson insertion is multiplication by the normalized character of the loop holonomy,

$$M_\epsilon = M_{w_R \circ \mathrm{Hol}_{C_\epsilon}}, \qquad \|M_\epsilon\| \le 1.$$

This is the Kogut–Susskind viewpoint at the level needed here: the lattice supplies the edge variables and their holonomy. It does not yet say which vectors are physical or how this particular representative is compared with a representative on another lattice.

Klauder–Shabanov projection supplies the first passage. The bounded physical operator is

$$\mathcal{W}_{R,\epsilon} = \mathbb{E}_\epsilon H_\epsilon M_\epsilon H_\epsilon \mathbb{E}_\epsilon, \qquad H_\epsilon = e^{-a_\epsilon Q_\epsilon/2}.$$

The two copies of $\mathbb{E}_\epsilon$ record input and output projection; the two heat factors regularize the insertion across the positive collar. Because the character is conjugation invariant, M_ϵ preserves the Gauss-law sector. The transported Ward identity then says more than gauge invariance at one scale: it says that the identification of physical sectors is compatible with the chosen refinement arrow.

8.2 The naturality defect and what each term means

Let $S = S_{\epsilon'\epsilon}$ be an admissible comparison map and abbreviate target-fiber objects by primes. Define the Wilson naturality defect

$$\mathfrak{D}_{\epsilon'\epsilon} = S\mathcal{W}_{R,\epsilon} - \mathcal{W}_{R,\epsilon'} S.$$

A direct telescoping insertion gives the exact identity

$$\begin{aligned}\mathfrak{D}_{\epsilon'\epsilon} = {} & (S\mathbb{E} - \mathbb{E}'S)HMH\mathbb{E} \\ & + \mathbb{E}'(SH - H'S)MH\mathbb{E} \\ & + \mathbb{E}'H'(SM - M'S)H\mathbb{E} \\ & + \mathbb{E}'H'M'(SH - H'S)\mathbb{E} \\ & + \mathbb{E}'H'M'H'(S\mathbb{E} - \mathbb{E}'S). \end{aligned} \tag{4}$$

TECHNIQUE: TELESCOPING A COMPOSITE DEFECT

To derive (4), move S from the far left of $\mathbb{E}HMH\mathbb{E}$ to the far right one factor at a time. At each move, add and subtract the expression in which only that factor has been replaced by its primed counterpart. This is the operator analogue of inserting $a - b = (a - c) + (c - b)$ repeatedly. The method is elementary, but it is powerful: every summand names the exact stage at which naturality may fail.

This formula is the comparison architecture in miniature. The first and last terms are projection defects; they vanish under projector naturality. The second and fourth are heat-transport defects; their control belongs to form compatibility, Mosco convergence, and the simultaneous small-collar estimate. The middle term is the holonomy-representative defect; it is controlled by compatibility of the selected Wilson lineage. A calculation made only after all carriers have been identified would replace these five questions by the single statement that the limiting operator exists. Equation (4) records why it exists and where it could have failed.

For contractions $\mathbb{E}, H, M$ and bounded S, the same identity yields the transparent estimate

$$\|\mathfrak{D}_{\epsilon'\epsilon}\| \le 2\|S\mathbb{E} - \mathbb{E}'S\| + 2\|SH - H'S\| + \|SM - M'S\|,$$

whenever the displayed defects are bounded on the selected carrier. In the exact finite system the projection terms are zero. In the thin-limit problem, the remaining terms need only vanish along the admissible lineage; no global unitary identification of every lattice carrier is required.

COMMON MISTAKE: OPERATOR-NORM CONTROL IS STRONGER THAN LINEAGE CONTROL

The displayed norm estimate is useful when bounded operator defects are available. Mosco convergence generally supplies strong convergence on chosen vectors, not operator-norm convergence on the entire carrier. For the thick–thin theorem, it is enough to control the defect along a dense admissible family and then use the uniform contraction bound to extend by density.

8.3 Infrared recovery, incision, and the surviving comparison

After transverse nonleakage, the selected zero-mode lineage lands in the class-function carrier $\mathcal{K}_C = L^2(G)^{\mathrm{Ad}\,G}$. On the KKN infrared branch, propagation across the fixed disk D is represented by the two-dimensional heat kernel. Character orthogonality then turns the same label R into

$$\langle W_R(C)\rangle = \exp\left[-\frac{1}{2}g_{\mathrm{eff}}^2 C_2(R)A(D)\right].$$

This is a recovery statement, not an identification of formalisms. KKN supplies the controlled infrared target and the effective scale; two-dimensional Yang–Mills evaluates the selected amplitude; admissible transport certifies that the positive-thickness physical operator reaches that target without transverse leakage.

The same passage has a topological remainder. The two boundary faces of the incision cancel in the projected observable by the Gauss-law Ward identity, while the relative class

$$[B_\epsilon, \partial B_\epsilon] \in H^3(B_\epsilon, \partial B_\epsilon; \mathbb{Z})$$

survives radial refinement. A linked magnetic insertion also retains the binary center phase $\omega_R(z)$. What does not survive is a nontrivial selected associator: refinement carries the same binary phase on both faces of the comparison cube, so their quotient is one.

Thus the worked loop separates three conclusions that are often compressed into one: the operator becomes physical by projection, its thin representative exists by analytic transport, and its incision data survive only in the degree supported by the $2+1$ geometry. The established methods provide the content at each stage. GAT exposes and tests the passages between them.

9 Coercivity and transverse nonleakage

The reduced Wilson carrier is

$$\mathcal{K}_C = L^2(G,\, \mathrm{d}U)^{\mathrm{Ad}\,G}.$$

With

$$Q_C = \frac{g_{\mathrm{eff}}^2}{2}(-\Delta_G),$$

the Peter–Weyl character basis gives

$$Q_C\chi_R = \frac{g_{\mathrm{eff}}^2}{2}C_2(R)\chi_R.$$

If $c_G = \min_{R\neq\mathbf{1}} C_2(R)$, then

Proposition 9.1 (Reduced Casimir coercivity)**.** *For $f = \sum_R f_R\chi_R$ in the form domain,*

$$\mathfrak{q}_C[f] = \frac{g_{\mathrm{eff}}^2}{2}\sum_R C_2(R)|f_R|^2 \geq \frac{g_{\mathrm{eff}}^2 c_G}{2}\left\| f - \int_G f\,\mathrm{d}U\right\|^2.$$

This is a gap on the reduced Wilson sector. It is not, by itself, the full $2+1$ Yang–Mills mass gap.

TECHNIQUE: COERCIVITY AS A POINCARE INEQUALITY

The constant character is the zero mode of $-\Delta_G$. Removing its average leaves only nontrivial representations, whose Casimir eigenvalues are bounded below by c_G. Proposition 9.1 is therefore a compact-group Poincare inequality written in the character basis. It controls the distance from the zero mode; it does not yet control every excitation of the full Yang–Mills Hamiltonian.

Coercivity is the conceptual hinge. The projector identifies the physical states, but coercivity determines whether unwanted transverse structure can remain energetically present as the collar is removed.

Let A_ϵ identify the transverse zero-mode sector with $\mathcal{K}_C$ and let J_ϵ be a recovery embedding. The positive-thickness bridge is the following hypothesis:

$$q_\epsilon[\Psi] \geq \mathfrak{q}_C[A_\epsilon Z_\epsilon \Psi] + c_\perp \epsilon^{-2} \|(1 - Z_\epsilon)\Psi\|^2 - r_\epsilon \|\Psi\|^2, \tag{5}$$

where $c_\perp > 0$ and $r_\epsilon \to 0$.

Corollary 9.2 (Transverse nonleakage)**.** *If $\|\Psi_\epsilon\|$ and $q_\epsilon[\Psi_\epsilon]$ are uniformly bounded, then*

$$\|(1 - Z_\epsilon)\Psi_\epsilon\| = O(\epsilon).$$

Indeed, rearranging (5) gives

$$c_\perp \epsilon^{-2} \|(1 - Z_\epsilon)\Psi_\epsilon\|^2 \leq q_\epsilon[\Psi_\epsilon] + r_\epsilon \|\Psi_\epsilon\|^2,$$

after discarding the nonnegative reduced form. Uniform boundedness of the right-hand side forces the transverse norm to be $O(\epsilon)$. This is the entire nonleakage mechanism: transverse mass becomes too expensive at the rate ϵ^{-2}.

COMMON MISTAKE: A SPECTRAL GAP IS NOT AUTOMATICALLY A MASS GAP

The reduced Casimir lower bound and the transverse penalty certify the selected Wilson lineage. A physical mass-gap theorem would additionally require the full physical Hamiltonian, its vacuum sector, and a uniform lower bound on all orthogonal physical excitations. The notebook keeps that stronger claim outside its boundary.

10 Mosco–Riesz thick–thin transport

We use the varying-Hilbert-space formulation of Mosco convergence developed for spectral structures by Kuwae and Shioya [9, 10, 8]. The two required statements are:

$$\begin{aligned} &\text{(M1)} \quad \Psi_\epsilon \rightharpoonup \Psi \Longrightarrow \liminf_{\epsilon\to 0} q_\epsilon[\Psi_\epsilon] \geq q_0[\Psi], \\ &\text{(M2)} \quad \text{for each } \Psi \in \mathrm{Dom}(q_0), \text{ there are } \Psi_\epsilon \to \Psi \text{ with } \limsup_{\epsilon\to 0} q_\epsilon[\Psi_\epsilon] \leq q_0[\Psi]. \end{aligned}$$

TECHNIQUE: THE TWO HALVES OF MOSCO CONVERGENCE

Condition (M1) says energy cannot disappear in a weak limit. Condition (M2) says every limiting state can be approximated without paying excess energy. The first prevents false low-energy states; the second prevents the limit from being too small. Together they identify the limiting closed form and yield strong resolvent convergence of its generator.

Mosco convergence gives strong resolvent convergence and fixed-time semigroup convergence. Fixed-time convergence alone is insufficient here because the collar time a_ϵ also tends to zero. The needed uniform estimate follows from

$$(1 - e^{-ax})^2 \le ax,$$

namely

$$\|(I - e^{-aQ})\Psi\|^2 \le a\, q_Q[\Psi]. \tag{6}$$

The scalar inequality behind (6) can be checked from the spectral theorem. For each spectral value $x \ge 0$, $|1 - e^{-ax}|^2 \le ax$; integrating against the spectral measure of Q gives the vector estimate. This extra diagonal bound is needed because both the operator Q_ϵ and the time a_ϵ vary.

Theorem 10.1 (Admissible thick–thin reconstruction)**.** *Assume* (A1)–(A5)*,* (5)*, Mosco convergence, $a_\epsilon \to 0$, and uniform energy bounds on a dense family of admissible lineages. Then*

$$\mathcal{W}_{R,\epsilon} \xrightarrow[\epsilon \downarrow 0]{\mathrm{s}} \mathbb{E}_0 M_{w_R,0} \mathbb{E}_0$$

on the selected physical carrier.

Proof. Corollary 9.2 removes transverse leakage. Mosco convergence identifies the limiting closed form and its semigroup. Estimate (6), together with $a_\epsilon \to 0$ and uniform energy, shows that both collar semigroups converge to the identity on the chosen core. Projector naturality and multiplier compatibility then give convergence on that core. The uniform contraction bound from Theorem 7.1 extends the convergence by density. □

If, in addition, a spectral window about the ground state remains uniformly isolated, the corresponding spectral projections converge strongly. We treat this spectral separation as an additional hypothesis rather than as a consequence of strong resolvent convergence alone.

> **LEARNING CHECKPOINT: WHAT CONVERGES, AND IN WHICH TOPOLOGY**
>
> Identify three different statements: resolvents converge strongly; heat semigroups converge strongly at fixed positive time; the collar semigroups tend to the identity while their time also tends to zero. Only the third statement reconstructs the thin Wilson insertion, and it needs the uniform-energy estimate in addition to ordinary fixed-time convergence.

11 Recovery of the heat-kernel area mechanism

The shrinking collar and the physical spanning region must be kept distinct. The collar area satisfies $a_\epsilon \to 0$ and supplies the identity limit. A spanning disk D has fixed physical area $A(D)$ and carries the nontrivial propagation.

In plain language: the collar disappears; the disk does not.

> **COMMON MISTAKE: DO NOT COLLAPSE THE PHYSICAL AREA**
>
> The parameter a_ϵ regularizes the operator and is driven to zero. The area $A(D)$ labels physical propagation and remains finite. They may carry the same dimensions and enter related heat kernels, but they are not the same geometric variable.

On the two-dimensional heat-kernel branch,

$$K_A(U) = \sum_Q d_Q \chi_Q(U) \exp\left[-\frac{1}{2} g_{\text{eff}}^2 C_2(Q) A\right]. \tag{7}$$

The convolution law

$$K_{A_1} * K_{A_2} = K_{A_1+A_2}$$

is the area-enriched bordism gluing law. Character orthogonality then gives

$$\langle W_R(C)\rangle = \exp\left[-\frac{1}{2} g_{\text{eff}}^2 C_2(R) A(D)\right]. \tag{8}$$

For the conventional KKN infrared normalization,

$$g_{\text{eff}}^2 = \frac{e^4 c_A}{2\pi}, \qquad \sigma_R = \frac{e^4 c_A C_2(R)}{4\pi},$$

so that for the fundamental representation of $SU(N)$,

$$\sigma_F = \frac{e^4(N^2-1)}{8\pi}.$$

These formulas belong to the selected KKN infrared branch [6, 4]; they are not asserted here as exact formulas for the unrestricted continuum theory. Exact heat-kernel Wilson-loop formulas in two-dimensional Yang–Mills are reviewed and derived in, for example, [1].

> **LEARNING CHECKPOINT: RECOVERING THE AREA EXPONENT**
>
> Insert the character expansion (7) into the disk amplitude and multiply by the normalized character of R. Haar orthogonality kills all terms except the matching representation. The surviving Casimir eigenvalue produces the exponent in (8). This is the exact two-dimensional step; the identification of the $2+1$ infrared branch with this step is the conditional part.

12 The $2+1$ incision and its relative class

Let $C \subset \Sigma^2$ be a spatial loop and

$$W = C \times I_t \subset M^3$$

its history. This is codimension one. A rounded positive-thickness neighborhood is

$$B_\epsilon = D_\epsilon(\nu_W) \cong W \times [-\epsilon, \epsilon] \cong S^1 \times D^2.$$

Theorem 12.1 (Incision persistence)**.** *For every $\epsilon > 0$,*

$$H^3(B_\epsilon; \mathbb{Z}) = 0, \qquad H^3(B_\epsilon, \partial B_\epsilon; \mathbb{Z}) \cong \mathbb{Z}.$$

Admissible orientation-preserving radial refinements preserve the relative generator, and fiber integration identifies it with $[W, \partial W] \in H^2(W, \partial W; \mathbb{Z})$.

Proof. B_ϵ deformation-retracts to S^1, giving the absolute vanishing. Poincare–Lefschetz duality gives $H^3(B_\epsilon, \partial B_\epsilon; \mathbb{Z}) \cong H_0(B_\epsilon; \mathbb{Z}) \cong \mathbb{Z}$. The normalized radial map

$$(x, s) \longmapsto (x, (\epsilon'/\epsilon)s)$$

is orientation preserving and maps the relative fundamental class to the relative fundamental class. Finally, the Thom isomorphism for the oriented normal line sends $[B_\epsilon, \partial B_\epsilon]$ to $[W, \partial W]$. □

TECHNIQUE: WHY RELATIVE COHOMOLOGY CHANGES THE ANSWER

Absolute cohomology asks whether the neighborhood itself carries a closed three-dimensional class. Relative cohomology permits the boundary to be treated as already trivialized and asks whether the region carries an orientation class relative to that boundary. A compact oriented three-manifold with boundary has a relative top class even when it retracts to a one-dimensional core. Thus $H^3(B_\epsilon) = 0$ and $H^3(B_\epsilon, \partial B_\epsilon) \cong \mathbb{Z}$ are not competing answers; they answer different questions.

Writing the principal boundary faces as

$$F_\epsilon^\pm = W \times \{\pm\epsilon\},$$

their oriented difference occurs in the boundary of a representative relative chain. Gauge projection and the Gauss-law Ward identity cancel the two operator-valued face contributions after admissible gluing. This does not annihilate the relative class. The correct conclusion is therefore a Ward-invisible boundary difference with a nonzero relative regional class. The boundary can become invisible to the projected observable without the region becoming topologically trivial.

WHAT BECAME VISIBLE

A traditional observable-first presentation sees the canceled boundary and may report nothing further. The regional lineage remembers that the cancellation occurred on the boundary of a nonzero relative cycle. The datum survives; the chosen observable is not sensitive enough to display it.

13 Electric–magnetic comparison on the solid torus

For a closed spacetime curve $\gamma \subset M^3$, the neighborhood $N_\epsilon(\gamma) \cong S^1 \times D^2$ has boundary torus with longitude ℓ and meridian m. Let $W_R(\ell)$ be the electric Wilson insertion and $T_z(m)$ the magnetic center-flux insertion. Their order–disorder relation is

$$T_z(m)W_R(\ell) = \omega_R(z)^{\mathrm{Lk}(\ell,m)} W_R(\ell)T_z(m), \tag{9}$$

where $R(z) = \omega_R(z)I_R$ [14].

Define the binary comparison

$$\Xi_{R,z,\epsilon} = T_{z,\epsilon}W_{R,\epsilon}T_{z,\epsilon}^{-1}W_{R,\epsilon}^{-1}.$$

For unit linking,

$$\Xi_{R,z,\epsilon} = \omega_R(z)I.$$

If radial refinement intertwines both insertions, then

$$S_{\epsilon'\epsilon}\Xi_{R,z,\epsilon} = \Xi_{R,z,\epsilon'}S_{\epsilon'\epsilon},$$

so the scalar phase is constant through the thick–thin system. For $SU(N)$, with $z_k = e^{2\pi ik/N}I$ and N-ality q_R,

$$\Xi_{R,z_k} = e^{2\pi ikq_R/N}I.$$

Thus $SU(2)$ in the fundamental representation gives -1, and $SU(3)$ gives $e^{2\pi i/3}$.

LEARNING CHECKPOINT: THE CENTER PHASE IN TWO EXAMPLES

For the $SU(2)$ fundamental representation, the nontrivial center element is $-I$ and acts as the scalar -1. For the $SU(3)$ fundamental representation, $z = e^{2\pi i/3}I$ acts as $e^{2\pi i/3}$. In each case the phase is fixed by the representation's center charge, not by the collar thickness. Radial refinement therefore transports the same scalar rather than creating a new phase.

14 Why the selected gerbe is silent

Bundle gerbes over a manifold are classified, up to stable equivalence, by their Dixmier–Douady class in $H^3(-;\mathbb{Z})$ [11, 12]. Since $H^3(B_\epsilon;\mathbb{Z}) = 0$, the local incision neighborhood supports no nontrivial absolute Dixmier–Douady class. This statement does not rule out an ambient gerbe, nor does it identify spacetime cohomology with algebraic group cohomology.

COMMON MISTAKE: TWO DIFFERENT THIRD COHOMOLOGIES

$H^3(B_\epsilon;\mathbb{Z})$ classifies geometric gerbe data on the incision neighborhood. $H^3(G;U(1))$ can classify algebraic associator data for a group or fusion structure. The two groups may communicate in a particular model, but they are not interchangeable symbols. Both the base geometry and the comparison maps must be specified before one claims a gerbe witness.

Algebraic associators are instead represented by classes in $H^3(G;U(1))$, as in Dijkgraaf–Witten theory [2]. Let S denote radial refinement, W electric insertion, and T magnetic insertion. The T–W face has comparison $\omega_R(z)$, but radial refinement transports the same phase to the refined face. Hence the total comparison-cube holonomy is

$$\Omega(S,T,W) = \omega_R(z)\omega_R(z)^{-1} = 1.$$

Theorem 14.1 (Gerbe silence with surviving binary charge)**.** *On the canonical $2+1$ regional–refinement branch,*

$$\Xi_2 = \omega_R(z), \qquad \Xi_3 = 1.$$

The incision preserves a finite electric–magnetic pairing but does not carry a nontrivial selected associator.

The silence is calculated, not assumed. The construction retains enough of the comparison geometry to test the associator and obtains the identity.

WHAT BECAME VISIBLE

The answer is not "there is no topology." The answer is more selective: relative orientation and binary center charge remain, while the chosen three-direction curvature vanishes. GAT earns the negative associator result because it preserved enough pre-strictification data to ask the question.

A framing change does not alter this conclusion. If $U = W_R(\ell)$, $V = T_z(m)$, and $VU = \zeta UV$, a Dehn twist with $D_pUD_p^{-1} = UV^p$ and $D_pVD_p^{-1} = V$ still gives

$$V(UV^p) = \zeta(UV^p)V.$$

The framing changes the representative but not the center phase or its flatness under refinement.

Chern–Simons theory could supply additional framing and ribbon-category data, but adding a Chern–Simons level changes the theory. It is not required for the pure-Yang–Mills thin limit. The appropriate limiting object is the area-enriched cylinder e^{-aQ_C}, which tends to the identity as $a \downarrow 0$.

15 The single-carrier visibility theorem

Suppose all projected observables are placed in one algebra $\mathcal{B}(\mathcal{H}_{\mathrm{phys}})$. Ordinary multiplication is strictly associative. Even if

$$T_a T_b = \eta(a,b) T_{ab},$$

associativity forces

$$\eta(a,b)\eta(ab,c) = \eta(b,c)\eta(a,bc),$$

or $\delta\eta = 1$.

Proposition 15.1 (Single-carrier invisibility)**.** *A single fixed operator carrier may retain a nontrivial projective H^2 class, but ordinary operator multiplication on that carrier cannot itself realize a nontrivial inherited H^3 associator. Such an associator requires distributed sector carriers and specified binary comparison isomorphisms.*

Indeed, for transports

$$T_a : \mathcal{H}_x \longrightarrow \mathcal{H}_{x+a}$$

and comparisons $\mu_{b,a} : T_b T_a \Rightarrow T_{ba}$, the two bracketings of three transports may differ by

$$\alpha_{c,b,a} = \Omega(c,b,a) I.$$

The admissible construction retains enough pre-strictification data to ask for Ω. In the canonical $2+1$ branch it then calculates $\Omega = 1$. This is more informative than identifying every carrier first and thereby making the question unavailable.

> **TECHNIQUE: WHY STRICTIFICATION CAN HIDE THE TEST**
>
> Inside one operator algebra, multiplication already has a fixed associativity law. A nontrivial inherited associator can only be detected before the sector carriers and their comparison maps are collapsed into that single algebra. Distributed carriers are therefore not decorative category language; they are the minimum bookkeeping needed to keep the associator question available.

16 Dimension comparison

The dimension change is encoded in the normal link of the Wilson history.

Theory	$\dim W$	$\dim M$	codimension	normal link
$2+1$	2	3	1	$S^0 = \{+,-\}$
$3+1$	2	4	2	S^1

In $2+1$, the time history has two normal faces. Their oriented difference is canceled at the projected boundary by the Ward identity, while the relative orientation class survives. There is no angular normal cover.

In $3+1$, the normal fiber is D^2 and the punctured normal link is S^1. Angular patches, overlap comparisons, and an additional independent Cech direction can support nontrivial comparison curvature. Codimension two does not by itself prove the existence of a gerbe; it supplies enough comparison geometry for one to be inhabited. The long-proof associator and its binary comparison maps are still required to decide that question.

A closed spacetime curve $\gamma \subset M^3$ must not be confused with the two-dimensional history $C \times I_t \subset M^3$. The former is codimension two and has knot-like meridional data; the latter is codimension one and is the two-face incision used in the time-zero reconstruction.

> **LEARNING CHECKPOINT: DIMENSION BEFORE ANALOGY**
>
> Compute the codimension and normal link before importing intuition from another dimension. In $2+1$, the Wilson history has an S^0 normal link: two faces. In $3+1$, it has an S^1 normal link with angular overlap data. The change in available comparison directions is the geometric reason the associator test can have a different answer.

17 Technique glossary for students

CORE VOCABULARY

Term	Working meaning in this notebook
Carrier	The Hilbert space, module, or class-function space on which a stated operator actually acts.
Admissible arrow	A selected comparison map satisfying the compatibility needed for the theorem being transported.
Lineage	A family of objects or vectors related by the chosen admissible arrows; convergence is often required only along such families.
Naturality defect	The difference between transporting then acting and acting then transporting.
Projection	Group averaging that selects invariant quantum vectors without choosing a global classical gauge slice.
Ward identity	The infinitesimal constraint relation obtained from gauge equivariance on a common invariant core.
Coercivity	A lower energy bound that prevents specified unwanted directions from remaining at finite cost.
Nonleakage	The conclusion that bounded-energy lineages concentrate in the selected sector as thickness tends to zero.
Mosco convergence	The liminf and recovery-sequence criterion that identifies limits of closed quadratic forms.
Relative class	A class that survives when the boundary is treated as trivialized, even if the corresponding absolute class vanishes.
Binary comparison	A phase or isomorphism comparing two ordered operations.
Associator	A ternary comparison between two bracketings of three transported operations.
Observable-silent	Present in the regional or ambient structure but invisible to the final selected observable.

> **OPEN EDGE**
>
> The natural continuation is not to force a gerbe into $2+1$. It is to test whether a noncanonical enrichment can couple to the surviving relative class without altering pure Yang–Mills dynamics. If it cannot, the notebook closes with a direct bordism lift and a genuinely center-visible but associator-flat incision.

18 Result ledger and conclusion

The logical status of the principal claims is summarized below.

Result	Status	Boundary
Klauder–Shabanov projector	Exact at compact lattice level	No continuum measure claim
Thick Wilson boundedness	Exact for each $\epsilon > 0$	Normalized character and contraction construction
Ward transport	Exact under equivariance and Haar compatibility	Common invariant core for differentiation
Information dominance	Exact relative to the endpoint record	Does not claim absolute or maximal primitiveness
Casimir coercivity	Exact on reduced Wilson sector	Not the full mass gap
Transverse nonleakage	Conditional	Requires (5)
Thick–thin reconstruction	Conditional	Mosco convergence, compatibility, and uniform energy
Heat-kernel area law	Exact on selected two-dimensional branch	KKN identification remains a branch statement
Relative incision class	Exact topologically	Survives by Thom descent, not as an absolute 3-class on the collapsed worldsheet
Electric–magnetic phase	Exact sector relation	A binary H^2-type witness
Canonical associator	Trivial on selected branch	Does not exclude enriched or ambient gerbes
Limit ordering	Explicit and typed	Full continuum and infinite-volume limits are not taken

The admissible-fibration machinery therefore passes a demanding dimensional test. It reconstructs a thin physical Wilson operator from bounded positive-thickness operators, transports gauge projection and Ward identities, isolates the reduced coercive mechanism, and recovers the heat-kernel area law on the appropriate infrared branch. At the same time, it does not force the $3+1$ obstruction into a geometry that cannot carry it.

A good dimensional test must be allowed to say no. Here the gerbe selector says no, but not because nothing is present. The relative incision remains. The electric–magnetic phase remains. What fails to appear is specifically the curvature of their selected three-direction comparison.

The $2+1$ incision is neither empty nor secretly gerbal. It carries a relative orientation lineage and a persistent electric–magnetic center phase. Its canonical regional–refinement comparison is nevertheless flat. The final picture is therefore

$$\boxed{2+1:\ \text{relative bordism} +\ \text{binary center charge} +\ \text{trivial selected associator}.}$$

That distinction is precisely the kind of information lost by premature strictification and retained by admissible transport.

The resulting claim is therefore deliberately strong: on this branch, admissible transport provides a strictly finer mathematical invariant than the final Wilson observable alone. It distinguishes constructions that become identical after projection and collapse but differ in provenance, relative topology, or comparison curvature. What it does not establish is that no still-finer enrichment can distinguish more.

References

[1] A. Ashtekar, J. Lewandowski, D. Marolf, J. Mourão, and T. Thiemann, "Closed formula for Wilson loops for $SU(N)$ quantum Yang–Mills theory in two dimensions," arXiv:hep-th/9605128 (1996).

[2] R. Dijkgraaf and E. Witten, "Topological gauge theories and group cohomology," *Commun. Math. Phys.* **129**, 393–429 (1990).

[3] L. D. Faddeev and V. N. Popov, "Feynman diagrams for the Yang–Mills field," *Phys. Lett. B* **25**, 29–30 (1967).

[4] D. Karabali, "The Hamiltonian approach to Yang–Mills $(2+1)$," arXiv:0906.0783 (2009).

[5] D. Karabali, C. Kim, and V. P. Nair, "Planar Yang–Mills theory: Hamiltonian, regulators and mass gap," arXiv:hep-th/9705087 (1997).

[6] D. Karabali, C. Kim, and V. P. Nair, "On the vacuum wave function and string tension of Yang–Mills theories in $(2+1)$ dimensions," *Phys. Lett. B* **434**, 103–109 (1998), arXiv:hep-th/9804132.

[7] J. Kogut and L. Susskind, "Hamiltonian formulation of Wilson's lattice gauge theories," *Phys. Rev. D* **11**, 395–408 (1975).

[8] K. Kuwae and T. Shioya, "Convergence of spectral structures: a functional analytic theory and its applications to spectral geometry," *Commun. Anal. Geom.* **11**, 599–673 (2003).

[9] U. Mosco, "Convergence of convex sets and of solutions of variational inequalities," *Adv. Math.* **3**, 510–585 (1969).

[10] U. Mosco, "Composite media and asymptotic Dirichlet forms," *J. Funct. Anal.* **123**, 368–421 (1994).

[11] M. K. Murray, "Bundle gerbes," *J. Lond. Math. Soc.* **54**, 403–416 (1996), arXiv:dg-ga/9407015.

[12] M. K. Murray and D. Stevenson, "Bundle gerbes: stable isomorphism and local theory," arXiv:math/9908135 (1999).

[13] V. P. Nair, "Gauge-invariant variables, WZW models and (2+1)-dimensional Yang–Mills theory," arXiv:hep-th/9911061 (1999).

[14] H. Reinhardt, "The 't Hooft loop in the Hamiltonian approach to Yang–Mills theory in Coulomb gauge," *Phys. Rev. D* **77**, 085007 (2008), arXiv:0706.0175.

[15] S. V. Shabanov and J. R. Klauder, "Towards a nonperturbative path integral in gauge theories," *Phys. Lett. B* **456**, 38–44 (1999), arXiv:hep-th/9902017.

[16] S. V. Shabanov, "Geometry of the physical phase space in quantum gauge systems," *Phys. Rep.* **326**, 1–163 (2000), arXiv:hep-th/0002043.

[17] E. Witten, "Quantum field theory and the Jones polynomial," *Commun. Math. Phys.* **121**, 351–399 (1989).

[18] E. Witten, "On quantum gauge theories in two dimensions," *Commun. Math. Phys.* **141**, 153–209 (1991).

Public release and research-assistance acknowledgment

Recommended citation. J. Scott Little, "Thick Wilson Loops in $2+1$ Yang–Mills: Projection, Admissible Thin Transport, and the Silence of the Gerbe," GAT Research Notebook RN-YM-01, version 1.1, September 2026.

Research-assistance acknowledgment. The physical and mathematical direction of this research program, its conjectures, conceptual prompts, and final mathematical judgments are the author's. OpenAI language-model tools materially assisted during development with framework simplification, structural organization, theorem staging, bookkeeping and ledger discipline, cross-framework recognition, source concordance, and editorial refinement. Responsibility for every statement and proof remains with the author.