

Yang–Mills Theory in the Single Monad Model and Duality of Time Framework:

Projected Lattice Gauge Theory, Dangerous Clusters, and Full Yang–Mills Identification

A full SMM/DTT-based rewrite of the projected lattice-gauge synthesis

Mohamed Haj Yousef

Physics Department, United Arab Emirates University, Al Ain, UAE

`mhajyousef@uaeu.ac.ae`

Abstract

This paper rewrites the projected lattice-gauge program as the Yang–Mills sector of the Single Monad Model (SMM) and Duality of Time Theory (DTT), under the standing assumption that the SMM/DTTtemporal framework is already established and accepted. In this reading, the compact auxiliary cycle is not merely a technical regulator but the finite-regulator realization of compact inner time, while the physical Euclidean-time direction used in Osterwalder–Schrader reconstruction models outer time. Relative to the earlier comprehensive draft, the rewrite now absorbs two additional theorem packages into the manuscript itself. First, the finite-dimensional low-window shadow sector is organized intrinsically by the Real commutant of the compact dangerous-cluster symmetry representation: the shadow complex structure is treated as a skew-adjoint square root of -1 in that real C^* -algebra, with operator-valued orientations on complex-type multiplicity blocks and Frobenius–Schur real/quaternionic data on self-conjugate blocks. Second, the analytic transport layer is rewritten in producer–consumer form: high-window positivity from Hamiltonian–mixed-operator coercivity, thermodynamic persistence of projected lower bounds, gap-preserving continuum transfer along a renormalization trajectory, and norm-resolvent stability of finite-rank dangerous-cluster projections on a common comparison space are recorded as abstract exports of named theorem targets in the concrete lattice-gauge realization. Part I develops the fixed-regulator Euclidean lattice theory, the projected low-window gap mechanism, the producer–consumer coercivity/transport layer, and the continuum/interface reduction that lowers Yang–Mills identification to equality on a countable Wilson-loop/curvature core. Part II imports the canonical-complexification and compact nonabelian spectral-orientation toolkit into the dangerous cluster, rewrites the oriented-shadow layer in intrinsic Real-commutant language, identifies the one-band sector as the multiplicity-one reduction of the general model-plus-error criterion and the two-band sector as its first genuinely operator-valued case, and records Yang–Mills matching on those first closed subcores. Part III proves finite-band exhaustion and explicit off-block decay directly from the band-effective spectral data, packages bounded-curvature extension through constructive band-adapted Wilson-loop approximation, and completes the full-core identification/gap-transfer chain. The result is a stronger synthesis in which the analytically tractable frontier is visible inside the paper itself: projected-sector/common-core invariance, volume-uniform coercivity, thermodynamic embeddings, scale-wise continuum trajectories, common-comparison-space band

realizations, and uniform isolation of the low discrete cluster remain explicit theorem targets, while the late-stage closure and identification steps are organized as theorem-level consequences of those packages.

Contents

I Part I: Fixed-regulator construction, projected gap, and continuum/interface theory	3
1 Introduction	4
1.1 Scope and stance	4
1.2 Standing SMM/DTT framework	5
1.3 Translation into projected lattice gauge language	5
1.4 Spectral-topological bridge and temporal decimation heuristic	6
1.5 Claim-status discipline and current frontier	7
1.6 Three-part organization	7
1.7 Global strategy	8
1.8 Relation to earlier drafts and sources	9
1.9 Main endpoint theorem	9
1.10 Reader’s guide	10
2 Fixed-regulator constructive core	11
2.1 Finite-regulator lattice geometry and configuration space	11
2.2 Reflection positivity and Osterwalder–Schrader reconstruction	12
2.3 Transfer matrix and Hamiltonian	12
2.4 Exact covariant inner-time operator	13
3 Projected low-window reduction and the finite-volume gap	15
3.1 Projected four-dimensional sector and coercivity targets	15
3.2 Concrete coercivity program on the projected core	16
3.3 Shifted Feshbach reduction	20
3.4 Packetwise phase counting and weighted packet multiplicity	21
3.5 Band-effective model and dangerous cluster	21
3.6 Rank-one and rank-two endgames	23
4 Large-volume and continuum transfer	25
4.1 Uniform large-volume theorem	25
4.2 Scale-uniform trajectory package	26
5 Countable cores and interface theorems	28
5.1 Projected observable packages	28
5.2 Countable Wilson-loop/curvature cores	29
5.3 Density theorems	30
5.4 Euclidean interface theorem	30

5.5	Operator-algebraic interface theorem	31
II Part II: Oriented shadows, closed sectors, localization, and curvature approximation		32
6	Oriented shadows, Real commutants, and intrinsic matching	33
6.1	Dangerous clusters with compact symmetry	33
6.2	Block symbols and leakage control	36
6.3	From oriented-shadow rigidity to Yang–Mills matching	37
7	Closure of the one-band and two-band Wilson-loop sectors	38
7.1	One-band sector	38
7.2	Two-band sector	39
7.3	Restricted Yang–Mills identification on the closed subcores	41
8	Finite-band exhaustion from band-effective data	41
8.1	Imported dangerous-cluster and band-effective data	41
8.2	Tuple-adapted resonance graphs and finite-band covers	42
8.3	Explicit off-block decay and reduction to closed sectors	45
9	Constructive curvature approximation by Wilson-loop polynomials	46
9.1	Admissible bounded curvature generators and finite-band compatibility	46
9.2	Constructive bounded functional calculus	48
10	Final synthesis: full-core Yang–Mills identification and gap transfer	49
10.1	Bandwise gluing on finite-band sectors	49
10.2	Curvature extension and mixed-word matching	49
10.3	Final identification and gap transfer	50
Status and scope boundary		52
A	Reflection-positivity boundary factorization and transfer kernel	53
A.1	Boundary factorization	53
A.2	One-slab kernel and gauge covariance	53
B	Twisted Fourier calculation and a toy packet model	54
B.1	Twisted Fourier modes	54
B.2	Packetwise toy phase-count bound	54
B.3	Toy dyadic decimation on the auxiliary circle	54

Part I

Part I: Fixed-regulator construction, projected gap, and continuum/interface theory

1 Introduction

1.1 Scope and stance

This paper rewrites the projected lattice-gauge program as the Yang–Mills sector of the Single Monad Model (SMM) and Duality of Time Theory (DTT). The governing stance is therefore different from the earlier v15 manuscript. There, compact inner time was deliberately bracketed as background motivation while the proof spine was presented in purely lattice, spectral, and operator-theoretic language. Here, by contrast, the SMM/DTTframework is taken as accepted background structure. The compact auxiliary cycle is read as the regulator-level manifestation of compact inner time, and the physical Euclidean-time direction used for Osterwalder–Schrader reconstruction is read as the regulator-level representative of outer time [19, 15, 16]. Within that stance, the central mathematical object of the paper remains the exact holonomy-twisted covariant inner-time operator generated by the mixed physical/inner-time plaquettes. Its low covariant-energy window isolates a finite-dimensional dangerous cluster, and the later stages of the argument are organized around that explicit finite-dimensional sector.

The present rewrite absorbs two distinct mathematical upgrades into the comprehensive manuscript itself. First, whenever the dangerous cluster carries compact symmetry, the shadow complex structure is no longer treated as a convenient auxiliary choice on a finite-dimensional block. Instead, following the canonical-complexification and compact nonabelian spectral-orientation papers, it is treated intrinsically as a point of the parameter space $X(U_n)$ of skew-adjoint square roots of -1 in the Real commutant $\mathcal{A}_R(U_n)$ of the cluster representation [13, 14]. This imports three tools directly into the Yang–Mills synthesis: the Real commutant as the universal invariant controlling symmetry-compatible complex structures, operator-valued orientation on complex-type multiplicity spaces, and the Frobenius–Schur real/quaternionic decomposition on self-conjugate blocks.

Second, the analytic transport layer is now written in producer–consumer form. The fixed-regulator projected-gap section remains the place where the low-window obstruction is isolated, but the consequences of the desired comparison estimates and convergence packages are no longer left as diffuse background assumptions. Instead, high-window positivity from a Hamiltonian–mixed-operator coercive inequality, thermodynamic persistence of projected lower bounds, orthogonalized recovery sequences against the vacuum, gap-preserving continuum transfer, and norm-resolvent stability of finite-rank dangerous-cluster projections on a common comparison space are recorded here as abstract theorem-level exports of named theorem targets in the concrete projected lattice-gauge realization [7, 8]. This separates the genuinely model-specific tasks from the functional-analytic transport machinery and makes the analytically tractable frontier visible inside the paper itself.

Two structural decisions therefore govern the rewrite. First, the SMM/DTTbackground is recorded explicitly rather than hidden behind motivational remarks. Second, the proof spine remains lattice, spectral, and operator-theoretic: accepted temporal ontology is not allowed to substitute

for constructive mathematical statements. The point of the rewrite is therefore not to weaken mathematical discipline, but to place the same mathematical chain inside the stronger foundational reading for which it was originally designed. A third, purely mathematical, decision is that both the finite-dimensional shadow sector and the coercivity/transport layer are now written in intrinsic language—Real-commutant theory on the dangerous cluster, and producer–consumer interface theorems on the transport side—rather than in the looser bookkeeping language of unspecified finite-dimensional choices and unnamed persistence packages.

1.2 Standing SMM/DTT framework

The accepted SMM/DTTbackground used throughout the paper is summarized in the following standing framework statement.

Assumption 1.1 (Standing SMM/DTT foundation). Throughout the paper we treat the following background principles as established input from the accepted SMM/DTTframework.

- (i) Temporal process, not pre-given spatial extension, is the primitive generative structure of the effective physical theory.
- (ii) Time has a dual organization: a compact inner-time direction governing cyclic renewal of microscopic configurations, and an extended outer-time direction governing effective macroscopic evolution.
- (iii) The finite auxiliary lattice cycle $\tau \in \mathbb{Z}/N_\tau\mathbb{Z}$ is the regulator-level realization of compact inner time, while the physical Euclidean-time coordinate entering reflection positivity and transfer-matrix reconstruction is the regulator-level representative of outer time.
- (iv) Nontrivial dependence on inner time carries genuine dynamical cost. In the lattice model this cost is measured by the mixed physical/inner-time plaquette sector and its associated covariant inner-time operator.
- (v) The projected four-dimensional sector is the effective outer-time shadow of the dual-time theory; ordinary Yang–Mills observables are recovered precisely when the explicit inner-time microstructure becomes invisible on the relevant Wilson-loop/curvature core.

The present manuscript does not rederive [theorem 1.1](#). Instead it uses that accepted framework to determine what the auxiliary cycle means, why the mixed plaquette term should be interpreted as an inner-time renewal cost, and why the projected four-dimensional sector is the mathematically relevant shadow seen by the effective Yang–Mills observer.

1.3 Translation into projected lattice gauge language

Under [theorem 1.1](#), the present lattice model should be read through the following translation dictionary.

- (i) The auxiliary lattice cycle $\tau \in \mathbb{Z}/N_\tau\mathbb{Z}$ is the finite-regulator avatar of compact inner time.
- (ii) The mixed physical/inner-time plaquette form is the exact gauge-covariant measure of failure of inner-time renewal to be covariantly constant.

- (iii) The dangerous cluster is the finite-dimensional low-window sector singled out from the inner-time operator after projection and Feshbach reduction; physically, it is the finite remnant of the nontrivial inner-time spectrum that survives in the effective four-dimensional shadow.
- (iv) The oriented-shadow / Real-commutant machinery is the compact, finite-dimensional specialization of the canonical-complexification program: the shadow complex structure is a point of the Real-commutant parameter space, and the shadow blocks are its operator-valued orientation and Frobenius–Schur components rather than an external add-on.
- (v) The producer–consumer coercivity/transport layer is the precise functional-analytic statement that once the projected-sector, coercivity, embedding, continuum-trajectory, and common-comparison-space theorem targets are discharged in the concrete lattice-gauge model, the same low-window mechanism persists through large volume and continuum passage and arrives at the identification stage as a preserved lower bound together with an eventually constant finite-rank dangerous-cluster family.
- (vi) Auxiliary invisibility is the mathematical statement that the accepted inner-time microstructure leaves no extra observable residue once the theory is expressed on the countable Wilson-loop/curvature core and compared with ordinary Yang–Mills data.

The notation DTTtherefore no longer functions merely as a historical label. It names the accepted dual-time framework whose compact inner-time / extended outer-time architecture is represented at regulator level by the lattice model analyzed below.

1.4 Spectral-topological bridge and temporal decimation heuristic

Within the accepted SMM/DTTframework, compact inner time predicts a discrete mode tower, and the shorter companion notes isolated two mathematical consequences of that picture [17, 18]. First, compactification of inner time should produce a spectral separation between the covariantly constant inner-time sector and the higher twisted modes. Second, decomposition along the compact cycle suggests an intrinsic temporal decimation procedure in which successive packets of inner-time modes are integrated out or organized into effective band data.

The present manuscript absorbs those ideas in a disciplined way. At fixed regulator, the auxiliary cycle is not metaphorical but an explicit lattice direction, and the mixed plaquettes generate the positive operator K_{τ,a,L,N_τ} whose twisted Fourier fibers are computed exactly. Later, the dangerous-cluster and finite-band layers realize the temporal-decimation idea through packetwise phase counting, spectral cutoffs, finite-band covers, and explicit off-block estimates. At the representation-theoretic stage, the same compactification principle reappears as conjugation-paired low spectral data on the dangerous cluster, whose atomic organization is handled by the Real commutant, operator-valued orientation, and Frobenius–Schur decomposition imported from [13, 14]. At the transport stage, the same low-window geometry is carried through large volume and continuum passage by explicit coercivity, Mosco-type convergence, and norm-resolvent stability packages [7, 8]. In that sense the paper may be read as the gauge-covariant operator-theoretic realization of the inner-time spectral picture already suggested by the accepted SMM/DTTbackground.

Two boundaries remain important even in this stronger foundational stance. First, no theorem below deduces the Yang–Mills gap from a bare formula for the first inner-time mode alone; the gap

still enters through the coercivity, Feshbach, localization, continuum-transfer, and interface packages displayed explicitly later on. Second, the reconstruction of the Hamiltonian remains anchored in reflection positivity for the physical Euclidean-time direction. The compact inner-time cycle governs spectral organization, but the Osterwalder–Schrader machinery still runs through the outer-time reflection framework.

1.5 Claim-status discipline and current frontier

The present rewrite keeps a strict logical status ladder.

- *Standing framework statement* refers to accepted background structure imported from SMM/DTT and used here as foundational input.
- *Theorem*, *proposition*, *lemma*, and *corollary* are statements proved in the present manuscript from named earlier sections and explicit assumptions.
- *Conditional theorem*, *conditional proposition*, and *conditional corollary* are continuation statements derived from explicit standing hypotheses displayed in the statement itself.
- *Theorem target*, *proposition target*, and *corollary target* are exact unresolved local tasks recorded in the notation later consumed; they are never cited as though already proved.
- *Interface theorem* is a producer–consumer statement: once the named theorem targets are discharged, the corresponding downstream imports become unconditional.
- *Assumption* packages a theorem target or imported companion result needed to move the mathematical argument forward.
- *Interpretation* records the physical or ontological reading of a proved mathematical fact.

This discipline matters even more in the present rewrite. Accepting the SMM/DTT framework does not erase the need to isolate the analytic frontier. What remains genuinely package-dependent is still visible rather than hidden: the projected-sector/common-core and coercivity inputs of Section 3, the thermodynamic embeddings, scale-uniform trajectory, common-comparison-space band realization, and isolated-contour stability inputs of Section 4, and the explicit pair-isolation/model-tail and admissible-precursor estimates that trigger the late-stage localization and curvature machinery of Sections 8 and 9.

1.6 Three-part organization

The manuscript is divided as follows.

- Part I contains Sections 2–5. It develops the fixed-regulator constructive core, the projected low-window dangerous-cluster mechanism, the producer–consumer coercivity/transport layer, and the continuum/interface reduction to an explicit countable Wilson-loop/curvature core.
- Part II contains Sections 6–9. It develops the Real-commutant and oriented-shadow formalism on the dangerous cluster, records the one-band and two-band closed sectors, and formulates the localization and bounded-curvature approximation packages needed for the full observable synthesis.

- (iii) Part III contains Section 10 and the appendices. It develops the end-stage gluing, mixed-word extension, vacuum-theory identification, and gap-transfer mechanism in a self-contained conditional form.

1.7 Global strategy

The logical chain is the following.

- (i) At fixed regulator the lattice Boltzmann weight gives a normalized gauge-invariant Gibbs measure; reflection positivity yields an Osterwalder–Schrader Hilbert space, transfer matrix, and Hamiltonian.
- (ii) The mixed physical/inner-time plaquettes define an exact positive covariant inner-time operator K ; after inner-time axial gauge it becomes a holonomy-twisted finite-difference operator on the compact inner-time cycle, and its twisted spectrum identifies the genuine low-energy obstruction. This is the regulator-level spectral core of the accepted dual-time picture.
- (iii) A projected four-dimensional sector is isolated. High covariant energies are controlled by a coercive estimate relative to K , while the low window is reduced to a finite-rank dangerous cluster through packetwise phase counting, band-effective comparison, and spectral cutoff.
- (iv) The finite-volume projected gap is then reduced to explicit low-window data: a strictly positive high-energy block, a finite-rank dangerous cluster with positive remainder, and a finite-dimensional model-plus-error positivity test. The one-band and two-band criteria are the multiplicity-one and first operator-valued specializations of that general finite-dimensional step.
- (v) The resulting projected lower bound is transported through large volume and continuum passage by a producer–consumer package: thermodynamic persistence under Mosco-form convergence, vacuum-compatible recovery sequences, and stable transport of dangerous clusters on a common comparison space with norm-resolvent convergence and an isolated contour.
- (vi) A countable Wilson-loop/curvature core is dense in the continuum Euclidean and operator-algebraic local structures. Equality on that core therefore suffices for full vacuum-theory identification.
- (vii) Inner-time invisibility is reformulated intrinsically on the dangerous cluster through the Real commutant $\mathcal{A}_R(U_n)$ of the compact symmetry representation, its parameter space $X(U_n)$ of skew-adjoint square roots of $-\mathbf{1}$, operator-valued orientations on complex-type blocks, Frobenius–Schur real/quaternionic factors on self-conjugate blocks, block symbols, and vanishing J -antilinear leakage.
- (viii) Section 8 upgrades those closed sectors to arbitrary finite Wilson-loop tuples by working directly with the band-effective spectrum: spectral atoms, tuple-adapted resonance graphs, pair-isolation inequalities, and model-tail estimates produce a finite-band cover, explicit off-block decay, and reduction to finitely many already-closed one-band and two-band sectors.

- (ix) Section 9 then extends the Wilson-loop sector to bounded curvature generators by constructive small-loop reconstruction, explicit bounded functional calculus, and band-adapted Wilson-loop approximation. Section 10 glues the finite-band Wilson-loop sectors, transfers the approximation to mixed words, upgrades full-core matching to vacuum-theory identification, and transfers the projected continuum spectral gap by unitary equivalence.

1.8 Relation to earlier drafts and sources

The present three-part manuscript consolidates earlier modular Yang–Mills papers, but it now does so inside a sharper producer–consumer architecture. The fixed-regulator constructive core remains the reflection-positivity and exact mixed-operator package of [5]. The fixed-scale projected gap and band-model reduction are supplied by the finite-volume gap paper together with its extracted projected-gap-to-continuum variant [6, 9]. The coercivity / thermodynamic / continuum transport layer is sharpened by the producer–consumer papers on high-energy coercivity, thermodynamic persistence, continuum transfer, and stable transport of dangerous clusters [7, 8]. The countable-core, density, extension, and reconstruction theorems reducing Yang–Mills identification to auxiliary invisibility on a countable core are the subject of [10]. The canonical complexification paper and the compact nonabelian spectral-orientation paper supply the intrinsic representation-theoretic toolkit used here on the dangerous cluster: the Real commutant as the universal invariant, operator-valued orientation on complex-type multiplicity blocks, the Frobenius–Schur real/quaternionic decomposition on self-conjugate blocks, and the multiplicity-one reduction underlying the one-band closure [13, 14]. The oriented-shadow/closure and final synthesis layers remain [11, 12]. Finally, the accepted SMM/DTTbackground supplies the foundational temporal ontology in which compact inner time, effective outer time, and the emergence of a four-dimensional shadow sector are no longer merely motivational language [19, 15, 16].

The earlier drafts therefore remain useful in four different ways. The modular Yang–Mills papers are part of the proof-production history for the present synthesis. The broader SMM/DTTpapers are part of the accepted foundational setting that explains why the projected lattice gauge program takes the form it does. The canonical-complexification and compact-orientation papers provide the intrinsic finite-dimensional language in which the dangerous-cluster shadow sector is now written. The new coercivity/transport papers provide the exact theorem targets and interface statements by which the tractable analytic frontier is exposed rather than hidden. The present text is nevertheless intended to be readable on its own.

1.9 Main endpoint theorem

The next theorem should be read as the master dependency theorem for the comprehensive SMM/DTT-based version.

Theorem 1.2 (Integrated endpoint / master dependency theorem). *Assume the standing SMM/DTTframework of theorem 1.1 together with the standing hypotheses introduced and then used in Sections 3–10, in particular the projected-sector/common-core and coercivity hypotheses of Section 3; the thermodynamic embedding, scale-uniform trajectory, common-comparison-space band realization, and isolated-contour stability hypotheses of Section 4; the observable continuity package of Section 5; the compact-symmetry and oriented-shadow hypotheses governing the closed one-band and two-band*

sectors in Sections 6 and 7; the finite-band localization package of Section 8; and the bounded-curvature approximation package of Section 9. Then the projected continuum theory and standard pure Yang–Mills theory are unitarily equivalent in their vacuum representations. If moreover the projected continuum lower bound of Section 4

$$H_{\text{cont}}^{\text{DTT}} \geq \delta \quad \text{on} \quad \mathcal{H}_{\text{cont}}^{\text{DTT}} \ominus \mathbb{C}\Omega_{\text{cont}}^{\text{DTT}}$$

holds for some $\delta > 0$, then

$$H_{\text{YM}} \geq \delta \quad \text{on} \quad \mathcal{H}_{\text{YM}} \ominus \mathbb{C}\Omega_{\text{YM}}.$$

In particular, within the bounded-core framework adopted throughout the paper, no additional end-stage matching principle remains once the localization, approximation, reconstruction, and transport packages displayed explicitly in Sections 8–10 are verified.

Proof. The proof is distributed across the present manuscript. Sections 2 and 3 construct the fixed-regulator model, the exact mixed operator, the low-window decomposition, and the projected dangerous-cluster gap mechanism. Section 4 transports not only the projected lower bound but also the dangerous-cluster package through the thermodynamic and continuum interfaces. Section 5 reduces identification to equality on an explicit countable core. Sections 6 and 7 localize inner-time invisibility and Yang–Mills matching on the dangerous cluster and verify the one-band and two-band closed sectors. Sections 8 and 9 then provide the localization and approximation steps required to pass from those closed sectors to arbitrary finite Wilson-loop/curvature words. Section 10 combines those ingredients with the Euclidean and operator-algebraic interface theorems to identify the full vacuum theories and transfer the spectral lower bound. The logical dependencies needed for the theorem therefore appear explicitly in the present text. $\square$

1.10 Reader’s guide

Sections 2 and 3 develop the fixed-regulator and projected-gap parts of the argument, now with explicit theorem targets for projected-sector/common-core invariance, coercivity, a concrete coercivity program with a first restricted local target, and finite-dimensional model-plus-error closure. Section 4 isolates the large-volume and continuum transport layer in producer–consumer form, including stable transport of dangerous clusters. Section 5 develops the countable-core, density, extension, and reconstruction infrastructure that reduces Yang–Mills identification to auxiliary invisibility on the explicit core. Section 6 develops the Real-commutant and oriented-shadow machinery on the finite-dimensional dangerous cluster. Section 7 then gives the first actual closures: one-band and two-band Wilson-loop sectors are explicitly closed and matched with the Yang–Mills target on those subcores. Section 8 derives finite-band exhaustion and off-block decay directly from the band-effective spectrum. Section 9 packages bounded curvature generators by constructive band-adapted Wilson-loop approximation. Section 10 glues those pieces into the final full-core identification and gap-transfer chain. The appendices record fixed-regulator auxiliary material and a compact audit checklist for the end-stage localization/approximation chain.

2 Fixed-regulator constructive core

This section develops the fixed-regulator constructive backbone of the theory inside the accepted SMM/DTTframework. The auxiliary cycle is treated here as the explicit compact lattice realization of inner time, and its mixed plaquettes generate an exact positive operator encoding the cost of nontrivial covariant renewal along that cycle. We define the finite-regulator lattice model, prove reflection positivity in physical Euclidean time, reconstruct the Osterwalder–Schrader Hilbert space, identify the transfer matrix and Hamiltonian, and compute the exact holonomy-twisted covariant inner-time operator generated by the mixed physical/inner-time plaquettes.

2.1 Finite-regulator lattice geometry and configuration space

Fix a compact gauge group G and a faithful finite-dimensional unitary representation $\rho: G \rightarrow U(m)$. Let $a > 0$ be the lattice spacing, let L and T be the physical spatial and Euclidean-time regulators, and let $N_\tau \in \mathbb{N}$ be the length of the auxiliary inner-time cycle. The physical base lattice is

$$\Lambda_{a,L,T} := (\mathbb{Z}/L\mathbb{Z})^3 \times (\mathbb{Z}/T\mathbb{Z}),$$

with coordinates $y = (x, t)$, $x \in (\mathbb{Z}/L\mathbb{Z})^3$, $t \in \mathbb{Z}/T\mathbb{Z}$. The auxiliary circle is $\mathbb{Z}/N_\tau\mathbb{Z}$, with coordinate τ .

The auxiliary cycle is an explicit finite-regulator ingredient of the model. Its role in the analysis is operator-theoretic: the later covariant operator K_{τ,a,L,N_τ} is a gauge-covariant finite-difference operator along this cycle, and the dangerous-cluster reduction isolates the finite-dimensional low covariant-energy sector relevant to the projected gap problem.

For $\mu \in \{0, 1, 2, 3\}$ let $U_\mu(y, \tau) \in G$ be the gauge variable on the physical link from (y, τ) to $(y + \hat{\mu}, \tau)$, and let $V(y, \tau) \in G$ be the auxiliary link variable from (y, τ) to $(y, \tau + 1)$. The full finite-regulator configuration space $\mathcal{U}_{a,L,T,N_\tau}$ is the product of all such link copies of G , hence compact. We denote by $d\mu_H$ the corresponding normalized Haar product measure.

Physical plaquettes are the usual Wilson loops inside the four physical directions. Mixed plaquettes involve one physical direction $\mu \in \{0, 1, 2, 3\}$ and the auxiliary cycle. Their holonomy is denoted by $W_{\mu\tau}(y, \tau)$. The DTT action is

$$S_{a,L,T,N_\tau}^{\text{DTT}}(U, V) = \kappa(a) \sum_{p \in \Lambda_{a,L,T} \times \mathbb{Z}/N_\tau\mathbb{Z}} \Re \text{Tr}(I - \rho(W_p)), \quad (2.1)$$

where the sum runs over physical and mixed plaquettes, and $\kappa(a) > 0$ is the scale-dependent coupling.

Lemma 2.1 (Gauge invariance and positivity of the action). *The action (2.1) is continuous, gauge invariant, and bounded below by 0.*

Proof. Continuity is immediate from continuity of group multiplication, inversion, and matrix trace in the finite-dimensional representation ρ . If $M \in U(m)$ has eigenvalues $e^{i\theta_1}, \dots, e^{i\theta_m}$, then

$$\Re \text{Tr}(I - M) = \sum_{j=1}^m (1 - \cos \theta_j) \geq 0.$$

Applying this to each plaquette term yields nonnegativity. Gauge invariance follows from conjugation invariance of the trace. $\square$

Theorem 2.2 (Finite-regulator Gibbs measure). *For every finite regulator (a, L, T, N_τ) the formula*

$$d\mu_{a,L,T,N_\tau}^{\text{DTT}}(U, V) := Z_{a,L,T,N_\tau}^{-1} e^{-S_{a,L,T,N_\tau}^{\text{DTT}}(U,V)} d\mu_H(U, V) \quad (2.2)$$

defines a normalized gauge-invariant probability measure on $\mathcal{U}_{a,L,T,N_\tau}$.

Proof. Because the action is continuous on a compact space, the Boltzmann weight is bounded and integrable. The partition function is finite and strictly positive, and the rest follows from [theorem 2.1](#) and Haar invariance. $\square$

2.2 Reflection positivity and Osterwalder–Schrader reconstruction

Let Θ be reflection in the physical Euclidean-time coordinate. Let $\mathcal{A}$ be the $*$ -algebra of bounded cylinder functions on $\mathcal{U}_{a,L,T,N_\tau}$ and $\mathcal{A}_+$ the positive-time subalgebra. Reflection positivity is proved by factorizing the boundary plaquette contribution into a finite sum of reflected squares; see [Appendix A](#).

Theorem 2.3 (Reflection positivity). *For every $F \in \mathcal{A}_+$,*

$$\int_{\mathcal{U}_{a,L,T,N_\tau}} (\Theta F) F d\mu_{a,L,T,N_\tau}^{\text{DTT}} \geq 0. \quad (2.3)$$

Equivalently, the finite-regulator DTT measure is reflection positive with respect to physical Euclidean-time reflection.

Proof. Using the decomposition of the action into positive-time, negative-time, and boundary contributions, the boundary factorization shows that the integrand becomes a finite sum of squares of positive-time functions. Each summand is nonnegative after Haar integration. $\square$

Define the reflection-positive form on $\mathcal{A}_+$ by

$$\langle F, G \rangle_{\text{OS}} := \int (\Theta F) G d\mu_{a,L,T,N_\tau}^{\text{DTT}}.$$

Let $\mathcal{N}_{\text{OS}} = \{F \in \mathcal{A}_+ : \langle F, F \rangle_{\text{OS}} = 0\}$, and set

$$\mathcal{H}_{a,L,N_\tau}^{\text{OS}} := \overline{\mathcal{A}_+ / \mathcal{N}_{\text{OS}}}^{\|\cdot\|_{\text{OS}}}.$$

The class of the constant function 1 is denoted by Ω . The residual gauge group acting on a fixed physical-time slice acts unitarily on $\mathcal{H}_{a,L,N_\tau}^{\text{OS}}$; the gauge-invariant physical Hilbert space $\mathcal{H}_{a,L,N_\tau}^{\text{phys}}$ is the fixed-point subspace.

2.3 Transfer matrix and Hamiltonian

Fix one outer-time slab. Let Σ_{a,L,N_τ} be the slice configuration space consisting of spatial links and auxiliary links on the slice $t = 0$, and let η denote the temporal links between slices $t = 0$ and $t = 1$.

Writing the action as a sum of horizontal slice terms $h(\sigma)$ and one-slab coupling terms $v(\sigma, \eta, \sigma')$, define the one-slab kernel

$$K_{a,L,N_\tau}(\sigma, \sigma') := \int e^{-v(\sigma, \eta, \sigma')} d\mu_H(\eta), \quad (2.4)$$

with half-horizontal weight $q_{a,L,N_\tau}(\sigma) := e^{-h(\sigma)/2}$. The transfer operator on $L^2(\Sigma_{a,L,N_\tau}, d\mu_\Sigma)$ is

$$(T_{a,L,N_\tau} f)(\sigma) := \int q_{a,L,N_\tau}(\sigma) K_{a,L,N_\tau}(\sigma, \sigma') q_{a,L,N_\tau}(\sigma') f(\sigma') d\mu_\Sigma(\sigma'). \quad (2.5)$$

Theorem 2.4 (Transfer matrix). *The operator T_{a,L,N_τ} defined by (2.5) is bounded, compact, positive, and self-adjoint. It commutes with the residual slice gauge action and therefore preserves $\mathcal{H}_{a,L,N_\tau}^{\text{phys}}$. After normalization by its top eigenvalue, it satisfies $\|T_{a,L,N_\tau}\| = 1$ and $T_{a,L,N_\tau}\Omega = \Omega$. The finite-regulator Hamiltonian is the nonnegative self-adjoint operator*

$$H_{a,L,N_\tau} := -\log T_{a,L,N_\tau}$$

on the gauge-invariant physical Hilbert space.

Proof. Appendix A records the kernel representation, positivity, symmetry, and gauge covariance. Compactness comes from continuity of the kernel on a compact slice space, and nonnegativity of the Hamiltonian follows from the normalization $0 \leq T \leq I$. $\square$

2.4 Exact covariant inner-time operator

For a physical link $e = (y, \mu)$ define the matrix-valued field

$$A_{y,\mu}(\tau) := \rho(U_\mu(y, \tau)),$$

and the covariant inner-time difference

$$(D_{\tau,V}^{\text{cov}} A_{y,\mu})(\tau) := \frac{\rho(V(y, \tau)) A_{y,\mu}(\tau+1) \rho(V(y + \hat{\mu}, \tau))^{-1} - A_{y,\mu}(\tau)}{\delta_\tau}, \quad (2.6)$$

where $\delta_\tau := 1/N_\tau$ and the index $\tau+1$ is taken modulo N_τ .

Theorem 2.5 (Exact mixed-plaquette quadratic identity). *For every outer-space-time point y , physical direction $\mu \in \{0, 1, 2, 3\}$, and inner-time index τ ,*

$$\Re \text{Tr}(I - \rho(W_{\mu\tau}(y, \tau))) = \frac{\delta_\tau^2}{2} \left\| (D_{\tau,V}^{\text{cov}} A_{y,\mu})(\tau) \right\|_{\text{HS}}^2. \quad (2.7)$$

Consequently,

$$\kappa(a) \sum_{y,\mu,\tau} \Re \text{Tr}(I - \rho(W_{\mu\tau}(y, \tau))) = \frac{\kappa(a) \delta_\tau^2}{2} \sum_{y,\mu,\tau} \left\| (D_{\tau,V}^{\text{cov}} A_{y,\mu})(\tau) \right\|_{\text{HS}}^2. \quad (2.8)$$

Proof. Set

$$M := \rho(V(y, \tau)) A_{y,\mu}(\tau+1) \rho(V(y + \hat{\mu}, \tau))^{-1} A_{y,\mu}(\tau)^{-1}.$$

Then $M = \rho(W_{\mu\tau}(y, \tau))$ and M is unitary. Hence

$$\|I - M\|_{\text{HS}}^2 = \text{Tr}((I - M)^*(I - M)) = 2\Re \text{Tr}(I - M).$$

Right multiplication by the unitary $A_{y,\mu}(\tau)^{-1}$ preserves the Hilbert–Schmidt norm, and (2.7) follows. Summing over (y, μ, τ) gives (2.8). $\square$

For each fixed y define the ordered inner-time holonomy

$$H_y := V(y, 0)V(y, 1) \cdots V(y, N_\tau - 1) \in G. \quad (2.9)$$

A gauge transform along the auxiliary cycle trivializes all auxiliary links except the closing link, which becomes H_y . In this axial gauge the covariant difference becomes a twisted discrete difference with unitary twist

$$H_e(B) := \rho(H_y)B\rho(H_{y+\hat{\mu}})^{-1}, \quad B \in \text{End}(\mathbb{C}^m). \quad (2.10)$$

On $\ell^2(\mathbb{Z}/N_\tau\mathbb{Z}; \text{End}(\mathbb{C}^m))$ the corresponding difference operator is

$$(D_{\tau, H_e}A)(\tau) = \frac{A(\tau+1) - A(\tau)}{\delta_\tau}, \quad 0 \leq \tau \leq N_\tau - 2, \quad (2.11)$$

$$(D_{\tau, H_e}A)(N_\tau - 1) = \frac{H_e A(0) - A(N_\tau - 1)}{\delta_\tau}. \quad (2.12)$$

Set $\Delta_{\tau, H_e} := D_{\tau, H_e}^* D_{\tau, H_e}$.

Proposition 2.6 (Fiberwise twisted spectrum). *Let $e = (y, \mu)$ be a physical link, and let $e^{i\theta_1}, \dots, e^{i\theta_r}$, $r = m^2$, be the eigenvalues of the unitary twist H_e on $\text{End}(\mathbb{C}^m)$, counted with multiplicity. Then*

$$\text{spec}(\Delta_{\tau, H_e}) = \left\{ \frac{4}{\delta_\tau^2} \sin^2\left(\frac{2\pi n + \theta_j}{2N_\tau}\right) : n = 0, \dots, N_\tau - 1, j = 1, \dots, r \right\}. \quad (2.13)$$

The kernel consists exactly of covariantly constant sequences in the H_e -invariant directions.

Proof. Choose an orthonormal eigenbasis of H_e . On each eigenspace the closing condition becomes a scalar twisted periodicity condition, and the twisted Fourier modes diagonalize the difference operator. The eigenvalue formula is the usual discrete Fourier calculation with phase shift θ_j . $\square$

Definition 2.7 (Covariant inner-time operator and covariant zero sector). Let K_{τ, a, L, N_τ} be the positive operator whose quadratic form is the mixed part of the action in (2.8). Its exact zero sector is

$$Z_{a, L, N_\tau}^{\text{cov}} := \ker K_{\tau, a, L, N_\tau}.$$

Theorem 2.8 (Twisted covariant inner-time spectrum package). *The mixed physical/auxiliary plaquette contribution to the DTT action is exactly the quadratic form of the positive operator K_{τ, a, L, N_τ} . After inner-time axial gauge, K_{τ, a, L, N_τ} decomposes fiberwise into holonomy-twisted discrete Laplacians of the form Δ_{τ, H_e} . Consequently, the exact zero sector is the covariant kernel determined by the holonomy twists rather than the untwisted Fourier-zero sector.*

Proof. This is the combination of theorems 2.5 and 2.6 and the direct-sum assembly over physical links. $\square$

Theorem 2.9 (Auxiliary-cycle interpretation of the low-window operator). *After inner-time axial gauge, the fiber operator D_{τ, H_e} is a holonomy-twisted finite-difference operator on a discrete circle. Accordingly, K_{τ, a, L, N_τ} is the gauge-covariant positive operator associated with that twisted difference, with the untwisted zero mode replaced by the covariant kernel selected by the holonomy twists. The low-window projection $P_\varepsilon = \mathbf{1}_{[0, \varepsilon]}(K)$ therefore isolates the corresponding low covariant-energy sector.*

Proof. Ignoring the gauge twist, Equations (2.11) and (2.12) is the standard finite-difference model on a circle of N_τ sites. The axial-gauge holonomy H_e changes only the closing condition and therefore produces the twisted covariant variant relevant to the auxiliary cycle. By theorem 2.5, the mixed plaquette form is exactly the quadratic form of $D_{\tau, H_e}^* D_{\tau, H_e}$, and theorem 2.6 identifies the resulting spectrum and the corresponding covariant kernel. $\square$

Remark 2.10 (Spectral-topological interpretation of the twisted fibers). In the untwisted abelian toy case the eigenvalues (2.13) reduce to the familiar discrete Fourier tower

$$\frac{4}{\delta_\tau^2} \sin^2\left(\frac{\pi n}{N_\tau}\right), \quad n = 0, \dots, N_\tau - 1.$$

This is the rigorous lattice avatar of the spectral-topological mechanism emphasized in [17]: compactification of the auxiliary cycle creates a discrete mode structure. In the nonabelian model, however, the relevant low-energy object is not the raw $n = 1$ mode but the holonomy-twisted covariant low window selected by $P_\varepsilon = \mathbf{1}_{[0, \varepsilon]}(K)$. Near-trivial holonomy can push positive twisted modes arbitrarily low, which is why the later gap mechanism proceeds through packetwise counting, Feshbach reduction, and dangerous-cluster analysis rather than through an immediate bare-mode estimate.

3 Projected low-window reduction and the finite-volume gap

This section develops the decisive finite-volume projected gap mechanism. The strategy is unchanged at the highest level: control the high-covariant-energy region by a coercive form estimate, reduce the low window by a shifted Feshbach map, isolate a finite-rank dangerous cluster through a compact-resolvent band model, and solve the remaining positivity problem on that finite-dimensional sector. What is new is the presentation. The section is now written in producer–consumer form. Two explicit theorem targets record the model-specific projected-sector/common-core and coercivity inputs; the next subsection expands them into a concrete coercivity program and a first restricted local target; the later theorems package their abstract consequences; and the finite-regulator gap mechanism is reduced to a finite-dimensional model-plus-error test whose rank-one and rank-two cases are the first closed sectors used later in the synthesis.

3.1 Projected four-dimensional sector and coercivity targets

The fixed-regulator theory of Section 2 supplies a nonnegative Hamiltonian on the gauge-invariant Hilbert space. The gap mechanism is formulated not on the full space but on a projected four-dimensional sector, denoted by $\mathcal{H}_{a, L, N_\tau}^{4D} \subset \mathcal{H}_{a, L, N_\tau}^{\text{phys}}$. The precise projection is part of the program input; at the level of the present paper we require only that the sector be compatible with the

dynamics and with the covariant inner-time operator. Accordingly we write

$$H := H_{a,L,N_\tau}|_{\mathcal{H}_{a,L,N_\tau}^{4D}}, \quad K := K_{\tau,a,L,N_\tau}|_{\mathcal{H}_{a,L,N_\tau}^{4D}},$$

and denote the vacuum vector in the projected sector by Ω . The excited sector is

$$\mathcal{H}_{\text{exc}} := \mathcal{H}_{a,L,N_\tau}^{4D} \ominus \mathbb{C}\Omega.$$

For $\varepsilon > 0$ let

$$P_\varepsilon := \mathbf{1}_{[0,\varepsilon]}(K), \quad Q_\varepsilon := I - P_\varepsilon.$$

The point of the low-window split is that the twisted spectrum of [theorem 2.8](#) allows arbitrarily small positive eigenvalues in the presence of near-trivial holonomy, so coercivity can only be expected on the high-energy region $Q_\varepsilon \mathcal{H}_{\text{exc}}$.

Theorem target 3.1 (Projected-sector invariance and common form core). *For each fixed regulator (a, L, N_τ) , prove that $\mathcal{H}_{a,L,N_\tau}^{4D}$ is invariant under H and under the spectral projections of K . Moreover, construct a dense form core $\mathcal{C} \subset \mathcal{H}_{\text{exc}}$ stable under P_ε and Q_ε and common to the quadratic forms of H and K .*

Theorem target 3.2 (Volume-uniform high-energy coercivity). *Establish constants $\alpha > 0$ and $\beta \geq 0$, independent of L in the large-volume regime, such that on the core from [theorem 3.1](#),*

$$q_H[\psi] \geq \alpha q_K[\psi] - \beta \|\psi\|^2. \quad (3.1)$$

Assumption 3.3 (Projected-sector invariance and common form core). The conclusion of [theorem 3.1](#) holds.

Assumption 3.4 (High-energy coercivity). The conclusion of [theorem 3.2](#) holds.

Theorem 3.5 (High-window positivity from coercivity). *Assume [theorems 3.3](#) and [3.4](#). Then for every $\varepsilon > 0$,*

$$H \geq \alpha\varepsilon - \beta \quad \text{on } Q_\varepsilon \mathcal{H}_{\text{exc}}. \quad (3.2)$$

In particular, if $0 < \delta < \alpha\varepsilon - \beta$, then $H - \delta$ is strictly positive on $Q_\varepsilon \mathcal{H}_{\text{exc}}$.

Proof. For $\psi \in Q_\varepsilon \mathcal{C}$, the spectral theorem gives $q_K[\psi] \geq \varepsilon \|\psi\|^2$. Inserting this into [\(3.1\)](#) gives

$$q_H[\psi] \geq (\alpha\varepsilon - \beta) \|\psi\|^2.$$

Density of $Q_\varepsilon \mathcal{C}$ in the form domain of $Q_\varepsilon H Q_\varepsilon$ yields the operator inequality [\(3.2\)](#). $\square$

Corollary 3.6 (Admissible low-window regime). *Any candidate projected-gap threshold δ must be paired with a low-window threshold ε satisfying*

$$\varepsilon > \beta/\alpha, \quad 0 < \delta < \alpha\varepsilon - \beta.$$

3.2 Concrete coercivity program on the projected core

The present subsection expands Theorem Targets [3.1–3.2](#) into a concrete proof program. Its role is to isolate the remaining model-specific analytic work required to establish high-energy

coercive separation in the projected sector. The producer–consumer coercivity papers show that once this front-end package is discharged, the high-window positivity bound and all downstream thermodynamic/continuum transport results follow abstractly; the real front-end burden is therefore to make the projected core explicit and to compare the mixed-direction part of the Hamiltonian to the exact mixed operator on that core.

Definition 3.7 (Projected Wilson-loop cylinder core). Let $\mathcal{C}_{a,L,N_\tau}^{\text{cyl}}$ denote the algebra of bounded gauge-invariant cylinder functions generated by finite linear combinations of projected Wilson-loop observables and bounded auxiliary-dressed local cylinder functions at fixed regulator (a, L, N_τ) . Let

$$\dot{\mathcal{C}}_{a,L,N_\tau}^{\text{cyl}} := \{F - \langle \Omega, F \Omega \rangle \mathbf{1} : F \in \mathcal{C}_{a,L,N_\tau}^{\text{cyl}}\}$$

be the vacuum-subtracted cylinder core, and let

$$D_{a,L,N_\tau}^{\text{cyl}} := \overline{\dot{\mathcal{C}}_{a,L,N_\tau}^{\text{cyl}}}^{\|\cdot\|_{\text{form}}} \subset \mathcal{H}_{a,L,N_\tau}^{4\text{D}} \ominus \mathbb{C}\Omega$$

denote its closure in the common form topology whenever the latter exists.

Theorem target 3.8 (Concrete projected-sector/common-core target). *For each fixed regulator (a, L, N_τ) , the projected sector $\mathcal{H}_{a,L,N_\tau}^{4\text{D}}$ is invariant under H_{a,L,N_τ} and under the spectral projections of $K_{\tau,a,L,N_\tau} \upharpoonright_{\mathcal{H}_{a,L,N_\tau}^{4\text{D}}}$. Moreover, the vacuum-subtracted projected Wilson-loop cylinder space $\dot{\mathcal{C}}_{a,L,N_\tau}^{\text{cyl}}$ is dense in $\mathcal{H}_{a,L,N_\tau}^{4\text{D}} \ominus \mathbb{C}\Omega$ and is stable under the low-window projections P_ε and Q_ε . Its form closure $D_{a,L,N_\tau}^{\text{cyl}}$ is a common form core for H and K .*

Proof route for Theorem Target 3.8. The proof should be organized into four local steps. First choose the projected four-dimensional sector as the closed span of the OS classes of the projected Wilson-loop/curvature cylinder family and verify that Ω lies in it. Second use the transfer-matrix construction and gauge-invariant cylinder stability to show that T_{a,L,N_τ} preserves the projected sector and hence, by functional calculus, that $H = -\log T$ preserves it as well. Third use the exact mixed-plaquette quadratic identity and compatibility of the projected cylinder family with the auxiliary direction to prove invariance under the spectral calculus of K . Fourth verify density and stability of the vacuum-subtracted projected cylinder family under P_ε and Q_ε , then identify its common form closure.

Proposition target 3.9 (Exact projected mixed-form realization). *On the core $D_{a,L,N_\tau}^{\text{cyl}}$, the projected mixed operator $K := K_{\tau,a,L,N_\tau} \upharpoonright_{\mathcal{H}_{a,L,N_\tau}^{4\text{D}}}$ has quadratic form*

$$q_K[\psi] = \frac{\kappa(a)\delta_\tau^2}{2} \sum_{y,\mu,\tau} \left\| (D_{\tau,V}^{\text{cov}} A_{y,\mu})(\tau) \psi \right\|^2$$

in the sense induced by the exact mixed-plaquette identity after projection to the four-dimensional sector.

Proof route for Proposition Target 3.9. Restrict the mixed-plaquette identity of Section 2 to the projected sector and pass to the OS representation. The only genuine point to verify is

compatibility of the projection with the mixed-direction covariant difference; vacuum subtraction does not alter the quadratic form on the excited sector.

Definition 3.10 (Mixed and remainder Hamiltonian forms). Assume Theorem Target 3.8. A decomposition

$$q_H = q_H^{\text{mix}} + q_H^{\text{rem}} \quad \text{on } D_{a,L,N_\tau}^{\text{cyl}}$$

is called *admissible* if:

- (i) q_H^{mix} is nonnegative and depends only on the projected mixed physical/auxiliary plaquette sector;
- (ii) q_H^{rem} is symmetric and lower semibounded;
- (iii) both forms are gauge invariant and stable on the projected cylinder core.

Theorem target 3.11 (Mixed-direction lower comparison). *There exists an admissible decomposition as in Definition 3.10 and a constant $\alpha_{\text{mix}} > 0$, independent of L in the large-volume regime, such that*

$$q_H^{\text{mix}}[\psi] \geq \alpha_{\text{mix}} q_K[\psi] \quad (\psi \in D_{a,L,N_\tau}^{\text{cyl}}).$$

Proof route for Theorem Target 3.11. The proof should proceed by isolating inside the projected Hamiltonian form the exact contribution generated by mixed physical/auxiliary plaquettes and comparing it termwise with the projected mixed quadratic form of Proposition Target 3.9. Concretely, one seeks a decomposition of the one-slab transfer kernel into a mixed-direction contribution plus a remainder and then uses the transfer-matrix/Hamiltonian reconstruction to show that the mixed part of q_H preserves a positive multiple of the same covariant inner-time renewal cost already measured by K . The difficult model-specific bridge is therefore not the abstract functional analysis but the identification of a transfer-kernel lower estimate of the form

$$q_H^{\text{mix}}[\psi] \geq c_0 \sum_{y,\mu,\tau} \left\| (D_{\tau,V}^{\text{cov}} A_{y,\mu})(\tau) \psi \right\|^2,$$

with $c_0 > 0$ independent of L .

Theorem target 3.12 (Uniform lower bound for the remainder). *There exists $\beta_{\text{rem}} \geq 0$, independent of L in the large-volume regime, such that*

$$q_H^{\text{rem}}[\psi] \geq -\beta_{\text{rem}} \|\psi\|^2 \quad (\psi \in D_{a,L,N_\tau}^{\text{cyl}}).$$

Proof route for Theorem Target 3.12. Bound the non-mixed plaquette, horizontal-slice, and normalization contributions from below on the projected cylinder core. The main point is to prove uniformity in L after projection and vacuum subtraction, so that every term not contributing to inner-time renewal is absorbed into a single scalar defect β_{rem} .

Corollary 3.13 (Concrete coercive comparison on the projected core). *Assume Theorem Target 3.8, Proposition Target 3.9, Theorem Target 3.11, and Theorem Target 3.12. Set*

$$\alpha := \alpha_{\text{mix}}, \quad \beta := \beta_{\text{rem}}.$$

Then

$$q_H[\psi] \geq \alpha q_K[\psi] - \beta \|\psi\|^2 \quad (\psi \in D_{a,L,N_\tau}^{\text{cyl}}),$$

with $\alpha > 0$ and $\beta \geq 0$ independent of L in the large-volume regime. In particular, Theorem Target 3.2 is discharged.

Proof. By Definition 3.10,

$$q_H[\psi] = q_H^{\text{mix}}[\psi] + q_H^{\text{rem}}[\psi].$$

Apply Theorem Target 3.11 to the first term and Theorem Target 3.12 to the second. $\square$

First restricted coercivity target. A realistic first closure of Theorem Target 3.2 is a restricted coercivity theorem on a projected local Wilson-loop cylinder core with bounded auxiliary dressing. In that form the proof burden reduces to a mixed-direction comparison between the transfer-matrix Hamiltonian form and the exact mixed-plaquette quadratic form, together with a uniform lower bound on the complementary remainder. The abstract high-window positivity and transport results of Sections 3–4 would then apply on that restricted projected sector without further functional-analytic input.

Fix a bounded rational region $O \subset \mathbb{R}^4$ and let $\mathcal{W}_{a,L,N_\tau}^{\text{loc}}(O)$ be the finite linear span of projected Wilson-loop cylinder observables supported in O with bounded auxiliary dressing. Let

$$\mathcal{W}_{a,L,N_\tau}^{\text{loc}}(O) := \{F - \langle \Omega, F\Omega \rangle \mathbf{1} : F \in \mathcal{W}_{a,L,N_\tau}^{\text{loc}}(O)\},$$

and let $D_{a,L,N_\tau}^{\text{loc}}(O)$ denote its closure in the common form topology whenever that topology is defined.

Theorem target 3.14 (Restricted coercivity on a projected local Wilson-loop core). *For each fixed regulator (a, L, N_τ) and bounded rational region O , there exist constants $\alpha_O > 0$ and $\beta_O \geq 0$, independent of L in the large-volume regime, such that*

$$q_H[\psi] \geq \alpha_O q_K[\psi] - \beta_O \|\psi\|^2 \quad (\psi \in D_{a,L,N_\tau}^{\text{loc}}(O) \cap (\mathcal{H}_{a,L,N_\tau}^{4D} \ominus \mathbb{C}\Omega)).$$

Reduction of the restricted target. The proof of Theorem Target 3.14 is reduced to the following two local estimates.

Proposition target 3.15 (Local mixed-form identification). *There exists a constant $c_K(a, N_\tau) > 0$ such that on $D_{a,L,N_\tau}^{\text{loc}}(O)$,*

$$q_K[\psi] = c_K(a, N_\tau) \sum_{(y,\mu,\tau) \in \text{Mix}(O)} \left\| (D_{\tau,V}^{\text{cov}} A_{y,\mu})(\tau) \psi \right\|^2,$$

where $\text{Mix}(O)$ denotes the set of mixed physical/auxiliary plaquettes whose physical support lies in O .

Proposition target 3.16 (Local Hamiltonian mixed-direction comparison). *There exists a decomposition*

$$q_H = q_H^{\text{mix},O} + q_H^{\text{rem},O} \quad \text{on } D_{a,L,N_\tau}^{\text{loc}}(O),$$

and constants $\alpha_O > 0$, $\beta_O \geq 0$, independent of L in the large-volume regime, such that

$$q_H^{\text{mix},O}[\psi] \geq \alpha_O q_K[\psi], \quad q_H^{\text{rem},O}[\psi] \geq -\beta_O \|\psi\|^2$$

for all $\psi \in D_{a,L,N_\tau}^{\text{loc}}(O) \cap (\mathcal{H}_{a,L,N_\tau}^{4D} \ominus \mathbb{C}\Omega)$.

Proof route for Proposition Target 3.16. This is the sharpest local bottleneck of the coercivity program. One must decompose the one-slab transfer kernel into a mixed-direction piece and a complementary remainder, then show that after OS reconstruction the mixed part of the Hamiltonian form satisfies a lower estimate of the same type as Proposition Target 3.15. In practice this means proving a local transfer-kernel lower bound generated by mixed physical/auxiliary plaquettes and showing that every other contribution is absorbed into a uniform scalar defect. If this proposition target can be verified on a sufficiently rich local projected core, then the abstract coercivity-to-high-window theorem already proved in Section 3 becomes available on that restricted sector.

Proposition 3.17 (Restricted coercivity from the local comparison lemmas). *Assume Proposition Targets 3.15 and 3.16. Then Theorem Target 3.14 holds.*

Proof. For $\psi \in D_{a,L,N_\tau}^{\text{loc}}(O) \cap (\mathcal{H}_{a,L,N_\tau}^{4D} \ominus \mathbb{C}\Omega)$,

$$q_H[\psi] = q_H^{\text{mix},O}[\psi] + q_H^{\text{rem},O}[\psi] \geq \alpha_O q_K[\psi] - \beta_O \|\psi\|^2.$$

□

Immediate export. Once Theorem Target 3.14 is available on a projected local core, the abstract coercivity-to-high-window theorem applies on that restricted sector. In particular, for every $\varepsilon > \beta_O/\alpha_O$ one obtains

$$H \geq \alpha_O \varepsilon - \beta_O \quad \text{on} \quad Q_\varepsilon D_{a,L,N_\tau}^{\text{loc}}(O),$$

so the remaining obstruction is localized to the low covariant-energy window already isolated by the dangerous-cluster mechanism.

3.3 Shifted Feshbach reduction

Write the block decomposition of $H - \delta$ with respect to $\mathcal{H}_{\text{exc}} = P_\varepsilon \mathcal{H}_{\text{exc}} \oplus Q_\varepsilon \mathcal{H}_{\text{exc}}$ as

$$H - \delta = \begin{pmatrix} A_\varepsilon(\delta) & B_\varepsilon \\ B_\varepsilon^* & D_\varepsilon(\delta) \end{pmatrix}.$$

When $D_\varepsilon(\delta)$ is strictly positive, the shifted Feshbach map is

$$F_\varepsilon(\delta) := A_\varepsilon(\delta) - B_\varepsilon D_\varepsilon(\delta)^{-1} B_\varepsilon^*. \quad (3.3)$$

Proposition 3.18 (Shifted Feshbach criterion). *Assume theorems 3.3 and 3.4 and let $0 < \delta < \alpha\varepsilon - \beta$, where α, β are the coercivity constants from theorem 3.4. Then*

$$H \geq \delta \quad \text{on} \quad \mathcal{H}_{\text{exc}} \quad \Longleftrightarrow \quad F_\varepsilon(\delta) \geq 0 \quad \text{on} \quad P_\varepsilon \mathcal{H}_{\text{exc}}.$$

Proof. By [theorem 3.5](#), $D_\varepsilon(\delta)$ is strictly positive on $Q_\varepsilon \mathcal{H}_{\text{exc}}$. The equivalence is the standard Schur-complement identity. $\square$

3.4 Packetwise phase counting and weighted packet multiplicity

The low-window problem is governed by the twisted eigenvalue formula [\(2.13\)](#). Low modes occur only when the holonomy phases θ_j are close to integer multiples of 2π . The first task is to count such phases packetwise.

Definition 3.19 (Packetwise phase-stable low-window reduction). Fix $\varepsilon > 0$. A packetization of the holonomy base is a partition of the relevant holonomy parameter set into finitely or countably many packets $\mathfrak{P}$ such that within each packet the phase set entering [\(2.13\)](#) varies by $O(\sqrt{\varepsilon})$, and the low twisted modes are stable under replacement by a packet representative. The corresponding packetized low window is the direct sum of the low eigenspaces associated with all packets.

Theorem 3.20 (Packetwise low-window rank bound). *Suppose the twisted fiber eigenvalue phases are packetized as above and that each packet supports at most $M_{\text{fib}}(\varepsilon)$ twisted low modes. If the number of relevant packets is at most $M_{\text{ph}}(\varepsilon)$, then*

$$\dim \text{Ran } P_\varepsilon \leq M_{\text{ph}}(\varepsilon) M_{\text{fib}}(\varepsilon).$$

In the nonabelian compact setting the natural refinement is the weighted packet-count bound

$$\sum_{\pi} \dim(M_{\pi}) N_{\pi, \text{low}}(\varepsilon) \leq M_*(\varepsilon),$$

where M_{π} are Peter–Weyl multiplicity spaces and $N_{\pi, \text{low}}(\varepsilon)$ counts packetized low modes in the π -block.

Proof. The abelian packet bound is immediate from the packet decomposition. In the nonabelian case one counts low modes on multiplicity spaces rather than irreducible factors; the weighted sum is exactly the natural dimension count in the Peter–Weyl decomposition. $\square$

Remark 3.21 (The first scale-sensitive quantity). The rank bound does not yet prove a gap. It isolates the first genuinely scale-sensitive quantity that must remain uniformly controlled along a continuum trajectory. In the rewritten producer–consumer architecture, collapse of packet number, near-trivial holonomy phases, or weighted packet multiplicity is therefore one of the earliest explicit failure modes to track in the renormalization-scale analysis.

3.5 Band-effective model and dangerous cluster

The low window is next compared with a compact-resolvent band model.

Definition 3.22 (Band-effective reduction data). Fix $\varepsilon > 0$ and $\delta > 0$ such that $F_\varepsilon(\delta)$ is defined. A band-effective reduction consists of:

- (i) a closed subspace $E_\varepsilon \subset P_\varepsilon \mathcal{H}_{\text{exc}}$ and a unitary map $W_\varepsilon: E_\varepsilon \rightarrow \mathcal{K}_\varepsilon$ onto a Hilbert space $\mathcal{K}_\varepsilon$;
- (ii) a self-adjoint compact-resolvent operator $L_{\varepsilon, \delta}$ on $\mathcal{K}_\varepsilon$;

(iii) constants $\eta_{\text{band}} \geq 0$, $c_{\text{band}} \geq 0$, and $r_{\text{out}} > 0$;

with the following properties. With respect to the orthogonal decomposition $P_\varepsilon \mathcal{H}_{\text{exc}} = E_\varepsilon \oplus E_\varepsilon^\perp$,

$$F_\varepsilon(\delta) = \begin{pmatrix} M_\varepsilon(\delta) & C_\varepsilon(\delta) \\ C_\varepsilon(\delta)^* & R_\varepsilon(\delta) \end{pmatrix}, \quad \|C_\varepsilon(\delta)\| \leq c_{\text{band}}, \quad R_\varepsilon(\delta) \geq r_{\text{out}} I, \quad (3.4)$$

and

$$\langle u, M_\varepsilon(\delta)u \rangle \geq \langle W_\varepsilon u, L_{\varepsilon, \delta} W_\varepsilon u \rangle - \eta_{\text{band}} \|u\|^2 \quad (u \in E_\varepsilon \cap \text{Dom}(M_\varepsilon(\delta))). \quad (3.5)$$

Theorem 3.23 (Constructive dangerous cluster). *Assume [theorem 3.22](#). Let $E_* > 0$ satisfy*

$$E_* > \eta_{\text{band}} + \frac{c_{\text{band}}^2}{r_{\text{out}}}. \quad (3.6)$$

Define the spectral projection

$$\Pi_{\varepsilon, E_*} := W_\varepsilon^* \mathbf{1}_{(-\infty, E_*]}(L_{\varepsilon, \delta}) W_\varepsilon,$$

extended by zero on $E_\varepsilon^\perp$. Then:

- (1) Π_{ε, E_*} has finite rank.
- (2) The complementary block

$$R_{\Pi_{\varepsilon, E_*}}(\delta) := (P_\varepsilon - \Pi_{\varepsilon, E_*}) F_\varepsilon(\delta) (P_\varepsilon - \Pi_{\varepsilon, E_*})$$

is strictly positive.

- (3) If in fact

$$E_* > \eta_{\text{band}} + \frac{2c_{\text{band}}^2}{r_{\text{out}}}, \quad (3.7)$$

then

$$R_{\Pi_{\varepsilon, E_*}}(\delta) \geq q_* I, \quad q_* := \min \left\{ \frac{1}{2}(E_* - \eta_{\text{band}}), \quad r_{\text{out}} - \frac{2c_{\text{band}}^2}{E_* - \eta_{\text{band}}} \right\} > 0. \quad (3.8)$$

Proof. Compact resolvent gives the finite rank of the cutoff spectral projection. On the orthogonal complement of Π_{ε, E_*} inside E_ε , (3.5) implies a lower bound by $(E_* - \eta_{\text{band}})I$. Together with (3.4) and a Schur-complement estimate using the off-block bound c_{band} , one obtains strict positivity and then the explicit margin (3.8). $\square$

The projection Π_{ε, E_*} is the dangerous cluster. Its range captures all genuinely low modes that can threaten positivity; the complement already has a positive remainder.

Theorem 3.24 (Dangerous cluster as finite-dimensional low-window sector). *Under the hypotheses of [theorems 2.9](#) and [3.23](#), the dangerous cluster $\text{Ran } \Pi_{\varepsilon, E_*}$ is the finite-dimensional part of the low covariant-energy window that remains after the uniformly positive tail has been removed by the band-effective Feshbach cutoff. Every later oriented-shadow and block-symbol analysis therefore takes place on an explicit finite-dimensional spectral sector selected from the projected model.*

Proof. By [theorem 2.9](#), $P_\varepsilon \mathcal{H}_{\text{exc}}$ is the low-window sector of the covariant operator K . [Theorem 3.23](#) defines Π_{ε, E_*} as the spectral cutoff on the band-effective model associated with the shifted Feshbach

map and proves that its complement already has a strictly positive remainder. Hence $\text{Ran } \Pi_{\varepsilon, E_*}$ captures precisely the finite-dimensional low-window modes left after the uniformly positive tail has been removed. $\square$

Definition 3.25 (Cluster effective operator). With respect to

$$P_\varepsilon \mathcal{H}_{\text{exc}} = \Pi_{\varepsilon, E_*} \mathcal{H}_{\text{exc}} \oplus (P_\varepsilon - \Pi_{\varepsilon, E_*}) \mathcal{H}_{\text{exc}},$$

write

$$F_\varepsilon(\delta) = \begin{pmatrix} M_\Pi(\delta) & C_\Pi(\delta) \\ C_\Pi(\delta)^* & R_\Pi(\delta) \end{pmatrix}.$$

The finite-rank cluster effective operator is

$$F_{\Pi, \varepsilon}^{\text{eff}}(\delta) := M_\Pi(\delta) - C_\Pi(\delta) R_\Pi(\delta)^{-1} C_\Pi(\delta)^*. \quad (3.9)$$

Proposition 3.26 (Cluster Schur complement). *Under the hypotheses of [theorem 3.23](#),*

$$F_\varepsilon(\delta) \geq 0 \text{ on } P_\varepsilon \mathcal{H}_{\text{exc}} \iff F_{\Pi, \varepsilon}^{\text{eff}}(\delta) \geq 0 \text{ on } \Pi_{\varepsilon, E_*} \mathcal{H}_{\text{exc}}.$$

Proof. Because $R_\Pi(\delta)$ is strictly positive, positivity of the full operator is equivalent to positivity of its Schur complement. $\square$

Theorem 3.27 (General model-plus-error criterion on the dangerous cluster). *Assume [theorems 3.3](#), [3.4](#) and [3.23](#). Fix $\varepsilon > \beta/\alpha$ and $0 < \delta < \alpha\varepsilon - \beta$. Suppose there exist a self-adjoint model operator $L_{\Pi, \varepsilon}(\delta)$ on the finite-dimensional space $\Pi_{\varepsilon, E_*} \mathcal{H}_{\text{exc}}$, an error bound $\eta_{\text{eff}} \geq 0$, and a margin $\mu_{\text{eff}} > 0$ such that*

$$\|F_{\Pi, \varepsilon}^{\text{eff}}(\delta) - L_{\Pi, \varepsilon}(\delta)\| \leq \eta_{\text{eff}}, \quad L_{\Pi, \varepsilon}(\delta) \geq (\eta_{\text{eff}} + \mu_{\text{eff}})I.$$

Then $H \geq \delta$ on $\mathcal{H}_{\text{exc}}$.

Proof. For every unit vector $x \in \Pi_{\varepsilon, E_*} \mathcal{H}_{\text{exc}}$ one has

$$\langle x, F_{\Pi, \varepsilon}^{\text{eff}}(\delta)x \rangle \geq \langle x, L_{\Pi, \varepsilon}(\delta)x \rangle - \eta_{\text{eff}} \geq \mu_{\text{eff}} > 0.$$

Hence $F_{\Pi, \varepsilon}^{\text{eff}}(\delta) \geq 0$ on the dangerous cluster. Proposition [3.26](#) yields $F_\varepsilon(\delta) \geq 0$ on the full low window, and Proposition [3.18](#) then gives $H \geq \delta$ on the full excited sector. $\square$

3.6 Rank-one and rank-two endgames

Theorem 3.28 (Multiplicity-one scalar model criterion / decisive one-band gap theorem). *Assume [theorem 3.27](#) and suppose the dangerous cluster has rank one and is spanned by a unit vector u_ε . Write*

$$m_{\text{eff}}(\varepsilon, \delta) := \langle u_\varepsilon, F_{\Pi, \varepsilon}^{\text{eff}}(\delta)u_\varepsilon \rangle.$$

Assume there exist a scalar model $\ell_\varepsilon(\delta) \in \mathbb{R}$, an error bound $\eta_{\text{eff}} \geq 0$, and a margin $\mu > 0$ such that

$$|m_{\text{eff}}(\varepsilon, \delta) - \ell_\varepsilon(\delta)| \leq \eta_{\text{eff}}, \quad \ell_\varepsilon(\delta) \geq \eta_{\text{eff}} + \mu.$$

Then $H \geq \delta$ on $\mathcal{H}_{\text{exc}}$.

Proof. Apply [theorem 3.27](#) with $L_{\Pi,\varepsilon}(\delta) = \ell_\varepsilon(\delta)I$ on the one-dimensional dangerous cluster. $\square$

Theorem 3.29 (First operator-valued case / decisive two-band gap theorem). *Assume the hypotheses of [theorem 3.27](#), except that the dangerous cluster now has rank two with orthonormal basis u_1, u_2 . Let*

$$M^{\text{eff}}(\varepsilon, \delta) := \left(\left\langle u_i, F_{\Pi,\varepsilon}^{\text{eff}}(\delta) u_j \right\rangle \right)_{i,j=1}^2.$$

Assume there exist a Hermitian model matrix

$$L_\varepsilon(\delta) = \begin{pmatrix} \ell_{11} & \ell_{12} \\ \ell_{12}^* & \ell_{22} \end{pmatrix},$$

an error bound $\eta_{\text{eff}} \geq 0$, and a number $\kappa > \eta_{\text{eff}}$ such that

$$\left\| M^{\text{eff}}(\varepsilon, \delta) - L_\varepsilon(\delta) \right\| \leq \eta_{\text{eff}} \quad (3.10)$$

and

$$\ell_{11} \geq \kappa, \quad \ell_{22} \geq \kappa, \quad (\ell_{11} - \kappa)(\ell_{22} - \kappa) \geq |\ell_{12}|^2. \quad (3.11)$$

Then $H \geq \delta$ on $\mathcal{H}_{\text{exc}}$.

Proof. The inequalities (3.11) assert that $L_\varepsilon(\delta) - \kappa I$ is positive semidefinite. Hence $L_\varepsilon(\delta) \geq \kappa I$, and (3.10) gives $M^{\text{eff}}(\varepsilon, \delta) \geq (\kappa - \eta_{\text{eff}})I \geq 0$. Apply [theorem 3.27](#). $\square$

Theorem 3.30 (Main finite-volume projected gap theorem). *Assume [theorems 3.3](#) and [3.4](#). Suppose that for some $\varepsilon > \beta/\alpha$ and $0 < \delta < \alpha\varepsilon - \beta$ the shifted Feshbach operator admits band-effective reduction data as in [theorem 3.22](#), the spectral cutoff of [theorem 3.23](#) produces a dangerous cluster of rank one or two, and the corresponding scalar or matrix positivity test of [theorems 3.28](#) and [3.29](#) holds with positive margin. Then*

$$H \geq \delta \quad \text{on } \mathcal{H}_{\text{exc}}.$$

Proof. Apply [theorem 3.28](#) in the one-dimensional case and [theorem 3.29](#) in the two-dimensional case. $\square$

Interface theorem 3.31 (Exports from the fixed-regulator gap layer). *Once the theorem targets [theorems 3.1](#) and [3.2](#) are discharged and the band-effective hypotheses of [theorems 3.22](#) and [3.23](#) are verified, the fixed-regulator gap layer exports four exact data to the transport and identification papers:*

- (1) *the admissible low-window regime of [theorem 3.6](#);*
- (2) *the high-window positivity bound of [theorem 3.5](#);*
- (3) *the finite-rank dangerous-cluster decomposition with positive remainder of [theorem 3.23](#);*
- (4) *the finite-dimensional model-plus-error closure mechanism of [theorem 3.27](#), together with its multiplicity-one and two-band specializations [theorems 3.28](#) and [3.29](#).*

These are precisely the low-window data consumed later by the thermodynamic/continuum transport layer and by the oriented-shadow identification layer.

Proof. The four statements are just the proved outputs of the present section, read as downstream imports once the named theorem targets and band-effective hypotheses have been established in the concrete lattice-gauge realization. $\square$

4 Large-volume and continuum transfer

This section records the large-volume and continuum-transport packages downstream of the fixed-regulator reduction. In the earlier comprehensive draft these layers were visible only as assumptions. They are now rewritten in producer–consumer form. The large-volume and continuum theorems themselves remain abstract functional-analytic consequences of explicit embedding and convergence packages; the concrete verification problem is separated into named theorem targets. The same rewrite also transports not merely a scalar lower bound but the dangerous-cluster package itself, via common-comparison-space realizations of the scale-wise band-effective operators and norm-resolvent stability of their Riesz projections.

4.1 Uniform large-volume theorem

Fix $a > 0$ and N_τ , and let $L \uparrow \infty$. Denote by H_L , K_L , Ω_L , and $\mathcal{H}_{L,\text{exc}}$ the corresponding projected objects.

Theorem target 4.1 (Local embeddings and vacuum compatibility). *Construct a Hilbert space $\mathcal{H}_\infty^{4\text{D}}$, a vacuum vector Ω_∞ , and isometric embeddings*

$$\iota_L: \mathcal{H}_L^{4\text{D}} \rightarrow \mathcal{H}_\infty^{4\text{D}}$$

compatible with the local observable net and the projected vacua, in the sense that $\iota_L \Omega_L \rightarrow \Omega_\infty$ in norm.

Theorem target 4.2 (Gap-preserving thermodynamic Mosco convergence). *Let q_L be the quadratic forms of the projected Hamiltonians H_L . Prove that q_L Mosco-converges, along the directed volume system, to a closed quadratic form q_∞ on $\mathcal{H}_\infty^{4\text{D}}$ associated with a nonnegative self-adjoint operator H_∞ , and that the recovery sequences can be orthogonalized against the projected vacua without increasing the limiting energy.*

Interface theorem 4.3 (Thermodynamic persistence interface). *Assume the fixed-scale exports of [theorem 3.31](#) together with [theorems 4.1](#) and [4.2](#). Suppose there exists $\delta > 0$ such that*

$$H_L \geq \delta \quad \text{on } \mathcal{H}_L^{4\text{D}} \ominus \mathbb{C}\Omega_L$$

for all sufficiently large L . Then

$$H_\infty \geq \delta \quad \text{on } \mathcal{H}_\infty^{4\text{D}} \ominus \mathbb{C}\Omega_\infty.$$

In particular, the projected infinite-volume theory at fixed lattice scale retains the same lower bound above the vacuum.

Proof. This is the standard thermodynamic lower-bound transfer under vacuum-compatible Mosco convergence. Let $\psi \in \text{Dom}(q_\infty)$ with $\psi \perp \Omega_\infty$. Choose a recovery sequence and orthogonalize it

against the finite-volume vacua; the orthogonalization does not raise the limiting energy by the second theorem target. The finite-volume lower bound then passes to the limit by the Mosco liminf inequality. $\square$

Theorem 4.4 (Uniform low-rank large-volume theorem). *Suppose there exist constants*

$$\alpha > 0, \beta \geq 0, \varepsilon > \beta/\alpha, 0 < \delta < \alpha\varepsilon - \beta, c_{\text{band}}, \eta_{\text{band}}, r_{\text{out}} > 0,$$

independent of L , such that for all sufficiently large L :

- (U1) *the projected-sector/common-core and coercivity package of [theorems 3.3](#) and [3.4](#) holds with the same α and β ;*
- (U2) *the shifted Feshbach operator $F_{L,\varepsilon}(\delta)$ admits band-effective data with the same constants $c_{\text{band}}, \eta_{\text{band}}, r_{\text{out}}$;*
- (U3) *the corresponding dangerous cluster has rank one or two;*
- (U4) *in the rank-one case a uniform scalar margin as in [theorem 3.28](#) holds;*
- (U5) *in the rank-two case a uniform matrix margin as in [theorem 3.29](#) holds.*

Then

$$H_L \geq \delta \quad \text{on } \mathcal{H}_L^{4\text{D}} \ominus \mathbb{C}\Omega_L$$

for all sufficiently large L .

Proof. Apply [theorem 3.30](#) volume by volume with the same constants. $\square$

4.2 Scale-uniform trajectory package

Let $a_n \downarrow 0$ be a renormalization trajectory. For each n assume a projected infinite-volume theory

$$H_n^{\text{proj}}, \quad \Omega_n^{\text{proj}} \in \mathcal{H}_n^{\text{proj}}, \quad H_n^{\text{proj}} \geq 0,$$

with nonnegative closed form q_n^{proj} . Let $\mathcal{H}_{\text{cont}}^{4\text{D}}$ be the projected continuum Hilbert space with nonnegative closed form q_{cont} . The convergence package is formulated through isometric embeddings $I_n: \mathcal{H}_n^{\text{proj}} \rightarrow \mathcal{H}_{\text{cont}}^{4\text{D}}$.

Assumption 4.5 (Common-core Mosco package). There exists a dense subspace $\mathcal{C}_{\text{loc}} \subset \text{Dom}(q_{\text{cont}})$ such that:

- (M1) $\mathcal{C}_{\text{loc}}$ is a form core for q_{cont} ;
- (M2) for every $\phi \in \mathcal{C}_{\text{loc}}$ there exist recovery vectors $\phi_n \in \mathcal{H}_n^{\text{proj}}$ with $I_n \phi_n \rightarrow \phi$ strongly and $q_n^{\text{proj}}[\phi_n] \rightarrow q_{\text{cont}}[\phi]$;
- (M3) whenever $I_{n_k} \psi_k \rightharpoonup \psi$ weakly in $\mathcal{H}_{\text{cont}}^{4\text{D}}$ and $\sup_k (q_{n_k}^{\text{proj}}[\psi_k] + \|\psi_k\|^2) < \infty$, one has

$$q_{\text{cont}}[\psi] \leq \liminf_{k \rightarrow \infty} q_{n_k}^{\text{proj}}[\psi_k];$$

(M4) $I_n \Omega_n^{\text{proj}} \rightarrow \Omega_{\text{cont}}^{\text{proj}}$ strongly;

(M5) if $\psi_n \perp \Omega_n^{\text{proj}}$, $I_n \psi_n \rightharpoonup \psi$, and the energies stay bounded, then $\psi \perp \Omega_{\text{cont}}^{\text{proj}}$.

Theorem target 4.6 (Scale-wise projected continuum trajectory). *Produce the scale-wise projected infinite-volume theories, mixed operators, embeddings, and vacuum compatibility needed to realize [theorem 4.5](#) along a renormalization trajectory $a_n \downarrow 0$, including compatibility of the projected mixed operators with the chosen auxiliary scaling $N_\tau(a_n)$.*

Theorem target 4.7 (Common-comparison-space band realization). *For fixed ε and δ , realize the scale-wise band-effective comparison operators on a common Hilbert space $\mathcal{K}$ as self-adjoint operators $L_{n,\varepsilon,\delta}$ together with a limiting operator $L_{\infty,\varepsilon,\delta}$. The realization must be compatible with the projected low-window band package produced at each fixed scale and with the transport of dangerous-cluster data used by the identification layer.*

Theorem target 4.8 (Norm-resolvent transport and isolated contour). *For the realization of [theorem 4.7](#), prove norm-resolvent convergence*

$$\left\| (L_{n,\varepsilon,\delta} - z)^{-1} - (L_{\infty,\varepsilon,\delta} - z)^{-1} \right\| \rightarrow 0$$

for each z in the common resolvent set of a contour Γ isolating the low discrete spectrum, and verify that Γ remains disjoint from the spectrum along the trajectory.

Theorem 4.9 (Continuum transfer of a projected spectral lower bound). *Assume [theorem 4.5](#). Suppose there exists $\delta > 0$ such that for all sufficiently large n ,*

$$H_n^{\text{proj}} \geq \delta \quad \text{on } \mathcal{H}_n^{\text{proj}} \ominus \mathbb{C} \Omega_n^{\text{proj}}.$$

Then

$$H_{\text{cont}}^{\text{proj}} \geq \delta \quad \text{on } \mathcal{H}_{\text{cont}}^{4\text{D}} \ominus \mathbb{C} \Omega_{\text{cont}}^{\text{proj}}.$$

Proof. Let $\phi \in \text{Dom}(q_{\text{cont}})$ with $\phi \perp \Omega_{\text{cont}}^{\text{proj}}$. Choose a recovery sequence ϕ_n from [theorem 4.5](#). Orthogonalize it by setting

$$\psi_n := \phi_n - \langle \phi_n, \Omega_n^{\text{proj}} \rangle \Omega_n^{\text{proj}}.$$

Then $\psi_n \perp \Omega_n^{\text{proj}}$, $I_n \psi_n \rightarrow \phi$ strongly, and $q_n^{\text{proj}}[\psi_n] = q_n^{\text{proj}}[\phi_n]$. Hence for large n ,

$$q_n^{\text{proj}}[\psi_n] \geq \delta \|\psi_n\|^2.$$

Passing to the limit along the recovery sequence gives $q_{\text{cont}}[\phi] \geq \delta \|\phi\|^2$. This is exactly the claimed lower bound. $\square$

Corollary 4.10 (Continuum persistence of the projected gap). *Assume the hypotheses of [theorems 4.4](#) and [4.5](#). Then the projected continuum Hamiltonian has spectral gap at least δ above the vacuum.*

Proof. Combine [theorem 4.4](#) and [theorem 4.9](#). $\square$

Theorem 4.11 (Stable transport of dangerous clusters). *Assume the conclusions of [theorems 4.7](#) and [4.8](#). Define the Riesz projections*

$$\Pi_n := \frac{1}{2\pi i} \int_{\Gamma} (\zeta - L_{n,\varepsilon,\delta})^{-1} d\zeta, \quad \Pi_{\infty} := \frac{1}{2\pi i} \int_{\Gamma} (\zeta - L_{\infty,\varepsilon,\delta})^{-1} d\zeta.$$

Then:

- (a) $\Pi_n \rightarrow \Pi_\infty$ in operator norm;
- (b) $\text{rank } \Pi_n = \text{rank } \Pi_\infty$ for all sufficiently large n .

Consequently, the dangerous-cluster package stabilizes along the renormalization trajectory and arrives at the identification layer as an eventually constant finite-rank family.

Proof. Norm-resolvent convergence is uniform on the compact contour Γ , hence the Cauchy integrals defining the Riesz projections converge in operator norm. A sufficiently small norm perturbation of a projection cannot change its rank, so the ranks stabilize. $\square$

Corollary 4.12 (Trajectory interface for the identification layer). *Assume the hypotheses of [theorems 4.9](#) and [4.11](#). Then the continuum/interface and oriented-shadow layers may begin from the projected continuum Hamiltonian, its vacuum, a preserved projected lower bound, and an eventually constant finite-rank dangerous-cluster family transported on a common comparison space.*

Proof. The continuum projected lower bound is provided by [theorem 4.9](#). The stable dangerous-cluster package follows from [theorem 4.11](#). These are exactly the two continuum-level inputs consumed downstream by the countable-core reduction and the dangerous-cluster shadow analysis. $\square$

Remark 4.13 (The most tractable verification target inside the present transport layer). Among the analytic theorem targets exposed by the present rewrite, the common-comparison-space realization and isolated-contour norm-resolvent transport of the band-effective models are the closest to theorem-level closure. Their abstract consequence, [theorem 4.11](#), is already fully proved once the comparison-space realization, norm-resolvent convergence, and contour isolation are supplied in the concrete model. By contrast, the coercive estimate itself remains a front-end lattice comparison problem, and auxiliary invisibility remains the final observable obstruction.

5 Countable cores and interface theorems

This section develops the countable-core and interface machinery needed to convert finite-dimensional matching statements into full vacuum-theory identification. In the rewritten producer–consumer architecture it begins only after the continuum lower bound and the stable dangerous-cluster family of [theorem 4.12](#) are available. The countable Wilson-loop/curvature core isolates the observable content of the projected continuum theory after auxiliary dressing is forgotten. We prove density of that core in both the Euclidean and operator-algebraic settings and then show that auxiliary invisibility on it is the only remaining obstruction to Yang–Mills identification.

5.1 Projected observable packages

Fix a bounded rational region $O \subset \mathbb{R}^4$. The projected continuum theory is assumed to carry a Euclidean local algebra $\mathcal{E}_{\text{cont}}^{4\text{D}}(O)$ and a bounded operator algebra $\mathcal{A}_{\text{cont}}^{4\text{D}}(O) \subset B(\mathcal{H}_{\text{cont}}^{4\text{D}})$. The intended generators are projected Wilson-loop observables and projected smeared curvature observables. The only properties required in the present section are separability, continuity of the parameter maps, and compatibility with forgetting the auxiliary dressing.

Assumption 5.1 (Continuity and separability package). For every bounded rational $O \subset \mathbb{R}^4$ the following data are available.

- (O1) A separable metric space $\Gamma(O)$ of admissible piecewise C^1 loops in O .
- (O2) A countable set $\widehat{G}$ of irreducible finite-dimensional unitary representations of the compact gauge group.
- (O3) A separable metric space $D_{\text{aux}}(O)$ of admissible neutral auxiliary dressings.
- (O4) A locally convex test-function space $\mathcal{T}(O)$ of curvature smearings, together with a countable dense subset.
- (O5) For each $\gamma \in \Gamma(O)$, $\pi \in \widehat{G}$, $d \in D_{\text{aux}}(O)$, and $f \in \mathcal{T}(O)$ there are projected observables

$$W_{\gamma,\pi,d}^{\text{proj}} \in \mathcal{E}_{\text{cont}}^{4\text{D}}(O), \quad F_{f,d}^{\text{proj}} \in \mathcal{E}_{\text{cont}}^{4\text{D}}(O),$$

with bounded representatives

$$W_{\gamma,\pi,d}^{\text{proj}} \in \mathcal{A}_{\text{cont}}^{4\text{D}}(O), \quad B_{f,d}^{\text{proj}} := b(F_{f,d}^{\text{proj}}) \in \mathcal{A}_{\text{cont}}^{4\text{D}}(O),$$

where $b(x) = x(1 + x^2)^{-1/2}$ is the bounded transform.

- (O6) The parameter maps into the Euclidean and operator algebras are continuous.
- (O7) Forgetting the auxiliary dressing defines a four-dimensional content map

$$q_O(W_{\gamma,\pi,d}^{\text{proj}}) = W_{\gamma,\pi}^{4\text{D}}, \quad q_O(F_{f,d}^{\text{proj}}) = F_f^{4\text{D}}.$$

- (O8) The local Euclidean algebra is the completion of the algebraic span of the projected Wilson-loop and curvature generators, and the local operator algebra is the norm closure of the algebra generated by their bounded representatives.

5.2 Countable Wilson-loop/curvature cores

Choose countable dense subsets

$$\Gamma_{\text{rat}}(O) \subset \Gamma(O), \quad \mathcal{T}_{\text{rat}}(O) \subset \mathcal{T}(O), \quad D_{\text{aux,rat}}(O) \subset D_{\text{aux}}(O).$$

Definition 5.2 (Projected and four-dimensional countable cores). Let $\mathcal{C}^{\text{proj}}(O)$ be the unital $*$ -algebra over $\mathbb{Q} + i\mathbb{Q}$ generated by the symbols

$$W_{\gamma,\pi,d}^{\text{proj}}, \quad F_{f,d}^{\text{proj}},$$

with $\gamma \in \Gamma_{\text{rat}}(O)$, $\pi \in \widehat{G}$, $d \in D_{\text{aux,rat}}(O)$, and $f \in \mathcal{T}_{\text{rat}}(O)$. Let $\mathcal{C}^{4\text{D}}(O)$ be the analogous algebra generated by the four-dimensional content symbols $W_{\gamma,\pi}^{4\text{D}}$ and $F_f^{4\text{D}}$.

Lemma 5.3 (Countability and quotient map). *For every bounded rational region O :*

- (i) $\mathcal{C}^{\text{proj}}(O)$ and $\mathcal{C}^{4\text{D}}(O)$ are countable sets;

(ii) the forgetful map q_O extends to a surjective unital $*$ -homomorphism

$$q_O: \mathcal{C}^{\text{proj}}(O) \twoheadrightarrow \mathcal{C}^{4\text{D}}(O).$$

Proof. Each parameter set is countable, hence the set of finite $*$ -words with rational complex coefficients is countable. The map on generators respects the $*$ -algebra operations and is obviously surjective. $\square$

Definition 5.4 (Auxiliary ideal). Two elements $X, Y \in \mathcal{C}^{\text{proj}}(O)$ are *auxiliary equivalent* if $q_O(X) = q_O(Y)$. Let $I_{\text{aux}}(O)$ be the two-sided $*$ -ideal generated by all differences $X - Y$ with $X \sim_{\text{aux}} Y$.

Lemma 5.5 (Quotient core). *The canonical map induced by q_O gives a $*$ -isomorphism*

$$\mathcal{C}^{\text{proj}}(O)/I_{\text{aux}}(O) \cong \mathcal{C}^{4\text{D}}(O).$$

Proof. By construction $I_{\text{aux}}(O) \subset \ker q_O$. Conversely, if $Z \in \ker q_O$, write Z as a finite linear combination of core monomials, group the terms with the same four-dimensional content, and observe that each grouped difference lies in the ideal generated by auxiliary-equivalent differences. Hence $\ker q_O = I_{\text{aux}}(O)$. $\square$

5.3 Density theorems

Theorem 5.6 (Euclidean density of the countable core). *Assume [theorem 5.1](#). For every bounded rational region O , the image of $\mathcal{C}^{\text{proj}}(O)$ is dense in $\mathcal{E}_{\text{cont}}^{4\text{D}}(O)$, and the image of $\mathcal{C}^{4\text{D}}(O)$ is dense in the corresponding four-dimensional Euclidean local algebra.*

Proof. By construction the full local Euclidean algebra is the completion of the span of all projected Wilson-loop and curvature generators. The continuity of the parameter maps and the density of the rational parameter sets show that every generator is approximable by generators from the countable subset. Rational complex coefficients approximate arbitrary complex coefficients, and density follows. $\square$

Theorem 5.7 (Operator-algebraic density of the countable core). *Assume [theorem 5.1](#). For every bounded rational region O , the local C^* -algebra $\mathcal{A}_{\text{cont}}^{4\text{D}}(O)$ is the norm closure of the $*$ -algebra generated by the bounded core observables*

$$W_{\gamma, \pi, d}^{\text{proj}}, \quad B_{f, d}^{\text{proj}},$$

with rational loop, smearing, and dressing parameters.

Proof. Exactly the same density argument works in operator norm because the generating maps into the bounded local algebra are norm-continuous. $\square$

5.4 Euclidean interface theorem

Assume there exists a normalized reflection-positive Euclidean functional

$$S^{\text{proj}}: \bigcup_O \mathcal{E}_{\text{cont}}^{4\text{D}}(O) \rightarrow \mathbb{C}$$

whose Osterwalder–Schrader reconstruction yields the projected continuum vacuum sector $(\mathcal{H}_{\text{cont}}^{4\text{D}}, \Omega_{\text{cont}}^{\text{proj}})$.

Theorem 5.8 (Euclidean extension from the countable core). *Assume [theorem 5.1](#). Let S_1 and S_2 be normalized continuous Euclidean functionals on the local Euclidean algebra such that*

$$S_1(A) = S_2(A) \quad \text{for all } A \in \mathcal{C}^{\text{proj}}(O) \text{ and all bounded rational } O.$$

Then $S_1 = S_2$ on the completed Euclidean algebra generated by the local spaces $\mathcal{E}_{\text{cont}}^{4\text{D}}(O)$.

Proof. By [theorem 5.6](#), the countable core is dense in each local algebra, and both functionals are continuous. Hence equality on the dense subset extends uniquely to equality on the completion. $\square$

Theorem 5.9 (Reconstruction uniqueness from core agreement). *Assume the hypotheses of [theorem 5.8](#). If in addition S_1 and S_2 are reflection positive and satisfy the same Osterwalder–Schrader regularity assumptions, then their Osterwalder–Schrader reconstructions are unitarily equivalent by a vacuum-preserving unitary.*

Proof. By [theorem 5.8](#), the two functionals agree on the full Euclidean algebra. The Osterwalder–Schrader pre-Hilbert structures therefore coincide, as do their null spaces, completed Hilbert spaces, and cyclic vacuum vectors. The time-translation semigroups and Hamiltonians reconstructed from the common functional are intertwined by the canonical unitary. $\square$

Definition 5.10 (Auxiliary invisibility on the Euclidean core). The projected continuum theory is *auxiliary invisible on the explicit Euclidean core* if for every bounded rational region O and every $A \in \mathcal{C}^{\text{proj}}(O)$ one has

$$S^{\text{proj}}(A) = S^{4\text{D}}(q_O(A)),$$

where $S^{4\text{D}}$ is the Euclidean functional of the four-dimensional shadow or Yang–Mills target theory.

Theorem 5.11 (Euclidean identification reduced to auxiliary invisibility). *Assume [theorem 5.1](#) and the common Osterwalder–Schrader regularity hypotheses. If the projected continuum theory is auxiliary invisible on the explicit Euclidean core and if the shadow or Yang–Mills target functional agrees with the projected one on that core, then the projected continuum theory and the target theory are unitarily equivalent in their vacuum representations.*

Proof. Auxiliary invisibility identifies the projected functional with the four-dimensional content functional on the countable core. By [theorem 5.8](#) the equality extends to the completed Euclidean algebra, and [theorem 5.9](#) then yields the vacuum-preserving unitary equivalence. $\square$

5.5 Operator-algebraic interface theorem

Let $\mathfrak{A}^{\text{proj}}$ be the quasi-local C^* -algebra generated by the local algebras $\mathcal{A}_{\text{cont}}^{4\text{D}}(O)$ and let ω^{proj} be the projected vacuum state. The four-dimensional or Yang–Mills target theory yields a corresponding quasi-local algebra $\mathfrak{A}^{\text{tar}}$ and vacuum state ω^{tar} .

Theorem 5.12 (Vacuum-preserving extension from the operator core). *Assume [theorem 5.1](#). Let*

$$\Phi_0: \bigcup_O \mathcal{C}^{\text{proj}}(O) \rightarrow \mathfrak{A}^{\text{tar}}$$

be a family of unital $*$ -homomorphisms compatible on overlaps and satisfying

$$\omega^{\text{proj}}(A) = \omega^{\text{tar}}(\Phi_0(A)) \quad (A \in \mathcal{C}^{\text{proj}}(O), O \text{ bounded rational}).$$

Then Φ_0 extends uniquely to a vacuum-preserving local $*$ -homomorphism

$$\Phi: \mathfrak{A}^{\text{proj}} \rightarrow \mathfrak{A}^{\text{tar}}.$$

If the same holds for the inverse direction, then the two quasi-local vacuum nets are isomorphic.

Proof. The countable core generates a dense local $*$ -subalgebra by [theorem 5.7](#). Because Φ_0 preserves the vacuum state, its local pieces are contractive with respect to the GNS seminorms and therefore extend by continuity. Compatibility on overlaps yields a global quasi-local map. If both directions exist, they are inverse isometric $*$ -isomorphisms. $\square$

Theorem 5.13 (Operator-algebraic identification reduced to auxiliary invisibility). *Under the assumptions of [theorem 5.12](#), auxiliary invisibility and core-level matching on the countable operator core imply vacuum-preserving equivalence of the projected continuum net and the target vacuum net.*

Proof. Auxiliary invisibility identifies auxiliary-equivalent core observables. The induced map on the quotient core $\mathcal{C}^{\text{proj}}(O)/I_{\text{aux}}(O) \cong \mathcal{C}^{4\text{D}}(O)$ preserves the vacuum states. Apply [theorem 5.12](#) to extend this local identification to the completed quasi-local nets. $\square$

Theorem 5.14 (Combined interface theorem). *Assume the projected continuum gap of [theorem 4.10](#), the countable-core continuity package of [theorem 5.1](#), and either the Euclidean or operator-algebraic interface hypotheses of [theorems 5.11](#) and [5.13](#). Then the only remaining obstruction to Yang–Mills identification is auxiliary invisibility on the explicit countable Wilson-loop/curvature core.*

Proof. The gap statement has already been transferred to the projected continuum theory. [Theorems 5.11](#) and [5.13](#) show that once equality is known on the explicit core, the completed vacuum theories coincide. Thus the remaining task is precisely the core-level auxiliary-invisibility and matching problem. $\square$

Remark 5.15 (Explicit failure modes and first verification targets). The value of the countable-core reduction is partly negative: it makes the remaining failure modes exact rather than implicit. Collapse of the scale-uniform packet-count or dangerous-cluster constants can destroy the projected continuum gap before any identification step is reached; the Mosco-type convergence scheme can fail to preserve the vacuum or the orthogonalized recovery-sequence property; observable continuity or separability can fail on the projected continuum core; and even after all extension/reconstruction machinery is in place, auxiliary invisibility can still fail on that core. The most realistic first entry points are therefore now visible in theorem-level language: a one-band Wilson-loop subcore, where the problem reduces to scalar cluster-level matching; a restricted local-curvature subcore, where bounded substitutes simplify the cylinder-functional estimates; or a two-band resonant window, where the obstruction is compressed to a common 2×2 shadow matrix.

Part II

Part II: Oriented shadows, closed sectors, localization, and curvature approximation

6 Oriented shadows, Real commutants, and intrinsic matching

This section rewrites the shadow layer of the argument in the intrinsic language of canonical complexification and compact nonabelian spectral orientation, now consuming the transported dangerous-cluster family provided by [theorem 4.12](#). The key imported point is that the dangerous cluster is not merely a finite-dimensional space on which one may choose a convenient complex structure: once compact symmetry is present, the relevant datum is the Real commutant of the cluster representation, and every admissible shadow complex structure is a skew-adjoint square root of $-\mathbf{1}$ in that real C^* -algebra [\[13, 14\]](#). Accordingly, the shadow sector is organized by operator-valued orientation on complex-type multiplicity blocks together with real/quaternionic data on self-conjugate blocks. The anti-linear leakage appearing later is precisely the defect of a compressed observable from being linear for the chosen point of this Real-commutant parameter space. Auxiliary invisibility and Yang–Mills matching are therefore reduced to an intrinsic matrix-algebra problem on the dangerous cluster.

6.1 Dangerous clusters with compact symmetry

From this point onward we work along the scale-uniform projected continuum trajectory. For each scale n the low window yields a finite-dimensional dangerous-cluster space

$$K_n := \text{Ran } \Pi_{n,\varepsilon} \subset \mathcal{H}_n^{\text{proj}},$$

where $\Pi_{n,\varepsilon} = \mathbf{1}_{[0,\varepsilon]}(K_n^{\text{cov}})$ denotes the scale-wise low-window projection. The projected continuum observables and their shadow counterparts compress to bounded self-adjoint operators

$$X_{n,j} := \Pi_{n,\varepsilon} G_j^{(n)} \Pi_{n,\varepsilon}|_{K_n}, \quad Y_{n,j} := \Pi_{n,\varepsilon} \tilde{G}_j^{(n)} \Pi_{n,\varepsilon}|_{K_n}.$$

Assumption 6.1 (Compact symmetry on the dangerous cluster). There exists a second-countable compact group $\mathbf{G}$ and, for each n , a continuous orthogonal representation

$$U_n : \mathbf{G} \rightarrow O(K_n)$$

such that every $X_{n,j}$ and every $Y_{n,j}$ commutes with $U_n(g)$ for all $g \in \mathbf{G}$.

Complexifying K_n gives a finite-dimensional unitary representation with Peter–Weyl decomposition

$$K_n^{\mathbb{C}} \cong \bigoplus_{\pi \in \widehat{\mathbf{G}}_n} H_\pi \otimes M_{n,\pi},$$

where H_π carries the irreducible representation and $M_{n,\pi}$ is the multiplicity space. Let C_n denote the canonical conjugation on $K_n^{\mathbb{C}}$, and let V_n be the complexified unitary representation. Canonical

conjugation transports the π -block to the $\bar{\pi}$ -block, and self-conjugate blocks split into real and quaternionic types according to the Frobenius–Schur sign.

Definition 6.2 (Complex commutant, Real commutant, and parameter space). Define the complex commutant and Real commutant of the dangerous-cluster representation by

$$\mathcal{A}(U_n) := \{T \in B(K_n^{\mathbb{C}}) : TV_n(g) = V_n(g)T \text{ for all } g \in \mathbf{G}\},$$

$$\mathcal{A}_R(U_n) := \{T \in \mathcal{A}(U_n) : TC_n = C_nT\},$$

and write

$$X(U_n) := \{J \in B(K_n) : J^* = -J, J^2 = -\mathbf{1}, JU_n(g) = U_n(g)J \text{ for all } g \in \mathbf{G}\}.$$

Proposition 6.3 (Intrinsic commutant formulation on the dangerous cluster). *Complexification induces an isometric real $*$ -isomorphism*

$$\text{Comm}_{\mathbb{R}}(U_n) := \{S \in B(K_n) : SU_n(g) = U_n(g)S \text{ for all } g \in \mathbf{G}\} \longrightarrow \mathcal{A}_R(U_n), \quad S \longmapsto S^{\mathbb{C}}.$$

In particular, $\mathcal{A}_R(U_n)$ is a real C^* -subalgebra of $B(K_n^{\mathbb{C}})$, and $X(U_n)$ is canonically identified with the set of skew-adjoint square roots of $-\mathbf{1}$ in $\mathcal{A}_R(U_n)$.

Proof. If $S \in \text{Comm}_{\mathbb{R}}(U_n)$, then $S^{\mathbb{C}}$ is complex-linear, commutes with $V_n(g)$ for every g , and commutes with the canonical conjugation because S is real-linear. Thus $S^{\mathbb{C}} \in \mathcal{A}_R(U_n)$. Conversely, if $T \in \mathcal{A}_R(U_n)$, then T preserves the fixed real subspace $K_n = \text{Fix}(C_n)$, and its restriction is a bounded real-linear operator commuting with U_n . The two constructions are inverse to one another and preserve the operator norm and adjoint. The final assertion is immediate from the defining relations of an orthogonal complex structure. $\square$

Corollary 6.4 (Functoriality under cluster intertwiners). *If $W : K_n \rightarrow K'_n$ is a real unitary intertwiner between compact dangerous-cluster representations U_n and U'_n , then conjugation by $W^{\mathbb{C}}$ induces an isometric real $*$ -isomorphism*

$$\text{Ad}(W^{\mathbb{C}}) : \mathcal{A}_R(U_n) \xrightarrow{\cong} \mathcal{A}_R(U'_n)$$

and carries $X(U_n)$ bijectively onto $X(U'_n)$.

Proof. Because $W^{\mathbb{C}}$ intertwines the complexified representations and commutes with canonical conjugation, conjugation by $W^{\mathbb{C}}$ preserves both the commuting and Real conditions. The statement about parameter spaces follows from [theorem 6.3](#). $\square$

Let $P_n \subset \widehat{\mathbf{G}}_{n,\text{cx}}$ be a representative set of complex-type conjugate pairs, let $\widehat{\mathbf{G}}_{n,\text{re}}$ and $\widehat{\mathbf{G}}_{n,\text{qu}}$ denote the real-type and quaternionic-type self-conjugate classes, and choose the usual antiunitary transports on the self-conjugate multiplicity spaces.

Theorem 6.5 (Finite-dimensional Real-commutant structure theorem). *Fix n and a Peter–Weyl realization of $K_n^{\mathbb{C}}$. Then there is an isometric real $*$ -isomorphism*

$$\mathcal{A}_R(U_n) \cong \prod_{\pi \in P_n} B(M_{n,\pi}) \times \prod_{\rho \in \widehat{\mathbf{G}}_{n,\text{re}}} B(\text{Fix}(c_{n,\rho})) \times \prod_{\sigma \in \widehat{\mathbf{G}}_{n,\text{qu}}} \{A \in B(M_{n,\sigma}) : Ac_{n,\sigma} = c_{n,\sigma}A\},$$

where the first factor is viewed as a real C^* -algebra, the second factor consists of bounded real-linear operators on the fixed real multiplicity spaces, and the third factor is the quaternionic-linear part commuting with the quaternionic transport. Consequently, $X(U_n)$ is the corresponding product of skew-adjoint square roots of $-\mathbf{1}$ on these factors.

Proof. This is the finite-dimensional compact-group specialization of the Real-commutant structure theorem proved in [13, 14]. Because K_n is finite-dimensional, the Peter–Weyl decomposition is atomic and no measurable transport issue remains. Schur’s lemma identifies the complex commutant blockwise as $\mathbf{1}_{H_\pi} \otimes B(M_{n,\pi})$. Imposing commutation with the canonical conjugation forces the displayed real/quaternionic restrictions on the multiplicity spaces, which is exactly the Frobenius–Schur decomposition of the Real commutant. $\square$

Corollary 6.6 (Artin–Wedderburn form and multiplicity-one reduction). *If all multiplicity spaces are finite-dimensional, then $\mathcal{A}_R(U_n)$ is a finite product of matrix algebras over $\mathbb{C}$, $\mathbb{R}$, and $\mathbb{H}$ as a real $*$ -algebra. In particular, if a complex-type block satisfies $\dim_{\mathbb{C}} M_{n,\pi} = 1$, then the corresponding factor of $X(U_n)$ consists of the two sign choices $\pm \mathbf{1}$.*

Proof. The first claim is the finite-dimensional Artin–Wedderburn form of the product decomposition in theorem 6.5. For a complex-type multiplicity space of dimension one, every bounded operator is scalar, and the relations $J^* = -J$ and $J^2 = -\mathbf{1}$ force that scalar to be $\pm i$. $\square$

Definition 6.7 (Oriented shadow structure). An *oriented shadow structure* on K_n is an element $J_n^{\text{sh}} \in X(U_n)$. Equivalently, it is a choice of operator-valued orientation on each complex-type multiplicity block together with real/quaternionic complex-structure data on the self-conjugate blocks. Relative to such a choice, every bounded operator T on K_n admits the canonical decomposition

$$T = T^{\text{lin}} + \Gamma_J(T), \quad T^{\text{lin}} := \frac{1}{2}(T - J_n^{\text{sh}} T J_n^{\text{sh}}), \quad \Gamma_J(T) := \frac{1}{2}(T + J_n^{\text{sh}} T J_n^{\text{sh}}). \quad (6.1)$$

Then T^{lin} commutes with J_n^{sh} , $\Gamma_J(T)$ anticommutes with J_n^{sh} , and the decomposition is unique. For the compressed generators we write

$$X_{n,j}^{\text{lin}} := \frac{1}{2}(X_{n,j} - J_n^{\text{sh}} X_{n,j} J_n^{\text{sh}}), \quad \Gamma_{n,j} := \frac{1}{2}(X_{n,j} + J_n^{\text{sh}} X_{n,j} J_n^{\text{sh}}).$$

Proposition 6.8 (Finite-dimensional classification of symmetry-compatible complex structures on the dangerous cluster). *Fix n . To specify an oriented shadow structure $J_n^{\text{sh}} \in X(U_n)$ is equivalent to specifying:*

- (a) *for each complex-type pair $\{\pi, \bar{\pi}\}$, an orthogonal complex structure on one multiplicity space $M_{n,\pi}$;*
- (b) *for each real-type self-conjugate block ρ , an orthogonal complex structure on the real multiplicity space $\text{Fix}(c_{n,\rho})$;*
- (c) *for each quaternionic-type self-conjugate block σ , a quaternionic-linear orthogonal complex structure on $M_{n,\sigma}$.*

In multiplicity one on a complex-type block, the only possibilities are the two sign choices from theorem 6.6. Thus the one-band sector is the multiplicity-one reduction of the general compact classification, while the two-band sector is its first genuinely operator-valued case.

Proof. This is the direct translation of [theorem 6.5](#) into the parameter-space language of [theorem 6.3](#). On complex-type pairs one chooses a skew-adjoint square root of -1 on one multiplicity space from each pair; on the self-conjugate blocks one obtains the real and quaternionic restrictions recorded above. $\square$

Remark 6.9 (Atomic simplification of the general canonical-complexification program). The general direct-integral classification of [\[13\]](#) involves measurable multiplicity fields, antiunitary transport systems, and a residual gauge-equivalence problem for those transports. None of that ambiguity survives on the dangerous cluster: the relevant representation is finite-dimensional and atomic, so the transport data are explicit and the Real commutant is already canonically determined up to finite-dimensional unitary equivalence.

Remark 6.10 (Failure of commutant rigidity and the meaning of leakage). In multiplicity one on a complex-type block, every commuting operator is automatically complex-linear for the chosen sign, which is why auxiliary leakage vanishes identically in the one-band sector. In higher multiplicity, however, an operator may commute with the compact symmetry yet fail to commute with the chosen J_n^{sh} . The anti-linear term $\Gamma_{n,j}$ is exactly this failure of commutant rigidity for the compressed observable $X_{n,j}$.

6.2 Block symbols and leakage control

Let $\mathcal{B}_n$ denote the oriented blocks selected by J_n^{sh} . For T in the commutant of U_n and J_n^{sh} , write $\sigma_{n,\alpha}(T)$ for its block symbol on the corresponding multiplicity-space factor of the Real-commutant decomposition in [theorem 6.5](#).

Definition 6.11 (Symbol defect and leakage size). For each generator G_j define the oriented symbol defect of $X_{n,j}$ relative to the shadow observable by

$$\Delta_{n,\alpha}(j) := \sigma_{n,\alpha}(X_{n,j}^{\text{lin}}) - \sigma_{n,\alpha}(Y_{n,j}),$$

and set

$$d_n(j) := \max_{\alpha \in \mathcal{B}_n} \|\Delta_{n,\alpha}(j)\|, \quad \ell_n(j) := \|\Gamma_{n,j}\|.$$

Theorem 6.12 (Cluster operator control by leakage and symbols). Assume [theorem 6.1](#) and fix an oriented shadow structure $J_n^{\text{sh}} \in X(U_n)$. Then for every scale n and every core generator G_j ,

$$\|X_{n,j} - Y_{n,j}\| \leq d_n(j) + \ell_n(j).$$

In particular, if $d_n(j) \rightarrow 0$ and $\ell_n(j) \rightarrow 0$, then $X_{n,j} - Y_{n,j} \rightarrow 0$ in operator norm on the dangerous cluster.

Proof. By [\(6.1\)](#),

$$X_{n,j} - Y_{n,j} = (X_{n,j}^{\text{lin}} - Y_{n,j}) + \Gamma_{n,j}.$$

Because $X_{n,j}$, $Y_{n,j}$, and J_n^{sh} all commute with the compact symmetry, the first summand belongs to the commutant of U_n and J_n^{sh} . By the block decomposition of [theorem 6.5](#), its operator norm is the maximum of the norms of its multiplicity-space symbols, namely $d_n(j)$. Add the leakage term $\Gamma_{n,j}$. $\square$

Corollary 6.13 (Ordered-exponential control on the dangerous cluster). *Fix a finite tuple $(j_1, \dots, j_m)$ and a compact set $K \subset \mathbb{R}^m$. There exists a constant $C_{K,m}$ such that for every scale n , every unit vector $\psi \in K_n$, and every $\lambda = (\lambda_1, \dots, \lambda_m) \in K$,*

$$\left| \left\langle \psi, e^{i\lambda_1 X_{n,j_1}} \dots e^{i\lambda_m X_{n,j_m}} \psi \right\rangle - \left\langle \psi, e^{i\lambda_1 Y_{n,j_1}} \dots e^{i\lambda_m Y_{n,j_m}} \psi \right\rangle \right| \leq C_{K,m} \sum_{r=1}^m (d_n(j_r) + \ell_n(j_r)).$$

Proof. For bounded self-adjoint operators A, B and real λ one has

$$\left\| e^{i\lambda A} - e^{i\lambda B} \right\| \leq |\lambda| \|A - B\|.$$

Apply this factor by factor to the ordered product and use [theorem 6.12](#). □

6.3 From oriented-shadow rigidity to Yang–Mills matching

Assumption 6.14 (Vanishing leakage and shadow symbol convergence). For every explicit core generator G_j ,

$$\ell_n(j) \rightarrow 0, \quad d_n(j) \rightarrow 0 \quad (n \rightarrow \infty).$$

Assumption 6.15 (Yang–Mills block-symbol convergence). For every explicit core generator G_j there exists a family of Yang–Mills block symbols $a_\alpha^{\text{YM}}(j)$, compatible with the Frobenius–Schur type of each oriented block, such that

$$\max_{\alpha \in \mathcal{B}_n} \left\| \sigma_{n,\alpha}(Y_{n,j}) - a_\alpha^{\text{YM}}(j) \right\| \rightarrow 0.$$

Assume also the shared analyticity and Osterwalder–Schrader regularity hypotheses required by the interface theorems of [Section 5](#).

Theorem 6.16 (Auxiliary invisibility from oriented-shadow rigidity). *Assume [theorems 6.1](#) and [6.14](#). Then the projected continuum theory is auxiliary invisible on the explicit countable Wilson-loop/curvature core. Equivalently, for every finite tuple from the explicit core the projected and shadow continuum cylinder functionals agree on a neighborhood of the origin.*

Proof. Low-window dominance reduces the continuum cylinder functionals to cluster expectations. [Theorem 6.13](#) shows that the cluster difference tends to zero uniformly on compact parameter sets because the symbol defects and leakage sizes vanish. Passing to the continuum limits gives equality of the projected and shadow cylinder functionals on the explicit core. □

Theorem 6.17 (Shadow-to–Yang–Mills matching from block symbols). *Assume [theorems 6.1](#) and [6.15](#). Then the shadow continuum theory and the Yang–Mills target theory agree on the explicit countable Wilson-loop/curvature core.*

Proof. Within the commutant of U_n and J_n^{sh} each shadow observable is determined blockwise by its multiplicity-space symbols. The block-symbol convergence hypothesis therefore identifies the finite-dimensional shadow model with the Yang–Mills block model on every oriented block. Continuity of ordered exponentials in finitely many matrix coefficients and the shared analyticity package then imply equality of the ordered cylinder functionals on the core. □

Theorem 6.18 (Intrinsic replacement theorem). *Assume [theorems 6.16](#) and [6.17](#). Then auxiliary invisibility and Yang–Mills matching are encoded intrinsically by compact dangerous-cluster symmetry, a shadow complex structure $J_n^{\text{sh}} \in X(U_n)$ selected inside the Real commutant, vanishing anti-linear leakage, and blockwise Yang–Mills symbol convergence. Consequently, once the closed-sector and finite-band hypotheses are verified, the Euclidean and operator-algebraic interface theorems of [Section 5](#) identify the projected continuum theory with standard pure Yang–Mills theory.*

Proof. The two theorems cited provide exactly the two pieces required by [theorem 5.14](#): auxiliary invisibility on the explicit core and matching of the shadow core with the Yang–Mills target. The interface theorems then apply verbatim. $\square$

7 Closure of the one-band and two-band Wilson-loop sectors

This section incorporates the closure results of [\[11\]](#) into the comprehensive manuscript, but now interprets them through the compact Real-commutant classification of [Section 6](#). The one-band Wilson-loop subcore W_{loop}^{1b} is the multiplicity-one complex-type case: the relevant factor of the dangerous-cluster Real commutant is scalar, operator-valued orientation collapses to a sign, and leakage vanishes by commutant rigidity. The two-band Wilson-loop subcore W_{loop}^{2b} is the first genuinely non-scalar case, in which the compact factor is two-dimensional and the verification problem becomes an explicit matrix problem for operator-valued orientation and its anti-linear defect. In both cases the last observable obstruction is therefore a finite-dimensional Real-commutant calculation on the dangerous cluster itself rather than an external matching postulate.

7.1 One-band sector

Let W_{loop}^{1b} denote the one-band Wilson-loop subcore. The one-band island is the selected low-window component on which the projected vacuum concentrates and whose couplings to the complement are negligible in the sense of the imported one-band low-window package.

Theorem 7.1 (One-band low-window dominance). *For every finite tuple from W_{loop}^{1b} , the full scale-wise projected and shadow cylinder functionals are asymptotically governed by their one-band dangerous-cluster compressions. If the corresponding shadow one-band scalars converge, the one-band continuum cylinder limits follow as well.*

Proof. This is the one-band specialization of the low-window dominance theorem proved in the closure paper: one-band vacuum concentration and small off-band Wilson couplings allow a direct Duhamel/telescoping reduction of the full ordered products to the selected one-band compressions. The shadow scalar-limit input then upgrades the scale-wise reduction to the corresponding continuum cylinder limits. $\square$

Theorem 7.2 (One-band compact symmetry and Schur rigidity). *On the one-band island the dangerous cluster carries a continuous orthogonal compact-group representation under which all compressed projected and shadow Wilson-loop observables are equivariant. If the retained shadow block is of complex type with multiplicity one, every such compressed self-adjoint observable is scalar, the oriented shadow complex structure is rigid up to sign, and auxiliary leakage vanishes identically.*

Proof. Compact symmetry is inherited by restricting the ambient gauge action through the gauge-equivariant one-band selector. Once the retained block is complex type of multiplicity one, [theorem 6.6](#) shows that the shadow complex structure on that factor is uniquely determined up to the two sign choices $\pm i1$. The corresponding Real-commutant factor is scalar, so every commuting self-adjoint endomorphism is automatically J_n^{sh} -linear. Hence the anti-linear leakage term in the decomposition (6.1) disappears exactly. $\square$

Theorem 7.3 (One-band closure theorem). *On W_{loop}^{1b} the following items of the oriented-shadow verification package hold internally:*

- (i) *low-window dominance and the reduction of full cylinder functionals to one-band expectations;*
- (ii) *compact dangerous-cluster symmetry and gauge equivariance;*
- (iii) *exact vanishing of auxiliary leakage and vanishing shadow symbol defects;*
- (iv) *automatic existence of the oriented shadow complex structure and rigidity of the shadow symbols.*

Equivalently, on the one-band Wilson-loop subcore all verification items except the imported continuum-trajectory package and Yang–Mills scalar-symbol convergence are settled internally.

Proof. Combine [theorems 7.1](#) and [7.2](#) with the multiplicity-one calibration argument of the one-band closure analysis. The projected/shadow discrepancy collapses to a single scalar defect, auxiliary leakage vanishes exactly, and the one-band continuum cylinder functionals reduce to one-band expectations. This is precisely the closure statement proved on W_{loop}^{1b} . $\square$

7.2 Two-band sector

Let W_{loop}^{2b} denote the two-band Wilson-loop subcore. The pair island is the first genuinely coupled low-window sector and supports an explicit 2×2 oriented symbol calculus.

Theorem 7.4 (Two-band low-window dominance). *For every finite tuple from W_{loop}^{2b} , the full scale-wise projected and shadow cylinder functionals are asymptotically governed by their two-band dangerous-cluster compressions. If the relevant shadow matrix limits and normalized two-band vacuum coordinates converge, the corresponding two-band continuum cylinder limits follow.*

Proof. This is the coupled-band analogue of [theorem 7.1](#). Two-band vacuum concentration and small couplings to the complement reduce the full cylinder functionals to the selected pair island. The additional matrix-limit and vacuum-coordinate inputs then yield the continuum cylinder limits on that two-band subcore. $\square$

Theorem 7.5 (Two-band oriented symbols and leakage matrices). *On the two-band island - the first genuinely operator-valued orientation sector - the compressed projected and shadow Wilson-loop observables admit oriented Hermitian 2×2 multiplicity-space symbols together with complex-symmetric 2×2 anti-linear leakage matrices. Vanishing auxiliary leakage and vanishing shadow symbol defects reduce to explicit entrywise decay of finitely many matrix coefficients.*

Proof. Restrict the compact dangerous-cluster symmetry to the pair island and choose the oriented two-band shadow datum. In this first non-scalar complex-type case, the relevant Real-commutant factor is a 2×2 complex matrix algebra viewed as a real C^* -algebra. Relative to the chosen orientation, each compressed observable splits into a J_n^{sh} -linear Hermitian 2×2 symbol and a complex-symmetric anti-linear leakage matrix. The finite list of calibration coordinates arising from this decomposition controls both the symbol defect and the leakage norm, so entrywise decay yields the stated vanishing properties. $\square$

Theorem 7.6 (Two-band closure theorem). *On W_{loop}^{2b} the following items of the oriented-shadow verification package hold internally:*

- (i) *coupled-band low-window dominance and reduction of the full cylinder functionals to two-band expectations;*
- (ii) *compact dangerous-cluster symmetry and gauge equivariance on the pair island;*
- (iii) *oriented Hermitian 2×2 block symbols and complex-symmetric leakage matrices with vanishing auxiliary leakage and vanishing shadow symbol defects;*
- (iv) *intrinsic closure of the projected and shadow Wilson-loop data on the pair island.*

Equivalently, on the two-band Wilson-loop subcore all verification items except the imported continuum-trajectory package and Yang–Mills matrix-symbol convergence are settled internally.

Proof. Combine [theorems 7.4](#) and [7.5](#) with the explicit two-band calibration argument. The remaining discrepancy is reduced to finitely many coordinate defects and the two-band vacuum coordinates. Once those defects are shown to vanish, the projected and shadow two-band data close intrinsically on the selected pair island. $\square$

Theorem 7.7 (Yang–Mills symbol convergence on the closed subcores). *On W_{loop}^{1b} the projected and shadow scalar symbols of each loop converge to the same Yang–Mills scalar target. On W_{loop}^{2b} the projected and shadow Hermitian 2×2 symbols converge entrywise, hence in operator norm, to the same Yang–Mills target matrix. Consequently the projected, shadow, and Yang–Mills cylinder functionals agree on both closed Wilson-loop subcores.*

Proof. In the one-band case, [theorem 7.3](#) reduces the projected/shadow discrepancy to a single scalar calibration defect and makes auxiliary leakage trivial. Yang–Mills matching is therefore exactly the convergence of that scalar to the Yang–Mills target, and the ordered cylinder functionals are then exponentials of the corresponding limiting scalar sums. In the two-band case, [theorem 7.6](#) expresses each compressed observable through an oriented Hermitian 2×2 symbol plus a vanishing leakage matrix. Entrywise convergence of the projected and shadow symbols to the same Yang–Mills target matrix implies operator-norm convergence, and the same factor-by-factor exponential estimate used in [Section 6](#) upgrades this to equality of the resulting cylinder functionals on compact parameter sets. $\square$

7.3 Restricted Yang–Mills identification on the closed subcores

Theorem 7.8 (One-band and two-band subcore identification). *Assume the imported continuum-trajectory package of Section 4. Then the projected continuum theory, the shadow theory, and the Yang–Mills target theory agree on both closed Wilson-loop subcores W_{loop}^{1b} and W_{loop}^{2b} at the level of continuum cylinder functionals. In particular, the old external matching input is eliminated on those first closed subcores.*

Proof. The closure theorems 7.3 and 7.6 supply the intrinsic oriented-shadow verification items on the one-band and two-band subcores, while theorem 7.7 supplies Yang–Mills scalar/matrix-symbol convergence on exactly those same subcores. The oriented-shadow replacement theorem theorem 6.18 therefore applies on both W_{loop}^{1b} and W_{loop}^{2b} , yielding the stated restricted identification. $\square$

8 Finite-band exhaustion from band-effective data

This section incorporates the localization and finite-band synthesis results of [12]. The closed one-band and two-band subcores of Section 7 are now treated as the first atomic Real-commutant factors resolved explicitly. The remaining task is to show that an arbitrary finite Wilson-loop tuple only sees finitely many low spectral atoms of the dangerous cluster and that the omitted part is both spectrally separated and weakly coupled back to the retained blocks. In the rewritten language, finite-band exhaustion says that every finite tuple reduces asymptotically to a finite product of complex, real, and quaternionic matrix factors carried by the compact dangerous-cluster symmetry, rather than to an undifferentiated finite-dimensional tail. The crucial gain over the earlier drafts is that finite-band exhaustion and off-block decay are tied directly to the band-effective spectral data rather than postulated as an additional end-stage miracle.

8.1 Imported dangerous-cluster and band-effective data

Assumption 8.1 (Imported packetized low-window and band-effective package). There exists a renormalization trajectory $a_n \downarrow 0$ with scale-wise projected continuum theories and a transported dangerous-cluster family whose rank is eventually constant by theorem 4.11:

$$\mathcal{H}_n^{\text{proj}}, \quad \Omega_n^{\text{proj}} \in \mathcal{H}_n^{\text{proj}}, \quad K_n^{\text{cov}} \geq 0,$$

and a packetwise phase-stable low-window reduction producing projections

$$\Pi_{n,\varepsilon} \leq \mathbf{1}_{[0,\varepsilon]}(K_n^{\text{cov}}), \quad \dim \text{Ran } \Pi_{n,\varepsilon} \leq M_*,$$

of uniformly bounded rank $M_* < \infty$. Write

$$K_n := \text{Ran } \Pi_{n,\varepsilon} \subset \mathcal{H}_n^{\text{proj}}.$$

For every rational Wilson loop γ used below there are bounded self-adjoint projected and shadow dangerous-cluster observables

$$X_{n,\gamma} := \Pi_{n,\varepsilon} L_\gamma^{(n)} \Pi_{n,\varepsilon}|_{K_n}, \quad Y_{n,\gamma} := \Pi_{n,\varepsilon} \tilde{L}_\gamma^{(n)} \Pi_{n,\varepsilon}|_{K_n}.$$

In addition, the band-effective reduction of Section 3 provides, for each scale n :

- (i) a finite-dimensional model space E_n and a unitary map $W_n: E_n \rightarrow K_n$;
- (ii) a self-adjoint compact-resolvent effective operator M_n on E_n ;
- (iii) a low spectral cutoff $E_n^* > 0$ and the corresponding low model projection

$$Q_n := \mathbf{1}_{(-\infty, E_n^*]}(M_n),$$

with rank at most M_* ;

- (iv) a positive tail margin $r_n \geq r_* > 0$ separating the retained low spectrum $Q_n E_n$ from the discarded model tail $(I - Q_n)E_n$.

The one-band and two-band closure results of Section 7 are available on the corresponding closed Wilson-loop subcores.

Definition 8.2 (Spectral atoms and tuple couplings). Fix a finite tuple of rational Wilson loops $J = (\gamma_1, \dots, \gamma_m)$. Let

$$\text{spec}(M_n|_{Q_n E_n}) = \{\lambda_{n,\nu} : \nu \in V_n\}$$

be the distinct low eigenvalues of the effective model below E_n^* , and let $q_{n,\nu}$ be the corresponding spectral projections on E_n . The associated dangerous-cluster spectral atoms are

$$P_{n,\nu} := W_n q_{n,\nu} W_n^*, \quad \nu \in V_n.$$

For distinct $\nu, \mu \in V_n$, define the projected/shadow coupling amplitude

$$c_{n,\nu\mu}(J) := \max_{1 \leq r \leq m} \left\{ \|P_{n,\nu} X_{n,\gamma_r} P_{n,\mu}\|, \|P_{n,\nu} Y_{n,\gamma_r} P_{n,\mu}\| \right\},$$

and the effective spectral gap

$$g_{n,\nu\mu} := |\lambda_{n,\nu} - \lambda_{n,\mu}|.$$

For a vertex ν , define its vacuum weight

$$\omega_{n,\nu} := \left\| P_{n,\nu} \Omega_n^{\text{proj}} \right\|^2.$$

8.2 Tuple-adapted resonance graphs and finite-band covers

Definition 8.3 (Tuple-adapted resonance graph). Fix $0 < \kappa_J < 1/4$ and define the tuple-dependent threshold

$$\tau_n(J) := \kappa_J \min \left\{ r_*, \inf_{\nu \neq \mu} \{g_{n,\nu\mu} : c_{n,\nu\mu}(J) > 0\} \right\} + \sup_{\nu \neq \mu} c_{n,\nu\mu}(J).$$

Let $\rho_n(J) \downarrow 0$ be a relevance threshold. The tuple-adapted resonance graph associated with J has vertex set V_n and edge set determined by

$$\{\nu, \mu\} \in E_n(J) \iff c_{n,\nu\mu}(J) > \frac{1}{2} \tau_n(J).$$

Write $\mathcal{C}_n(J)$ for the set of connected components of this graph. A connected component $C \in \mathcal{C}_n(J)$ is called *vacuum relevant* if

$$\|P_{n,C}\Omega_n^{\text{proj}}\| > \rho_n(J), \quad P_{n,C} := \sum_{\nu \in C} P_{n,\nu}.$$

Assumption 8.4 (Pair-isolation and model-tail estimate). For every finite tuple J there exists n_J such that for all $n \geq n_J$ and all distinct $\nu, \mu, \sigma \in V_n$:

(i) if $c_{n,\nu\mu}(J) > \frac{1}{2}\tau_n(J)$, then

$$\max\{c_{n,\nu\sigma}(J), c_{n,\mu\sigma}(J)\} \leq \frac{1}{2}\tau_n(J);$$

(ii) if $c_{n,\nu\mu}(J) > \frac{1}{2}\tau_n(J)$, then the corresponding model gap is small only for that pair,

$$g_{n,\nu\mu} \leq \Gamma_J \tau_n(J), \quad \min\{g_{n,\nu\sigma}, g_{n,\mu\sigma}\} \geq 4\Gamma_J \tau_n(J),$$

for some tuple-dependent constant $\Gamma_J > 0$ independent of n ;

(iii) there is a model-tail quantity $t_n(J) \rightarrow 0$ such that

$$\|(I - W_n Q_n W_n^*)\Omega_n^{\text{proj}}\| \leq t_n(J).$$

Proposition 8.5 (Finite-band cover from band-effective data). Assume [theorems 8.1 and 8.4](#). Then for all sufficiently large n :

(i) every connected component of the tuple-adapted resonance graph has size at most two;

(ii) the set of vacuum-relevant components is finite, with

$$|\mathcal{C}_n^{\text{rel}}(J)| \leq \min\{M_*, \rho_n(J)^{-2}\};$$

(iii) the retained projections

$$\{P_{n,\alpha} : \alpha \in A_J\} := \{P_{n,C} : C \in \mathcal{C}_n^{\text{rel}}(J)\}$$

form a finite cover by one-band and two-band islands.

Proof. If a connected component contained three distinct vertices, it would contain a path of length two, say $\{\nu, \mu\}$ and $\{\mu, \sigma\}$. The first pair-isolation inequality then forces $c_{n,\mu\sigma}(J) \leq \frac{1}{2}\tau_n(J)$, contradicting the edge relation. Thus every connected component has size at most two. The bound on the number of vacuum-relevant components follows from orthogonality of the corresponding projections and the uniform rank bound $\dim K_n \leq M_*$. $\square$

Definition 8.6 (Finite-band cover attached to a tuple). Define the finite-band cover projection and finite-band remainder by

$$P_{n,J}^{\text{fin}} := \sum_{\alpha \in A_J} P_{n,\alpha}, \quad R_{n,J}^{\text{fin}} := \Pi_{n,\varepsilon} - P_{n,J}^{\text{fin}}.$$

The omitted components are exactly those connected components of the resonance graph that are not vacuum relevant at the threshold $\rho_n(J)$.

Theorem 8.7 (Finite-band exhaustion from band-effective data). *Assume [theorems 8.1](#) and [8.4](#). Then for every finite Wilson-loop tuple $J = (\gamma_1, \dots, \gamma_m)$ and all sufficiently large n :*

- (i) *each retained block in the finite-band cover is either a one-band island or a two-band island inside $\Pi_{n,\varepsilon}$;*
- (ii) *the vacuum remainder obeys the explicit estimate*

$$\left\| R_{n,J}^{\text{fin}} \Omega_n^{\text{proj}} \right\| \leq t_n(J) + \sqrt{M_*} \rho_n(J);$$

- (iii) *every omitted connected component carries vacuum mass at most $\rho_n(J)^2$;*
- (iv) *the cover depends only on the tuple J and the thresholds $\tau_n(J)$ and $\rho_n(J)$.*

In particular, every finite Wilson-loop tuple is asymptotically supported on finitely many one-band and two-band islands.

Proof. The first item is exactly the size bound from [theorem 8.5](#). For the remainder, split

$$R_{n,J}^{\text{fin}} \Omega_n^{\text{proj}} = (\Pi_{n,\varepsilon} - W_n Q_n W_n^*) \Omega_n^{\text{proj}} + \sum_{C \notin \mathcal{C}_n^{\text{rel}}(J)} P_{n,C} \Omega_n^{\text{proj}}.$$

The first term is bounded by $t_n(J)$ from [theorem 8.4](#). For the second term, orthogonality gives

$$\left\| \sum_{C \notin \mathcal{C}_n^{\text{rel}}(J)} P_{n,C} \Omega_n^{\text{proj}} \right\|^2 = \sum_{C \notin \mathcal{C}_n^{\text{rel}}(J)} \left\| P_{n,C} \Omega_n^{\text{proj}} \right\|^2 \leq M_* \rho_n(J)^2,$$

because each omitted component contributes at most $\rho_n(J)^2$ and there are at most M_* such components. $\square$

Remark 8.8 (Artin–Wedderburn reading of the finite-band cover). Because every retained block inherits compact symmetry, the finite-band cover decomposes the tuple into finitely many atomic Real-commutant factors of complex, real, and quaternionic type. The one-band and two-band closures are therefore the first matrix-sized instances of the general compact classification from [Section 6](#), and finite-band exhaustion says that every finite Wilson-loop word asymptotically sees only finitely many such factors.

Remark 8.9 (Intrinsic temporal-RG reading of the finite-band package). The shorter continuum note [\[18\]](#) proposed integrating out odd auxiliary Fourier modes as an intrinsic temporal renormalization step. The present manuscript uses a more conservative and gauge-covariant version of the same idea. After diagonalizing the twisted auxiliary fibers, the low-window problem is organized not by raw Fourier labels but by band-effective spectral atoms and tuple-adapted resonance graphs. Passing from the full low window $\Pi_{n,\varepsilon}$ to the finite-band cover $P_{n,J}^{\text{fin}}$ should therefore be read as the rigorous replacement of that temporal decimation heuristic: one discards only those low spectral atoms that are both vacuum-irrelevant and weakly coupled to the chosen Wilson-loop tuple, and the quantitative error is tracked explicitly by $t_n(J)$, $\rho_n(J)$, and the off-block defects.

8.3 Explicit off-block decay and reduction to closed sectors

Definition 8.10 (Total off-block defects). Fix a tuple $J = (\gamma_1, \dots, \gamma_m)$ and its finite-band cover. For each loop γ_r define

$$\eta_{n,J}^{\text{proj}}(\gamma_r) := \sum_{\alpha \neq \beta} \|P_{n,\alpha} X_{n,\gamma_r} P_{n,\beta}\|, \quad \eta_{n,J}^{\text{sh}}(\gamma_r) := \sum_{\alpha \neq \beta} \|P_{n,\alpha} Y_{n,\gamma_r} P_{n,\beta}\|.$$

Proposition 8.11 (Explicit off-block decay estimate). Assume [theorems 8.1](#) and [8.4](#). Fix a tuple J and its finite-band cover. Then for every loop γ_r in J ,

$$\eta_{n,J}^{\text{proj}}(\gamma_r) + \eta_{n,J}^{\text{sh}}(\gamma_r) \leq 2M_*^2 \tau_n(J).$$

In particular,

$$\eta_{n,J}^{\text{proj}}(\gamma_r) \rightarrow 0, \quad \eta_{n,J}^{\text{sh}}(\gamma_r) \rightarrow 0.$$

Proof. Distinct selected blocks come from distinct connected components of the resonance graph, so no vertex in one selected block is connected by an edge to a vertex in another. By definition of the edge set, every corresponding projected or shadow cross-block matrix element is therefore at most $\frac{1}{2}\tau_n(J)$. There are at most M_*^2 ordered pairs of selected blocks, and summing these bounds yields the stated estimate. $\square$

Corollary 8.12 (Ordered-word stability on the finite-band cover). Let $J = (\gamma_1, \dots, \gamma_m)$ be a finite tuple of Wilson-loop generators and let $K \subset \mathbb{R}^m$ be compact. There exists a constant $C_{J,K} > 0$ such that

$$\sup_{\lambda \in K} \left| \Xi_{J,n}^{\text{DTT}}(\lambda) - \Xi_{J,n}^{\text{DTT,blk}}(\lambda) \right| + \sup_{\lambda \in K} \left| \Xi_{J,n}^{\text{sh}}(\lambda) - \Xi_{J,n}^{\text{sh,blk}}(\lambda) \right| \leq C_{J,K} \left(t_n(J) + \sqrt{M_*} \rho_n(J) + \sum_{r=1}^m \eta_{n,J}^{\text{proj}}(\gamma_r) + \sum_{r=1}^m \eta_{n,J}^{\text{sh}}(\gamma_r) \right),$$

where $\Xi_{J,n}^{\text{DTT,blk}}$ and $\Xi_{J,n}^{\text{sh,blk}}$ denote the corresponding block-diagonal finite-band cylinder functionals.

Proof. Split the vacuum as $\Omega_n^{\text{proj}} = P_{n,J}^{\text{fin}} \Omega_n^{\text{proj}} + R_{n,J}^{\text{fin}} \Omega_n^{\text{proj}}$ and use [theorem 8.7](#) to bound the vacuum remainder. On the retained finite-band sector, decompose each compressed Wilson-loop observable into its block-diagonal part plus its off-block part and telescope the ordered product. Each off-block replacement is controlled by [Proposition 8.11](#), and the standard bound $\|e^{i\lambda A} - e^{i\lambda B}\| \leq |\lambda| \|A - B\|$ converts the one-factor estimates into the displayed ordered-word estimate. $\square$

Corollary 8.13 (Tupewise localization of the dangerous cluster). For every finite Wilson-loop tuple J , the part of the dangerous cluster seen by J is, up to the errors controlled by [theorems 8.7](#) and [8.12](#), a finite direct sum of one-band and two-band components. Thus each finite observable probe accesses only a finite-dimensional localized sector of the dangerous cluster.

Proof. Combine the finite-band cover of [theorem 8.7](#) with the block-diagonal reduction of [theorem 8.12](#). The retained finite-band sector is a finite direct sum of one-band and two-band islands, and the omitted part contributes only the vanishing remainder terms already estimated there. $\square$

Corollary 8.14 (Off-block decay package for finite-band reduction). Every finite Wilson-loop tuple from the explicit core is asymptotically governed, at the level of ordered cylinder functionals, by finitely many one-band and two-band blocks indexed by A_J , with total error controlled by the

right-hand side of [theorem 8.12](#). In particular, the finite-band cover is admissible for the one-band and two-band closure theorems of [Section 7](#).

Proof. This is a direct restatement of [theorem 8.12](#) together with the fact that every retained block has size one or two by [theorem 8.7](#). Those are exactly the cases handled by the closed-sector theorems. $\square$

Theorem 8.15 (Reduction of finite Wilson-loop tuples to closed sectors). *Assume [theorem 7.8](#) together with the standing hypotheses of the present section. Then every finite tuple of Wilson-loop generators from the explicit core is asymptotically governed by finitely many already-closed one-band and two-band sectors.*

Proof. By [theorems 8.7](#) and [8.14](#), the ordered products attached to the tuple reduce to a finite direct sum of one-band and two-band block-diagonal sectors, up to the explicit vacuum-remainder and off-block errors estimated above. By [theorem 7.8](#), each retained block is already matched with the Yang–Mills target on the corresponding closed subcore. Summing over the finitely many retained blocks yields the reduction theorem. $\square$

9 Constructive curvature approximation by Wilson-loop polynomials

This section incorporates the constructive curvature-extension results of [\[12\]](#). At this stage the Wilson-loop side of the late-stage program is already closed on one-band and two-band islands and has been propagated to arbitrary finite Wilson-loop tuples through finite-band exhaustion. What remains is to treat bounded curvature generators on the same explicit core. The point of the present section is that this step is now constructive rather than schematic: small-loop reconstruction, explicit polynomial/functional-calculus bounds, and a concrete approximation schedule produce band-adapted Wilson-loop approximants compatible with every finite-band cover.

9.1 Admissible bounded curvature generators and finite-band compatibility

Definition 9.1 (Band-adapted Wilson-loop approximation scheme). A sequence $Q_{N,j}$ of finite self-adjoint $*$ -polynomials in rational Wilson loops is a *band-adapted Wilson-loop approximation scheme* for a bounded curvature generator C_j if:

- (i) $Q_{N,j} \rightarrow C_j$ in the seminorm topology of the explicit core algebra W ;
- (ii) for every finite-band cover associated with any finite tuple containing C_j ,

$$\sup_n \max_{\alpha \in A_j} \max_{\bullet \in \{\text{DTT, sh, YM}\}} \left\| P_{n,\alpha}(Q_{N,n,j}^\bullet - C_{n,j}^\bullet) P_{n,\alpha} \right\| \rightarrow 0;$$

- (iii) the corresponding ordered cylinder functionals are approximated uniformly on compact parameter sets when C_j is replaced by $Q_{N,j}$.

Assumption 9.2 (Admissible precursor/common-core package). Every bounded curvature generator C_j on the explicit countable core is of the form $C_j = f_j(A_j)$, where A_j is a self-adjoint gauge-invariant precursor on the common core and f_j is one of the bounded symbols used in the program,

in particular either an exponential $f_s(u) = e^{isu}$ or the bounded transform $\chi(u) = u(1 + u^2)^{-1/2}$. The common-core regularity and locality assumptions needed for small-loop reconstruction hold uniformly in the projected, shadow, and Yang–Mills models.

Definition 9.3 (Small rational plaquettes and explicit Wilson precursors). Fix a rational side length $h > 0$. For a rational base point x and directions $\mu < \nu$, let $\square_{\mu\nu}(x; h)$ be the positively oriented rational plaquette of side length h based at x , and let $W_{\mu\nu}(x; h)$ be its renormalized Wilson holonomy. Define the anti-Hermitian plaquette defect

$$D_{\mu\nu,h}(x) := \frac{W_{\mu\nu}(x; h) - W_{\mu\nu}(x; h)^*}{2ih^2},$$

and for a direction ρ the centered finite difference

$$\Delta_{\rho,h}B(x) := \frac{B(x + he_\rho) - B(x - he_\rho)}{2h}.$$

The explicit Wilson precursor $A_{h,j}$ is obtained by replacing each curvature component in A_j by the corresponding plaquette defect, each derivative by the corresponding centered finite difference, and each spatial integral by the rational Riemann sum on mesh $h\mathbb{Q}^3$. The result is a finite self-adjoint $*$ -polynomial in rational Wilson loops.

Proposition 9.4 (Local small-loop reconstruction). *For every admissible precursor there exists a constant $c_j^{\text{loc}} < \infty$ such that, uniformly on the common core and uniformly in the projected, shadow, and Yang–Mills models,*

$$\|D_{\mu\nu,h}(x) - F_{\mu\nu}(x)\| \leq c_j^{\text{loc}}h, \quad \|\Delta_{\rho,h}D_{\mu\nu,h}(x) - D_\rho F_{\mu\nu}(x)\| \leq c_j^{\text{loc}}h,$$

for all sufficiently small rational h .

Proof. Expand the plaquette holonomy around the base point by the nonabelian Taylor/Baker–Campbell–Hausdorff formula. The leading term is the curvature, the next term is linear in h , and gauge-covariant finite differences recover the derivative terms with the same order of accuracy. The resulting estimates are exactly the local precursor estimates needed for the constructive curvature step. $\square$

Proposition 9.5 (Finite-band precursor estimate). *For every admissible bounded curvature generator $C_j = f_j(A_j)$ there exists a constant $C_j < \infty$ such that for the Wilson precursors $A_{h,j}$,*

$$\sup_n \max_{\alpha \in A_J} \max_{\bullet \in \{\text{DTT, sh, YM}\}} \|P_{n,\alpha}(A_{h,n,j}^\bullet - A_{n,j}^\bullet)P_{n,\alpha}\| \leq C_j h$$

on every finite-band cover associated with a finite tuple containing C_j .

Proof. Combine the local small-loop reconstruction from Proposition 9.4 with the finite-band localization of Section 8. Each retained block is either one-band or two-band, so the blockwise compression of the precursor error is controlled by the same $O(h)$ estimate uniformly across the finite-band cover. $\square$

9.2 Constructive bounded functional calculus

For the two principal bounded symbols one may choose explicit polynomial approximants.

- (i) For $f_s(u) = e^{isu}$, use the truncated Taylor polynomial

$$p_{m,s}^{\text{exp}}(u) := \sum_{q=0}^m \frac{(isu)^q}{q!}.$$

If $\sigma(A_j) \subset [-M_j, M_j]$, the remainder satisfies

$$E_{m,j}^{\text{exp}} \leq e^{|s|M_j} \frac{(|s|M_j)^{m+1}}{(m+1)!}.$$

- (ii) For $\chi(u) = u(1+u^2)^{-1/2}$, let p_m^χ be the degree- m Chebyshev approximant on $[-M_j, M_j]$. For every $1 < \rho < \rho_\chi(M_j)$,

$$E_{m,j}^\chi \leq \frac{2M_\rho(\chi, M_j)}{\rho - 1} \rho^{-m},$$

where $M_\rho(\chi, M_j)$ is the maximum of $|\chi(M_j z)|$ on the Bernstein ellipse E_ρ .

Define $E_{m,j}$ to be the relevant approximation error for the chosen symbol, and let $L_{m,j}$ denote the Lipschitz constant of the polynomial approximant p_m on the relevant spectral interval.

Theorem 9.6 (Constructive band-adapted approximation theorem). *Let $C_j = f_j(A_j)$ be an admissible bounded curvature generator. Choose a rational mesh sequence $h_N \downarrow 0$ and polynomial degrees $m_N \uparrow \infty$ such that*

$$2E_{m_N,j} + L_{m_N,j} C_j h_N \leq 2^{-N},$$

and define

$$Q_{N,j} := p_{m_N}(A_{h_N,j}).$$

Then:

- (i) $Q_{N,j} \rightarrow C_j$ in the seminorm topology of W ;
(ii) for every finite-band cover associated with a finite tuple containing C_j ,

$$\sup_n \max_{\alpha \in A_j} \max_{\bullet \in \{\text{DTT}, \text{sh}, \text{YM}\}} \|P_{n,\alpha}(Q_{N,n,j}^\bullet - C_{n,j}^\bullet)P_{n,\alpha}\| \leq 2^{-N};$$

- (iii) for every compact interval $[-\Lambda, \Lambda]$,

$$\sup_{|\lambda| \leq \Lambda} \sup_n \max_{\alpha \in A_j} \max_{\bullet \in \{\text{DTT}, \text{sh}, \text{YM}\}} \|P_{n,\alpha}(e^{i\lambda Q_{N,n,j}^\bullet} - e^{i\lambda C_{n,j}^\bullet})P_{n,\alpha}\| \leq \Lambda 2^{-N}.$$

In particular, the sequence $Q_{N,j}$ is a band-adapted Wilson-loop approximation scheme for C_j .

Proof. The blockwise precursor estimate from Proposition 9.5 contributes the term $L_{m_N,j} C_j h_N$, while the polynomial approximation of the bounded symbol contributes $E_{m_N,j}$. The defining choice of h_N and m_N therefore gives the displayed 2^{-N} blockwise estimate. Seminorm convergence in W

is proved by the same small-loop and polynomial approximation bounds on the common core. The one-factor exponential estimate follows from the elementary bound $\|e^{i\lambda A} - e^{i\lambda B}\| \leq |\lambda| \|A - B\|$ for bounded self-adjoint operators. $\square$

Corollary 9.7 (Mixed-word stability under curvature replacement). *Assume [theorems 8.7](#) and [9.6](#). Let J be a finite tuple from the explicit core allowing both Wilson-loop and bounded curvature generators. Replacing each curvature slot by its band-adapted Wilson polynomial approximant changes the corresponding ordered cylinder functional by a quantity that can be made arbitrarily small uniformly in n on compact parameter sets.*

Proof. Replace one curvature slot at a time and telescope the ordered product. Each replacement is controlled by the one-factor exponential estimate of [theorem 9.6](#). Since the tuple length is finite and the parameter set is compact, the total error is bounded by a finite sum of terms tending to zero. $\square$

10 Final synthesis: full-core Yang–Mills identification and gap transfer

This section incorporates the synthesis results of [\[12\]](#). The late-stage chain is now explicit and continuous: the closed one-band and two-band sectors of [Section 7](#) are glued through the finite-band cover of [Section 8](#), so arbitrary finite tuples are reduced to finitely many atomic Real-commutant factors; bounded curvature generators are inserted by the constructive approximation scheme of [Section 9](#); and the resulting full-core matching is upgraded to vacuum-theory identification by the density and extension theorems of [Section 5](#). Within the modular framework of the program, there is no longer any additional end-stage matching or approximation input beyond these visible packages.

10.1 Bandwise gluing on finite-band sectors

Theorem 10.1 (Bandwise gluing on finite-band sectors). *Assume the standing hypotheses of [Sections 3](#) and [4](#). Let $J = (\gamma_1, \dots, \gamma_m)$ be a finite tuple of Wilson-loop generators from the explicit core. Then the projected, shadow, and Yang–Mills cylinder functionals associated with J agree on compact parameter sets after passage to the finite-band cover. Equivalently, finite-band Wilson-loop tuples are reconstructed from the already-closed one-band and two-band sectors carried by the components in the finite-band cover.*

Proof. By [theorem 8.15](#), the ordered products attached to J are asymptotically equivalent to a finite direct sum of ordered products on one-band and two-band sectors. On each retained block, [theorem 7.8](#) gives equality of the projected, shadow, and Yang–Mills cylinder functionals. Summing over the finitely many retained components and using [theorem 8.12](#) to control the vacuum remainder and off-block errors gives the claimed bandwise gluing statement. $\square$

10.2 Curvature extension and mixed-word matching

Proposition 10.2 (Approximation transfer on a finite-band cover). *Assume [theorems 8.7](#) and [9.6](#). Let C_j be a bounded curvature generator and let $Q_{N,j}$ be the corresponding band-adapted Wilson-loop*

approximation sequence. Then for every finite-band cover associated with a finite tuple containing C_j ,

$$\sup_n \max_{\alpha \in A_J} \left\| P_{n,\alpha} (Q_{N,n,j}^\bullet - C_{n,j}^\bullet) P_{n,\alpha} \right\| \rightarrow 0$$

for $\bullet \in \{\text{DTT}, \text{sh}, \text{YM}\}$, and the same convergence holds for one-factor exponentials uniformly on compact parameter sets.

Proof. This is the direct blockwise content of [theorem 9.6](#). The one-factor exponential convergence follows from the same bounded-operator perturbation estimate used in [Section 9](#). $\square$

Theorem 10.3 (Curvature extension theorem to the full explicit core). *Assume [theorems 9.6](#) and [10.1](#). Let $J = (j_1, \dots, j_m)$ be any finite tuple from the explicit core, allowing both Wilson-loop and bounded curvature generators. Then the projected, shadow, and Yang–Mills cylinder functionals for J agree asymptotically on compact parameter sets. Equivalently, the matching proved on the one-band and two-band Wilson-loop subcores extends to all bounded mixed words in the explicit Wilson-loop/curvature core.*

Proof. Replace each bounded curvature slot by its band-adapted Wilson polynomial approximant and apply [theorem 9.7](#). The resulting tuple is a finite Wilson-loop tuple, so [theorem 10.1](#) applies. Finally remove the approximation using [Proposition 10.2](#). Since the number of curvature slots is finite, the total perturbation can be made arbitrarily small uniformly on compact parameter sets. $\square$

Corollary 10.4 (Full-core word matching). *Under the hypotheses of [theorem 10.3](#), the projected continuum DTT, shadow, and Yang–Mills cylinder functionals agree on every finite word in the explicit Wilson-loop/curvature core. In particular,*

$$S^{\text{DTT}}(A) = S^{\text{sh}}(A) = S^{\text{YM}}(A) \quad (A \in W_0^{\text{wc}}).$$

Proof. Apply [theorem 10.3](#) to the ordered cylinder functionals and differentiate at the origin using the shared analyticity hypothesis. $\square$

10.3 Final identification and gap transfer

Theorem 10.5 (Full-core auxiliary invisibility and Yang–Mills matching). *Assume [theorem 10.3](#). Then the projected continuum theory is auxiliary invisible on the full explicit Wilson-loop/curvature core, and the shadow core agrees with the Yang–Mills target core on the same set. Equivalently,*

$$S^{\text{DTT}}(A) = S^{\text{YM}}(A) \quad (A \in W_0^{\text{wc}}).$$

Proof. The statement is exactly the content of [theorem 10.4](#), interpreted simultaneously as auxiliary invisibility of the projected continuum theory and as agreement of the shadow core with the Yang–Mills target core. $\square$

Theorem 10.6 (Final full-core identification theorem). *Assume [theorem 10.5](#) and the countable-core continuity package of [theorem 5.1](#). Then:*

- (i) *the Euclidean functionals S^{DTT} and S^{YM} agree on the completed common algebra generated by the explicit core;*

- (ii) the Osterwalder–Schrader Hilbert spaces reconstructed from those functionals are unitarily equivalent by a vacuum-preserving unitary U ;
- (iii) the corresponding Hamiltonians are intertwined by the same unitary,

$$UH_{\text{cont}}^{\text{DTT}}U^{-1} = H_{\text{YM}}, \quad U\Omega_{\text{cont}}^{\text{DTT}} = \Omega_{\text{YM}}.$$

In particular, the projected continuum theory and standard pure Yang–Mills theory are unitarily equivalent in their vacuum representations.

Proof. By [theorem 10.5](#), the projected and Yang–Mills Euclidean data agree on the explicit countable Wilson-loop/curvature core. The density theorems of [Section 5](#) extend that equality to the completed Euclidean and operator-algebraic local structures. Applying [theorems 5.11](#) and [5.13](#) to the extended equality gives the vacuum-preserving unitary equivalence and the Hamiltonian intertwining. $\square$

Corollary 10.7 (Transfer of the projected continuum spectral gap). *Assume the hypotheses of [theorem 10.6](#) and the projected continuum lower bound*

$$H_{\text{cont}}^{\text{DTT}} \geq \delta \quad \text{on } \mathcal{H}_{\text{cont}}^{\text{DTT}} \ominus \mathbb{C}\Omega_{\text{cont}}^{\text{DTT}}$$

for some $\delta > 0$. Then the Yang–Mills Hamiltonian satisfies

$$H_{\text{YM}} \geq \delta \quad \text{on } \mathcal{H}_{\text{YM}} \ominus \mathbb{C}\Omega_{\text{YM}}.$$

Proof. The Hamiltonians are unitarily equivalent by [theorem 10.6](#), so the projected continuum lower bound transfers unchanged to the Yang–Mills Hamiltonian. $\square$

Corollary 10.8 (Exponential Euclidean-time clustering from the transferred gap). *Assume the hypotheses of [theorem 10.7](#). Let A and B be bounded local operators in the Yang–Mills vacuum representation, and write*

$$A_0 := A - \langle \Omega_{\text{YM}}, A\Omega_{\text{YM}} \rangle \mathbf{1}, \quad B_0 := B - \langle \Omega_{\text{YM}}, B\Omega_{\text{YM}} \rangle \mathbf{1}.$$

Then for every $s \geq 0$,

$$\left| \langle \Omega_{\text{YM}}, A_0 e^{-sH_{\text{YM}}} B_0 \Omega_{\text{YM}} \rangle \right| \leq \|A_0 \Omega_{\text{YM}}\| \|B_0 \Omega_{\text{YM}}\| e^{-\delta s}.$$

In particular, the connected Euclidean-time two-point functions of bounded local observables decay exponentially with rate at least δ .

Proof. Because $B_0 \Omega_{\text{YM}} \perp \Omega_{\text{YM}}$, the spectral theorem together with [theorem 10.7](#) implies

$$\left\| e^{-sH_{\text{YM}}} B_0 \Omega_{\text{YM}} \right\| \leq e^{-\delta s} \|B_0 \Omega_{\text{YM}}\|.$$

Pairing this vector with $A_0 \Omega_{\text{YM}}$ gives the stated estimate. $\square$

Remark 10.9 (Program-level meaning after the present rewrite). Taken together with the earlier sections, the late-stage chain is now materially stronger than in the earlier comprehensive draft. One-band and two-band sectors are explicitly closed; finite-band exhaustion and off-block decay are derived

directly from the band-effective spectral data; bounded curvature generators admit constructive band-adapted Wilson-loop approximation schemes; and the final full-core identification/gap-transfer theorem follows by a transparent gluing-and-extension argument. Within the bounded-core framework, there is no remaining end-stage obstacle in the form of auxiliary invisibility, Yang–Mills matching, finite-band exhaustion, or curvature approximation on the explicit core itself.

Remark 10.10 (What remains if one wants a completely unconditional single-paper proof). Sections 7–10 now absorb the late-stage closure and synthesis results of Papers 4 and 5 into the comprehensive manuscript. What remains to be proved directly inside a single manuscript is therefore upstream and explicit: the projected-sector/common-core and coercivity inputs of Section 3, the scale-uniform trajectory / Mosco package of Section 4, and the concrete pair-isolation/model-tail and admissible-precursor estimates that trigger the localization and curvature machinery of Sections 8 and 9. No hidden terminal matching principle remains beyond those visible packages.

Status and scope boundary

The present manuscript is the single retained SMM/DTT-based rewrite of the projected lattice-gauge program. Earlier companion drafts are kept as proof-production history, while the broader SMM/DTT papers are now treated as part of the accepted background framework rather than as mere historical motivation. At the level of notation, theorem statements, logical dependencies, and proof-production order, the present text is intended to be readable on its own.

In particular, the notation DTT no longer functions only as a model label. It denotes the accepted dual-time background whose compact inner-time / extended outer-time architecture is represented at regulator level by the auxiliary cycle and the physical Euclidean-time reflection direction. Likewise, the SMM background is not used as a substitute for proof, but as the accepted foundational setting in which sequential renewal of effective space, the inner-time spectral tower, and the existence of a projected four-dimensional shadow sector are taken as genuine structural input [19, 15, 16].

The mathematical scope boundary nevertheless remains sharp. The argument itself is still formulated directly in terms of the lattice model, its mixed plaquette operator, the projected low-window sector, the dangerous cluster, the producer–consumer coercivity/transport layer, and the countable Wilson-loop/curvature core. The manuscript now contains a self-contained fixed-regulator backbone, an explicit theorem-target and interface-theorem formulation of the projected-gap and continuum/transport packages, a self-contained treatment of the Real-commutant/oriented-shadow layer, and a late-stage theorem chain in which the first closed subcores, finite-band localization, constructive curvature approximation, full-core matching, identification, and gap transfer are all recorded explicitly.

What remains assumption-driven is concentrated in six visible places:

- (i) the projected-sector/common-core theorem target of [theorem 3.1](#);
- (ii) the volume-uniform high-energy coercivity theorem target of [theorem 3.2](#);
- (iii) the thermodynamic embedding and vacuum-compatibility theorem target of [theorem 4.1](#);
- (iv) the scale-wise projected continuum trajectory theorem target of [theorem 4.6](#);

- (v) the common-comparison-space band realization and isolated-contour norm-resolvent transport theorem targets of [theorems 4.7](#) and [4.8](#);
- (vi) the explicit pair-isolation/model-tail and admissible-precursor estimates underlying [Sections 8](#) and [9](#) if one wishes to derive them here from first principles.

Within the bounded-core framework adopted throughout the paper, there is therefore no hidden late-stage frontier. The unresolved burden is exactly the finite list of visible analytic packages named above, while the closure, transport, localization, approximation, and full-core synthesis steps themselves are written as an integrated theorem chain inside the accepted SMM/DTTbackground.

The rewrite also makes the near-term theorem-level entry points explicit. The most tractable analytic target is the common-comparison-space stability package behind [theorem 4.11](#); the most realistic observable-level entry points are the one-band Wilson-loop subcore, a restricted local-curvature subcore, and the two-band resonant window singled out in [theorem 5.15](#). These targets do not close the full Yang–Mills proof by themselves, but they now sit inside the comprehensive manuscript as precise theorem-level problems rather than as diffuse programmatic aspirations.

A Reflection-positivity boundary factorization and transfer kernel

This appendix supplies two fixed-regulator ingredients repeatedly used in the main text.

A.1 Boundary factorization

Let the action be decomposed as

$$S = S_+ + \Theta S_+ + S_\partial,$$

where S_+ contains only positive-time plaquettes and S_∂ the plaquettes meeting the reflection plane. Each boundary plaquette weight admits a positive character expansion

$$e^{-\kappa \Re \operatorname{Tr}(I - \rho(W))} = \sum_{\lambda} c_{\lambda} \chi_{\lambda}(W), \quad c_{\lambda} \geq 0.$$

Grouping the boundary plaquettes and using Haar factorization across the reflection plane, one obtains

$$(\Theta F) F e^{-S} = \sum_{\beta} (\Theta(F e^{-S_+/2} G_{\beta})) (F e^{-S_+/2} G_{\beta}),$$

with absolute convergence in the sup norm. Integrating against the Haar product measure gives a sum of nonnegative squares, proving reflection positivity.

A.2 One-slab kernel and gauge covariance

The decomposition of the action into horizontal slice terms and one-slab interaction terms leads to the kernel [\(2.4\)](#) and the transfer operator [\(2.5\)](#). Positivity and symmetry of the kernel are inherited from positivity and reality of the Boltzmann factors, while gauge covariance follows because the slice gauge action acts simultaneously on both kernel arguments and leaves the slab action invariant. Compactness is a consequence of continuity of the kernel on the compact slice configuration space.

B Twisted Fourier calculation and a toy packet model

B.1 Twisted Fourier modes

Let $\mathcal{H}_\theta = \ell^2(\mathbb{Z}/N_\tau\mathbb{Z})$ with twisted periodicity $a(N_\tau) = e^{i\theta}a(0)$. The difference operator

$$(D_\theta a)(\tau) = \frac{a(\tau+1) - a(\tau)}{\delta_\tau}, \quad 0 \leq \tau \leq N_\tau - 2,$$

with closing condition

$$(D_\theta a)(N_\tau - 1) = \frac{e^{i\theta}a(0) - a(N_\tau - 1)}{\delta_\tau}$$

is diagonalized by the twisted Fourier basis

$$a_n(\tau) = \frac{1}{\sqrt{N_\tau}} \exp(i(2\pi n + \theta)\tau/N_\tau), \quad n = 0, \dots, N_\tau - 1.$$

The eigenvalues of $D_\theta^* D_\theta$ are

$$\frac{4}{\delta_\tau^2} \sin^2\left(\frac{2\pi n + \theta}{2N_\tau}\right).$$

Low modes therefore occur only when $2\pi n + \theta$ is small modulo 2π .

B.2 Packetwise toy phase-count bound

Suppose the phase parameter θ ranges in a compact set and is partitioned into packets of diameter $O(\sqrt{\varepsilon})$. Within one packet there are at most finitely many integers n such that

$$\frac{4}{\delta_\tau^2} \sin^2\left(\frac{2\pi n + \theta}{2N_\tau}\right) \leq \varepsilon.$$

Hence the total number of low modes is bounded by the number of packets times the maximal number of low modes per packet. This toy model is the abelian shadow of the packetwise phase-count mechanism used in the main text before Peter–Weyl refinement.

B.3 Toy dyadic decimation on the auxiliary circle

Assume N_τ is even and work in the untwisted scalar model of the previous subsection. Writing

$$a(\tau) = \sum_{n=0}^{N_\tau-1} \hat{a}_n a_n(\tau),$$

with respect to the Fourier basis a_n , define the even-mode decimation operator by

$$(\mathcal{R}a)(\tau) := \sum_{m=0}^{N_\tau/2-1} \hat{a}_{2m} \frac{1}{\sqrt{N_\tau/2}} e^{2\pi i m \tau / (N_\tau/2)}.$$

This simply discards the odd Fourier labels and rescales the circle length by a factor of two. In the toy model the low-energy shells of $D_\theta^* D_\theta$ are therefore naturally organized by dyadic mode blocks. The main text does not use $\mathcal{R}$ directly, because the gauge-covariant problem is twisted

and the relevant low sector is selected by holonomy-dependent spectral atoms rather than by raw Fourier labels. Nevertheless, the band-effective cutoffs, finite-band cover, and off-block estimates of Sections 3 and 8 should be read as the rigorous replacement of this toy temporal decimation heuristic.

References

- [1] K. G. Wilson, *Confinement of quarks*, Phys. Rev. D **10** (1974), 2445–2459.
- [2] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green’s functions. I*, Commun. Math. Phys. **31** (1973), 83–112.
- [3] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green’s functions. II*, Commun. Math. Phys. **42** (1975), 281–305.
- [4] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Ann. Physics **110** (1978), 440–471.
- [5] M. Haj Yousef, *Reflection Positivity and Covariant Inner-Time Structure in a Fixed-Regulator Lattice Gauge Model*, preprint, April 2026.
- [6] M. Haj Yousef, *Finite-Volume Spectral Gap via Low-Rank Resonance Reduction in a Projected Lattice Gauge Sector*, preprint, April 2026.
- [7] M. Haj Yousef, *High-Energy Coercivity, Thermodynamic Persistence, and Continuum Transfer in the Projected DTT Sector*, preprint, April 2026.
- [8] M. Haj Yousef, *Projected Gaps and Gap-Preserving Continuum Limits for the DTT Program*, preprint, March 2026.
- [9] M. Haj Yousef, *Projected Gaps from Uniform Band Models to the Continuum*, preprint, March 2026.
- [10] M. Haj Yousef, *Continuum Transfer, Core Density, and Reduction of Yang–Mills Identification to Auxiliary Invisibility*, preprint, April 2026.
- [11] M. Haj Yousef, *Auxiliary Invisibility via Band-Localized Closure and Symbol Convergence*, preprint, April 2026.
- [12] M. Haj Yousef, *Finite-Band Exhaustion, Curvature Extension, and Full Yang–Mills Identification*, preprint, April 2026.
- [13] M. Haj Yousef, *Canonical Complexification of Real Hilbert Representations*, preprint, April 2026.
- [14] M. Haj Yousef, *Nonabelian Spectral Orientation and Complex Structures Commuting with Compact Group Representations*, preprint, April 2026.
- [15] M. Haj Yousef, *A Semi-Discrete Yang–Mills Model on a Compact Manifold: Existence of Minimisers, Fluctuations, and Spectral Gap*, preprint, February 2026.

- [16] M. Haj Yousef, *Spectral Mass and Temporal Ontology: A Complex-Time Approach to Gauge Theory Foundations*, preprint.
- [17] M. Haj Yousef, *Mass Gap Generation in Yang–Mills Theory from Inner-Time Spectral Topology: A Duality of Time Approach*, Progress of Theoretical and Experimental Physics (submitted), 2025.
- [18] M. Haj Yousef, *Existence of Four-Dimensional Yang–Mills Theory with a Non-Zero Mass Gap via the Duality of Time Theory*, preprint, June 2025.
- [19] M. Haj Yousef, *Mathematical Foundations of the Single Monad Model and Duality of Time Theory: Compact Inner Time, Stabilized Events, Effective Spacetime Dynamics, and Lorentzian Geometry from Order, Volume, and Discrete Action*, preprint, April 2026.