

Wilson Transport and the Vacuum Module

J. Scott Little

Website notes October 22, 2026

A Wilson loop is easy to write down. Follow a closed path, multiply the parallel transports in their prescribed order, and take a trace. On a finite graph with compact gauge group, this gives a bounded observable. There is already something satisfying about that construction: the intermediate choices of frame disappear, and what remains belongs to the closed path.

The question that interests me begins when we open the loop again. What must we keep in order to cut the geometry, move between local descriptions, or pass to a finer scale and still recognize the observable we started with? The trace has forgotten the intermediate frames. The construction that compares two descriptions cannot afford to forget them quite so early.

This is where the Hilbert module becomes physically important. Its left and right actions remember the two sides of a cut. Its composition tells us how to join those sides again. The electric–magnetic graph frustum determines this structure through its boundaries, seams, and ordered holonomy words. If we introduce the module only after the calculation, as an abstract space satisfying a list of conditions, it becomes much harder to see why it has the shape it does.

There is a second distinction I want to keep visible. Comparing an observable across scale and comparing the vacua at those scales are different operations. In the finite example below, an inherited Wilson operator is transported exactly. The vacuum is transported with a small, quantified error. Both statements hold at once, and together they tell us more than either statement alone.

What the cut asks us to retain

Think of the graph frustum as the finite region connecting two graph presentations. The magnetic part uses ordered products around faces and paths. The electric part acts on the edge variables and on the boundary data exposed by cutting those paths. These actions meet at the same banks. Their compatibility is part of the geometry we have to preserve.

Suppose an open piece has an incoming representation space V_L and an outgoing representation space V_R . A transport across that piece takes an incoming vector to an outgoing vector, so its natural home is

$$E_{RL} = \text{Hom}(V_L, V_R).$$

An operator on V_R acts from the left; an operator on V_L acts from the right. This is the elementary bimodule structure. It says which boundary each operation belongs to. The right Hilbert-module inner product is

$$\langle T, S \rangle_{B(V_L)} = T^*S.$$

Its value is still an operator on the incoming boundary space. Taking a scalar inner product would give a number; keeping this operator-valued inner product retains information that another piece of the geometry may need when it is attached.

Now join $S : V_0 \rightarrow V_1$ to $T : V_1 \rightarrow V_2$. The shared space V_1 is the seam, and the composite is TS . For finite nonzero bank spaces, this is expressed by

$$E_{21} \hat{\otimes}_{B(V_1)} E_{10} \cong E_{20}, \quad T \otimes S \longmapsto TS.$$

The word “balanced” has a concrete meaning here. For a seam operator A , the descriptions $(TA) \otimes S$ and $T \otimes (AS)$ represent the same joined operation. In matrix components, the shared index is summed: $(TS)_{ac} = \sum_b T_{ab} S_{bc}$. The module tells us which index is shared before we perform the sum.

The same structure appears in the infinitesimal electric action:

$$\delta T = \rho_{R*}(X_R)T - T\rho_{L*}(X_L).$$

Here ρ_* is the infinitesimal representation of the gauge action. At a common bank, the contribution from one piece cancels the contribution from its neighbor:

$$(-T\rho_*(X))S + T(\rho_*(X)S) = 0.$$

This little calculation is worth pausing over. The signs, the order of multiplication, and the spaces on which the operators act all matter. We can see the cancellation because both sides of the cut are still present.

Once the whole path closes, we may take a trace. Taking separate traces earlier would generally change the result: even for square matrices, $\text{Tr}(TS)$ is not generally $\text{Tr}(T) \text{Tr}(S)$. For an open piece between different boundary spaces, a trace may not be available at all. The shape of the module is what makes the eventual closure meaningful.

Crossing a physical spacetime cover

A cover introduces a related problem. We describe the same physical region in overlapping patches and ask whether those descriptions agree where they meet. The Čech nerve organizes the comparisons: double overlaps carry transition data, triple overlaps compare compositions, and higher overlaps record further coherence.

That organization is useful, but the physical cancellation still has to be proved. To cross the physical spacetime cover, we need the actual Gauss–Ward identities, applied on domains where the relevant operators and boundary terms are defined. The bank calculation above shows the algebraic pattern of one cancellation. The physical identity must also include the terms and seam identifications supplied by the construction.

I want to be careful here because an appealing diagram can make this step look easier than it is. Two scalar formulas may agree while the lines in which they take values, or the maps identifying those lines across an overlap, have not yet been compared. A selector’s line and intertwiner carry information that disappears when we retain only a scalar answer.

This also explains the role of a scale bordism filling the transgression layer. The filling participates in the comparison of the boundary coefficient data. Square and cube calculations check that the

resulting local comparisons compose coherently. To use that coherence physically, its boundary maps must be identified with the actual module and seam maps. This is the point at which the geometry of the filling has to meet the Ward identities of the physical cover.

Following one Wilson loop through a refinement

We can make part of this discussion quite concrete. Consider the native cap-8 to cap-16 root-star refinement in the selected scalar $PSU(3)$ branch. These names specify one finite construction. Write s for its coarse Hilbert space and t for its fine Hilbert space. I will call these the two carriers, meaning the full selected spaces on which the observables act.

Each oriented coarse edge $e = (v, w)$ is represented by a prescribed two-child word z_e in the fine graph. At a promoted endpoint, a native collar holonomy k_v supplies the endpoint comparison; at an unchanged endpoint, $k_v = I$. The map from fine configurations to coarse configurations is

$$p(g_t)_e = k_v^{-1} z_e k_w.$$

The collars record how the endpoint frames are related. They are part of the actual geometry of this refinement. The corresponding map on wavefunctions is pullback:

$$(J\psi)(g_t) = \psi(p(g_t)).$$

The finite source proves that p pushes the fine Haar measure forward to the coarse Haar measure. Thus pulling back preserves the Hilbert norm, and $J^*J = I$. It also proves compatibility with the scalar Gauss/copy selection. Every sector of the selected coarse carrier is retained with its multiplicity. Spectator factors, which this comparison leaves inert, are carried by the identity.

Now follow a closed coarse path $\gamma = e_1 \cdots e_N$. At each intermediate vertex, the outgoing collar from one edge meets the inverse collar on the next. The product therefore telescopes:

$$\prod_{r=1}^N p(g_t)_{e_r} = k_{v_0}^{-1} \left(\prod_{r=1}^N z_{e_r} \right) k_{v_0}.$$

Only a conjugation at the starting vertex remains. This is exactly the information the trace can forget without changing the closed observable.

Use the adjoint character, normalized by its dimension eight:

$$w_s = \frac{1}{8} \chi_{\text{ad}} \left(\prod_r g_{e_r} \right), \quad w_t = \frac{1}{8} \chi_{\text{ad}} \left(\prod_r z_{e_r} \right).$$

These functions are real, bounded in absolute value by one, and defined on $PSU(3)$. Conjugation invariance gives $w_t = w_s \circ p$. If $W_j = M_{w_j}$ denotes multiplication by w_j , then

$$\boxed{K_a = W_t J - J W_s = 0.}$$

There are two routes in this equation. We may apply the coarse Wilson observable and then transport the state, or transport the state and apply the inherited fine Wilson observable. The two routes agree on the entire selected coarse carrier. The fine word was specified by the child paths before this calculation, so the equality tests the geometric comparison.

The same calculation works for the adjoint observable. It is useful to distinguish that comparison from taking the literal adjoint of K_a :

$$K_a^\sharp = W_t^* J - J W_s^* = 0, \quad K_a^* = J^* W_t^* - W_s^* J^* = 0.$$

The first map goes from coarse to fine; the second goes from fine to coarse. In this example $W_j = W_j^*$, but the distinction between the two map types remains.

Let $P_J = JJ^*$ project onto the inherited coarse range inside the fine carrier. The two intertwining identities imply

$$[P_J, W_t] = 0, \quad J^* W_t J = W_s, \quad J W_s J^* = W_t P_J.$$

Thus the inherited range reduces this Wilson operator. The last equation also tells us precisely where the comparison lives: $J W_s J^*$ is supported on that range. The full fine carrier can contain additional states and additional observables.

What becomes an algebra and what the trace forgets

The loops supply generators. On the selected Hilbert-space representation, form the unital $*$ -algebra generated by their multiplication operators and take its weak operator closure. The result is a von Neumann algebra, also called a W^* -algebra. For these equal-time multiplication observables it is commutative. Electric operators introduce further structure and are not included in this statement.

The Haar-preserving pullback gives a faithful normal unital $*$ -homomorphism of the gauge-invariant multiplication algebras. Faithfulness means that a nonzero coarse multiplication observable is not lost under this map. Normality ensures compatibility with the increasing bounded limits used in the von Neumann algebra. Restricting the map to the Wilson-generated algebra identifies it with its inherited fine image. New fine loops need not belong to that image.

There is another kind of information that ordinary observable transport can forget. Suppose the unitary overlap maps satisfy

$$u_{ij} u_{jk} = \zeta_{ijk} u_{ik}, \quad \zeta_{ijk} \in U(1).$$

The phase ζ_{ijk} records a possible failure of strict composition at the level of the coefficient data. In the adjoint action on an operator, however, that phase cancels against its complex conjugate. Consequently,

$$\text{Ad}(u_{ij}) \text{Ad}(u_{jk}) = \text{Ad}(u_{ik}).$$

An ordinary observable can therefore be well defined while the coefficient band remains nontrivial. Here “coefficient band” refers to the phase data retained by the line-valued comparisons. Its invisibility under conjugation does not erase it from those comparisons.

This gives an operational way to discuss choices. Rephasing an intertwiner leaves its adjoint action unchanged. Multiplying a Wilson operator itself by a phase generally changes that operator. Changing a cover requires comparison through the overlap maps on a common refinement. Changing a prism or filling requires comparing its boundary data and the residual closed class. If that residual acts trivially in the ordinary realization, it gives the same ordinary observable; the coefficient system may still retain flat information or closed-surface holonomy.

Local square and cube coherence helps us make these comparisons consistently. It does not say that every closed residual is trivial. The remaining physical task is to identify the coefficient system and

its maps with those of the actual charged selector construction. The scalar Wilson identity above provides one evaluated part of that task.

Why the vacuum needs its own comparison

A geometric embedding of carriers tells us how to carry a wavefunction to the finer graph. The fine Hamiltonian determines which wavefunction is its vacuum. There is no reason to identify those two operations before examining the Hamiltonians and their ground states.

For this finite refinement, the composed geometric Paper I supplies normalized positive scalar ground vectors obeying

$$\|J\phi_s - \phi_t\| < \delta_*, \quad \delta_* = 0.003089852, \quad \langle \phi_t, J\phi_s \rangle > 0.999995226.$$

This is the numerical input used below. It says that the transported coarse vacuum and the independently determined fine vacuum are close in the fine Hilbert norm. Positivity fixes their relative phase for the comparison.

The complete ground projections retain the unchanged spectator multiplicity:

$$P_j = |\phi_j\rangle\langle\phi_j| \otimes I_{\text{spectator}}.$$

Compress J between these ground spaces and take its polar map. The positive overlap gives

$$V = |\phi_t\rangle\langle\phi_s| \otimes I_{\text{spectator}}.$$

On the ground module, V takes the coarse vacuum to the fine vacuum while preserving each spectator factor. Extend it by zero on the source ground complement. Its purpose and domain differ from those of J , which transports the full coarse carrier.

We can now see why substituting V for J in the Wilson equation changes the question. Suppress the spectator identity and write

$$W_j\phi_j = \mu_j\phi_j + \chi_j, \quad \chi_j \perp \phi_j, \quad \mu_j = \langle W_j\phi_j, \phi_j \rangle.$$

The first term is the vacuum expectation multiplied by the vacuum. The second is the response orthogonal to the ground. Its squared norm is the variance of the Wilson observable in that vacuum.

Scalar inner products here are linear in the first argument, and $|u\rangle\langle v|$ means $x \mapsto \langle x, v \rangle u$. Expanding the two compositions gives

$$\begin{aligned} W_t V - V W_s &= |\chi_t\rangle\langle\phi_s| - |\phi_t\rangle\langle\chi_s| \\ &\quad + (\mu_t - \mu_s)|\phi_t\rangle\langle\phi_s|. \end{aligned}$$

The terms have a physical interpretation. The fine Wilson operator can take the transported ground into a fine excitation. A coarse Wilson insertion can have a component returning an excited source vector to the coarse ground. Finally, the two vacuum expectations may differ. A map defined only on the ground module has no reason to intertwine all these processes.

Acting on the coarse vacuum makes the issue particularly clear:

$$\|(W_t V - V W_s)\phi_s\|^2 = \text{Var}_t(W_t) + |\mu_t - \mu_s|^2.$$

Even if the expectations agree, the fine variance can remain. That term measures the response to the observable. It is not, by itself, a failure of vacuum transport.

If we ask only for the ground-compressed comparison, the excited terms disappear:

$$P_t W_t P_t V - V P_s W_s P_s = (\mu_t - \mu_s) V, \quad |\mu_t - \mu_s| < 2\delta_*.$$

The estimate follows by comparing the expectation of the same contraction W_t in $J\phi_s$ and ϕ_t . Each occurrence of the state contributes at most δ_* . We have used exact Wilson transport and approximate vacuum transport together, each for the part it controls.

Turning the overlap into a response estimate

The overlap certificate reaches beyond a single expectation value. It controls how the ground projections themselves move. Put

$$P_a = J P_s J^*, \quad E = P_t J - J P_s = (P_t - P_a) J.$$

The operator E measures the disagreement between projecting after transport and transporting after projection. In this rank-one scalar ground, with the same full spectator identity on both sides,

$$\varepsilon := \|E\| = \sqrt{1 - |\langle \phi_t, J\phi_s \rangle|^2} < \delta_*.$$

Geometrically, this is the sine of the angle between the two scalar ground lines. The distance between their orthogonal projections gives the upper bound, and applying E to the source ground vector attains it. Tensoring with the spectator identity preserves the norm, so no multiplicity is discarded.

Now consider the operator

$$B_j = (1 - P_j) W_j P_j.$$

Read it from right to left: retain the ground component, insert the Wilson observable, and retain the part outside the ground. This is precisely the part of the observable responsible for the centered response. Using $W_t J = J W_s$, we obtain

$$B_t J - J B_s = (1 - P_t) W_t E - E W_s P_s.$$

Both terms are controlled by the projection defect because the remaining factors are contractions. Therefore

$$\|B_t J - J B_s\| \leq 2\varepsilon < 0.006179704.$$

The excited response need not be small. What is small is the disagreement between its two transported descriptions. This is the comparison we need if Wilson insertion is to remain a meaningful probe as the vacuum changes.

We can also compare the centered response vectors directly. Let $a = J\phi_s$ and $b = \phi_t$, suppressing spectators, and write P_a, P_b for their rank-one projections. Then

$$\chi_t - J\chi_s = (1 - P_b) W_t (b - a) + (P_a - P_b) W_t a.$$

The first term comes from changing the vacuum vector. The second comes from changing the projection used to remove its ground component. Since $\|P_a - P_b\| \leq \|a - b\|$ and $\|W_t\| \leq 1$, both terms have the same available bound. It follows that

$$\begin{aligned}\|\chi_t - J\chi_s\| &< 2\delta_*, \\ |\sqrt{\text{Var}_t(W_t)} - \sqrt{\text{Var}_s(W_s)}| &< 2\delta_*, \\ |\text{Var}_t(W_t) - \text{Var}_s(W_s)| &< 4\delta_* = 0.012359408.\end{aligned}$$

The last estimate uses that each response norm is at most one. These inequalities control the size and fluctuation of the response in Hilbert norm. To attach an energy to it, we would also need the Hamiltonian comparison and its domain.

Several insertions and the next physical step

A single loop is only the beginning of what we may want to measure. We can insert several inherited Wilson observables and retain or remove the ground component between insertions. Such an operation is a finite ordered word in Wilson contractions, their adjoints, and the projections P_j or $1 - P_j$.

Let A_s, A_t be corresponding words containing r projection factors. Move J through the word one factor at a time. Every inherited Wilson factor passes through exactly. Each ground projection contributes E , and each complementary projection contributes $-E$. All the other factors have norm at most one, so

$$\|A_t J - J A_s\| \leq r\epsilon.$$

This explains why the number of projection factors appears in the bound. The error counts the places where we ask the two ground decompositions to agree. There is no additional contribution from the exactly transported Wilson factors.

For corresponding ground vectors with the same unit spectator vector, comparing the endpoint states adds two more terms:

$$|\langle A_t \phi_t, \phi_t \rangle - \langle A_s \phi_s, \phi_s \rangle| < (r + 2)\delta_*.$$

We therefore have a finite comparison of whole sequences of insertions, uniform over the retained spectator multiplicity. The original positive-overlap estimate has become a usable statement about observables and their responses.

For me, the value of this calculation is that it lets us see what each piece of the construction is doing. The open module retains the boundary actions. The seam determines their composition. The Wilson trace closes the word. The carrier map compares the inherited observable, and the polar ground map compares the vacua. Keeping those roles visible makes the estimates easier to interpret and the remaining physical question more precise.

The next realization step is to match the selector's actual line, intertwiner, and lateral-seam data, using the Ward identities across the physical cover. An unbounded refinement system also needs estimates along that system and a specified normal receiving realization. The finite result gives us a concrete place to begin those comparisons: we know which Wilson operator survives exactly, how its vacuum changes, and how that change enters a finite sequence of measurements.

Companion material

J. Scott Little, *Scale Bordisms and Wilson Observable Algebras: Relative Filling, Čech Descent, and Normal Realization in $3 + 1$ Dimensions*, consolidated draft v0.27.0, especially Section 6.1. [Download the companion paper](#).

Finite construction and numerical certificate: J. Scott Little, *Geometric Physical Realization of Sources: Yang–Mills at Finite Scale*, composed manuscript, supplied revision 108, Section 14.61.1; provenance to the geometric source’s Development Notes V–AJ. The numerical input is the certificate quoted there. The residual identities and finite-word estimates above are deductions from that input and the exact inherited Wilson comparison.