The Geometry of Admissible Transport V

Mass Identification and Invariant Spectral Transfer

Orbit thresholds, same-carrier realization, and the mass-square spectral shadow

J. Scott Little

Independent Researcher

Hostile-proof and prose-integrity audited authorial manuscript

Primary-construction firewall. The five-paper Wilson–Tube Yang–Mills sequence is frozen and non-superseded. This paper extracts a general invariant-transfer theorem, recognizes its realization in the frozen construction, and combines it with downstream vacuum lineage. It does not replace, reorganize, or reconstruct Papers I–V.

Scope and status. The general transfer theorem, source concordance, typed carrier dictionary, proof-boundary ledger, and sharp countermodels are complete. No unresolved proof obligation remains inside the declared domain. The Wilson–Tube construction remains frozen and non-superseded.

Authorial manuscript v1.0.2 — August 14, 2026

Abstract

An energy gap is not, by notation alone, an invariant spectral gap. The two estimates live on different observables, and a reconstruction may also change carriers before either observable is available. This paper isolates a general transfer mechanism. A covariant projection-valued measure on a symmetry space carries a height observable and an invariant orbit-floor observable. If the height operator is identified, on one Hilbert carrier, with the positive square root of a reconstructed coercive operator, then covariance transports the height exclusion around every symmetry orbit. The orbit infimum becomes an invariant lower threshold. For the proper orthochronous Lorentz action on the forward cone, the height is energy and the orbit floor is relativistic mass; nonzero null orbits are excluded and the energy-square certificate becomes a mass-square certificate without numerical renaming. A separate admissible-transport calculus tracks certificates, sharp spectral edges, compression, products, infinite assembly, Mosco limits, and source provenance. Explicit boundary models show that the same-carrier comparison, exact zero fiber, full orbit covariance, target-complement control, uniform infinite floor, domain declarations, and lineage data cannot be silently removed. The Wilson–Tube Yang–Mills sequence is treated as a frozen primary construction instantiating the general interface, while vacuum identity is supplied downstream by route-independent lineage.

1 Construction status and direction of use

The status hierarchy is fixed before any import:

Wilson–Tube Papers I–V	:	`FROZEN -- PRIMARY CONSTRUCTION`,
Admissible Transport Paper I	:	`EXISTING -- DO NOT REWRITE FROM SCRATCH`,
Admissible Transport downstream program	:	`ACTIVE; COMPLETED MANUSCRIPTS INDIVIDUALLY FROZEN`,
Measurement/interface theory	:	`EXPORTED DOWNSTREAM PROGRAM`.

Within that active downstream program, GAT IV is frozen at authorial manuscript v1.0.0 and the present GAT V manuscript is frozen only after the audits recorded here. This current paper-level status refines, rather than alters, the bootstrap hierarchy. The governing direction is

$$\boxed{\text{original proof construction} \longrightarrow \text{later recognition and explanation} \longrightarrow \text{new consequences.}}$$

GAT V therefore does not prove the frozen Yang–Mills theorem again. Its general contribution is to identify the exact mechanism by which a coercive spectral exclusion for a reconstructed generator becomes a symmetry-invariant exclusion on the same retained physical sector.

That distinction determines the prose as well as the theorem architecture. The frozen construction already reaches its physical conclusion by its own sequence of finite, cofinal, thin, and spectral passages. The present paper asks a later question: what mathematical fact made the last passage possible, and which hypotheses prevent an energy estimate from being misread as an invariant-mass estimate? The answer must be abstract enough to expose the mechanism, but narrow enough not to retroactively reorganize the construction that discovered it.

The generality is downstream rather than revisionary. The abstract theorem applies whenever a covariant spectral support has a meaningful orbit floor. Relativistic mass is the principal example, not part of the abstract definition. Yang–Mills is the terminal recognition corollary, not the ambient category of the theorem.

The argument therefore proceeds in the order

$$\boxed{\begin{array}{c}\text{physical obstruction} \longrightarrow \text{orbit-covariant spectral geometry} \longrightarrow \text{same-carrier closure}\\ \longrightarrow \text{general transfer} \longrightarrow \text{Poincaré specialization} \longrightarrow \text{Yang–Mills recognition}\end{array}}$$

The abstract orbit-floor theorem therefore precedes the Poincaré corollary, but only after the physical energy-to-mass obstruction has motivated it. The order records the logical dependence of the argument: the obstruction earns the abstraction, the abstraction isolates the transferable statement, and the Yang–Mills result enters only after the general interface has been closed.

2 The object changes twice

The chain

$$H_{\text{thin}} \longrightarrow P^0 \longrightarrow M^2 = P_\mu P^\mu$$

does not carry one spectrum through three names. A closed form first produces a nonnegative self-adjoint operator K^{thin} and its positive square root H_{thin}. A strongly continuous translation representation separately produces strongly commuting Stone generators P^μ and their joint PVM E_P. Invariant mass square is then the joint-functional-calculus operator

$$M^2 = \int p_\mu p^\mu \, dE_P(p).$$

The objects are related, but they are not interchangeable:

$$\boxed{K^{\mathrm{thin}} \neq H_{\mathrm{thin}} \neq P^0 \neq E_P \neq M^2.}$$

This is the precise point at which a correct construction can become unreadable to an external inference engine. Each arrow may be valid while the composite remains uncertified because the intermediate identifications have been suppressed. A coercive form controls the operator associated with that form. A translation representation supplies a joint spectral measure. Lorentz-invariant mass is a function of the joint spectral variable. Moving from one sentence to the next therefore changes both the observable and, potentially, the carrier on which it acts.

The first missing datum is same-carrier realization: H_{thin} and P^0 must be proved equal as self-adjoint operators on one physical Hilbert space. The second missing datum is orbit geometry: positivity of P^0 does not compare M^2 with $(P^0)^2$ at fixed momentum. The third missing datum is lineage: the carrier on which the analytic theorem holds must be the retained terminal carrier of the selected source packet.

These three data answer different objections. Same-carrier realization answers whether the reconstructed dynamics is the physical time translation. Orbit geometry answers whether a frame-dependent exclusion controls the invariant spectrum. Lineage answers whether the Hilbert space on which those statements hold is still the physical sector selected at the beginning of the reconstruction. None can be replaced by the other two.

3 Covariant spectral systems and orbit floors

The relativistic argument is a special case of a more general spectral fact.

The relevant invariant is not obtained by comparing two operators pointwise. It is obtained by moving a spectral point through every symmetry-related description and asking for the lowest height visible along that orbit. Covariance is load-bearing because it forces the entire orbit of an admitted spectral point to remain admitted. A lower height exclusion can then be tested at every point of that orbit. The missing datum is therefore an orbit floor.

Definition 3.1 (Orbit-covariant spectral system)**.** An *orbit-covariant spectral system* is a tuple

$$\mathfrak{O} = (\mathcal{H}, X, G, V, E, h, Z, \Pi)$$

with the following data:

(O1) X is a locally compact second-countable Hausdorff space, G is a second-countable topological group, and G acts continuously on X by homeomorphisms;

(O2) V is a strongly continuous unitary representation of G on $\mathcal{H}$;

(O3) E is a sharp Borel PVM on X satisfying

$$V(g)E(B)V(g)^* = E(gB);$$

(O4) $h : X \to [0, \infty)$ is a continuous *height function*;

(O5) $Z \subseteq X$ is closed and G-invariant;

(O6) Π records the carrier and physical provenance of the system.

Write $S = \operatorname{supp} E$ and

$$H_E := \int_X h(x) \, dE(x).$$

The height is frame- or observer-dependent. The previous data determine its values on every transformed spectral point, but they do not yet package the lowest height along an orbit.

Definition 3.2 (Orbit floor and invariant shadow)**.** For $x \in X$, set

$$\rho_G(x) := \inf_{g \in G} h(gx).$$

Define the *orbit-floor spectral shadow*

$$Q_G := \int_X \rho_G(x)^2 \, dE(x).$$

The square is a normalization choice adapted to energy-square coercivity.

Lemma 3.3 (Support exhaustion and orbit-floor measurability)**.** *For an orbit-covariant spectral system:*

(i) $E(X \setminus S) = 0$*; consequently, Borel sets* B, C *satisfying* $B \cap S = C \cap S$ *have* $E(B) = E(C)$*;*

(ii) ρ_G *is upper semicontinuous, hence Borel, and is constant on* G*-orbits;*

(iii) the spectral shadow is invariant:

$$V(g) Q_G V(g)^* = Q_G, \qquad g \in G,$$

with equality of the corresponding maximal operator domains.

Proof. Let $\mathcal{B}$ be a countable base for X. The union of the members of $\mathcal{B}$ that are E-null is the largest open E-null set: every open null set is a union of null basis elements. Countable additivity therefore gives $E(X \setminus S) = 0$ and the first conclusion.

Choose a countable dense subset $D \subseteq G$. For fixed x, continuity of $g \mapsto h(gx)$ gives

$$\rho_G(x) = \inf_{g \in D} h(gx).$$

This is a countable infimum of continuous functions of x, hence upper semicontinuous. Replacing x by $g_0 x$ only right-translates the set over which the infimum is taken, proving orbit invariance. Covariance of the PVM and orbit invariance of ρ_G then give the final assertion by unbounded Borel functional calculus, including equality of maximal domains. ☐

Lemma 3.4 (Support invariance)**.** *For an orbit-covariant spectral system,* $S = \operatorname{supp} E$ *is* G*-invariant.*

Proof. If an open set O is E-null, covariance makes gO null. Applying the same statement to g^{-1} shows that the complement of the largest open null set is invariant. ☐

Theorem 3.5 (Orbit-threshold transfer)**.** *Let* $\mathfrak{O}$ *be an orbit-covariant spectral system. Assume*

$$S \cap h^{-1}(0) = S \cap Z, \qquad P_Z := E(Z),$$

and, for some $\eta > 0$*,*

$$H_E^2 \geq \eta(1 - P_Z).$$

The inequality is understood in the closed-form sense on $\operatorname{Dom} H_E$*. Then*

$$S \subseteq Z \cup \{x : \rho_G(x)^2 \geq \eta\}, \qquad Q_G \geq \eta(1 - P_Z).$$

In particular, no nonzero spectral orbit whose height infimum is below $\sqrt{\eta}$ *can occur.*

Proof. Set $A = \{x \notin Z : h(x)^2 < \eta\}$. For every sufficiently large integer n, let

$$A_n = \{x \notin Z : h(x)^2 \leq \eta - 1/n\}.$$

If $\psi \in \operatorname{Ran} E(A_n)$, then $\psi \in \operatorname{Dom} H_E$, $P_Z \psi = 0$, and spectral calculus gives

$$\|H_E \psi\|^2 \leq (\eta - 1/n)\|\psi\|^2.$$

The assumed form inequality gives the reverse lower bound $\|H_E \psi\|^2 \geq \eta \|\psi\|^2$, so $E(A_n) = 0$. Since A is the increasing union of these A_n, PVM continuity gives $E(A) = 0$. The set A is open, hence it is disjoint from S.

Let $x \in S \setminus Z$. Since Z is invariant, $gx \notin Z$ for every g. Lemma 3.4 gives $gx \in S$, so $h(gx)^2 \geq \eta$ for every g. Taking the infimum yields $\rho_G(x)^2 \geq \eta$. Lemma 3.3 removes the complement of S, and functional calculus gives the closed-form inequality for Q_G. ☐

The theorem does not require the orbit infimum to be attained. That point is decisive for nonzero null Lorentz orbits, whose energy can approach zero under boosts without ever reaching the zero momentum.

The mechanism should be read in this order. The form inequality first removes a strict height sublevel from the support. Covariance then removes every orbit that would enter that sublevel after a symmetry transformation. Only after that support statement has been proved does functional calculus convert the orbit floor into an operator inequality. The invariant threshold is therefore a consequence of spectral support rigidity, not a formal manipulation of the symbols H_E and Q_G.

Boundary 3.6 (Effective symmetry is part of the theorem)*.* Replacing G by a smaller realized subgroup replaces ρ_G by a larger orbit floor and may weaken the invariant interpretation. A time-axis estimate becomes a relativistic mass estimate only when the represented covariance group is large enough to see the relevant Lorentz orbit infimum.

4 Same-carrier realization and exact zero

The orbit theorem consumes an inequality for H_E. Reconstruction supplies an inequality for K^{thin}. The missing criterion is equality of the two time evolutions on one recovery representation.

Unitary equivalence would preserve spectra, but it would not by itself say that the reconstructed ground projection, the translation vacuum, and the orbit-covariant PVM belong to one physical realization. The theorem needs a stronger statement: the two groups act on the same Hilbert space, implement the same recovered dynamics, and fix the same cyclic vector. Under those hypotheses the apparent comparison collapses to equality by density and Stone uniqueness.

Definition 4.1 (Thin dynamical packet)**.** A *thin dynamical packet* on $\mathcal{H}$ is

$$\mathfrak{D} = (q, K, H_{\text{thin}}, P_\Omega, \kappa, \Pi_{\text{thin}}),$$

where q is a densely defined closed nonnegative form, K is its associated self-adjoint operator, $H_{\text{thin}} = K^{1/2}$, $\operatorname{Ran} P_\Omega = \ker K$, and

$$K \geq \kappa(1 - P_\Omega), \qquad \kappa > 0.$$

Theorem 4.2 (Cyclic same-carrier identification)**.** *Let $\mathfrak{A}_{\text{rec}} \subseteq \mathcal{B}(\mathcal{H})$ be a represented $*$-algebra and let Ω be cyclic. Suppose $V_1(t)$ and $V_2(t)$ are strongly continuous unitary groups such that*

$$V_j(t)\Omega = \Omega, \qquad V_j(t) A V_j(t)^* = \alpha_t(A)$$

for every $A \in \mathfrak{A}_{\text{rec}}$, $t \in \mathbb{R}$, and $j = 1, 2$. Then $V_1(t) = V_2(t)$ for all t, and their self-adjoint Stone generators are equal, including their maximal domains.

Proof. For $A \in \mathfrak{A}_{\text{rec}}$,

$$V_j(t) A\Omega = \alpha_t(A)\Omega.$$

The groups agree on the dense subspace $\mathfrak{A}_{\text{rec}}\Omega$ and hence on $\mathcal{H}$. Stone uniqueness identifies the generators. ☐

Boundary 4.3 (Equality is stronger than common implementation)*.* Two unitary groups may implement the same algebra automorphisms while differing by a unitary in the commutant. The common invariant cyclic vector removes that ambiguity. Abstract similarity of dynamics does not.

Proposition 4.4 (General exact-zero identification)**.** *Let $\mathfrak{O}$ and $\mathfrak{D}$ act on the same $\mathcal{H}$. Assume, for an owned scalar $a_E > 0$,*

$$H_E = a_E H_{\text{thin}} = a_E K^{1/2}$$

as self-adjoint operators, including equality of maximal domains after scalar rescaling, and assume $S \cap h^{-1}(0) = S \cap Z$. Then

$$\boxed{P_\Omega = E(Z).}$$

Proof. The spectral theorem gives $\mathbf{1}_{\{0\}}(H_E) = E(h^{-1}(0)) = E(Z)$. The last equality follows from Lemma 3.3 and the zero-fiber condition. On the other hand, positive scalar rescaling preserves the kernel, so

$$\mathbf{1}_{\{0\}}(H_E) = \mathbf{1}_{\{0\}}(a_E K^{1/2}) = P_\Omega.$$

☐

The conclusion identifies the full complete-ground support. It does not assert rank one, factoriality, or uniqueness of a vacuum vector.

This is also where zero energy and zero momentum cease to be interchangeable shorthand. Equality of generators identifies their kernels. The zero-fiber condition then identifies the kernel of the height operator with the declared invariant zero set. In the relativistic specialization the forward-cone condition makes that set the single point $p = 0$. Without the zero-fiber statement, an additional zero-height spectral component could survive even when the same-carrier dynamical identification is exact.

5 The general invariant-transfer theorem

The previous results now close one analytic packet.

Closure here has two layers. Analytic closure means that the coercive operator, the height generator, the covariant PVM, and the orbit-floor observable all live on one carrier and satisfy the hypotheses needed for functional calculus. Physical closure means that this carrier and its distinguished projection retain the source marking and provenance of the selected reconstruction branch. The first layer proves the inequality. The second determines whose inequality it is.

Definition 5.1 (Closed invariant-transfer packet)**.** A *closed invariant-transfer packet* is a pair $(\mathfrak{O}, \mathfrak{D})$ on one Hilbert space such that:

(IT1) $S \cap h^{-1}(0) = S \cap Z$;

(IT2) $H_E = a_E H_{\text{thin}} = a_E K^{1/2}$ as self-adjoint operators, with an owned normalization $a_E > 0$ and $c_E = a_E^2$;

(IT3) $K \geq \kappa(1 - P_\Omega)$ with an owned $\kappa > 0$;

(IT4) the source, marking, normalization, branch, and lineage record is retained in $\Pi = (\Pi_{\text{thin}}, \Pi_{\text{spectral}})$.

The packet is *physically closed* when P_Ω, K, and E lie in one certified source-provenance component.

Theorem 5.2 (General invariant spectral transfer)**.** *Every closed invariant-transfer packet satisfies*

$$P_\Omega = E(Z), \quad S \subseteq Z \cup \{x : \rho_G(x)^2 \geq c_E \kappa\}, \quad Q_G \geq c_E \kappa (1 - P_\Omega).$$

Moreover, P_Ω reduces Q_G and $P_\Omega \leq \mathbf{1}_{\{0\}}(Q_G)$, so (Q_G, P_Ω) is an invariant certificate packet in the sense of Section 7. If the packet is physically closed, these conclusions belong to the terminal spectral shadow of the same retained source datum.

Proof. Proposition 4.4 gives $P_\Omega = E(Z)$. Same-carrier identification turns the thin inequality into

$$H_E^2 = c_E K \geq c_E \kappa (1 - E(Z)).$$

Theorem 3.5 gives the support and invariant-shadow conclusions. Because $E(Z)$ and Q_G belong to the same functional calculus, $E(Z)$ reduces Q_G; on $S \cap Z$ one has $h = 0$ and hence $\rho_G = 0$, so Lemma 3.3 gives $Q_G E(Z) = 0$. Physical closure supplies identity of provenance; it is not needed for the analytic inequality. ☐

This separation is structural. Covariant spectral calculus proves an analytic result. Lineage proves which physical object the result is about.

The theorem consequently has no need to encode the entire history of admissible transport inside its spectral proof. Provenance enters as a typed interface condition after the analytic packet has closed. Conversely, analytic equality cannot be inferred from provenance alone. The two parts meet at one terminal carrier, but they retain different theorem ownership.

6 Relativistic mass as the orbit-floor specialization

Let $X = \overline{V}_+$, let $G = SO^+(1, d)$, take $Z = \{0\}$, and set $h(p) = p^0$. The forward cone is a closed, locally compact, second-countable, Lorentz-invariant space, and $h : X \to [0, \infty)$ now has exactly the type required by Definition 3.1.

For $p \in \overline{V}_+$,

$$\rho_G(p) = \begin{cases} \sqrt{p_\mu p^\mu}, & p^2 > 0, \\ 0, & p^2 = 0. \end{cases}$$

The first identity is obtained by boosting to the rest frame. The second follows by boosts along the null direction, which rescale the positive energy toward zero.

This computation isolates the no-light-cone-escape statement. A nonzero null point may have energy bounded below in one chosen frame or on one truncated ray. Full Lorentz covariance nevertheless generates boosts for which that energy approaches zero. If a positive energy exclusion holds on the entire covariant support, the null orbit cannot be present. Timelike orbits behave differently: their lowest attainable energy is their invariant mass. The same support argument therefore excludes the null orbit and assigns the retained timelike spectrum a uniform invariant floor.

Corollary 6.1 (Mass identification and invariant transfer)**.** *Let U be a strongly continuous positive-energy representation of $\mathcal{P}_+^\uparrow$ on $\mathcal{H}$, with sharply covariant joint translation PVM E_P and* $\operatorname{supp} E_P \subseteq \overline{V}_+$*. Let $(q, K, H_{\mathrm{thin}}, P_\Omega, \kappa)$ be a thin dynamical packet on the same carrier. If*

$$P^0 = H_{\mathrm{thin}} = K^{1/2}, \qquad K \geq \kappa(1 - P_\Omega),$$

then

$$\begin{aligned} P_\Omega &= E_P(\{0\}), \\ \operatorname{supp} E_P &\subseteq \{0\} \cup \{p \in \overline{V}_+ : p^2 \geq \kappa\}, \\ E_P(\{p \neq 0 : p^2 = 0\}) &= 0, \\ M^2 &\geq \kappa(1 - P_\Omega). \end{aligned}$$

Thus κ is a certified mass-square threshold and $\sqrt{\kappa}$ is the corresponding mass threshold.

Proof. The spectrum condition gives $E_P(\overline{V}_+) = 1$, so E_P may be regarded as a PVM on the Borel space $X = \overline{V}_+$ without changing its functional calculus. On this spectral space, p^0 is nonnegative and the orbit system satisfies Definition 3.1. On the forward cone, the zero-energy face is $\{0\}$. The invariant orbit-floor shadow is

$$Q_G = \int p^2 \, dE_P(p) = M^2.$$

Theorem 5.2 gives the conclusions. □

No pointwise inequality $M^2 \geq c(P^0)^2$ is used. The transfer is possible because the joint support is Lorentz invariant. In particular, the theorem owns the constant: orbit transfer does not silently replace an energy threshold by a different mass threshold.

The distinction answers the central audit objection. An energy gap by itself controls distance from the zero-energy hyperplane in one frame. A mass gap controls distance from the entire light cone in the Lorentz-invariant geometry of joint momentum space. The orbit theorem supplies exactly the missing passage between those statements, and it supplies it only when the represented covariance is strong enough to move through the relevant orbits.

7 Certificates and sharp edges

The previous theorem produces a valid lower certificate. It does not prove that spectrum occurs at that value.

Definition 7.1 (Invariant certificate packet)**.** Let $Q \geq 0$ be a self-adjoint invariant spectral shadow and let P_0 be an orthogonal projection that reduces Q and satisfies $P_0 \leq \mathbf{1}_{\{0\}}(Q)$. For finite scalars define

$$\operatorname{Cert}(Q, P_0) := \{\eta \in [0, \infty) : Q \geq \eta(1 - P_0)\}$$

in the closed-form sense, and define

$$\kappa_{\mathrm{sharp}}(Q, P_0) := \inf \operatorname{Spec}\Big(Q\big|_{(1-P_0)\mathcal{H}}\Big),$$

with $+\infty$ when the nonzero subspace is absent. An owned certificate is an element $\kappa_{\mathrm{cert}} \in \operatorname{Cert}(Q, P_0)$ with its theorem, normalization, carrier, and provenance record.

The two quantities should remain separate in every later ledger. A certificate belongs to a proof route and may carry conservative losses. The sharp edge belongs to the terminal operator and records the actual bottom of its nonvacuum spectrum. Improving an estimate can raise a certificate without moving the spectrum; changing the selected sector can move the sharp edge without invalidating an inherited certificate.

Proposition 7.2 (Certificate interval). *If $(1-P_0)\mathcal{H} \neq 0$, then*

$$\mathrm{Cert}(Q, P_0) = [0, \kappa_{\mathrm{sharp}}(Q, P_0)].$$

If $(1-P_0)\mathcal{H} = 0$, then

$$\mathrm{Cert}(Q, P_0) = [0, \infty), \qquad \kappa_{\mathrm{sharp}}(Q, P_0) = +\infty.$$

In both cases, in the extended-real sense,

$$\kappa_{\mathrm{sharp}}(Q, P_0) = \sup \mathrm{Cert}(Q, P_0).$$

Proof. On the reducing nonzero subspace, the certificate inequality is precisely $Q \geq \eta$. The spectral theorem identifies its largest valid finite constant with the bottom of the restricted spectrum. If that subspace is absent, every finite scalar is a certificate, but $+\infty$ is not itself treated as an operator coefficient. □

A valid certificate can decrease because a proof uses a conservative comparison. The sharp edge can decrease only because the actual assembled spectrum approaches zero. These are different failure modes.

8 General transport calculus

The general invariant shadow Q_G admits a transport calculus independent of its relativistic specialization.

The need for this calculus is not cosmetic. Reconstruction changes carriers by unitary seams, isometric embeddings, reducing selections, infinite assemblies, and limiting passages. These operations do not act identically on thresholds. Exact unitary transport preserves the whole spectral packet, compression can delete low spectrum, and an infinite family can accumulate at zero even when every member is individually positive. A threshold statement is therefore incomplete until the operation carrying it has been named.

Definition 8.1 (Spectral and lineage-spectral arrows). For covariant spectral systems, a bounded map $W : \mathcal{H}_1 \to \mathcal{H}_2$ is a *spectral arrow* when

$$W E_1(B) = E_2(B) W$$

for every Borel set B. When the PVMs are the joint spectral measures of abelian unitary representations, it is enough to intertwine those representations. A *lineage-spectral arrow* is a spectral arrow that also carries a certified passage between the provenance records.

Theorem 8.2 (All-Borel naturality). *If W is a spectral arrow, then*

$$W\, \mathrm{Dom}\, f(E_1) \subseteq \mathrm{Dom}\, f(E_2), \qquad W f(E_1)\psi = f(E_2) W\psi$$

for every Borel f and every $\psi \in \mathrm{Dom}\, f(E_1)$. If W is unitary, the maximal domains are carried bijectively. If the PVMs are the joint spectral measures of abelian unitary representations, an intertwiner of those representations is a spectral arrow and therefore satisfies the same conclusions.

Proof. Simple-function calculus follows from the defining PVM identity, and bounded Borel calculus follows by monotone-class closure. For an arbitrary Borel f, take the canonical truncations

$$f_n = f\, \mathbf{1}_{\{|f| \leq n\}}.$$

If $\psi \in \mathrm{Dom}\, f(E_1)$, then

$$\|W(f_n(E_1) - f_m(E_1))\psi\| \leq \|W\|\, \|(f_n(E_1) - f_m(E_1))\psi\|.$$

Thus $f_n(E_2)W\psi = W f_n(E_1)\psi$ converges. Because W intertwines the spectral projections, $E_2(\{|f| \leq n\})W\psi$ converges to $W\psi$. Closedness of $f(E_2)$ now gives $W\psi \in \mathrm{Dom}\, f(E_2)$ and the asserted equality. Apply the same argument to W^{-1} in the unitary case. In the abelian representation case, the PVM identity itself follows from equality of the matrix-coefficient Fourier transforms. □

Theorem 8.3 (Elementary certificate rules)**.** *Let (Q_i, P_i) be invariant certificate packets.*

(i) If W is a unitary spectral arrow with $WP_1 = P_2W$, then the full certificate sets and sharp edges agree.

(ii) If W is unitary, $P_2 = WP_1W^*$, and $Q_2 = cWQ_1W^*$ with $c > 0$, then certificates and sharp edges are multiplied by c; unsquared thresholds are multiplied by $\sqrt{c}$.

(iii) If J is an isometric spectral arrow and $P_1 = J^*P_2J$, its range reduces the target calculus. Compression preserves every target certificate and can only raise the sharp edge.

(iv) A source estimate on $\operatorname{Ran} J$ gives no estimate on $(\operatorname{Ran} J)^\perp$.

Every physical identity assertion additionally requires a lineage-spectral arrow.

Proof. All statements follow from Theorem 8.2 and the spectral theorem on the relevant reducing subspaces. The complement statement follows by adjoining an arbitrary nonnegative operator on $(\operatorname{Ran} J)^\perp$. □

Theorem 8.4 (Infinite assembly)**.** *For a direct sum of invariant certificate packets,*

$$\kappa_{\mathrm{sharp}}\left(\bigoplus_i Q_i, \bigoplus_i P_i\right) = \inf_i \kappa_{\mathrm{sharp}}(Q_i, P_i).$$

For a measurable decomposable direct integral in which $x \mapsto P_x$ is a measurable projection field and the restricted operators on $(1 - P_x)\mathcal{H}_x$ form a measurable self-adjoint field,

$$\kappa_{\mathrm{sharp}}\left(\int_X^{\oplus} Q_x \, d\mu(x), \int_X^{\oplus} P_x \, d\mu(x)\right) = \operatorname*{ess\,inf}_x \kappa_{\mathrm{sharp}}(Q_x, P_x).$$

The analogous infimum or essential-infimum rule holds for recorded certificates.

Proof. The restricted operator is the direct sum or decomposable integral of the restricted fiber operators. The bottom of its spectrum is the infimum or essential infimum of the fiber bottoms. □

Pointwise positivity is insufficient for an infinite assembly. If the sharp fiber edges are $1/n$, every fiber is gapped but the sum has spectrum accumulating at zero. If only the chosen nonsharp certificates are $1/n$ while all sharp edges equal one, the failure is merely in the ledger.

9 Relativistic assembly rules

Finite products use additional geometry that is special to the forward cone.

Lemma 9.1 (Future-cone sum rigidity)**.** *For $p_1, \ldots, p_N \in \overline{V}_+$,*

$$\left(\sum_i p_i\right)^2 \geq \left(\sum_i \sqrt{p_i^2}\right)^2.$$

Moreover, $\sum_i p_i = 0$ if and only if every $p_i = 0$.

Proof. Future-causal vectors satisfy $p_i \cdot p_j \geq \sqrt{p_i^2 p_j^2}$. Expanding the square proves the first statement. The nonnegative time components prove the second. □

Theorem 9.2 (Relativistic threshold assembly)**.** *For positive-energy Poincaré packets with nonzero vacuum spaces:*

(i) the full finite tensor product has vacuum projection $\bigotimes_i P_{\Omega,i}$ and sharp mass-square edge $\min_i \kappa_{\mathrm{sharp},i}$;

(ii) a translation-reducing product selection, central quotient, or sector compression inherits the minimum certificate, while its sharp edge may rise;

(iii) a corridor family has the uniform certificate $\inf_c \kappa_{\mathrm{cert},c}$ and uniform sharp edge $\inf_c \kappa_{\mathrm{sharp},c}$;

(iv) exact unitary seams with one normalization preserve the entire certificate interval and sharp edge.

Proof. The product PVM is pushed forward by momentum addition. Lemma 9.1 gives the zero fiber and the minimum lower bound. Tensoring one near-edge nonvacuum vector with vacuum vectors in all other factors proves sharp equality for the full product. Reducing selection can delete but cannot add low spectrum. The corridor and unitary claims follow from the general calculus. □

The full-product equality must not be transferred to a selected branch. Selection may remove every single-factor excitation at the minimum edge.

10 From source coercivity to the thin carrier

The invariant theorem begins on the terminal carrier. Admissible reconstruction begins with finite source data. The intervening passages must retain more than a number.

A positive seed constant has no intrinsic destination. It becomes a terminal certificate only through a chain that controls newly created directions, retains the chosen packet, and survives the limiting form construction. This is why source coercivity is recorded together with marking and provenance. The number measures a lower bound; the packet determines which physical branch that lower bound belongs to.

Definition 10.1 (Marked coercive source packet)**.** A *marked coercive source packet* is

$$\mathfrak{C}_{\mathrm{src}} = (\mathbf{V}_0, \mathbf{P}_0, \mathbf{H}, \kappa, \mathcal{Z}, \mathfrak{m}, \mathfrak{h}, \mathfrak{s}, \Pi_{\mathrm{src}}),$$

where $\mathbf{V}_0$ is the source datum, $\mathbf{P}_0$ its complete-ground family,

$$\mathbf{H} \geq E_0 + \kappa(1 - \mathbf{P}_0), \qquad \kappa > 0,$$

$\mathcal{Z}$ and $(\mathfrak{m}, \mathfrak{h})$ record marking and source fingerprint, $\mathfrak{s}$ records seam and target-complement certification, and Π_{src} records provenance.

Equal constants do not identify source packets. Marking, holonomy, quotient, regional carrier, branch, and provenance are constitutive inputs.

Proposition 10.2 (One-sided seam firewall)**.** *Let $J : \mathcal{H}_\alpha \to \mathcal{H}_\beta$ be an isometry. An estimate*

$$q_\beta[Jx] \geq \kappa \|(1 - P_\beta)Jx\|^2$$

controls only $\operatorname{Ran} J$*. It implies no lower bound on* $(\operatorname{Ran} J)^\perp$*.*

Proof. On the orthogonal complement, adjoin an arbitrary nonnegative closed form. Choosing the zero form leaves the displayed range estimate unchanged while creating a zero-cost target complement. □

Proposition 10.3 (Independent target-complement criterion)**.** *Suppose*

$$\mathcal{H}_\beta = J\mathcal{H}_\alpha \oplus \mathcal{K}_{\mathrm{new}}, \qquad P_\beta = JP_\alpha J^*,$$

the form has no cross term, the old image has lower certificate κ, and

$$q_\beta[y] \geq \kappa \|y\|^2, \qquad y \in \mathcal{K}_{\mathrm{new}}.$$

Then $q_\beta \geq \kappa(1 - P_\beta)$ on the full target form domain.

Proof. Write $\Psi = Jx + y$ and add the two orthogonal form estimates. □

Hypothesis 10.4 (Comparison–Mosco packet)**.** Let $(\mathcal{H}_n, q_n, P_n)$ approach $(\mathcal{H}, q, P)$ through comparison maps J_n. Assume every $\psi \in \operatorname{Dom} q$ has a recovery family $\psi_n \in \operatorname{Dom} q_n$ such that

$$J_n\psi_n \to \psi, \quad q_n[\psi_n] \to q[\psi], \quad \|\psi_n\| \to \|\psi\|, \quad \|(1 - P_n)\psi_n\| \to \|(1 - P)\psi\|,$$

and

$$q_n[\phi] \geq \kappa_n \|(1 - P_n)\phi\|^2, \qquad \liminf_n \kappa_n \geq \kappa_* > 0.$$

Theorem 10.5 (Comparison–Mosco coercive transfer)**.** *Under Hypothesis 10.4,*

$$q[\psi] \geq \kappa_* \|(1-P)\psi\|^2, \qquad \psi \in \operatorname{Dom} q.$$

Proof. For $\varepsilon > 0$, eventually $\kappa_n \geq \kappa_* - \varepsilon$. Apply the endpoint inequality to a recovery family, pass to the limit, and let ε decrease to zero. □

This theorem transfers a lower certificate. Its proof uses only the recovery-side energy convergence and ground-complement norm convergence displayed in Hypothesis 10.4; the Mosco lower condition is not needed for this scalar inequality. Full Mosco convergence is needed upstream to certify that q is the actual limiting closed form, and Riesz stability remains separately necessary to identify an isolated limiting zero projection. The theorem does not prove occupancy at the edge or equality of sharp positive spectral values.

11 Threshold lineage and route comparison

Vacuum lineage identifies a physical object across changing carriers. Threshold lineage performs the analogous bookkeeping for a certified lower bound, but the two notions are not identical. The physical destination may be route-independent while the numerical strength of available estimates depends on the chosen comparisons. The missing datum is a route ledger that records losses without multiplying physical sectors.

Definition 11.1 (Certificate passage and threshold lineage)**.** A *certificate passage* consists of a certified carrier passage, preservation of marking and provenance, a declared numerical rule, and a theorem validating the destination certificate. A *threshold lineage* is a composable chain

$$\boxed{\Delta_{\text{seed}} \rightsquigarrow \kappa_{\text{cof}} \xrightarrow{\text{Mosco}} \kappa_{\text{thin}} \xrightarrow{c_E} \eta_E \xrightarrow{\text{orbit}} \kappa_{\text{inv}} \leq \kappa_{\text{sharp}}.}$$

Every arrow records its carrier, theorem owner, normalization, numerical loss rule, marking, and source-provenance component.

The orbit arrow changes interpretation without numerical loss. The factor c_E belongs earlier, in the same-carrier comparison. If

$$(P^0)^2 \geq c_E K^{\text{thin}}, \qquad K^{\text{thin}} \geq \kappa_{\text{thin}}(1-P_\Omega),$$

then the energy-square certificate entering orbit transfer is $\eta_E = c_E \kappa_{\text{thin}}$.

Definition 11.2 (Lineage-relative route envelope)**.** Let Γ_0 be a nonempty family of certified routes in one source-provenance component, all terminating in the same invariant packet (Q, P_0). If route γ gives $\eta_\gamma \in \operatorname{Cert}(Q, P_0)$, define

$$\kappa_{\text{rec}}(\Gamma_0) := \sup_{\gamma \in \Gamma_0} \eta_\gamma.$$

Theorem 11.3 (Route-envelope theorem)**.** *The route envelope satisfies, in the extended-real order,*

$$\kappa_{\text{rec}}(\Gamma_0) \leq \kappa_{\text{sharp}}(Q, P_0).$$

If $\kappa_{\text{rec}}(\Gamma_0) < +\infty$, it is a valid terminal certificate. If $\kappa_{\text{rec}}(\Gamma_0) = +\infty$, then $(1-P_0)\mathcal{H} = 0$ and every finite scalar is a valid terminal certificate; $+\infty$ itself is not used as an operator coefficient. Thus physical route independence is compatible with route-dependent proof strength.

Proof. Every route gives $Q \geq \eta_\gamma(1-P_0)$, hence $\eta_\gamma \leq \kappa_{\text{sharp}}(Q, P_0)$ by Proposition 7.2. Taking the supremum gives the extended-real inequality. When the supremum is finite, monotone passage in the quadratic-form inequality proves that it is a certificate. If the supremum is infinite, then $\kappa_{\text{sharp}}(Q, P_0) = +\infty$; Proposition 7.2 forces the nonzero subspace to be absent. □

Certificates from different provenance components may not be pooled merely because their terminal operators are isomorphic.

12 The source-to-invariant bridge

The preceding components now answer the complete interface question. Source recognition supplies a marked positive packet. Destination classification says whether it is retained. Comparison–Mosco passage places its coercive content on the thin carrier. Same-carrier realization identifies the physical height generator. Orbit covariance turns that height exclusion into an invariant threshold. Lineage certifies that the endpoint is still the selected source sector. Removing any one step changes the conclusion that can be stated.

Theorem 12.1 (General source-recognition-to-invariant-threshold theorem)**.** *Assume:*

(B1) a marked coercive source packet with a positive finite certificate;

(B2) a retained cofinal family with independent target-complement certification and a uniform floor $\kappa_{\mathrm{cof}} > 0$;

(B3) destination classification placing the family in one retained provenance component;

(B4) comparison–Mosco convergence and complete-ground projection transport to a thin closed form;

(B5) an orbit-covariant spectral system on the same terminal Hilbert carrier;

(B6) a same-carrier comparison

$$H_E = a_E H_{\mathrm{thin}}, \qquad H_{\mathrm{thin}} = (K^{\mathrm{thin}})^{1/2}, \qquad c_E = a_E^2 > 0;$$

(B7) the zero-fiber condition $S \cap h^{-1}(0) = S \cap Z$;

(B8) source-to-terminal lineage for the distinguished projection.

Then

$$P_\Omega = E(Z), \qquad Q_G \geq c_E \kappa_{\mathrm{cof}}(1 - P_\Omega), \qquad \kappa_{\mathrm{sharp}}(Q_G, P_\Omega) \geq c_E \kappa_{\mathrm{cof}}.$$

The conclusion belongs to the terminal spectral shadow of the same selected source packet.

Proof. The source and complement hypotheses give a uniform endpoint certificate on the retained branch. Theorem 10.5 carries it to the thin form. The same-carrier comparison contributes c_E and turns it into a height-square exclusion. Theorem 5.2 produces the orbit-invariant threshold. The lineage hypothesis identifies the projection and carrier with the selected source packet. Proposition 7.2 gives the sharp-edge inequality. ☐

The theorem is the general bridge. Neither Yang–Mills nor Poincaré geometry appears in its statement. Those structures enter when one chooses the source packet and identifies the orbit-floor observable.

13 Sector-resolved invariant shadows

Let (Q, P_0) be an invariant certificate packet with $P_0 = E(Z)$, and let a finite terminal atomic abelian algebra be

$$\mathcal{Z}_0 = \bigoplus_s \mathbb{C}\widehat{z}_s.$$

Assume every $\widehat{z}_s$ commutes with the covariant PVM and the terminal invariant shadow. Define

$$\mathcal{H}_s = \widehat{z}_s \mathcal{H}, \qquad P_{0,s} = P_0 \widehat{z}_s P_0, \qquad Q_s = Q|_{\mathcal{H}_s}.$$

Theorem 13.1 (Sector-resolved invariant threshold)**.** *If $Q \geq \eta(1 - P_0)$, then on every sector*

$$\boxed{Q_s \geq \eta(\widehat{z}_s - P_{0,s}).}$$

Moreover $P_{0,s} = \widehat{z}_s E(Z) = E_s(Z)$.

Proof. Compress the global functional-calculus identity and inequality by the reducing projection $\widehat{z}_s$. ☐

If a connected symmetry group acts continuously and preserves the finite atomic algebra, it fixes the atoms pointwise: the induced continuous homomorphism lands in the discrete finite automorphism group and is therefore constant. Disconnected admitted symmetries may still act by an explicitly permitted finite permutation. None of this implies sector uniqueness, vector uniqueness, or factoriality.

14 Yang–Mills recognition corollary

The general theorem now receives the selected Wilson–Tube data. The five obligations are

$$\boxed{\begin{array}{ll} \mathrm{MI}_1: & P^0 = H_{\text{thin}} = (K^{\text{thin}})^{1/2}, \\ \mathrm{MI}_2: & P_\Omega = E_P(\{0\}), \\ \mathrm{MI}_3: & \operatorname{supp} E_P \setminus \{0\} \subseteq \{p \in \overline{V}_+ : p^2 \geq m_{\text{cert}}^2\}, \\ \mathrm{MI}_4: & M^2 \geq m_{\text{cert}}^2(1 - P_\Omega), \\ \mathrm{MI}_5: & \text{the same retained source sector carries } \mathrm{MI}_1\text{–}\mathrm{MI}_4. \end{array}}$$

The list is not a replacement decomposition of frozen Paper V. It is an audit interface imposed afterward. Its purpose is to expose theorem ownership at the exact places where an automated or skeptical reader might otherwise compose unlike objects. The frozen paper owns the analytic realization and invariant transfer. The admissible-transport sequence explains how the same selected sector reaches that terminal representation.

Theorem 14.1 (Yang–Mills mass-identification recognition theorem)**.** *On each frozen Paper-V branch satisfying its declared exact-dilation Sector-I or scalar Sector-III hypotheses, the Wilson–Tube data instantiate the Poincaré specialization of Theorem 12.1. On the normalized exact branch,*

$$c_E = 1, \qquad \kappa_{\text{inv}} = \kappa_{\text{thin}}, \qquad M^2 \geq \kappa_{\text{thin}}(1 - P_\Omega).$$

The frozen construction owns MI_1*–*MI_4 *[5, Theorems 8.10, 8.13, 12.19, and 12.22–12.25]. GAT IV supplies the route-independent source-to-terminal provenance needed for* MI_5 *[4, Theorems 7.4 and 8.1]. Its terminal equation in Theorem 9.1 imports Paper V's exact-zero theorem and is therefore a lineage interpretation, not an independent analytic premise of the present proof. If the terminal marking algebra is descent-compatible with connected Poincaré covariance, Theorem 13.1 gives the sectorwise mass-square inequality.*

Proof. Recognition supplies the exact complete-ground source packet, positive cofinal floor, full marking, and source fingerprint [1, Proposition 14.1 and Theorems 14.2 and 14.4]. GAT I classifies the destination and separates retention, migration, and escape [2, Theorems 8.3–8.5 and Corollary 9.1]. GAT III supplies the thin complete-ground carrier [3, Theorems 7.1 and 8.1]. GAT IV identifies one route-independent source-to-terminal vacuum lineage [4, Theorems 7.4 and 8.1]. Its Theorem 9.1 then interprets the Paper-V exact-zero identity as the terminal Hilbert shadow of that lineage; it is not used to re-prove exact zero here.

Frozen Paper V constructs the terminal Poincaré representation and joint PVM, identifies its time generator with the positive thin square root, proves forward-cone support, and performs Lorentz-orbit transfer [5, Theorem 4.3, Definition 4.5, Theorem 8.10, Theorem 8.13, and Theorems 12.19 and 12.22–12.25]. These inputs instantiate the general bridge with $X = \overline{V}_+$, $G = SO^+(1,3)$, $h(p) = p^0$, $Z = \{0\}$, and $Q_G = M^2$. The proof recognizes the interface among those results; it does not replace their construction. □

For the selected SU(3) recognition packet,

$$\Delta_\triangle^{SU(3)} = 1.2482172840978913294\ldots, \qquad \kappa_{\text{cof}} = 0.3374044376020202719\ldots.$$

The numerical drop is an owned conservative source-side comparison [1, Proposition 14.1 and Theorem 14.2]. It is not a Lorentz-transfer loss. After physical normalization, the κ quantities are mass-square certificates and the certified mass is their square root.

Sector II remains conditional on a physically verified corridor realization and a positive uniform corridor infimum. A non-scalar correspondence-valued branch requires its own spectral-correspondence theorem. Neither case is silently absorbed into the scalar Hilbert-space theorem.

The conclusion is therefore exact on its declared branches and silent elsewhere. It neither weakens the frozen Yang–Mills result nor promotes conditional branches into unconditional ones. It identifies, in general form, why the successful scalar branches support the final invariant-mass statement.

15 Nonclaims and domain of applicability

The conclusion is deliberately narrower than a universal reconstruction theorem.

N1. A source coercive certificate is not itself a terminal invariant gap.

N2. Equal thresholds, ranks, or zero energies do not establish physical lineage.

N3. An old-image estimate does not control a new target complement.

N4. Mosco transfer preserves a lower certificate, not occupancy at the sharp edge.

N5. An energy-square bound becomes an invariant bound only under sufficient covariance to compute the required orbit floor.

N6. Compression can raise a sharp edge; infinite assembly can lower it to an infimum or essential infimum.

N7. Vacuum completeness, sector uniqueness, vector uniqueness, and factoriality remain distinct.

N8. Arbitrary correspondence or Morita equivalence does not automatically provide a Hilbert-space unitary or a transported PVM.

N9. GAT V does not modify the frozen theorem statements or discovery order of Papers I–V.

Outside these hypotheses the theorem provides no certification. This is a boundary of the interface, not a negative statement about possible model-specific arguments.

A Exact source concordance

Every entry below was checked against its native manuscript rather than inherited from an earlier status ledger. The source versions are Recognition dated August 6, 2026; GAT I v2.6.0; GAT III release candidate v1.0; GAT IV v1.0.0; and frozen Paper V journal manuscript v1.0.12. The distinction between Paper V's internal PHIGT layer and its later mobile-incision global closure is retained explicitly.

Interface	Exact owner	Imported conclusion
Exact finite seed	Recognition [1], Proposition 14.1	Complete-ground projection and exact three-sector seed gap.
Cofinal uniform floor	Recognition [1], Theorem 14.2	Common positive floor; endpoint gaps may be larger.
One-sided seam firewall	Recognition [1], Proposition 14.3	Old-image comparison plus independent target-complement certification.
Full source fingerprint	Recognition [1], Theorem 14.4	Marking, holonomy, quotient, profiles, provenance, and late-tail demand.
Recognition boundary	Recognition [1], Warning 16.7	Source recognition is not downstream reconstruction or invariant mass positivity.
Destination classification	GAT I [2], Theorems 8.3–8.5; Corollary 9.1	Retention/migration/escape, Mosco–Riesz stability, and selected-branch no migration.
Thin complete ground	GAT III [3], Theorems 7.1 and 8.1	Thin physical carrier and exact complete-ground support.
Vacuum composition and spatialization	GAT IV [4], Theorems 6.2 and 6.5	Coherent composition of retained passages and its graph-reducible spatial shadow.
Route-independent lineage	GAT IV [4], Theorems 7.4 and 8.1	One source-certified, route-independent terminal vacuum lineage.
Terminal spectral interpretation	GAT IV [4], Theorem 9.1	Interprets Paper V's imported exact-zero identity as the terminal lineage shadow; it is not an independent analytic proof.
Sector zero shadows	GAT IV [4], Proposition 9.2 and Theorem 9.4	Descended marking atoms reduce the PVM and carry sectorwise zero-momentum shadows.
Joint spectral system	Frozen Paper V [5], Theorem 4.3; Definition 4.5	Joint PVM, strongly commuting momenta, and mass square by joint calculus.

Interface	Exact owner	Imported conclusion
Reduction before identification	Frozen Paper V [5], Corollary 4.16; Remark 4.17	The ground projection reduces the PVM; reduction alone does not identify $p = 0$.
Comparison–Mosco	Frozen Paper V [5], Hypothesis 7.2; Theorems 7.5, 7.9, and 7.13	Recovery, Glue–Mosco convergence, complete-ground Riesz stability, and uniform separation.
Internal same-carrier bridge	Frozen Paper V [5], Theorem 8.5; Lemma 8.6; Theorem 8.10	The physical time generator equals the positive thin square root; PHIGT identifies the zero-energy projection.
Internal invariant transfer	Frozen Paper V [5], Theorems 8.12–8.13	Energy-square threshold, exact joint zero, null-orbit exclusion, and invariant mass-square certificate.
All-Borel arrows	Frozen Paper V [5], Theorem 10.3	Joint and mass PVM naturality for exact translation intertwiners.
Unitary/isometric split	Frozen Paper V [5], Proposition 10.4; Theorems 10.8–10.9	Exact unitary support transport and range-only isometric transport.
Corridor families	Frozen Paper V [5], Theorems 10.12 and 13.12	Uniform-infimum rule and conditional physical realization.
Finite products	Frozen Paper V [5], Lemma 10.15; Theorem 10.16	Spectral convolution, tensor zero projection, and minimum bound.
Central quotients	Frozen Paper V [5], Theorem 10.18; Corollary 10.19	Reducing compression cannot lower the threshold.
Direct sums/integrals	Frozen Paper V [5], Theorems 10.21–10.22	Infimum and essential-infimum assembly rules.
Global same-carrier bridge	Frozen Paper V [5], Theorem 12.19	Descended positive energy and $P_0^2 = K^{\mathrm{thin}}$ on the mobile-incision carrier.
Global cone and orbit closure	Frozen Paper V [5], Theorems 12.22–12.25; Theorem 13.7	Forward cone, exact zero, Lorentz rigidity, joint-spectrum threshold, and mass-square inequality.
Final branch closure	Frozen Paper V [5], Theorem 13.18	Observable mass gap and physical no leakage on declared branches.

B Typed carrier dictionary

Stage	Native datum	Carrier	Certified relation
Source	$\mathbf{V}_0$, $\mathbf{P}_0$, κ_{src}	Finite recognition diagram	Realization interface; not operator equality.
Cofinal	Retained marked packet, κ_{cof}	Directed admissible diagram	Packet retention with marking and provenance.
Thin	q, K, P_Ω	Terminal physical Hilbert space	Mosco form limit and complete-ground transport.
Time flow	$H_{\mathrm{thin}} = K^{1/2}$	The same terminal carrier	Positive square root.
Spectral height	$H_E = \int h\, dE$	The same terminal carrier	Same-carrier generator identification.
Orbit shadow	$Q_G = \int \rho_G^2\, dE$	The same terminal carrier	Covariant orbit-threshold transfer.
Poincaré case	P^0, E_P, M^2	Physical translation representation	$h = p^0$, $\rho_G^2 = p^2$.
Vacuum lineage	P_Ω	Terminal Hilbert shadow	$P_\Omega = E(Z)$ after exact-zero identification.
Sector	$P_{0,s}$, Q_s	$\widehat{z}_s\mathcal{H}$	Reducing compression by a fixed marking atom.

Literal equality is reserved for one declared carrier. Realization crosses native types. Spectral covariance transports a PVM. Lineage identifies source provenance. Compression restricts a reducing system. A shadow is a carrier-dependent representative, not a universal operator.

C Threshold-operation dictionary

Operation	Certificate rule	Sharp-edge rule
Unitary seam	unchanged	equality
Normalization $c > 0$	multiply by c	multiply by c
Reducing compression	inherited	may rise
Corridor family	infimum	infimum
Full finite Poincaré product	minimum	exact minimum
Selected product/quotient	minimum inherited	at least the minimum
Direct sum	infimum	exact infimum
Direct integral	essential infimum	exact essential infimum

Every row is a spectral rule. To become a statement about one physical reconstruction, the row must be accompanied by its certified provenance passage.

D Proof-boundary ledger

Boundary joint	Failure mode	Resolution
PVM support	Equality on S used without proving that E is supported on S	Second countability yields a countable null base and $E(X \setminus S) = 0$.
Orbit floor	Infimum over an uncountable group may fail to be visibly measurable	Second countability of G reduces the infimum to a countable dense subset; ρ_G is upper semicontinuous.
Relativistic carrier	$h(p) = p^0$ is not nonnegative on all of $\mathbb{R}^{1+d}$	The spectrum condition restricts the PVM to the invariant Borel space $X = \overline{V}_+$, where $h : X \to [0, \infty)$ has the declared type.
Height exclusion	Strict sublevel set declared null directly from a form inequality	Exhaust by uniformly separated sublevels $h^2 \leq \eta - 1/n$ before invoking PVM continuity.
Exact zero	Positive normalization handled by an implicit rescaled packet	Proposition 4.4 includes the owned scalar $a_E > 0$ explicitly.
Certificate packet	Restricted operator used without a reducing hypothesis	P_0 is required to reduce Q and lie in its zero spectral projection.
Unbounded calculus	Domain transport hidden inside inclusion notation	Theorem 8.2 proves maximal-domain transport by bounded truncation and closedness.
Compression	Normalization and isometric rules omitted projection compatibility	Unitary and isometric rules now state the exact vacuum-projection relations.
Infinite assembly	Fiber restriction assumed measurable without declaration	The direct-integral theorem now requires measurable projection and restricted-operator fields.
Route envelope	$+\infty$ treated as an operator coefficient	Infinite envelope is separated from finite certificates and forces an absent nonzero subspace.
Mosco transfer	Full Mosco convergence appeared load-bearing in a recovery-only estimate	The lower-certificate lemma records its minimal recovery hypotheses; full Mosco and Riesz roles remain upstream.
Program circularity	GAT IV Theorem 9.1 could appear to re-prove Paper V exact zero	GAT IV supplies lineage only; its terminal equation is explicitly an interpretation of the imported Paper-V theorem.

The listed failure modes are excluded inside the declared domain. Each resolution narrows or types a passage; none strengthens the frozen Yang–Mills claims.

E Sharp boundary models

The preceding ledger establishes the proof inside its declared domain. Sharpness asks a different question: what fails when one crosses that domain? A proposed implication may be false, an expression may cease to be defined, or a stated regularity package may be sufficient without being minimal. These outcomes must not be conflated.

E.1 False implications

W1. Energy gap and exact vacuum do not imply invariant mass gap. Fix $\Delta > 0$ and set

$$\mathcal{H} = \mathbb{C}\Omega \oplus L^2([\Delta, \infty), d\varepsilon).$$

Let the four strongly commuting translation generators vanish on Ω and act on the second summand by multiplication with

$$p(\varepsilon) = (\varepsilon, \varepsilon, 0, 0).$$

Then the joint spectrum lies in $\overline{V}_+$,

$$P^0 \geq \Delta(1 - P_\Omega), \qquad P_\Omega = E_P(\{0\}),$$

but

$$M^2 = (P^0)^2 - \mathbf{P}^2 = 0$$

on the entire nonvacuum summand. Thus MI_1, MI_2, and an energy gap do not imply MI_3 or MI_4. The missing datum is full Lorentz-orbit covariance: the truncated massless ray is not invariant under $\mathcal{P}_+^\uparrow$.

W2. The zero-fiber condition cannot be inferred from same-carrier realization. Take $X = \{0, z\}$, $\mathcal{H} = \mathbb{C}^2$, and let E have one rank-one atom at each point. Put $Z = \{0\}$ and $h(0) = h(z) = 0$. Then

$$H_E = 0, \qquad \ker H_E = \mathcal{H}, \qquad E(Z) \neq \mathbf{1}_{\{0\}}(H_E).$$

The carrier is fixed and the calculus is exact, but the extra zero-height point creates an unlicensed zero sector.

W3. A thin certificate does not transfer without an owned comparison. On $\mathbb{C}^2$, let

$$P_0 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad K = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad H_E = \begin{pmatrix} 0 & 0 \\ 0 & \epsilon \end{pmatrix}, \qquad 0 < \epsilon < 1.$$

Then $K \geq 1 - P_0$ and $\ker K = \ker H_E = \operatorname{Ran} P_0$, yet for the trivial orbit action the invariant shadow is

$$Q_G = H_E^2 = \begin{pmatrix} 0 & 0 \\ 0 & \epsilon^2 \end{pmatrix}.$$

Its sharp edge can be arbitrarily small. Common carrier and common zero projection do not supply the missing dynamical comparison.

W4. The normalization is squared. In the preceding two-dimensional model, impose the genuine bridge

$$H_E = aK^{1/2}, \qquad a = \tfrac{1}{2}.$$

Then $Q_G = H_E^2 = \frac{1}{4}K$. The valid transferred coefficient is $a^2\kappa = \frac{1}{4}$, while silently using $a\kappa = \frac{1}{2}$ gives a false lower bound. An energy normalization cannot be renamed as an energy-square normalization.

W5. An old-image estimate does not control a new target complement. Let

$$Q_\alpha = \operatorname{diag}(0, 1), \qquad P_\alpha = \operatorname{diag}(1, 0),$$

and embed $\mathcal{H}_\alpha = \mathbb{C}^2$ isometrically into $\mathcal{H}_\beta = \mathbb{C}^2 \oplus \mathbb{C}$. Define

$$Q_\beta = Q_\alpha \oplus 0, \qquad P_\beta = P_\alpha \oplus 0.$$

The unit certificate holds on the old image, but the new basis vector lies in $(1 - P_\beta)\mathcal{H}_\beta \cap \ker Q_\beta$. No positive target certificate survives.

W6. Reducing compression preserves a certificate but may raise the sharp edge. For

$$Q = \operatorname{diag}(0, 1, 2), \qquad P_0 = \operatorname{diag}(1, 0, 0),$$

the sharp edge is one. Compression to the reducing subspace spanned by the first and third basis vectors preserves the unit certificate but has sharp edge two. Therefore compression gives an inequality, not sharp equality.

W7. Pointwise gaps do not give an infinite-assembly gap. For each $n \geq 1$, take

$$Q_n = \operatorname{diag}(0, 1/n), \qquad P_n = \operatorname{diag}(1, 0).$$

Every fiber has a positive sharp edge, but

$$\kappa_{\text{sharp}}\left(\bigoplus_{n\geq 1} Q_n, \bigoplus_{n\geq 1} P_n\right) = \inf_{n\geq 1} \frac{1}{n} = 0.$$

Uniform positivity, not pointwise positivity, is the load-bearing infinite-assembly hypothesis.

W8. A transferred certificate does not assert occupancy at its value. Take the constant sequence

$$Q_n = Q = \operatorname{diag}(0, 2), \qquad P_n = P_0 = \operatorname{diag}(1, 0), \qquad \kappa_n = 1.$$

All finite-dimensional Mosco and projection convergence requirements hold, and the unit certificate transfers. Nevertheless the restricted spectrum is $\{2\}$: the value one is not occupied and the sharp edge is two.

W9. Disconnected symmetry may permute marking atoms. Let $\mathcal{Z} = \mathbb{C} \oplus \mathbb{C}$ act diagonally on $\mathbb{C}^2$, and let the nontrivial element of $\mathbb{Z}_2$ be represented by the coordinate swap. Conjugation preserves $\mathcal{Z}$ globally but interchanges its two atoms. Connectedness is therefore essential to the pointwise-fixation conclusion, though not to global algebra preservation.

W10. Vacuum completeness does not imply the three uniqueness properties. Let translations act trivially on $\mathbb{C}^2$, so

$$P_\Omega = E_P(\{0\}) = \mathbf{1}.$$

With the two-atom marking algebra $\mathbb{C} \oplus \mathbb{C}$ and ambient algebra $\mathbb{C} \oplus \mathbb{C}$, the vacuum is spectrally complete while there are two vacuum sectors, a two-dimensional vacuum space, and a nonfactorial algebra. Vacuum completeness, sector uniqueness, vector uniqueness, and factoriality remain logically distinct.

E.2 Undefined or unlicensed passages

W11. Unbounded spectral calculus requires domains. On $L^2((1, \infty), dx)$ let T be multiplication by x and let $\psi(x) = x^{-1}$. Then $\psi \in L^2$ but $x\psi \notin L^2$, so $\psi \notin \operatorname{Dom} T$. Even for the identity spectral arrow, writing $T\psi$ as a Hilbert vector is undefined. The maximal-domain statement in Theorem 8.2 is not optional notation.

W12. A nonmeasurable fiber field does not define a direct integral operator. Let $A \subset [0, 1]$ be non-Lebesgue-measurable and take one-dimensional fibers. Set $P_x = \mathbf{1}$ for $x \in A$ and $P_x = 0$ otherwise. Then

$$x \longmapsto \langle P_x 1, 1 \rangle = \mathbf{1}_A(x)$$

is not measurable. Hence $\int^\oplus P_x\, dx$ is not a decomposable projection, and no essential-infimum theorem can be applied.

W13. An infinite sharp edge is not an operator coefficient. For $\mathcal{H} = \mathbb{C}$, $Q = 0$, and $P_0 = \mathbf{1}$, the nonzero complement is absent,

$$\operatorname{Cert}(Q, P_0) = [0, \infty), \qquad \kappa_{\text{sharp}}(Q, P_0) = +\infty.$$

Every finite inequality is valid, but the string $Q \geq +\infty(1 - P_0)$ contains the undefined product $+\infty \cdot 0$. The extended value belongs to the order ledger, not the operator inequality.

W14. Morita equivalence does not manufacture a Hilbert-space unitary. The algebras $\mathbb{C}$ and $M_2(\mathbb{C})$ are strongly Morita equivalent, while their defining irreducible Hilbert carriers $\mathbb{C}$ and $\mathbb{C}^2$ have different dimensions. The equivalence induces representations through a correspondence; it is not a unitary between the displayed carriers. A separate graph-reducibility or spatial-effectivity theorem is required.

W15. Spectral coincidence does not manufacture provenance. Take two copies of the same triple $(\mathcal{H}, Q, P_0)$ and attach distinct source records $\Pi_1 \neq \Pi_2$. After forgetting provenance, every rank, spectrum, certificate, and matrix coefficient agrees. No statement in the forgotten data can decide whether a certified source passage exists. Lineage is additional structure, not a spectral invariant.

W16. An empty route family carries no certificate. If $\Gamma_0 = \varnothing$, a convention for $\sup \varnothing$ cannot create a certified passage or a physical threshold. This is why Definition 11.2 requires a nonempty family before forming $\kappa_{\text{rec}}(\Gamma_0)$.

E.3 What is sharp and what is merely sufficient

The countermodels prove necessity of the corresponding logical or typing boundaries. They do not prove that every topological hypothesis in Definition 3.1 is minimal. In particular, second countability is used to earn the two precise facts

$$E(X \setminus \operatorname{supp} E) = 0, \qquad \rho_G \text{ is Borel.}$$

A more general version may assume these facts directly or derive them from another regularity package. The present theorem keeps the second-countable formulation because it is checkable, stable under the intended applications, and sufficient for the proof.

Boundary	Witness	Certified lesson
N1	W3, W5, W7	Source coercivity needs comparison, complement control, and a uniform assembly floor.
N2	W15	Spectral coincidence does not determine physical lineage.
N3	W5	Range control leaves a new complement uncontrolled.
N4	W8	Mosco transfer gives a lower certificate, not edge occupancy.
N5	W1	Energy exclusion needs full orbit covariance to become invariant exclusion.
N6	W6, W7	Compression may raise and infinite assembly may lower the sharp edge.
N7	W9, W10	Symmetry and vacuum completeness do not collapse sectorial distinctions.
N8	W14	Correspondence equivalence is not automatically spatial unitary transport.
N9	Firewall	This is a construction-status boundary, not an operator implication.

The boundary models do not attack the theorem. They explain why its hypotheses carry mathematical information.

F Freeze ledger and conclusion

F.1 Frozen theorem spine

The following results form the stable dependency spine of the paper.

Interface	Stable result	Role
Orbit geometry	Theorem 3.5	Height exclusion becomes an invariant orbit-floor threshold.
Same-carrier dynamics	Theorem 4.2	Equality of time evolutions identifies their generators.
Exact zero	Proposition 4.4	The complete ground equals the declared zero spectral fiber.
General transfer	Theorem 5.2	Coercivity transfers to the invariant spectral shadow.
Relativistic specialization	Corollary 6.1	Energy-square exclusion becomes mass-square exclusion.
Certificate semantics	Proposition 7.2	Valid certificates are separated from the sharp edge.
Spectral transport	Theorems 8.2–8.4	Domains, seams, compression, sums, and integrals are typed.
Mosco endpoint	Theorem 10.5	A uniform lower certificate reaches the thin form.
Route comparison	Theorem 11.3	Route-dependent proof strengths share one terminal packet.
Source bridge	Theorem 12.1	A retained source certificate reaches the invariant shadow.
Sector resolution	Theorem 13.1	Fixed marking atoms carry compressed invariant thresholds.
Yang–Mills recognition	Theorem 14.1	The frozen construction instantiates the general interface.

F.2 Conclusion

The integrated architecture is

$$\boxed{\begin{array}{c}\text{marked source coercivity} \longrightarrow \text{retained thin coercivity} \longrightarrow \text{same-carrier height generator} \\ \longrightarrow \text{covariant orbit-floor threshold} \longrightarrow \text{sector-resolved invariant shadow.}\end{array}}$$

The first line is general admissible transport. The orbit-floor theorem is general covariant spectral geometry. Relativistic mass is the Lorentz specialization. Wilson–Tube Yang–Mills is the frozen primary construction recognized at the endpoint.

The decisive passage is not $H \mapsto M$ by change of notation. It is the preservation of one excluded spectral neighborhood under every admitted symmetry transformation. Once same-carrier realization makes the reconstructed Hamiltonian the height observable of the covariant PVM, the orbit floor records what no change of frame can remove. In relativistic momentum space that irreducible remainder is mass.

The support, measurability, zero-fiber, unbounded-domain, compression, direct-integral, and extended-threshold joints are closed. The sharp-boundary models show why the same-carrier bridge, zero-fiber condition, orbit covariance, target-complement control, uniform infinite floor, measurable assembly, domain declarations, connected atom action, and lineage data cannot be silently deleted. The analytic proof does not consume GAT IV's imported terminal equation; GAT IV enters only at the lineage interpretation. Within the declared domain, the theorem spine has no unresolved proof obligation.

Vacuum identity is a statement of certified lineage before it is a statement of spectral equality. Mass identity adds a second demand: the spectral exclusion must survive the full geometry of the symmetry orbit. Together they explain how a source-certified coercive packet reaches a terminal, sector-resolved invariant shadow without turning later recognition into a replacement construction.

Acknowledgment of research assistance

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.

References

[1] J. Scott Little, *Cohomological and Coercive Recognition of Compact Simple Gauge Groups: Sectorial Certificates for the Wilson–Tube Yang–Mills Construction*, research manuscript, August 6, 2026.

[2] J. Scott Little, *The Geometry of Admissible Transport: Grassmannian Gluing, Spectral Trichotomy, and Sectorial Type III Local Algebras*, version 2.6.0, August 5, 2026.

[3] J. Scott Little, *The Geometry of Admissible Transport III: Admissible Yang–Mills Fibrations and Thin Haag–Kastler Reconstruction*, release candidate v1.0, August 12, 2026.

[4] J. Scott Little, *The Geometry of Admissible Transport IV: Vacuum Lineage and Spectral Identification*, authorial manuscript v1.0.0, August 14, 2026.

[5] J. Scott Little, *Spectral Transfer, Interaction, and No-Leakage for Thin Yang–Mills Observable Representations: Physical Zero-Mode Architecture, Hilbert-Gerbe Spectral Stability, Incision-Relative Mass-Gap Formation, and Haag-Type Free-Net Non-Equivalence*, journal manuscript, version 1.0.12, July 25, 2026.