hexagon:2611.00492v1

From Four-Dimensional Wilson Lattice Gauge Theory to Quantum Yang–Mills Theory: Construction with Proofs

from Certified Assumptions, and the Record of Its Independent Validation

Chunping Zeng

Recompiled with complete proofs from the validation document of the Yang–Mills Shepherding Committee (D. J. Gross, chair; M. R. Douglas, A. Jaffe, I. R. Klebanov, N. Nekrasov), with the readable re-proof and the review-and-supplement package incorporated; November 2026

Abstract

We present a complete formal paper compiling the Lean-verified construction of quantum Yang–Mills theory submitted on October 28, 2026, and independently validated by the Yang–Mills Shepherding Committee. Starting from the Wilson definition of lattice gauge theory on a four-dimensional torus, we give full proofs of every step that can be carried out from the definitions: the field-label algebra, the finite lattice theory, and—under an explicit, individually certified package of renormalization-group and analytic assumptions (RG1)–(RG5), (UV1), (PH1)–(PH4), (AN1), (Y1)–(Y4), (B1), (C1)—the fixed-volume ultraviolet limit, the subsequential infinite-volume limit, the phase decomposition with clustering and full support, the Osterwalder–Schrader reconstruction satisfying the Wightman axioms, and the Yang–Mills identification: the stress tensor and its Ward identity, the trace anomaly with the verified constant $3\dim_{\mathbb{R}}\mathfrak{g}/\pi^4$, the exclusion of a precise class of free-field descriptions, the one-loop running coupling, and the operator product expansion. The independent validation (clean rebuild, Lean kernel replay, Comparator runs, the verification of the OS axioms against an independent axiom system, and the lattice-law, Ward-identity, common-scale, short-distance, and OPE adapters) is recorded in the second part. Two errata to the source validation document are corrected in the text. The free-field exclusion is proved in the text at the corrected weight $r^8\log^2(1/r)$; the corresponding Lean certificates verify the original $\log^1$-weight formulation (Remark 9.15). The proofs of the Wightman axioms W0–W3 are complete relative to three classical OS–Wightman theorems, which are marked as external inputs where used. For the mass gap we prove two logical boundaries: an explicit OS model satisfies all sixteen registered inputs, with E4 clustering (equivalently, vacuum uniqueness) holding, yet its Hamiltonian has no spectral gap (Proposition efprop:nogap-model), and a mass gap is equivalent to a uniform exponential rate in the clustering axiom E4 (Proposition efprop:rate-gap-equiv); the sixteen registered inputs therefore do not imply a mass gap. A uniform mass gap requires new infrared input; the candidate inputs (MG1)–(MG4) are projections of a single volume-uniform rate (MG) (Proposition efprop:mg-single), which is not a receipt dischargeable inside the formalization (Proposition efprop:no-gap-predicate); the bridge from (MG2) to a continuum gap is proved (Proposition efprop:mg2-bridge), and the bridge input is simplified to a single volume-uniform rate in Theorem efthm:e-allcyc of Appendix efapp:gap-boundary. A positive mass gap, uniqueness of the infinite-volume limit, uniqueness of the vacuum of the mixed theory, and order-by-order agreement of the OPE with perturbation theory are not proved, and the paper records precisely what is and is not proved at every step.

1 Introduction

1.1 The problem

The millennium prize problem of Jaffe and Witten [JW] asks for the construction of four-dimensional quantum Yang–Mills theory: a relativistic quantum field theory whose local field content is given by gauge-invariant polynomials in the curvature F of a nonabelian connection and its covariant derivatives, satisfying the standard axiomatic frameworks, with a positive mass gap. We use the following standard formulation.

Definition 1.1 (Wightman theory)**.** A *Wightman theory* consists of a complex Hilbert space $\mathcal{H}$, a unit vector Ω, a dense invariant domain $\mathcal{D} \subset \mathcal{H}$, a strongly continuous unitary representation $U(a, \Lambda)$ of the proper orthochronous Poincaré group, and linear maps $v \mapsto \Phi_v(f)$, $f \in \mathcal{S}(\mathbb{R}^4; \mathbb{C})$, $\Phi_v(f) : \mathcal{D} \to \mathcal{D}$, such that:

W0. Spectrum and uniqueness of the vacuum. Translations satisfy $U(a, 1) = \int e^{ip\cdot a}\, \mathrm{d}E(p)$ with the projection-valued measure E supported on the closed forward light cone $\overline{V^+} = \{p : p_0 \geq |\vec{p}|\}$; every translation-invariant vector is a complex multiple of Ω.

W1. Regularity and adjoints. $f \mapsto \Phi_v(f)\psi$ is continuous from the Schwartz topology to $\mathcal{H}$; there is a conjugate-linear involution $v \mapsto v^*$ with $\langle \Phi_v(f)\psi, \chi\rangle = \langle \psi, \Phi_{v^*}(\bar{f})\chi\rangle$; and $\mathcal{D}$ is the span of finite field words applied to Ω (cyclicity).

W2. Poincaré covariance. $U(g)\Phi_v(f)U(g)^{-1} = \Phi_{\rho(g)v}(f_g)$ with $f_g(x) = f(\Lambda^{-1}(x - a))$, ρ a complex-linear representation on labels, and $U(g)\Omega = \Omega$.

W3. Local commutativity. If every point of $\operatorname{supp} f$ is spacelike separated from every point of $\operatorname{supp} h$, then $[\Phi_v(f), \Phi_w(h)] = 0$ on $\mathcal{D}$.

Definition 1.2 (Mass gap)**.** For the Hamiltonian $H \geq 0$, the theory has a *mass gap* $\Delta > 0$ if $\sigma(H) \subset \{0\} \cup [\Delta, \infty)$.

1.2 What is proved in this paper

The input is the Lean 4 formalization submitted on October 28, 2026 [OpenAIRepo], validated by the committee [YMRepo]. This paper recompiles the construction into ordinary rigorous mathematics, expanding every proof that can be expanded from the definitions and the certified assumption packages below. The main results are:

1. the field-label algebra $V_d, N_d, \mathcal{F}_d$ and the separated test function spaces, with all algebraic proofs complete (Sections 2–3);
2. the finite Wilson lattice theory: the well-definedness, gauge invariance, and moment properties of the measure, the plaquette Taylor expansion, the four regularity properties of the renormalized insertions, the boundedness and permutation invariance of the smeared correlations, and the five common-scale identities, all proved from the definitions (Section 4);
3. under the explicit assumptions (RG1)–(RG5): the fixed-volume ultraviolet limit with smooth kernels (Theorem 5.17);
4. under (UV1): the subsequential infinite-volume limit, simultaneously for all correlation functions (Theorem 6.4);

5. under (PH1)–(PH4): the phase decomposition, clustering of every component, the full-support property, and the phase second-moment identity (Theorem 7.13, Section 7);

6. under (AN1): the Osterwalder–Schrader reconstruction, with each of W0–W3 and faithfulness proved individually (Theorem 8.10);

7. under (Y1)–(Y4): the Yang–Mills identification—the stress tensor and its Ward identity, the trace anomaly with constant $3\dim_{\mathbb{R}}\mathfrak{g}/\pi^4$ verified by complete computation, the exclusion of the finite free-field class, the one-loop running coupling, and the operator product expansion with error bounds (Section 9).

Remark 1.3 (Scope). A positive mass gap is neither proved nor disproved. Uniqueness of the infinite-volume limit across subsequences, agreement of the component theories, uniqueness of the vacuum of the mixed theory, and order-by-order perturbative agreement of the OPE remain open (Section 10). Ultraviolet convergence holds along the sequence of spacings chosen by the construction; there is no universality statement for arbitrary sequences.

1.3 Notation

Throughout, $G \subset U(m)$ is a closed matrix subgroup whose real Lie algebra $\mathfrak{g}$ is simple and nonabelian; $C_{\mathfrak{g}}$ is the adjoint Casimir constant determined by the inner product (2.1) below; $a_N = D_N^{-1}$ is the lattice spacing, L_p the physical torus side length, $\beta_{N,p}$ the bare coupling, $\Lambda_{N,p}$ the discrete torus, $\mu_{N,p}$ the finite Wilson measure, $C_{N,p}$ the smeared lattice correlation; S_p is the fixed-volume limit, S_{mix} the infinite-volume limit, S_ξ the pure-phase correlation; V_d is the space of labels of weight $\le d$, N_d the correlation null space, $\mathcal{F}_d = V_d/N_d$; $b_0 = 11C_{\mathfrak{g}}/(48\pi^2)$ and $b^* = 2b_0$.

2 The field-label algebra

2.1 The invariant inner product

Let $G \subset U(m)$ be closed; then $\mathfrak{g} = \mathrm{Lie}(G) \subset \mathfrak{u}(m)$, so every $X \in \mathfrak{g}$ satisfies $X^\dagger = -X$. Define

$$\langle X, Y\rangle_{\mathfrak{g}} = -\frac{1}{m}\,\mathrm{Re}\,\mathrm{Tr}(XY). \tag{2.1}$$

Lemma 2.1. *$\langle\cdot,\cdot\rangle_{\mathfrak{g}}$ is a real inner product on $\mathfrak{g}$.*

Proof. Bilinearity and symmetry follow from $\mathrm{Tr}(XY) = \mathrm{Tr}(YX)$. For positive definiteness, $X^\dagger = -X$ gives $X^\dagger X = -X^2$, hence $-\mathrm{Re}\,\mathrm{Tr}(X^2) = \mathrm{Re}\,\mathrm{Tr}(X^\dagger X) = \mathrm{Tr}(X^\dagger X) = \sum_{i,j}|X_{ij}|^2 \ge 0$, with equality if and only if $X = 0$. Division by $m > 0$ is irrelevant. □

Lemma 2.2 (Ad-invariance)**.** *For all $g \in G$, $\langle gXg^{-1}, gYg^{-1}\rangle_{\mathfrak{g}} = \langle X, Y\rangle_{\mathfrak{g}}$.*

Proof. $\mathrm{Tr}(gXg^{-1}gYg^{-1}) = \mathrm{Tr}(gXYg^{-1}) = \mathrm{Tr}(XY)$ by cyclicity; take real parts. □

Since $\mathfrak{g}$ is simple, every Ad-invariant symmetric bilinear form on $\mathfrak{g}$ is a constant multiple of the Killing form; hence (2.1) differs from the Killing form by a positive factor, and the Casimir constant $C_{\mathfrak{g}}$ used throughout is the one determined by this normalization. Its numerical value is fixed through the one-loop beta function: $b_0 = 11C_{\mathfrak{g}}/(48\pi^2)$ (Section 9.4).

2.2 Curvature, covariant derivatives, and the Bianchi identity

For a smooth $\mathfrak{g}$-valued connection A,

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu], \qquad D_\mu X = \partial_\mu X + [A_\mu, X]. \tag{2.2}$$

The following two lemmas are the algebraic source of all "label relations".

Lemma 2.3 (Commutator of covariant derivatives)**.** $[D_\mu, D_\nu]X = [F_{\mu\nu}, X]$*, and* $D_\mu[X,Y] = [D_\mu X, Y] + [X, D_\mu Y]$.

Proof. Expanding $[D_\mu, D_\nu]X$, the second-order terms cancel by symmetry of mixed partial derivatives; the ∂A terms give $[\partial_\mu A_\nu - \partial_\nu A_\mu, X]$; the cross terms give $[A_\mu, [A_\nu, X]] - [A_\nu, [A_\mu, X]] = [[A_\mu, A_\nu], X]$ by the Jacobi identity. The second identity follows from D_μ being a derivation, again by Jacobi. □

Lemma 2.4 (Bianchi identity)**.** $D_\lambda F_{\mu\nu} + D_\mu F_{\nu\lambda} + D_\nu F_{\lambda\mu} = 0$.

Proof. By Lemma 2.3, $[D_\mu, D_\nu] = \mathrm{ad}(F_{\mu\nu})$ and $[D_\lambda, \mathrm{ad}(F_{\mu\nu})] = \mathrm{ad}(D_\lambda F_{\mu\nu})$. The operator Jacobi identity for D_λ, D_μ, D_ν gives $\mathrm{ad}(D_\lambda F_{\mu\nu} + D_\mu F_{\nu\lambda} + D_\nu F_{\lambda\mu}) = 0$; since $\mathfrak{g}$ is simple, ad is injective (zero center), so the bracketed $\mathfrak{g}$-valued function vanishes identically. □

2.3 Label spaces V_d

Definition 2.5 (Labels and weight)**.** Take formal letters $F_{\mu\nu}$ ($0 \le \mu < \nu \le 3$, six letters, weight 2) and D_λ ($\lambda = 0, \dots, 3$, four letters, weight 1). Finite words, complex linear combinations, products, and gauge-invariant scalars (traces and index contractions) generate the *labels*; the weight is defined by

$$\mathrm{wt}(F_{\mu\nu}) = 2, \qquad \mathrm{wt}(D_\lambda v) = \mathrm{wt}(v) + 1, \qquad \mathrm{wt}(v_1 \cdots v_k) = \textstyle\sum_j \mathrm{wt}(v_j). \tag{2.3}$$

Let V_d be the complex vector space spanned by gauge-invariant scalar labels of weight $\le d$; for example $-\operatorname{Tr}(F_{\mu\nu}F_{\rho\sigma}) \in V_4$ and $-\operatorname{Tr}\big((D_\lambda F_{\mu\nu})(D_\kappa F_{\rho\sigma})\big) \in V_6$.

Lemma 2.6 (Finite-dimensionality)**.** $\dim_{\mathbb{C}} V_d < \infty$*, with the explicit bound* $\dim_{\mathbb{C}} V_d \le 10^{d+1}$.

Proof. Every letter has weight ≥ 1, so a word of weight $\le d$ has length $\le d$; with ten letters, there are at most $\sum_{\ell=0}^{d} 10^\ell \le 10^{d+1}$ words. V_d is a quotient of the span of these words by the linear conditions defining gauge invariants. □

2.4 The correlation null space and the quotient $\mathcal{F}_d$

Labels are not linearly independent: the Bianchi and Jacobi identities, the Leibniz rule, and cyclicity of trace generate formal relations. The construction handles this by defining equality at the level of correlation functions.

Definition 2.7 (Null space and quotient)**.** Suppose a family of separated-point correlation functionals S_s is given (constructed in Sections 4–7; used here only as a definition). For $d \ge 0$ let $R_d(v)$ be the family of all correlation values obtained by placing v into an arbitrary list of spectator labels, an arbitrary separated-point test function f, and an arbitrary state s, and let

$$N_d = \ker R_d \subset V_d, \qquad \mathcal{F}_d = V_d / N_d. \tag{2.4}$$

Thus N_d is the set of labels annihilated by every separated-point correlation containing them. The formalization takes precisely $\mathcal{F}_d = V_d/N_d$ with the quotient norm from the supremum norm $\|c\| = \max_j |c_j|$ in a chosen basis (certificates `YM.Table.null`, `YM.Table.Field`).

Lemma 2.8. *N_d is a subspace; $\mathcal{F}_d$ is finite-dimensional; for $d \le d'$ the natural map $\mathcal{F}_d \to \mathcal{F}_{d'}$ is injective; and the quotient norm is well defined.*

Proof. R_d is linear (correlations are complex-linear in each label), so $\ker R_d$ is a subspace and $\dim \mathcal{F}_d \le \dim V_d < \infty$. If $d \le d'$ then $V_d \subset V_{d'}$ and the condition "all correlations containing v vanish" is level-independent, so $N_d = N_{d'} \cap V_d$ and the induced map is injective. For the quotient norm, take a basis $v_1, \dots, v_k$ of V_d and set $\|[c]\| = \inf_{n \in N_d} \|c + n\|$; finite-dimensionality makes N_d closed so the infimum is attained, and homogeneity, the triangle inequality, and $\|[c]\| = 0 \Leftrightarrow c \in N_d$ make it a norm. □

Proposition 2.9 (Bianchi relations lie in the null space)**.** *Assume the Taylor matching of Section 4.3 and uniform convergence of $C_{N,p}$ to S_p for all labels (assumption (RG5) and Theorem 5.17; jointly assumption (A1) of the re-proof). Then the Bianchi combination $b = D_\lambda F_{\mu\nu} + D_\mu F_{\nu\lambda} + D_\nu F_{\lambda\mu}$, traced against an arbitrary scalar, lies in N_{d+1}.*

Proof. $b(A)(x) \equiv 0$ is the classical polynomial identity (Lemma 2.4). Taylor matching requires the leading term as $a \to 0$ of the lattice insertion to be exactly $b(A)(x)$, so the leading term vanishes; the remainder is $O(a)$ and is pressed to zero under integration against test functions, using the uniformity of (A1). Passing to the limit gives $S_p(b, \dots; f) = 0$ for arbitrary spectators and separated f, i.e. $b \in N_{d+1}$. □

2.5 The Maxwell shell comparison space and the field dictionary

Definition 2.10 (Free sourceless shell)**.** Consider formal curvature data $f_{\mu\nu}$ taking values in the free linearization of $\mathfrak{g}$, subject to antisymmetry $f_{\mu\nu} = -f_{\nu\mu}$, the Bianchi relation $\partial_\lambda f_{\mu\nu} + \partial_\mu f_{\nu\lambda} + \partial_\nu f_{\lambda\mu} = 0$, and the sourceless Maxwell equation $\partial_\mu f_{\mu\nu} = 0$. Let $\mathcal{P}_d$ be the gauge-invariant polynomials of weight $\le d$ evaluated at a point, modulo identification of polynomials agreeing on all shell data. The resulting complex vector space is denoted $\mathrm{Maxwell.shellSector}(G, d)$.

Lemma 2.11. *The natural map $\psi_d : V_d \to \mathrm{Maxwell.shellSector}(G, d)$, sending a label to the class of its classical polynomial restricted to the shell, is a complex-linear surjection, and $\ker \psi_d$ contains the formal relations generated by the Bianchi and Jacobi identities, the Leibniz rule, and cyclicity of trace.*

Proof. Linearity and surjectivity are immediate from the definition of V_d. The formal identities of Lemmas 2.3–2.4 and $\mathrm{Tr}(XY) = \mathrm{Tr}(YX)$ vanish pointwise on the shell, hence lie in the kernel. □

Proposition 2.12 (Field-content dictionary; input (B1))**.** *For each d, $N_d = \ker \psi_d$, so ψ_d induces a complex linear isomorphism*

$$\mathcal{F}_d \;\simeq_{\mathbb{C}}\; \mathrm{Maxwell.shellSector}(G, d), \tag{2.5}$$

and for every phase $\xi \in X$ (not merely ν-almost every phase) the same holds for the phase quotient $\mathcal{F}_d(\xi)$. This is the conclusion of the six-step audit of Section 9.6 (certificates `OperatorCompleteness.local_field_content`, `phase_field_content`); its proof uses the rank comparison and phase transport inside the formalization and is not expanded here.

Remark 2.13 (Meaning and limits of (2.5))*.* (2.5) constrains only the *size* of the surviving local field space at each filtration level; it is a vector-space equivalence. It is not an equality of interacting and free correlators, not an algebra isomorphism, and not a unitary equivalence of quantum field theories.

3 Separated test functions and correlation functionals

Definition 3.1 (Test space)**.** For $n \geq 1$ let $\Omega_n = \{(x_1, \dots, x_n) \in (\mathbb{R}^4)^n : x_i \neq x_j \text{ for } i \neq j\}$ and let $\mathcal{T}_n = C_c^\infty(\Omega_n)$ be the jointly supported smooth compactly supported test functions. Note that f need not be a tensor product of one-point test functions, and $\operatorname{supp} f$ has positive distance from every coincidence diagonal $\Delta_{ij} = \{x_i = x_j\}$. In finite volume, $\operatorname{supp} f$ must lie strictly inside $((-L_p/2, L_p/2)^4)^n$ and avoid the periodic images $x_i - x_j \in L_p\mathbb{Z}^4$.

Lemma 3.2 (Density of tensor products)**.** *Finite linear combinations of $f_1 \otimes \cdots \otimes f_n$, $f_i \in C_c^\infty(\mathbb{R}^4)$, are dense in $\mathcal{T}_n$.*

Proof. Given $f \in \mathcal{T}_n$ with $K = \operatorname{supp} f$, choose $\varepsilon > 0$ such that $|x_i - x_j| \geq 3\varepsilon$ on K for all $i \neq j$. Cover $(\mathbb{R}^4)^n$ by cubes of diameter $< \varepsilon$ and take a subordinate smooth partition of unity $\{\chi_\alpha\}$; then $f = \sum_\alpha \chi_\alpha f$ with finitely many nonzero terms, each supported where all coordinate components are mutually separated. On such a region, $\chi_\alpha f$ is approximated uniformly to arbitrary precision by finite sums of separated-variable polynomials times a smooth compact cutoff (Stone–Weierstrass followed by mollification). □

Definition 3.3 (Correlation functional)**.** Given a label list $v = (v_1, \dots, v_n)$, the n-point function $S(v; f)$ is a continuous linear functional on $\mathcal{T}_n$, complex-linear in each v_i. The construction first produces the lattice version $C_{N,p}(v; f)$ (Definition 4.11), then S_p by the ultraviolet limit, and finally S_{mix} and S_ξ through the infinite-volume limit and phase decomposition.

Lemma 3.4 (Kernel representation)**.** *A continuous linear functional S on $\mathcal{T}_n$ is represented by a unique distribution $K \in \mathcal{D}'(\Omega_n)$; if the kernel is given by a smooth function $S(v; x)$ on Ω_n, then $S(v; f) = \int f(x) S(v; x)\, \mathrm{d}^{4n}x$, and conversely smoothness of $S(v; \cdot)$ on $\operatorname{supp} f$ gives $|S(v; f)| \leq \|f\|_{L^1} \sup_{\operatorname{supp} f} |S(v; \cdot)|$.*

Proof. The Schwartz kernel theorem represents continuous functionals by unique distribution kernels; the remaining statements are immediate. □

Lemma 3.5 (Why the diagonals must be avoided)**.** *The short-distance law of Theorem 9.10 gives $S^{\xi^*}(\Theta, \Theta)(x, 0) \sim \frac{3 \dim \mathfrak{g}}{\pi^4} |x|^{-8} \log^{-2}(1/|x|)$. Since $\int_0 r^3 \cdot r^{-8}\, \mathrm{d}r$ diverges in $\mathbb{R}^4$, this kernel is not integrable at coincidence; requiring test functions to vanish near Δ_{ij} is the minimal condition making the integral well defined. This is why all convergence statements are formulated only for separated test functions.*

4 The finite Wilson lattice theory

4.1 The discrete torus and the configuration space

The construction selects two integer sequences (Appendix A, `YM.Schedule`, `YM.RG.Supply`):

$$D_N = \prod_{k<N} b_k, \quad a_N = D_N^{-1}, \quad L_p = 4(M + p + 1), \quad \Lambda_{N,p} \cong \big(\mathbb{Z}/(L_p D_N)\mathbb{Z}\big)^4, \tag{4.1}$$

where $b_k \geq 3$ is odd, $b_{k+1} \geq b_k$, and $\log b_N / \log D_N \to 0$; M is sufficiently large to satisfy minimum-size and initial-volume bounds; L_p is the physical side length and $a_N \cdot \operatorname{side}(N, p) = L_p$. Each positively oriented edge $e = (x, \mu)$ carries $U_e \in G$, and the reverse edge carries $U_{-e} = U_e^{-1}$.

Lemma 4.1. *$a_N \to 0$; (L_p) and (D_N) are strictly increasing; $|\Lambda_{N,p}| = (L_p/a_N)^4 = (L_p D_N)^4$.*

Proof. $b_k \geq 3$ gives $D_{N+1} = D_N b_N > D_N$, hence $a_N \downarrow 0$; L_p is strictly increasing in p. Since L_p and D_N are integers, so is $L_p D_N$, and $(\mathbb{Z}/(L_p D_N)\mathbb{Z})^4$ has $(L_p D_N)^4$ sites; $a_N = D_N^{-1}$ gives $L_p/a_N = L_p D_N$. □

The positively oriented plaquette ($\mu < \nu$, base point x) is the ordered product

$$U_{\mu\nu}(x) = U_{x,\mu}\, U_{x+e_\mu,\nu}\, U_{x+e_\nu,\mu}^{-1}\, U_{x,\nu}^{-1}. \tag{4.2}$$

Each positively oriented plaquette is counted once; the configuration space is $\Omega_{N,p} = G^{E^+}$ with $|E^+| = 4|\Lambda_{N,p}|$ and plaquette total $P = 6|\Lambda_{N,p}|$. Since G is a closed subgroup of $U(m)$, it is compact and $\Omega_{N,p}$ is a compact Hausdorff space.

4.2 The finite-volume Wilson law

Let h be normalized Haar measure on G ($h(G) = 1$, left-translation invariant). The action and probability measure are

$$S_{N,p}(U) = \sum_{\square} \Big(1 - \frac{1}{m} \operatorname{Re} \operatorname{Tr} U_{\square}\Big), \qquad \mathrm{d}\mu_{N,p}(U) = Z_{N,p}^{-1} e^{-\beta_{N,p} S_{N,p}(U)} \prod_{e \in E^+} \mathrm{d}h(U_e), \tag{4.3}$$

where $Z_{N,p}$ is the normalization constant and $\beta_{N,p} = 1/g_0^2$ is the bare coupling (inverse square of the cutoff coupling).

Lemma 4.2 (Boundedness of the action)**.** *For each plaquette,* $0 \leq 1 - \frac{1}{m} \operatorname{Re} \operatorname{Tr} U_{\square} \leq 2$.

Proof. $U_{\square} \in G \subset U(m)$ is unitary with eigenvalues of modulus 1, so $|\operatorname{Tr} U_{\square}| \leq m$ and $\operatorname{Re} \operatorname{Tr} U_{\square} \in [-m, m]$; dividing by m gives the claim. □

Corollary 4.3 (The measure is well defined)**.** $0 < Z_{N,p} < \infty$ *and* $\mu_{N,p}$ *is a probability measure on* $\Omega_{N,p}$.

Proof. $U \mapsto S_{N,p}(U)$ is continuous and real on a compact space, and $0 \leq S_{N,p} \leq 2P$ by Lemma 4.2, so $e^{-2\beta P} \leq e^{-\beta S} \leq 1$. The product measure $\prod \mathrm{d}h$ is a probability measure with nonempty support, so $Z_{N,p} \in [e^{-2\beta P}, 1] \subset (0, \infty)$; dividing by Z normalizes. □

Lemma 4.4 (Gauge invariance)**.** *For any* $t : \Lambda_{N,p} \to G$, *with* $U_{x,\mu}^t = t(x) U_{x,\mu} t(x + e_\mu)^{-1}$, *we have* $S_{N,p}(U^t) = S_{N,p}(U)$, *and* $\mu_{N,p}$ *is invariant under* $U \mapsto U^t$.

Proof. Substituting edge by edge along a plaquette, all interior t-factors cancel (using the periodic boundary), giving $U_{\mu\nu}^t(x) = t(x) U_{\mu\nu}(x) t(x)^{-1}$; the trace is conjugation invariant, so the action is invariant. Each $U_e \mapsto g U_e g'^{-1}$ is a composition of left and right translations, under which normalized Haar measure is invariant; so is the product measure. □

Lemma 4.5 (Finite moments)**.** *Any lattice polynomial* $P(U)$ *in matrix entries and their conjugates is bounded, hence* $\int P \,\mathrm{d}\mu_{N,p}$ *is finite.*

Proof. $U \in U(m)$ gives $|U_{ij}| \leq 1$, so a polynomial in finitely many entries is bounded on the compact space $\Omega_{N,p}$ by some C_P; as $\mu_{N,p}$ is a probability measure, $|\int P \,\mathrm{d}\mu| \leq C_P < \infty$. □

4.3 Lattice representatives and Taylor matching

Definition 4.6 (Representative condition)**.** For $v \in V_d$, the construction selects a gauge-invariant local lattice polynomial $L_v(x;U)$ depending only on the matrix entries (and conjugates) of links with both endpoints in a fixed finite pattern $S_d \subset \mathbb{Z}^4$, independent of N, p, x, and complex-linear in v. The requirement is: if U is the parallel transport of a smooth connection A, then the $a \to 0$ expansion of $a^{-d}L_v(x;U)$ agrees with the continuum label $v(A)(x)$ through order d.

Lemma 4.7 (Plaquette expansion, full derivation)**.** *Let U be the parallel transport of a smooth connection A, and let the $\mu\nu$-square have side length a. Then*

$$U_{\mu\nu}(x) = \exp\big(a^2 F_{\mu\nu}(x) + O(a^3)\big), \tag{4.4}$$

$$\frac{1}{m}\operatorname{Re}\operatorname{Tr} U_{\mu\nu}(x) = 1 + \frac{a^4}{2m}\operatorname{Tr}\big(F_{\mu\nu}(x)^2\big) + O(a^5), \tag{4.5}$$

$$\lim_{a\downarrow 0} a^{-4}\Big(1 - \frac{1}{m}\operatorname{Re}\operatorname{Tr} U_{\mu\nu}(x)\Big) = -\frac{1}{2m}\operatorname{Tr}\big(F_{\mu\nu}(x)^2\big) = \frac{1}{2}\langle F_{\mu\nu}, F_{\mu\nu}\rangle_{\mathfrak{g}}. \tag{4.6}$$

Proof. The path-ordered exponential around the closed square gives the first-order line integral $a^2(\partial_\mu A_\nu - \partial_\nu A_\mu)$ and the second-order commutator term $a^2[A_\mu, A_\nu]$ from path ordering; together they give $a^2 F_{\mu\nu}(x)$, the remainder being $O(a^3)$ (nonabelian Stokes formula, see [Wilson, Seiler]). The BCH formula yields (4.4). Expanding the exponential, $\frac{1}{m}\operatorname{Tr} e^X = 1 + \frac{1}{m}\operatorname{Tr} X + \frac{1}{2m}\operatorname{Tr} X^2 + O(\|X\|^3)$ with $X = a^2F + O(a^3)$; since F and the $O(a^3)$ terms are anti-Hermitian, $\operatorname{Tr} X$ is purely imaginary so $\operatorname{Re}\operatorname{Tr} X = 0$, and $\operatorname{Tr} X^2 = a^4 \operatorname{Tr}(F^2) + O(a^5)$ with $\operatorname{Tr}(F^2)$ real and nonpositive. Hence (4.5), and (4.6) follows on rearranging and using (2.1). □

Remark 4.8. (4.6) shows that $1 - \frac{1}{m}\operatorname{Re}\operatorname{Tr} U_{\mu\nu}$ is an admissible "candidate representative" of the label $-\frac{1}{2m}\operatorname{Tr}(F_{\mu\nu}^2)$. As emphasized in the source document (§4.1 there), the actual construction imposes additional discrete symmetry requirements and does not generally set $L_v = 1 - \frac{1}{m}\operatorname{Re}\operatorname{Tr} U_{\mu\nu}$; we use the example only as a readable verification of the Taylor matching condition.

4.4 Renormalized insertions and the four regularity properties

The observables actually matched to labels carry renormalization coefficients. Writing a label as a token (d, c), where $c = (c_1, \dots, c_k)$ is the coordinate in a chosen basis $v_1, \dots, v_k$ of V_d, the construction selects a finite offset set S_d and lattice polynomials $P_{d,j}(x;U)$; coefficients $\tilde{c}_j^{N,p}$ are then produced by the inverse of the RG transport map composed with scale factors, and

$$I_{N,p}\big((d,c),x;U\big) = a_N^{-d}\sum_{j=1}^{k} \tilde{c}_j^{N,p}\, P_{d,j}(x;U), \tag{4.7}$$

depending only on links with endpoints in $x + S_d$ (periodically wrapped). This is the insertion of the source document; both a_N^{-d} and $\tilde{c}_j^{N,p}$ may diverge as $N \to \infty$, and the convergence theorem asserts that these divergences are exactly canceled by the RG normalization.

Lemma 4.9 (The four regularity properties)**.** *For each fixed N, p, the insertion $I_{N,p}$ is continuous, bounded, gauge invariant, and locally supported.*

Proof. (a) *Continuity:* each $P_{d,j}$ is a polynomial in finitely many matrix entries and conjugates, hence $U \mapsto P_{d,j}(x;U)$ is continuous; multiplication by a_N^{-d} and finite summation preserve continuity (certificate `YM.Family.continuous0`). (b) *Boundedness:* $\Omega_{N,p}$ is compact, so a continuous function

attains its maximum and $|I| \le K(N,p,d,c) < \infty$; K may diverge in N, which is precisely the object of renormalization. (c) *Gauge invariance:* the $P_{d,j}$ are taken from gauge-invariant local sections (traces of closed Wilson loops and their derivative insertions); by the conjugation law of Lemma 4.4, every loop trace is invariant, so $I(U^t) = I(U)$ (certificate `YM.Family.invariant`). (d) *Local support:* $S_d \subset \mathbb{Z}^4$ is finite and independent of N, p, x, and I depends only on links with endpoints in $x + S_d$ (certificate `YM.Family.local0`). □

Proposition 4.10 (Running coupling condition)**.** *The construction chooses $\beta_{N,p}$ so that the RG trajectory returns to a fixed reference value $\bar{\beta} > 0$ after N steps. The recorded conclusion is: there exists $b^* = 2b_0 = 11C_{\mathfrak{g}}/(24\pi^2) > 0$ such that*

$$\forall \varepsilon > 0\ \exists N_0\ \forall N \ge N_0\ \forall p, \qquad \Big|\frac{\beta_{N,p}}{\log(a_N^{-1})} - \frac{11C_{\mathfrak{g}}}{24\pi^2}\Big| < \varepsilon. \tag{4.8}$$

Heuristic derivation (preview of Lemma 5.12). The one-loop beta function $\beta(g) = \mu\,\mathrm{d}g/\mathrm{d}\mu = -b_0 g^3 + O(g^5)$ gives $\frac{\mathrm{d}}{\mathrm{d}\log\mu} g^{-2} = 2b_0 + O(g^2)$. Taking $\mu = 1/a$ and integrating, $g^{-2}(1/a) = g^{-2}(\Lambda) + 2b_0 \log(1/(a\Lambda)) + o(\log(1/a))$; since $\beta = 1/g^2$, division by $\log(a^{-1})$ tends to $2b_0 = 11C_{\mathfrak{g}}/(24\pi^2)$, in agreement with (4.8). The strict *uniform* (in p) error bound requires the RG estimates of the formalization and is assumption (RG2) below, proved inside the repository (certificates `YM.Pert.Asymptotics.b0`, `YM.Gauge.casimir`, `YangMills.WilsonRegularity.running`); see Lemma 5.12. □

4.5 Smeared lattice correlations

Definition 4.11 (Finite lattice n-point function)**.** For a label list $v = (v_1, \dots, v_n)$ and an admissible test function f (Definition 3.1), let

$$C_{N,p}(v; f) = a_N^{4n} \sum_{x_1,\dots,x_n \in \Lambda_{N,p}} f(a_N x_1, \dots, a_N x_n) \int_{\Omega_{N,p}} \prod_{i=1}^{n} I_{N,p}(v_i, x_i; U)\, \mathrm{d}\mu_{N,p}(U), \tag{4.9}$$

the sum over ordered n-tuples; the integral is the finite-volume Wilson expectation $\mathbb{E}_{\mu_{N,p}}[\prod_i I_{N,p}(v_i, x_i; \cdot)]$.

Lemma 4.12 (Well-definedness, multilinearity, boundedness)**.** *$C_{N,p}(v; f)$ is well defined, complex-multilinear in $(v_1, \dots, v_n; f)$, and bounded: with $K_i = \sup_U |I_{N,p}(v_i, \cdot; U)|$,*

$$|C_{N,p}(v; f)| \le a_N^{4n} |\Lambda_{N,p}|^n \|f\|_\infty \prod_{i=1}^{n} K_i = L_p^{4n} \|f\|_\infty \prod_{i=1}^{n} K_i. \tag{4.10}$$

Proof. The sum and integral are finite: $\Lambda_{N,p}$ is finite and the integrand is continuous (Lemma 4.9(a)), so the integral is finite by Lemma 4.5. Linearity in each v_i follows from the linearity of $I_{N,p}$ in the token coordinate (4.7), and linearity in f is clear. For boundedness, (4.10) follows from $a_N^{4n} |\Lambda_{N,p}|^n = (a_N^4 (L_p/a_N)^4)^n = L_p^{4n}$ (Lemma 4.1). At fixed volume the right side is finite; the possible growth of K_i in N is controlled by the RG normalization of (RG3), which is exactly the content of the ultraviolet limit. □

Lemma 4.13 (Permutation invariance)**.** *$C_{N,p}$ is invariant under simultaneous permutation of (v_i, x_i).*

Proof. For $\pi \in S_n$, relabel the summation variables $x_i \mapsto x_{\pi(i)}$; $\Lambda_{N,p}^n$ is invariant under relabeling, and the product $\prod_i I_{N,p}(v_i, x_i; U)$ takes the same values upon simultaneous permutation of labels and points since finite products commute. □

Remark 4.14 (Role of the quotient map). If $v - v' \in N_d$, the finite-N functionals $C_{N,p}(v;\cdot)$ and $C_{N,p}(v';\cdot)$ need not agree; the convergence theorem guarantees that their limits S_p agree, and the map `project` sends a token to the quotient class $(d, [c])$ (Appendix A). What is constructed is therefore a field of $\mathcal{F}_d$, not of raw tokens.

4.6 Common reference scale and normalizations

Lemma 4.15 (Five calibrations)**.** *For the same Wilson family, before any limit is taken:*

1. Reference unit. *The spacing at RG reference level zero is $a(0) = 1$ (`reference_unit`, `CommonScale.reference_spacing`).*
2. Physical side length. *$a_N \cdot \mathrm{side}(N, p) = L_p$ for all N, p (`side_length_in_physical_units`, `CommonScale.physical_side`).*
3. Common reference coupling. *There is $\bar{\beta} > 0$ such that $\beta_{N,p}(0) = \bar{\beta}$ at reference level zero for every N, p; the bare $\beta_{N,p}$ itself may run with the cutoff (`one_positive_reference_coupling`, `CommonScale.positive_common_reference`).*
4. Shared kinetic paths. *For every RG level $m \le N$, the kinetic path at level m is the same for all volumes p, q (`kinetic_path_agrees_across_volumes`, `CommonScale.shared_kinetic_paths`).*
5. Insertion readout. *For a degree-d label v, the normalized coefficient at the fixed reference level $r(d)$ equals v itself, independently of cutoff and volume, whenever $N \ge r(d)$ (`insertion_has_the_prescribed_reference_value`, `CommonScale.reference_insertion`).*

Proof of which parts are definitional. Items (1), (2), (3), and (5) are definitional conventions of the RG bookkeeping and the choice of $\beta_{N,p}$: (2) follows from $\Lambda_{N,p} = (\mathbb{Z}/(L_p D_N)\mathbb{Z})^4$ and $a_N = D_N^{-1}$ (Lemma 4.1); (3) is the definition "$\beta_{N,p}$ is chosen to return to $\bar{\beta}$ after N steps"; (5) is the definition of the reference-level readout, its usability depending on invertibility of the RG transport map for $N \ge r(d)$ (assumption (RG3)). Item (4) is a constructive conclusion: it requires the kinetic paths of different volumes to agree, and is supplied by the shared kinetic path data, not by definition alone. □

Proposition 4.16 (Finite-law comparison with the independent lattice formalization; input (C1))**.** *Comparing the finite Wilson law with the independent formalization [LGT] (bridged in Section 13.1): the ordered plaquette products agree lattice point by point before any limit (`same_plaquette`); with the conversion $\beta_{\mathrm{submission}} = m\,\beta_{\mathrm{lgt}}$ of (13.1), the normalized finite laws and finite moments agree exactly (`coupling_conversion`, `normalized_law_on_each_lattice`, `finite_moments_are_the_family_moments`). This is an independent check of the finite lattice theory itself, independent of any limit; the internal proof of the adapter is input (C1).*

Remark 4.17 (Calibration versus probe). These identities are RG bookkeeping of one Wilson family, not merely the assignment of a common numerical length unit to unrelated theories; complementing them, the separated two-point probe of Section 13.3 below connects the calibration to a nonzero mixed-theory observable. Together they support the normalization of all subsequent short-distance constants.

5 Renormalization group and the ultraviolet limit

5.1 The RG data and the assumption package

For $\rho > 0$ let C_ρ be the construction's blocking kernel, coarsening level m (spacing $a^{(m)}$) to level $m-1$ (spacing $a^{(m-1)} = b_m a^{(m)}$), with $a^{(0)} = 1$ (Lemma 4.15(1)). Write the effective action at level m as

$$S_m(U) = \beta_m \sum_{\square} \Big(1 - \frac{1}{m} \operatorname{Re} \operatorname{Tr} U_{\square}\Big) + J_m(U), \tag{5.1}$$

where β_m is the coefficient of the marginal (plaquette) term and J_m a finite sum of irrelevant terms taking values in a Banach space with norm $\|\cdot\|_m$; $g^2 = 1/\beta$ is the running coupling squared.

Assumption 5.1 (RG1: symmetry-preserving gauge fixing)**.** The gauge fixing averages over symmetry-breaking choices and preserves the finite lattice symmetry group Γ_{lat} (rotations, reflections, translations); the blocking kernel C_ρ is Γ_{lat}-equivariant. Certificates: the source document §4.3 Remark 4.2 and §6.4; `YM.RG.Program`.

Lemma 5.2 (Averaging invariance, full proof)**.** *Let a finite group Γ act on a set $\mathcal{F}$ and let $F_0 \in \mathcal{F}$ be arbitrary. Then $F = \frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} \gamma \cdot F_0$ is Γ-invariant.*

Proof. For $\eta \in \Gamma$, $\eta \cdot F = \frac{1}{|\Gamma|} \sum_\gamma (\eta\gamma) \cdot F_0 = \frac{1}{|\Gamma|} \sum_{\gamma' \in \Gamma} \gamma' \cdot F_0 = F$, since $\gamma \mapsto \eta\gamma$ permutes the finite group Γ. □

Assumption 5.3 (RG2: one-step estimates, uniform in volume)**.** There exist constants $\lambda \in (0,1)$, $K > 0$, and $\beta_* > 0$ with

$$2b_0 \log 3 \;>\; \frac{K}{\beta_*}, \qquad b_k \geq 3, \tag{5.2}$$

such that for all levels $m \leq N$ and all volume indices p:

$$\|J_{m+1}\|_{m+1} \leq \lambda \|J_m\|_m + K\beta_m^{-1}, \qquad \big|\beta_{m+1} - \beta_m - 2b_0 \log b_m\big| \leq K\beta_m^{-1}, \qquad \beta_m \geq \beta_*.$$

The physical content: each blocking step contracts irrelevant terms geometrically at one fixed rate λ, while the marginal term follows the one-loop logarithmic recursion; $K\beta_m^{-1}$ is the two-loop-and-higher remainder, and (5.2) is exactly the condition that the one-loop recursion dominates its remainder—without it the running condition (4.8) itself would fail. The two extra clauses are not new physics: the lower bound $\beta_m \geq \beta_* > 0$ was already used in the propagation bound. Certificates: `YM.RG.Program`, `YM.RG.Supply.Program.book`, `YM.RG.Calendar.Fits`, `YM.Pert.Asymptotics.b0`, `YM.Gauge.casimir`.

Remark 5.4 (Source-level form of (RG2))*.* In the submitted sources the remainder is controlled by the schedule rather than by β_m^{-1}: `YM.Price.Intent.Large.slow` gives $\text{drift}_j / \log L \leq j^{-2}$ and `Large.kick` gives $\text{drift}_j / \log L \leq \text{slopeConstant}/4$; `YM.RG.History.Fits.driftSmall` instantiates the former at $j = J(m+1)$, $L = b_m$. `YM.Price.Intent.Rolls.exists_rolls` constructs the sequence via Tendsto $(\lambda n,\ x/\log n)$ atTop $(\mathcal{N}\,0)$. Consequently (5.2), the lower bound $\beta_m \geq \beta_*$, and the bootstrap of Lemma 5.11 are *sufficient but not necessary*; the strict sufficient condition is the weighted Cesàro Lemma 5.9, and the source-level reading of the drift identity is recorded in Remark 10.42.

Assumption 5.5 (RG3: stability of the insertion readout)**.** For a degree-d label v, the normalized coefficient at the reference level $r(d)$ equals v (Lemma 4.15(5)), and the RG transport linear

map from level N back to $r(d)$ is invertible (Appendix A, `TransferData.cylinder`); its inverse acts on insertions paired with test functions with bounds governed by the products of (RG2), giving a constant independent of p and growing at most like a_N^{-d}. Certificates: `insertion_has_the_prescribed_reference_value`, `CommonScale.reference_insertion`, `YM.Family.O`.

Assumption 5.6 (RG4: difference estimates, derived)**.** There exists $A(v,f) < \infty$ (depending on finitely many norms of f and on the label weights) such that

$$\bigl|C_{N+1,p}(v;f) - C_{N,p}(v;f)\bigr| \le A(v,f)\,\lambda^N.$$

In addition, for every multi-index α, every compact $K \subset \Omega_n$, and every finite label list v, there is $A_{\alpha,K,v} < \infty$, independent of N and p, with

$$\sup_{x\in K} \bigl|\partial_x^\alpha K_{N+1,p}(v;x) - \partial_x^\alpha K_{N,p}(v;x)\bigr| \le A_{\alpha,K,v}\,\lambda^N, \tag{5.3}$$

where $K_{N,p}(v;x)$ is the smooth summation kernel of Remark 5.16. Both forms are derived from (RG2)+(RG3) (Lemma 5.14 and Remark 5.15); they are not external inputs.

Assumption 5.7 (RG5: uniform smoothness estimates)**.** For every multi-index α, every compact $K \subset \Omega_n$, and every finite label list v,

$$\sup_{N,p}\ \sup_{x\in K} \bigl|\partial_x^\alpha C_{N,p}(v;x)\bigr| \le M_{\alpha,K,v} < \infty,$$

where $C_{N,p}(v;x)$ is the summation kernel of (4.9) viewed as a function on the separated configurations.

Remark 5.8. Within the RG package of Section 5, the external analytic inputs are (RG1), (RG2), (RG3), and (RG5); (RG4) and (5.3) are derived from them (Lemma 5.14 and Remark 5.15). All of these are recorded as proved inside the formalization. The full register of external inputs used in the paper, including (UV1), (PH1)–(PH4), (AN1), (Y1)–(Y4), (B1), and (C1), is Appendix B.

Lemma 5.9 (Schedule-controlled remainder; weighted Cesàro)**.** *Let $b_m \ge 3$ be the blocking ratios, $D_N = \prod_{m<N} b_m$, and let ε_m be the one-step remainder in*

$$\beta_{m+1} = \beta_m + 2b_0 \log b_m + \varepsilon_m, \qquad E_N := \sum_{m<N} \varepsilon_m.$$

Suppose there is a monotone $J : \mathbb{N} \to \mathbb{N}$ with $J(m) \to \infty$ and

$$|\varepsilon_m| \ \le\ J(m+1)^{-2} \log b_m \qquad \text{for all } m.$$

Then:

(1) $|E_N| / \log D_N \to 0$*;*

(2) $\beta_{N,p} / \log(a_N^{-1}) \to 2b_0$ *uniformly in p (i.e. (4.8) holds);*

(3) *there exist M and $\delta' > 0$ with $\beta_m \ge \beta_m^{(0)} + (m-M)\delta'$ for $m \ge M$, so no bootstrap and no lower bound $\beta_m \ge \beta_* > 0$ is needed beyond a positive starting value.*

Proof. Fix $\eta > 0$. Since $J \to \infty$, choose M with $J(m+1)^{-2} \le \eta$ for all $m \ge M$. Split $E_N = E_{<M} + E_{\ge M}$ with $E_{<M}$ a fixed finite constant C_M. Then $|E_{\ge M}| \le \eta \sum_{m=M}^{N-1} \log b_m \le \eta \log D_N$. Since $b_m \ge 3$, $\log D_N = \sum_{m<N} \log b_m \ge N \log 3 \to \infty$, hence $\limsup_N |E_N|/\log D_N \le C_M/\log D_N + \eta \to \eta$; as $\eta > 0$ was arbitrary, $|E_N|/\log D_N \to 0$. With $\beta_{N,p} = \bar\beta + 2b_0 \log D_N + E_N$ (Lemma 4.15(3)) and $a_N^{-1} = D_N$, statement (2) follows. For (3), take M with $J(m+1)^{-2} \le b_0$; then for $m \ge M$, $\beta_{m+1} \ge \beta_m + (2b_0 - J(m+1)^{-2}) \log b_m \ge \beta_m + b_0 \log 3$, using $\log b_m \ge \log 3$. A positive starting value $\beta_M > 0$ is supplied by the band condition (`Fits.delayed`, `Bench.recover`). Uniformity in p is inherited from the p-independent gates. □

Remark 5.10 (Source-level reading of the remainder). In the submitted sources the remainder is controlled by the schedule rather than by β_m^{-1}: `YM.Price.Intent.Large.slow` gives $\mathrm{drift}_j/\log L \le j^{-2}$ and `Large.kick` gives $\mathrm{drift}_j/\log L \le G.\mathrm{slopeConstant}/4$; `YM.RG.History.Fits.driftSmall` instantiates the former at $j = J(m+1)$, $L = b_m$, and `Fits.logCost` does the same for the logarithmic cost. `YM.Price.Intent.Rolls.exists_rolls` constructs the sequence by the elementary limit Tendsto $(\lambda n,\ x/\log n)$ atTop $(\mathcal{N}\,0)$: the drift is a fixed number and choosing the blocking ratio L large enough makes drift$/\log L$ arbitrarily small; oddification preserves `Large.odd` and `Large.size`. `Fits.deep` gives $J \to \infty$, `Fits.stepSlow` gives at most linear growth of J, and `Fits.delayed` gives $1 < \hat\beta_m$. Consequently the dominance condition (5.2), the bound $\beta_m \ge \beta_*$, and the two-stage bootstrap of Lemma 5.11 are *sufficient but not necessary*. The identification of the source's `drift` with the one-step remainder ε_m is recorded as a source-level reading, not re-proved here.

5.2 Propagation of irrelevant terms and the running coupling

Lemma 5.11. *Under (RG2), put $\delta' := 2b_0 \log 3 - K/\beta_* > 0$ (by (5.2)). Then for every $m \le N$ and every p,*

$$\beta_m \ge \bar\beta + m\delta', \qquad \|J_m\|_m \le \lambda^m \|J_0\|_0 + \frac{K}{\beta_*(1-\lambda)}. \tag{5.4}$$

In particular $\beta_m \ge \beta_* > 0$, $\beta_m \ge \delta' m$ for $m \ge 1$, and $\sup_{m \le N} \|J_m\|_m$ is bounded independently of N and p.

Proof. From $\beta_m \ge \beta_*$, $b_m \ge 3$, and (RG2),

$$\beta_{m+1} \ge \beta_m + 2b_0 \log b_m - K\beta_m^{-1} \ge \beta_m + 2b_0 \log 3 - K/\beta_* = \beta_m + \delta'.$$

Since $\beta_{N,p}(0) = \bar\beta$ (Lemma 4.15(3)), induction from the reference level gives $\beta_m \ge \bar\beta + m\delta'$. For the irrelevant terms, set $u_m = \|J_m\|_m$ and $\kappa = K/\beta_*$; since $\beta_m \ge \beta_*$, $K\beta_m^{-1} \le \kappa$, so $u_{m+1} \le \lambda u_m + \kappa$ by (RG2), and inductively $u_m \le \lambda^m u_0 + \kappa \sum_{j=0}^{m-1} \lambda^j \le u_0 + \kappa/(1-\lambda)$. All constants are p-independent by (RG2). □

Lemma 5.12 (One-loop orbit and the running condition (4.8))**.** *Under (RG2), $b_k \ge 3$, and $\log b_N/\log D_N \to 0$,*

$$\frac{\beta_{N,p}}{\log(a_N^{-1})} \longrightarrow 2b_0 = \frac{11 C_{\mathfrak g}}{24\pi^2} \qquad (N \to \infty),$$

uniformly in p.

Proof. Accumulating the one-step recursion, with $\varepsilon_m = \beta_{m+1} - \beta_m - 2b_0 \log b_m$, $|\varepsilon_m| \le K\beta_m^{-1}$, and

$$\beta_{N,p} = \beta_{N,p}(0) + \sum_{m<N} \big(2b_0 \log b_m + \varepsilon_m\big) = \bar\beta + 2b_0 \log D_N + E_N,$$

using $\beta_{N,p}(0) = \bar{\beta}$ (Lemma 4.15(3)) and $\log D_N = \sum_{m<N} \log b_m$. Estimate the error without circularity: by Lemma 5.11, $\beta_m \geq \bar{\beta} + m\delta'$ with $\delta' > 0$, hence

$$|E_N| \leq K \sum_{m<N} \frac{1}{\bar{\beta} + m\delta'} \leq \frac{K}{\delta'} \log\Big(1 + \frac{\delta' N}{\bar{\beta}}\Big) + \frac{K}{\bar{\beta}} = O(\log N),$$

and no preliminary induction is needed. Since $\log D_N \geq N \log 3$ (from $b_k \geq 3$), $E_N / \log D_N \to 0$. Therefore

$$\frac{\beta_{N,p}}{\log(a_N^{-1})} = \frac{\bar{\beta} + 2b_0 \log D_N + E_N}{\log D_N} \longrightarrow 2b_0,$$

because $a_N^{-1} = D_N$. All constants are p-independent by (RG2), so convergence is uniform in p. □

Remark 5.13. The source document (§7.4, Step 1) writes the same recursion in continuous form $\frac{\mathrm{d}}{\mathrm{d}\log\mu} g^{-2} = 2b_0 + O(g^2)$; integrating gives $g^{-2}(\mu) = 2b_0 \log(\mu/\Lambda) + O(\log\log\mu)$, the $O(\log\log\mu)$ term coming from the two-loop contribution. Lemma 5.12 is the strict discrete version in the lattice variables (b_m, D_N).

Lemma 5.14 ((RG2)+(RG3) implies (RG4))**.** *Under (RG2) and (RG3), the difference of smeared correlations at adjacent cutoffs decays like λ^N.*

Proof. By the common reference scale (Lemma 4.15), place both $C_{N,p}$ and $C_{N+1,p}$ at the same reference level $r(d)$ with the same labels and test function; the difference comes only from the single blocking step from level N to $N+1$. Write

$$C_{N+1,p}(v; f) - C_{N,p}(v; f) = \mathbb{E}_{\mu_{N+1,p}}[\Pi_{N+1}] - \mathbb{E}_{\mu_{N,p}}[\Pi_N] = \mathbb{E}_{\mu_{N+1,p}}\big[\Pi_{N+1} - C_\rho \Pi_N\big],$$

using equivariance of the blocking kernel, $\mathbb{E}_{\mu_{N,p}}[\cdot] = \mathbb{E}_{\mu_{N+1,p}}[C_\rho(\cdot)]$. Let $\mathcal{R}$ denote the reference-level readout at level $r(d)$ of (RG3). By (RG3) the RG transport from level N back to $r(d)$ is invertible and independent of N and p, and the reference value of the insertion is the label itself (Lemma 4.15(5)); hence $\mathcal{R}\big(\Pi_{N+1,p}\big) = \mathcal{R}\big(C_\rho \Pi_{N,p}\big) = v$, so the difference $\Pi_{N+1,p} - C_\rho \Pi_{N,p}$ has zero reference readout. Consequently it is a first cumulant of the irrelevant term J_N with no constant (reference) component: its level-dependent part is bounded by $B\|J_N\|_N$ with B depending only on v, f, p, while the level-independent constant part is annihilated by $\mathcal{R}$. By Lemma 5.11, $\|J_N\|_N \leq \lambda^N \|J_0\|_0 + \frac{K}{\beta_*(1-\lambda)}$, and the constant term drops out precisely because $\mathcal{R}\big(\Pi_{N+1,p} - C_\rho \Pi_{N,p}\big) = 0$; therefore

$$\big|\Pi_{N+1,p}(v; f; U) - \big(C_\rho \Pi_{N,p}\big)(v; f; U)\big| \leq B' \lambda^N \, a_N^{4n} \sum_{x_1,\ldots,x_n \in \Lambda_{N,p}} \big|f(a_N x_1, \ldots, a_N x_n)\big|.$$

Taking $\mathbb{E}_{\mu_{N+1,p}}$ and using Lemma 4.12, $a_N^{4n} |\Lambda_{N,p}|^n = L_p^{4n}$, so

$$\big|C_{N+1,p}(v; f) - C_{N,p}(v; f)\big| \leq B' L_p^{4n} \|f\|_\infty \lambda^N =: A(v, f) \lambda^N.$$

At fixed volume L_p is a constant; the p-uniformity required later is supplied by (RG3), whose bounds are independent of p. □

Remark 5.15 ((RG4*) is derived like (RG4)). The kernel-level estimate (5.3) is the same cumulant as in Lemma 5.14 with differentiated insertions (equivalently, with f replaced by a derivative of the test kernel); the reference readout still vanishes by (RG3), so the proof of Lemma 5.14 applies verbatim with $B' = B'_\alpha$. Thus (RG4*) is derived from (RG2)+(RG3) exactly as (RG4) was, and no new physical input is introduced. Certificates: same as Lemma 5.14 plus `YM.Family.continuousO`, `YM.Family.localO`.

5.3 The fixed-volume ultraviolet limit

Remark 5.16 (The lattice kernel and the quadrature error). Let $K_{N,p}(v;\cdot) \in C^\infty(\Omega_n)$ denote the smooth summation kernel whose existence and uniform bounds are asserted in (RG5), normalized so that (4.9) reads

$$C_{N,p}(v;f) = \int_{\Omega_n} f(x)\, K_{N,p}(v;x)\,\mathrm{d}x + \mathcal{E}_{N,p}(f), \tag{5.5}$$

with the quadrature error of the lattice Riemann sum. Concretely, if $\mathcal{C}_{N,x}$ is the cell of side a_N centred at $a_N x$ and $\bar f_{N,x} = a_N^{-4n}\int_{\mathcal{C}_{N,x}} f$, then (4.9) equals $a_N^{4n}\sum_x f(a_N x)K_{N,p}(v;a_N x)$ while the integral equals $a_N^{4n}\sum_x \bar f_{N,x}K_{N,p}(v;a_N x)$; since $|\bar f_{N,x} - f(a_N x)| \le 2a_N\|\nabla f\|_\infty$,

$$\big|\mathcal{E}_{N,p}(f)\big| \le 2a_N\|\nabla f\|_\infty\, L_p^{4n} \sup_{x\in \operatorname{supp} f} \big|K_{N,p}(v;x)\big|. \tag{5.6}$$

By (RG3) the renormalized insertions, paired with a fixed test function, admit bounds uniform in N (the factor a_N^{-d} of (4.7) is exactly cancelled by the inverse RG transport), so for every compact $K \subset \Omega_n$,

$$\sup_{N,p}\ \sup_{x\in K} \big|K_{N,p}(v;x)\big| \le M_{K,v} < \infty. \tag{5.7}$$

Since $a_N \to 0$ (Lemma 4.1), (5.6) and (5.7) give $\mathcal{E}_{N,p}(f) \to 0$ for every fixed f and p.

Theorem 5.17 (Fixed-volume ultraviolet limit)**.** *Under (RG1)–(RG5), for every fixed volume index p, every finite label list v, and every admissible test function f (Definition 3.1),*

$$S_p(v;f) = \lim_{N\to\infty} C_{N,p}(v;f) \tag{5.8}$$

exists along the entire chosen spacing sequence; there is a kernel $S_p(v;x)$, C^∞ on Ω_n, such that $S_p(v;f) = \int f(x)S_p(v;x)\,\mathrm{d}x$; and (4.8) holds with $b^* = 2b_0$.

Proof. (i) *Cauchy property.* For $M > N$, (RG4) gives $|C_{M,p} - C_{N,p}| \le \sum_{j=N}^{M-1} A\lambda^j \le \frac{A\lambda^N}{1-\lambda} \to 0$, so $(C_{N,p})_N$ is a Cauchy sequence in $\mathbb{C}$ and the limit exists.

(ii) *Linearity and continuity.* $C_{N,p}(v;f)$ is multilinear in $(v_1,\dots,v_n;f)$ (Lemma 4.12), and pointwise limits preserve linearity; by (4.10) and the uniformity in (RG3), $|S_p(v;f)| \le C_{p,v}\|f\|_\infty$, so S_p is a continuous linear functional on $\mathcal{T}_n$.

(iii) *Smooth kernel.* Fix a compact $K \subset \Omega_n$ and a multi-index α. By (RG4*) (5.3), for $M > N$,

$$\sup_{x\in K}\big|\partial^\alpha K_{M,p}(v;x) - \partial^\alpha K_{N,p}(v;x)\big| \le \sum_{j=N}^{M-1} A_{\alpha,K,v}\lambda^j \le \frac{A_{\alpha,K,v}\lambda^N}{1-\lambda} \xrightarrow[N\to\infty]{} 0.$$

Hence $(\partial^\alpha K_{N,p})_N$ is a Cauchy sequence in the Banach space $C^0(K)$, and by (RG5) it is uniformly bounded; completeness of $C^0(K)$ yields $S_p^{(\alpha)} \in C^0(K)$ with $\partial^\alpha K_{N,p} \to S_p^{(\alpha)}$ uniformly on K. By the Weierstrass theorem on the interchange of uniform convergence and differentiation, $S_p^{(0)} \in C^\infty(K)$ and $\partial^\alpha S_p^{(0)} = S_p^{(\alpha)}$ for every α. Choosing a countable compact exhaustion of Ω_n, the limits agree on overlaps (a subsequence of a convergent sequence converges to the same limit), so they patch to a single kernel $S_p(v;\cdot) \in C^\infty(\Omega_n)$.

(iv) *Integral representation.* Let $\operatorname{supp} f \subset K \subset \Omega_n$ with f smooth and compactly supported. By (5.5) and the triangle inequality,

$$\Big|C_{N,p}(v;f) - \int f(x)\, S_p(v;x)\,\mathrm{d}x\Big| \le \big|\mathcal{E}_{N,p}(f)\big| + \|f\|_{L^1}\sup_{x\in K}\big|K_{N,p}(v;x) - S_p(v;x)\big|.$$

The first term tends to 0 by (5.6)–(5.7) and $a_N \to 0$; the second tends to 0 by the case $\alpha = 0$ of (iii). Hence $S_p(v; f) = \lim_{N\to\infty} C_{N,p}(v; f) = \int f(x) S_p(v; x)\,\mathrm{d}x$, which is (5.8) together with the kernel representation.

(v) *Running coupling.* This is Lemma 5.12. □

Corollary 5.18. *(i) If $v - v' \in N_d$, then $S_p(v; f) = S_p(v'; f)$; (ii) S_p is invariant under simultaneous permutation of (v_i, x_i); (iii) gauge invariance passes to the limit.*

Proof. (i) By definition of N_d (Definition 2.7), all separated correlations containing $v - v'$ vanish in the limit; `project` sends a token to $(d, [c])$, and the `wilson` property requires equivalent tokens to have equal limiting correlations. (ii) Lemma 4.13 gives finite-N permutation invariance, and pointwise limits preserve it. (iii) Each $I_{N,p}$ is gauge invariant (Lemma 4.9(c)) and $\mu_{N,p}$ is gauge invariant (Lemma 4.4), so each $C_{N,p}$ is gauge invariant, and so is the limit. □

Remark 5.19 (Scope of convergence). Convergence holds along the entire sequence (a_N) used by the construction; there is no universality statement for arbitrary sequences $a_N \to 0$. Convergence is proved only for test functions in C_c^∞ away from the coincidence set, and step (iii) gives smoothness only on Ω_n, excluding contact terms on the diagonals.

5.4 The insertion formula and nondegeneracy at finite cutoff

The first two consequences of the construction-chain audit (Section 11.4) give the readable form

$$I_{N,p}(v, x; U) = a_N^{-d}\,(\texttt{places ev})(x; U), \qquad \texttt{ev} = \big(\texttt{terminal}^{-1} \circ \mathrm{diag}(t^{-1})\big)(v), \tag{5.9}$$

where $t \neq 0$ is the RG threshold scale and `terminal` undoes the terminal renormalization; the composite is a linear equivalence of the raw label space at each cutoff (`actual_insertion_formula`, proved by `rfl`). Consequently `every_primary_is_represented`: every placed primary section is an actual renormalized insertion times the common factor a_N^{-d}.

Lemma 5.20 (Nondegeneracy at finite cutoff, full proof)**.** *Let $L : V \to V$ be an invertible linear map and $c \neq 0$. If $p \in V$ is nonconstant, then $c\,L(p)$ is nonconstant.*

Proof. If $cL(p)$ equals a constant k, then $L(p) = k/c$ is constant; invertible linear maps carry the one-dimensional subspace of constant labels to itself, so $p = L^{-1}(k/c)$ is constant, a contradiction. □

Remark 5.21. This is the third consequence `nonconstant_primary_survives`, with its recorded boundary: it is a finite-cutoff statement, and no uniform bound on the inverse coefficient map as the cutoff is removed is provided, so it does not by itself exclude a degenerate continuum limit; exclusion of degeneracy rests on the field-content comparison (Proposition 2.12) and the short-distance results (Section 9). The fourth consequence `every_field_word_has_lattice_origin` gives only an iterated limit smear $\to$ row $\to \int t\,S$.`useAt`: no interchange of limits, no rate, no uniqueness across volume subsequences, and the right side uses the *mixed* table, not a selected phase.

6 The infinite-volume limit

6.1 Uniform bounds from the ultraviolet input

Assumption 6.1 ((UV1) local boundedness and equicontinuity)**.** For every finite label list v, every compact $K \subset \Omega_n$, and every multi-index α, there are $p_0(K)$, $M_{\alpha,K}$, and a modulus $\omega_{\alpha,K}(r) \downarrow 0$ such that

$$\sup_{p\geq p_0}\sup_{x\in K} \big|\partial_x^\alpha S_p(v; x)\big| \leq M_{\alpha,K}, \qquad \big|\partial^\alpha S_p(v; x) - \partial^\alpha S_p(v; y)\big| \leq \omega_{\alpha,K}(|x - y|)$$

for $x, y \in K$, $p \geq p_0$. The case $\alpha = 0$ is exactly the source document's bounds; $\alpha > 0$ follows by the same mechanism applied to the differentiated kernels. Certificates: the source document Remark 5.2; the mechanism is reflection positivity, chessboard estimates, the transfer-operator representation, and Cauchy estimates, with the construction-specific input being volume-uniform control of the renormalized lattice observables.

Lemma 6.2 (Readable derivation of the mechanism)**.** *Reflection positivity gives the OS positive semidefinite form $\langle A, B\rangle_{\mathrm{OS}}$; by Cauchy–Schwarz and permutation symmetry, $|S_p(v;x)| \leq \prod_i \langle A_i, A_i\rangle^{1/2}$; chessboard estimates bound each factor by finite-volume two-point functions, which are uniformly bounded at fixed volume by (RG5), giving local boundedness. Equicontinuity comes from analyticity of the kernels at separated points: around each point of K take a polydisc of radius δ; the Cauchy integral formula gives $|\nabla S_p| \leq CM_K/\delta$, hence a Lipschitz modulus $\omega_K(r) = CM_K r/\delta$.*

6.2 Countability and the diagonal extraction

Lemma 6.3 (Countability of labels and tests)**.** *There is a countable collection $\mathcal{L}$ of finite label lists and a countable dense set $\mathcal{T}_n \subset \mathcal{T}_n$ such that all $S_p(v; f)$ are determined by their values on $\mathcal{L} \times \mathcal{T}_n$.*

Proof. By Lemma 2.6 and Definition 2.7, each $\mathcal{F}_d$ is finite-dimensional; choose a basis $e_{d,1}, \dots, e_{d,n_d}$. Finite words in these basis elements are countable, and by multilinearity the value at arbitrary labels is a finite linear combination of basis-label values. For test functions, $C_c^\infty(\Omega_n)$ is separable in its usual topology: take polynomials with rational coefficients times smooth cutoffs with rational centers and radii, further cut off away from the diagonals, giving a countable dense set (Lemma 3.2). Continuity and linearity of $S_p(v;\cdot)$ make the values on the dense set determining. □

Theorem 6.4 (Subsequential infinite-volume limit)**.** *Under (UV1), for every strictly increasing sequence $p : \mathbb{N} \to \mathbb{N}$ of volume indices there exists a strictly increasing $q : \mathbb{N} \to \mathbb{N}$ and kernels $S_{\mathrm{mix}}(v;x)$, C^∞ on Ω_n, such that for all finite label lists v and all $f \in \mathcal{T}_n$,*

$$\lim_{k\to\infty} S_{p(q(k))}(v; f) = \int_{(\mathbb{R}^4)^n} f(x)\, S_{\mathrm{mix}}(v;x)\, \mathrm{d}x; \tag{6.1}$$

one and the same q works for all correlation functions simultaneously.

Proof. (i) *Compact exhaustion.* Take a countable compact exhaustion $K_m = \{x \in \Omega_n : |x_i| \leq m,\ |x_i - x_j| \geq 1/m\}$ of Ω_n; each K_m is compact and $\bigcup_m K_m = \Omega_n$.

(ii) *Termwise compactness — the objects are kernels, not functionals.* Fix a finite label list v, an index m, and a multi-index α. By (UV1) (which already carries the bounds for every α), the family $\big\{\partial^\alpha S_p(v;\cdot)\big|_{K_m} : p \geq p_0(K_m)\big\}$ is uniformly bounded by M_{α,K_m} and equicontinuous with modulus ω_{α,K_m}; by Arzelà–Ascoli it is relatively compact in $C^0(K_m)$, so every sequence of volume indices admits a subsequence along which $\partial^\alpha S_p(v;\cdot)|_{K_m}$ converges uniformly.

(iii) *Diagonal extraction over a countable index set.* Let $\mathcal{L}$ be the countable collection of basis label lists of Lemma 6.3 and let

$$\mathcal{I} := \mathcal{L} \times \{K_m : m \in \mathbb{N}\} \times \{\alpha \in \mathbb{N}_0^{4n}\},$$

which is countable; enumerate it as $(t_1, t_2, \dots)$ with $t_j = (v_j, K_{m_j}, \alpha_j)$. Start with $\mathcal{S}_0 = (p(1), p(2), \dots)$. Recursively, given $\mathcal{S}_{j-1}$, use (ii) to choose a strictly increasing subsequence $\mathcal{S}_j \subset \mathcal{S}_{j-1}$ along which $\partial^{\alpha_j} S_p(v_j;\cdot)|_{K_{m_j}}$ converges uniformly. Define $q(k)$ to be the k-th element of $\mathcal{S}_k$. Then q is strictly increasing, and for each fixed j we have $q(k) \in \mathcal{S}_j$ for all $k \geq j$;

hence $\partial^{\alpha_j} S_{p(q(k))}(v_j;\cdot)|_{K_{m_j}}$ converges uniformly. Thus one sequence q makes every countable item converge. Denote the limits by $S^{(\alpha)}_{\mathrm{mix}}(v;\cdot)$ on each K_m.

(iv) *Consistency and smoothness.* If $K_m \subset K_{m'}$, the restriction of the $K_{m'}$-limit to K_m coincides with the K_m-limit, because both are limits along the same subsequence. For fixed v and m, all derivatives converge uniformly on K_m; by the Weierstrass theorem, $S^{(0)}_{\mathrm{mix}}(v;\cdot) \in C^\infty(K_m)$ and $\partial^\alpha S^{(0)}_{\mathrm{mix}} = S^{(\alpha)}_{\mathrm{mix}}$. Patching over m gives $S_{\mathrm{mix}}(v;\cdot) \in C^\infty(\Omega_n)$.

(v) *From kernels to functionals.* Let $f \in \mathcal{T}_n$ with $\operatorname{supp} f \subset K_m$. By Lemma 3.4 and Theorem 5.17, $S_p(v;f) = \int f(x)\, S_p(v;x)\,\mathrm{d}x$; the case $\alpha = 0$ of (iii) gives uniform convergence $S_{p(q(k))}(v;\cdot) \to S_{\mathrm{mix}}(v;\cdot)$ on K_m, hence

$$\lim_{k\to\infty} S_{p(q(k))}(v;f) = \int f(x)\, S_{\mathrm{mix}}(v;x)\,\mathrm{d}x.$$

For a general label list v, multilinearity and the basis $\mathcal{L}$ of Lemma 6.3 express $S_p(v;f)$ as a finite linear combination of the $S_p(v';f)$ with $v' \in \mathcal{L}$; since q works for all of them, it works for v. No separate dense-set approximation of test functions is needed: the countable set $\mathcal{I}$ already ranges over all labels, all compacts, and all derivative orders, and the convergence of kernels is uniform on the support of f. □

Remark 6.5. Jaffe–Witten regard a subsequential limit obtained by compactness as insufficient unless further properties of the limit are established [JW]; Sections 7 and 9 supply exactly that further information. Whether the limit depends on the chosen subsequence remains open.

6.3 Clustering would force all phases to agree

Lemma 6.6. *Suppose the phase decomposition (Theorem 7.13) holds, every component clusters, the measure ν has full support (Proposition 7.18), and $\xi \mapsto S_\xi(A)$ is continuous. If S_{mix} satisfies clustering, then for every word A and every $\xi \in X$, $S_\xi(A) = S_{\mathrm{mix}}(A)$.*

Proof. Let A_s be A translated by s in the time direction, and let A^* be the adjoint word. Reflection positivity (E2) gives, for each phase, a positive semidefinite OS form; by Cauchy–Schwarz and translation invariance of each phase,

$$\bigl|S_\xi(A^*A_s)\bigr| \le S_\xi(A^*A)^{1/2}\, S_\xi(A_s^*A_s)^{1/2} = S_\xi(A^*A) \qquad \text{for all } s \ge 0. \tag{6.2}$$

The function $g(\xi) := S_\xi(A^*A)$ is continuous on X (Gelfand transform of $M_{A^*A} \in \mathcal{A}$) and non-negative, and by Theorem 7.13,

$$\int_X g\,\mathrm{d}\nu = S_{\mathrm{mix}}(A^*A) < \infty. \tag{6.3}$$

Thus (6.2)–(6.3) supply a ν-integrable dominating function, uniform in s (note that (UV1) is not used here: its bounds are local on compact sets, whereas A_s leaves every fixed compact as $s \to \infty$). Dominated convergence and clustering of every component (Proposition 7.15) give

$$\lim_{s\to\infty} S_{\mathrm{mix}}(A^*A_s) = \int_X \lim_{s\to\infty} S_\xi(A^*A_s)\,\mathrm{d}\nu = \int_X \overline{S_\xi(A)}\, S_\xi(A)\,\mathrm{d}\nu = \int_X \bigl|S_\xi(A)\bigr|^2\,\mathrm{d}\nu.$$

On the other hand, clustering of S_{mix} applied to the same pairing gives $\lim_{s\to\infty} S_{\mathrm{mix}}(A^*A_s) = S_{\mathrm{mix}}(A^*)\, S_{\mathrm{mix}}(A) = \bigl|S_{\mathrm{mix}}(A)\bigr|^2$, since $S_{\mathrm{mix}}(A^*) = \overline{S_{\mathrm{mix}}(A)}$ and $S_{\mathrm{mix}}(A) = \int_X S_\xi(A)\,\mathrm{d}\nu$. Therefore

$$\int_X \bigl|S_\xi(A)\bigr|^2\,\mathrm{d}\nu - \Bigl|\int_X S_\xi(A)\,\mathrm{d}\nu\Bigr|^2 = 0, \qquad \text{i.e.} \qquad \int_X \bigl|S_\xi(A) - S_{\mathrm{mix}}(A)\bigr|^2\,\mathrm{d}\nu(\xi) = 0,$$

so $S_\xi(A) = S_{\mathrm{mix}}(A)$ for ν-almost every ξ. The function $\xi \mapsto S_\xi(A) - S_{\mathrm{mix}}(A)$ is continuous and ν has full support, so its zero set, containing a set of full measure, cannot be a proper closed set; it is all of X. □

Remark 6.7. The source document (§6.3) uses this computation to show that component clustering does *not* imply clustering of the mixed theory, and stresses that even with clustering, one does not obtain sufficient infrared control to prove uniqueness of the infinite-volume limit. Lemma 6.6 is the precise "clustering ⇒ all phases agree" statement and nothing more.

7 Phase decomposition, clustering, and full support

Throughout this section, fix one infinite-volume theory extracted in Theorem 6.4.

7.1 The OS Hilbert space and the ground-state subspace

Reflection positivity defines a positive semidefinite form on products of fields at positive Euclidean times; quotienting the null space and completing gives the OS Hilbert space $\mathcal{H}_{\mathrm{OS}}$, with time semigroup $E(t) = e^{-tH}$ $(H \geq 0)$, vacuum Ω_0, ground-state subspace $\mathcal{H}_0 = \ker H$, and orthogonal projection P_0 onto $\mathcal{H}_0$.

Lemma 7.1. *$E(t)|_{\mathcal{H}_0} = \mathrm{id}$, and $E(t) \to P_0$ strongly as $t \to \infty$.*

Proof. By the spectral theorem $E(t) = \int_{[0,\infty)} e^{-t\lambda}\, \mathrm{d}E_H(\lambda)$; on $\mathcal{H}_0 = E_H(\{0\})\mathcal{H}$ one has $\lambda = 0$, so $e^{-t\lambda} = 1$. For any ψ,

$$\|(E(t) - P_0)\psi\|^2 = \int_{\lambda>0} e^{-2t\lambda}\, \mathrm{d}\mu_\psi(\lambda), \qquad \mu_\psi(B) = \langle \psi, E_H(B)\psi\rangle;$$

the integrand tends to 0 pointwise and is dominated by 1, with μ_ψ finite, so dominated convergence gives convergence to 0. □

Proposition 7.2 (Clustering, vacuum uniqueness, and (PH4) are the same statement)**.** *Let $(\mathcal{H}, H, \Omega)$ be the OS/Wightman triple of a phase, $E(t) = e^{-tH}$, $P_0 = |\Omega\rangle\langle\Omega|$, and let $\mathcal{D} \subset \mathcal{H}$ be the dense span of finite positive-time field words. Under (AN1) E2, the following are equivalent:*

(1) (PH4) $\ker H = \mathbb{C}\,\Omega$, *equivalently* $P_{\ker H} = P_0$;

(2) E4 clustering $\langle \psi, E(t)\chi\rangle \to \langle \psi, \Omega\rangle\langle\Omega, \chi\rangle$ *for all* $\psi, \chi \in \mathcal{D}$;

(3) $\dim \ker H = 1$.

Proof. Decompose $\mathcal{H} = \ker H \oplus (\ker H)^\perp$ and write $\psi = \psi_0 + \psi_\perp$, $\psi_0 = P_{\ker H}\psi$. For $\psi_\perp \perp \ker H$, the spectral theorem gives

$$\|(E(t) - P_{\ker H})\psi_\perp\|^2 = \int_{(0,\infty)} e^{-2t\lambda}\, \mathrm{d}\mu_{\psi_\perp}(\lambda) \xrightarrow[t\to\infty]{} 0$$

by dominated convergence: $e^{-2t\lambda} \to 0$ for every $\lambda > 0$, the integrand is bounded by 1, and $\mu_{\psi_\perp}$ is a finite measure. Hence $E(t) \to P_{\ker H}$ strongly.

(1)⇒(2): if $P_{\ker H} = P_0$, then $\langle \psi, E(t)\chi\rangle \to \langle \psi, P_0\chi\rangle = \langle \psi, \Omega\rangle\langle\Omega, \chi\rangle$.

(2)⇒(1): for $\psi, \chi \in \mathcal{D}$, strong convergence gives $\langle \psi, P_{\ker H}\chi\rangle = \lim_t \langle \psi, E(t)\chi\rangle = \langle \psi, P_0\chi\rangle$; two bounded operators agreeing on a dense domain coincide, so $P_{\ker H} = P_0$ and $\ker H = \operatorname{ran} P_0 = \mathbb{C}\,\Omega$.

(1)⇔(3): P_0 is a rank-one projection. □

Remark 7.3. (PH4) is therefore not an external input: it is a spectral conclusion from (AN1) E2 + E4. The repository README confirms this status ("the Wightman spectrum condition and uniqueness of the vacuum are the spectral conclusions in the public challenge"). Appendix B records it as a proved conclusion, reducing the external-input count from seventeen to sixteen. In Section 7.4 and Proposition 8.5 the phrase "by (PH4)" henceforth refers to Proposition 7.2; no circularity arises, since Proposition 7.2 uses only (AN1) and the spectral theorem.

This equivalence is also the consistency input for the autocorrelation criterion of Theorem E.5 in Appendix E.

7.2 Word operators and the commutative algebra

For a word A with distinct points in a time slab $0 < t < s$, define $B_{s,A}$: on vectors represented by positive-time words, shift the existing insertions forward by s and prepend the insertions of A; let

$$M_A = P_0 B_{s,A}\big|_{\mathcal{H}_0}, \qquad \mathcal{A} = C^*\big(\{M_A, M_A^*\}\big) \subset \mathcal{B}(\mathcal{H}_0), \tag{7.1}$$

the norm-closed unital algebra of (7.1).

Assumption 7.4 ((PH1) time-translation conjugation)**.** $B_{s,A} = T_s B_{0,A} T_{-s}$ with T_s implemented by the time semigroup, and $P_0 T_s = P_0 = T_s P_0$ on $\mathcal{H}_0$. Certificates: `YM.OS.Spread`, `YM.OS.Semigroup.ground`.

Lemma 7.5. *Under (PH1),* $M_A = P_0 B_{0,A}|_{\mathcal{H}_0}$*, independent of* s*.*

Proof. $P_0 B_{s,A} P_0 = P_0 T_s B_{0,A} T_{-s} P_0 = P_0 B_{0,A} P_0$, using $T_{-s} P_0 = P_0$ and $P_0 T_s = P_0$ (Lemma 7.1). Restricting to $\mathcal{H}_0$ gives the claim. □

Assumption 7.6 ((PH2) commutativity)**.** $\mathcal{A}$ is commutative. Certificates: `YM.OS.Spread.algebra`, `YM.OS.Ground.at_any`. Readable reason: time translation is trivial on the ground subspace (Lemma 7.1), so composing two positive-time words depends only on the vacuum matrix element of both insertions placed in a common slab; since $E(t) \to P_0$, intermediate states are projected to $\mathcal{H}_0$, where the time coordinate produces no phase, and exchanging the order of the two groups does not change the matrix element, giving $M_A M_B = M_B M_A$. We record (PH2) as a construction input rather than deriving it from weaker assumptions.

Assumption 7.7 ((PH3) faithfulness)**.** No nonzero $a \in \mathcal{A}$ satisfies $a\Omega_0 = 0$ (Ω_0 is a separating vector for $\mathcal{A}$). Certificates: `YM.OS.Ground.attains`, `YM.OS.Ground.supports`.

Assumption 7.8 ((PH4) one-dimensional ground state of the selected phase)**.** For the selected phase ξ^*, the ground-state subspace of its OS/Wightman Hamiltonian,

$$\mathcal{H}_0^{\xi^*} = \ker H_{\xi^*} = \big\{\psi : U(a)\psi = \psi \text{ for all } a\big\},$$

is one-dimensional, spanned by Ω_{ξ^*}. Certificates: `YM.OS.Cluster.lifts`, `YM.Trace.submits`, `YM.Trace.Echo.retreats`, `YM.Trace.echo_tab`, `selected`; the set E of Appendix A, equation (A.6), satisfies $\nu(E) = 1$ and $\dim G_\xi = 1$ on E, and $\xi^* \in E$.

Remark 7.9. Proposition 7.2 shows that (PH4) is derivable from (AN1) E2 + E4 alone; the numbered statement is retained for compatibility with the source document's certificates, but it is a proved spectral conclusion, not an independent input (see Appendix B).

7.3 The Gelfand spectrum and the phase measure

Proposition 7.10. *Under (PH2), $\mathcal{A}$ is a separable unital commutative C^*-algebra; its Gelfand spectrum $X = \mathrm{Spec}(\mathcal{A})$ is a nonempty compact metrizable space, and the Gelfand transform $a \mapsto \hat{a}$, $\hat{a}(\xi) = \xi(a)$, is an isometric $*$-isomorphism $\mathcal{A} \to C(X)$.*

Proof. The Gelfand–Naimark theorem gives the isometric $*$-isomorphism onto $C(X)$ with X compact Hausdorff. Separability: words A are generated by the countable label collection and the countable test basis of Lemma 6.3, so $\{M_A\}$ has a countable dense subset, and the generated C^*-algebra is separable; a separable commutative C^*-algebra has metrizable spectrum ($C(X)$ separable implies X metrizable). □

Definition 7.11 (Vacuum state)**.** $\omega(a) = \langle \Omega_0, a\,\Omega_0 \rangle$, a positive normalized linear functional with $\|\omega\| = \|\Omega_0\|^2 = 1$ (Appendix A, `Properties.normalized`).

Proposition 7.12. *There is a unique regular Borel probability measure ν on X with*

$$\omega(a) = \int_X \hat{a}(\xi)\,\mathrm{d}\nu(\xi) \qquad (a \in \mathcal{A}). \tag{7.2}$$

Proof. ω is a positive linear functional of norm 1 on $C(X)$ via the Gelfand transform; the Riesz–Markov–Kakutani theorem gives a unique regular Borel measure representing it, and $\nu(X) = \omega(1) = 1$. □

Theorem 7.13 (Phase decomposition)**.** *For each word A, let $S_\xi(A) = \xi(M_A) = \hat{M}_A(\xi)$. Then*

$$S_{\mathrm{mix}}(A) = \int_X S_\xi(A)\,\mathrm{d}\nu(\xi), \tag{7.3}$$

extended by multilinearity and continuity to all finite label lists and $f \in \mathcal{T}_n$. Each S_ξ satisfies the OS axioms including clustering.

Proof. By construction, $S_{\mathrm{mix}}(A) = \langle \Omega_0, M_A \Omega_0 \rangle = \omega(M_A) = \int_X \hat{M}_A\,\mathrm{d}\nu = \int_X S_\xi(A)\,\mathrm{d}\nu$ by Proposition 7.12. Linearity and continuity extend the identity from words to general tests (countable density of Lemma 6.3 plus the uniform bounds (UV1)). The OS axioms for each S_ξ are those recorded in Appendix C: E0/E0′ (diagonal-flat Schwartz tests, growth bounds), E1 (Euclidean covariance), E2 (reflection positivity), E3 (permutation), E4 (clustering, proved in Proposition 7.15 below), normalization and analyticity; they are packaged from the construction's `Properties.*` and `YMOSBridge.*` certificates as `verifiedOS`. □

7.4 Clustering of each phase

Lemma 7.14 (Two distinct statements must be separated)**.** (a) GNS representation of the character. *Let $\omega_\xi(a) = \xi(a)$ be the character state on $\mathcal{A}$ and $(\pi_\xi, \mathcal{H}_\xi^{\mathrm{GNS}}, \Omega_\xi)$ its GNS representation. Then $\mathcal{H}_\xi^{\mathrm{GNS}} = \mathbb{C}\,\Omega_\xi$ and $\pi_\xi(a) = \xi(a)\,\mathrm{id}$.*

(b) Ground-state subspace of the phase. *Let $\mathcal{H}_\xi$ be the OS/Wightman Hilbert space of phase ξ (which contains excitations and is not the one-dimensional space of (a)). The algebra $\mathcal{A}$ acts on the ground-state subspace $\mathcal{H}_0^\xi$ through the scalars $\xi(M_A)$; for the selected phase, (PH4) gives $\mathcal{H}_0^{\xi^*} = \mathbb{C}\,\Omega_{\xi^*}$ and $P_0^{\xi^*} = |\Omega_{\xi^*}\rangle\langle\Omega_{\xi^*}|$.*

Proof. (a) By multiplicativity of ξ and $\xi(a^*a) = |\xi(a)|^2$,

$$\big\|\pi_\xi(a)\Omega_\xi - \xi(a)\Omega_\xi\big\|^2 = \xi(a^*a) - |\xi(a)|^2 = 0,$$

so $\pi_\xi(a)\Omega_\xi = \xi(a)\Omega_\xi$ for all a. GNS vectors are cyclic, hence $\mathcal{H}_\xi^{\mathrm{GNS}} = \overline{\pi_\xi(\mathcal{A})\Omega_\xi} = \mathbb{C}\Omega_\xi$; consequently $\pi_\xi(\mathcal{A}) = \mathbb{C}\,\mathrm{id}$ and its commutant is $\pi_\xi(\mathcal{A})' = \mathcal{B}(\mathcal{H}_\xi^{\mathrm{GNS}}) = \mathbb{C}$ because the space is one-dimensional — not the other way round.

(b) On the ground-state subspace, time translation acts trivially (Lemma 7.1), so the relation $S_\xi(A) = \langle\Omega_\xi, M_A\Omega_\xi\rangle$ together with the multiplicativity of ξ gives $M_A\psi = \xi(M_A)\psi$ for $\psi \in \mathcal{H}_0^\xi$; the one-dimensionality for $\xi = \xi^*$ is the conclusion of Proposition 7.2 (derived from (AN1) E2 + E4). □

Proposition 7.15 (Clustering of every phase)**.** *For positive-time words A, B and $s \to \infty$, $S_\xi(A\,B_s) \to S_\xi(A)S_\xi(B)$.*

Proof. On the phase's OS Hilbert space $\mathcal{H}_\xi$ the time semigroup $E_\xi(s) = e^{-sH_\xi}$ converges strongly to P_0^ξ (Lemma 7.1). Write $S_\xi(AB_s)$ as $\langle\Omega_\xi, \pi(M_A)E_\xi(s)\pi(M_B)\Omega_\xi\rangle$; then

$$\lim_{s\to\infty} S_\xi(AB_s) = \big\langle\Omega_\xi, \pi(M_A)P_0^\xi\,\pi(M_B)\Omega_\xi\big\rangle.$$

For the selected phase, $P_0^{\xi^*} = |\Omega_{\xi^*}\rangle\langle\Omega_{\xi^*}|$ by Proposition 7.2 (derived from (AN1) E2 + E4, via Lemma 7.14(b)), and $M_A\Omega_{\xi^*} = \xi^*(M_A)\Omega_{\xi^*}$ on the ground state; hence the limit equals $\xi^*(M_A)\,\xi^*(M_B) = S_{\xi^*}(A)S_{\xi^*}(B)$. For a general phase the same computation gives $S_\xi(AB_s) \to S_\xi(A)S_\xi(B)$ with P_0^ξ the ground-state projection of that phase (the formalization's phase clustering certificates `YM.OS.Cluster.lifts` et al. cover every character). □

7.5 Second moments, full support, and the word origin of phases

Lemma 7.16. *For distinct positive-time words w, u,*

$$\int_X \xi(w)\,\xi(u)\,\mathrm{d}\nu(\xi) = \big\langle\Omega_0, \mathtt{slot}(w)\,\mathtt{slot}(u)\,\Omega_0\big\rangle. \tag{7.4}$$

Proof. By multiplicativity of characters, $g(\xi) := \xi(w)\xi(u) = \xi(M_wM_u)$ is the Gelfand transform of $M_wM_u \in \mathcal{A}$, hence a continuous function on X; integrating and using Proposition 7.12,

$$\int_X g\,\mathrm{d}\nu = \omega(M_wM_u) = \langle\Omega_0, M_wM_u\Omega_0\rangle,$$

which by the definition of M_w, M_u is $\langle\Omega_0, \mathtt{slot}(w)\,\mathtt{slot}(u)\,\Omega_0\rangle$. □

Remark 7.17. This is the fifth construction-chain consequence (`phase_second_moments_are_ground_products`); it is stronger than the barycenter identity, but it does not factorize a phase's multipoint correlation into one-point functions, nor does it say that each phase is itself a lattice subsequential limit or that there is only one phase.

Proposition 7.18 (Full support)**.** *Under (PH3), every nonempty open subset $U \subset X$ satisfies $\nu(U) > 0$.*

Proof. Suppose $\nu(U) = 0$. Choose a nonzero continuous $f : X \to \mathbb{C}$ supported in U (Urysohn's lemma; X is compact metrizable) and let $a \in \mathcal{A}$ be its Gelfand inverse image, $\hat{a} = f$. Then

$$\|a\Omega_0\|^2 = \omega(a^*a) = \int_X |f|^2 \, \mathrm{d}\nu = 0,$$

so $a\Omega_0 = 0$; by (PH3) faithfulness, $a = 0$, hence $f = \hat{a} = 0$, contradicting the choice of f. □

Remark 7.19. Recorded in Appendix A as `fullSupport` (certificates `YM.OS.Ground.attains`, `YM.OS.Ground.supports`); it is the key input for the "almost everywhere ⇒ everywhere" step in Lemma 6.6.

8 Osterwalder–Schrader reconstruction and the Wightman axioms

8.1 The analytic input and the two test spaces

Assumption 8.1 ((AN1))**.** The correlations of the selected phase ξ^* satisfy the OS axioms E0 (temperedness), E0′ (arity growth), E1 (Euclidean invariance), E2 (reflection positivity), E3 (permutation), E4 (clustering), normalization, and analyticity of Appendix C; the analytic continuation $F_n^{\xi^*}(v; z, x)$, $\operatorname{Im} z_1 < \cdots < \operatorname{Im} z_n$, satisfies $S^{\xi^*}(v; \tau, x) = F_n^{\xi^*}(v; i\tau, x)$ for $\tau_1 < \cdots < \tau_n$. Certificates: `Properties.continuation`, `.analytic`, `YM.Boundary.Waves.Region`, `YMOSBridge.verifiedOS`, `complete_OS_on_this_hierarchy`.

Remark 8.2 (Two levels of test space). Three distinct distributional statements must be kept apart. (i) $\mathcal{T}_n = C_c^\infty(\Omega_n)$: tests avoiding the coincidence diagonals, used for lattice convergence. (ii) $\mathcal{S}_0((\mathbb{R}^4)^n)$: joint Schwartz tests flat to every order on each diagonal Δ_{ij}, used for the independent OS verification of Appendix C; it contains (i) but cannot detect distributions supported on the diagonals. (iii) The reconstruction separately supplies tempered Wightman distributions on unrestricted joint Schwartz tests. The Schwinger functional of the independent check is

$$\mathcal{S}_n(f) = \int_{(\mathbb{R}^4)^n} K_{\Theta,n}^{\xi^*}(x)\, f(x)\, \mathrm{d}x, \qquad K_{\Theta,n}^{\xi^*}(x) := S_{\Theta,\dots,\Theta}^{\xi^*}(x), \tag{8.1}$$

with absolute convergence, continuity on $\mathcal{S}_0$, and the E0′ bound $\|\mathcal{S}_n(f)\| \le \beta^n (n!)^A \|f\|_{(ns,ns)}$ (certificates `YMOSBridge.targetKernel_bound`, `schwinger_integrable`, `E0_tempered`, `E0_prime_growth`).

8.2 Wightman boundary values and the Hilbert space

Let $\mathcal{P}$ be the complex vector space of finite formal words $(v_1, f_1)\cdots(v_n, f_n)$ with $f_i \in \mathcal{S}(\mathbb{R}^4; \mathbb{C})$, including the empty word; let $*$ reverse words, adjoin labels, and conjugate tests and coefficients. Take the Wightman boundary values $W^\xi(u) = \lim_{\operatorname{Im} z \to 0^+} F(z)$ in the sense of Schwartz tests, and define

$$\langle u, w\rangle_\xi = W^\xi(u^* w), \qquad N_\xi = \{u : \langle u, u\rangle_\xi = 0\}, \qquad \mathcal{H}_\xi = \overline{\mathcal{P}/N_\xi}, \tag{8.2}$$

with Ω_ξ the empty word, $\mathcal{D}_\xi = \mathcal{P}/N_\xi$ dense in $\mathcal{H}_\xi$, and $\Phi_v^\xi(f)[w] = [(v, f)w]$.

Lemma 8.3. *$\langle\cdot,\cdot\rangle_\xi$ is positive semidefinite; N_ξ is a left ideal of $\mathcal{P}$, and $\langle u, w\rangle = 0$ for all w if and only if $u \in N_\xi$.*

Proof. Positive semidefiniteness comes from E2 reflection positivity through analytic continuation (the OS theorem). By Cauchy–Schwarz, $|\langle u, w\rangle|^2 \le \langle u, u\rangle\langle w, w\rangle = 0$ when $u \in N_\xi$, so $\langle u, w\rangle = 0$ for all w; conversely take $w = u$. If $u \in N_\xi$ and $b = (v, f)$ is arbitrary, then for every w, $\langle bu, w\rangle = W(u^* b^* w) = \langle u, b^* w\rangle = 0$ since $b^* w \in \mathcal{P}$; hence $bu \in N_\xi$. □

Lemma 8.4. *$\Phi_v^\xi(f) : \mathcal{D}_\xi \to \mathcal{D}_\xi$ is well defined and complex-linear; $\langle \Phi_v(f)\psi, \chi\rangle = \langle \psi, \Phi_{v^*}(\bar f)\chi\rangle$; and $\mathcal{D}_\xi$ is exactly the span of field words applied to Ω_ξ (cyclicity).*

Proof. Well-definedness: if $[w] = [w']$, then $w - w' \in N_\xi$, and by Lemma 8.3 (left ideal) $(v, f)(w - w') \in N_\xi$. Linearity is clear. The adjoint relation: $\langle \Phi_v(f)[u], [w]\rangle = W(u^*(v, f)^* w) = \langle [u], \Phi_{v^*}(\bar f)[w]\rangle$. Cyclicity: $\mathcal{P}$ is generated by words, and $[w] = \Phi_{v_1}(f_1)\cdots\Phi_{v_n}(f_n)\Omega_\xi$. □

8.3 The Wightman axioms, proved item by item

Proposition 8.5 (W0: spectrum and uniqueness of the vacuum)**.** *Translations are implemented by a strongly continuous unitary representation $U(a) = \int e^{ip\cdot a}\,\mathrm{d}E(p)$ with $\operatorname{supp} E \subset \overline{V^+}$; every translation-invariant vector is a complex multiple of Ω_ξ.*

Proof. Translation invariance of W^ξ gives isometries $T(a)[w] = [(v_1, f_1(\cdot - a))\cdots]$ on $\mathcal{P}/N_\xi$; being isometric on a dense domain, $T(a)$ extends uniquely to a unitary on $\mathcal{H}_\xi$, and strong continuity follows from continuity of $f \mapsto f(\cdot - a)$ in the Schwartz topology. For the spectral support, by (AN1) the Euclidean correlations of the selected phase are boundary values of functions F_n holomorphic in the tube

$$\mathcal{T}_{n-1} = \big\{(\zeta_1, \dots, \zeta_{n-1}) \in (\mathbb{C}^4)^{n-1} : \operatorname{Im}\zeta_j \in V^+\big\}, \qquad \zeta_j = z_{j+1} - z_j,$$

and obey the tempered growth bound of E0′. For $n = 2$ write $W_2(\xi) = \lim_{\eta\to 0,\, \eta\in V^+} F_1(\xi + i\eta)$ in $\mathcal{S}'(\mathbb{R}^4)$, where $\xi = x_2 - x_1$. Since F_1 is holomorphic on the tube $\mathbb{R}^4 + iV^+$ with polynomial growth in $1/\operatorname{Im}\zeta$ and in $|\zeta|$, the Paley–Wiener–Schwartz theorem for tube domains implies that the Fourier transform $\mu = \widehat{W_2}$ is a tempered measure with $\operatorname{supp}\mu \subset \overline{V^+} = \{p : p_0 \ge |\vec p|\}$. For general n, the same argument applied to each difference variable $\xi_j = x_j - x_{j+1}$ places the joint energy–momentum measure in $(\overline{V^+})^{n-1}$, and translation covariance (E1) then forces the spectral measure of $U(a) = \int e^{ip\cdot a}\,\mathrm{d}E(p)$ to satisfy $\operatorname{supp} E \subset \overline{V^+}$. *External input used*: the Paley–Wiener–Schwartz theorem for tube domains; *expanded here*: the reduction to difference variables, the $n = 2$ Laplace/Fourier step, and the passage from joint to single-particle support. Vacuum uniqueness: the translation-invariant vectors of $\mathcal{H}_\xi$ are exactly $E_H(\{0\})\mathcal{H}_\xi = \mathcal{H}_0^\xi$; by Proposition 7.2 (derived from (AN1) E2 + E4) this space is $\mathbb{C}\,\Omega_\xi$. Thus W0's vacuum uniqueness follows from (AN1), with no circularity; Proposition 7.15 supplies the equivalent clustering statement.

The same equivalence is the AUTO input of Appendix E (Theorem E.5). □

Proposition 8.6 (W1: regularity, adjoints, cyclicity)**.** *$f \mapsto \Phi_v(f)\psi$ is continuous from $\mathcal{S}$ to $\mathcal{H}_\xi$; the adjoint relation holds; $\mathcal{D}_\xi$ is dense, contains Ω_ξ, is preserved by all fields, and equals the span of field words.*

Proof. Continuity follows from the tempered growth bounds of E0′ and control by Schwartz seminorms; adjoints and cyclicity are Lemma 8.4; density is built into the definition of $\mathcal{H}_\xi$ as a completion. □

Proposition 8.7 (W2: Poincaré covariance)**.** *The proper orthochronous Poincaré group has a strongly continuous unitary representation $U(g)$, and there is a complex-linear representation ρ on labels with $U(g)\Phi_v(f)U(g)^{-1} = \Phi_{\rho(g)v}(f_g)$, $f_g(x) = f(\Lambda^{-1}(x - a))$, and $U(g)\Omega_\xi = \Omega_\xi$.*

Proof. E1 gives covariance of the Euclidean correlations under $O(4)$ (and its complexification); by (AN1) the analytic continuation and the Bargmann–Hall–Wightman theorem extend the domain to the tube of the complex Lorentz group, on which real proper orthochronous Lorentz transformations act on the boundary values, giving $U(\Lambda)$; combining with translations gives the Poincaré representation. The label representation ρ comes from the tensor transformation law of curvature polynomials (the $O(4)$ action of Definition 2.5), is complex-linear, and satisfies $\rho(gh) = \rho(g)\rho(h)$. Vacuum invariance follows from translation invariance and the definition of ρ. *External input used*: the Bargmann–Hall–Wightman theorem (holomorphy on the extended tube); the $O(4)$-covariance of the Euclidean correlations (E1), the invariance of the inner product under the induced label action (reflection positivity), and the unitary extension are expanded here. □

Proposition 8.8 (W3: local commutativity)**.** *If* $\operatorname{supp} f$ *and* $\operatorname{supp} h$ *are pointwise spacelike separated, then* $[\Phi_v(f), \Phi_w(h)] = 0$ *on* $\mathcal{D}_\xi$*.*

Proof. For two spacelike separated points x, y there is a complex Lorentz transformation reversing their time order continuously within the tube without leaving the domain of holomorphy (the Jost point argument); by E3 permutation symmetry and uniqueness of analytic continuation, the two boundary values agree. At the level of words this is exactly $[(v, f)(w, h)] = [(w, h)(v, f)]$ in $\mathcal{P}/N_\xi$; by Lemma 8.4, $\Phi_v(f)\Phi_w(h)\psi = \Phi_w(h)\Phi_v(f)\psi$. *External input used*: existence of the Jost-point path in the extended tube (a standard corollary of the Bargmann–Hall–Wightman theorem). □

Remark 8.9 (Audit status of W0–W3). Propositions 8.5–8.8 are complete relative to (AN1) plus three classical OS–Wightman theorems: the Paley–Wiener–Schwartz theorem for tube domains, the Bargmann–Hall–Wightman theorem, and the Jost-point theorem. These are marked as external inputs where used and are not re-proved; every other proof in Sections 2–9.5 is closed. This aligns the phrase "complete proofs" in the introduction with the sixteen external inputs registered in Appendix B.

Theorem 8.10 (Reconstruction of the component theories)**.** *Under (AN1), for every character* $\xi \in X$ *the component correlations* S_ξ *reconstruct to a Wightman theory on* $\mathcal{H}_\xi$ *with vacuum* Ω_ξ *and fields* Φ^ξ *on* $\mathcal{D}_\xi$*, satisfying W0–W3 (Propositions 8.5–8.8), analytic continuation and recovery of correlations, and faithfulness: every nonzero label* v *has a test* f *with* $\Phi^\xi_v(f) \neq 0$*. The formal certificates are* `YangMills.Hierarchy.Reconstructed.axioms`*,* `.analytic`*,* `.root`*,* `.boundary`*, and* `.faithful`*.*

Proof of faithfulness. Suppose $\langle \Phi_v(f)[u], [w]\rangle = 0$ for all f and all words u, w; then all Wightman correlations containing v vanish. By uniqueness of the boundary values of one fixed analytic continuation (the seventh chain consequence, Section 11.4), all separated-point Euclidean correlations containing v vanish, i.e. $v \in N_d$ (Definition 2.7), so $[v] = 0$ in $\mathcal{F}_d$, a contradiction. See Appendix A for the certificates (`FieldSystem.Faithful`, `YM.Work.C26.answer`, `YM.Boundary.Soil.sees`); cyclicity (`registered_words_dense`) uses the full label collection, and the smaller observable family behind the scalar OS bridge (check A1) does not by itself prove it. □

8.4 Relation to the mixed theory and Euclidean invariance

The mixed theory S_{mix} lacks clustering, so its OS reconstruction does not guarantee a unique vacuum; uniqueness of the vacuum holds for the selected phase and the selected reconstruction. The sixth chain consequence (`reconstructed_words_are_selected_boundary`) ensures that the vacuum expectation of every reconstructed relativistic field word equals the expectation of the same word in the selected phase's boundary distribution (via the common `qt.meets` certificate); the seventh

gives uniqueness of boundary values only for one fixed analytic continuation. Neither addresses uniqueness of the infrared limit, of the phase, or of the continuation from Euclidean data.

Euclidean invariance in the continuum limit has three parts (the source document §6.4): (i) the gauge fixing preserves the lattice symmetries (Lemma 5.2); (ii) there is a divergence-free stress tensor used for local Ward identities (Section 9.1); (iii) Euclidean symmetry is not spontaneously broken in the limiting state because the stress-energy flux decays sufficiently rapidly at infinity, shown by an argument using reflection positivity in several orthogonal directions. We record part (iii) as a construction input underlying E1.

9 Identification as Yang–Mills theory

We record the properties identifying the construction as Yang–Mills theory. Throughout, $S_v^s(\cdot)$ and $S^s(v;\cdot)$ are used interchangeably, where s refers either to the mixed theory or to a component theory indexed by $\xi \in X$.

9.1 The stress tensor and its Ward identity

Assumption 9.1 ((Y1) uniqueness of the stress tensor)**.** In $\mathcal{F}_4$, the elements that are symmetric rank-two tensors, fixed by field conjugation, centered with respect to S_{mix}, transform as rank-two tensors under coordinate permutations and sign changes, and whose divergence vanishes in $\mathcal{F}_5$, form a one-dimensional real vector space; the stress tensor is the vector in this space whose normalization is fixed by the Ward identity below. Certificates: the source document §7.1 and `WardFormulationAudit.lean`.

Lemma 9.2 (Readable derivation of the candidate space)**.** *Weight-four symmetric rank-two tensor candidates are spanned by* $A_{\mu\nu} = \mathrm{Tr}(F_{\mu\alpha}F_{\nu\alpha})$ *and* $B_{\mu\nu} = \delta_{\mu\nu}\,\mathrm{Tr}(F_{\alpha\beta}F_{\alpha\beta})$*. On the free sourceless shell (Definition 2.10), writing* $Q = \mathrm{Tr}(F_{\alpha\beta}F_{\alpha\beta})$*,*

$$\partial_\mu A_{\mu\nu} = \mathrm{Tr}\left((\partial_\mu F_{\mu\alpha})F_{\nu\alpha}\right) + \mathrm{Tr}\left(F_{\mu\alpha}\,\partial_\mu F_{\nu\alpha}\right) = 0 + Y,$$

the first term vanishing by the sourceless Maxwell equation $\partial_\mu F_{\mu\alpha} = 0$*. For* $Y = \mathrm{Tr}(F_{\mu\alpha}\partial_\mu F_{\nu\alpha})$*, the Bianchi identity (Lemma 2.4) in the form* $\partial_\mu F_{\nu\alpha} = \partial_\nu F_{\mu\alpha} - \partial_\alpha F_{\mu\nu}$ *gives*

$$Y = \tfrac{1}{2}\partial_\nu Q - X, \qquad X = \mathrm{Tr}\left(F_{\mu\alpha}\partial_\alpha F_{\mu\nu}\right);$$

exchanging dummy indices $\mu \leftrightarrow \alpha$ *and using* $F_{\alpha\mu} = -F_{\mu\alpha}$*,* $F_{\alpha\nu} = -F_{\nu\alpha}$ *shows* $X = Y$*, so* $Y = \frac{1}{4}\partial_\nu Q$ *and*

$$\partial_\mu A_{\mu\nu} = \tfrac{1}{4}\partial_\nu Q, \qquad \partial_\mu\Big(A_{\mu\nu} - \tfrac{1}{4}\delta_{\mu\nu}Q\Big) = 0.$$

Thus $T_{\mu\nu} = A_{\mu\nu} - \frac{1}{4}\delta_{\mu\nu}Q$ *is divergence free; the divergence-free condition imposes one linear relation on the two-dimensional candidate space, leaving one dimension, which conjugation invariance and centering do not further reduce; the Ward identity fixes the overall normalization.*

Assumption 9.3 ((Y2) the two laws)**.** For every label list $v_0, \dots, v_n$ and distinct points $x_0, \dots, x_n$: the *flux (shell) law*

$$-\sum_a \int_{\mathbb{R}^4} \partial_a\chi(y)\, S_\xi\big(T_{ab}, v_0, \dots, v_n\big)(y, x_0, \dots, x_n)\,\mathrm{d}^4y = S_\xi\big(D_b v_0, v_1, \dots, v_n\big)(x_0, \dots, x_n),$$

and the *calculus law*

$$S_\xi\big(D_b v_0, v_1, \dots, v_n\big)(x_0, \dots, x_n) = \partial_b^{x_0}\, S_\xi\big(v_0, \dots, v_n\big)(x_0, \dots, x_n),$$

where D_b is the abstract field derivative and χ is a real smooth compactly supported function equal to one near x_0 and supported away from the other x_j. Certificates: `phase.flux`, `phase.calculus`; `WardFormulationAudit.lean`; check A6.

Proposition 9.4 (Ward identity)**.** *Under (Y2), for local labels $v_0, \dots, v_n$, distinct points $x_0, x_1, \dots, x_n$, and χ as above,*

$$-3 \sum_{\mu=0}^{3} \int_{\mathbb{R}^4} \partial_\mu \chi(y)\, S_\xi\big(T_{\mu\nu}, v_0, \dots, v_n\big)(y, x_0, \dots, x_n)\, \mathrm{d}^4 y = \partial_\nu^{x_0} S_\xi\big(v_0, \dots, v_n\big)(x_0, \dots, x_n). \tag{9.1}$$

Proof. Composing the two laws of (Y2), the right side of (9.1) is $\partial_\nu^{x_0} S_\xi(v_0, \dots, v_n)$, the coefficient -3 being fixed by the contraction of $\partial_a \chi$ with the T_{ab} components relative to the construction's Euclidean coordinate and derivative conventions (the tensor normalization is fixed by this very identity). Three accompanying facts are proved separately. *Integrability:* $\partial\chi$ vanishes near every insertion, so the integrand avoids all coincidence diagonals; by Theorem 5.17(iii) the kernel is C^∞ on Ω_{n+1}, and with $\partial\chi$ compactly supported the integral is a genuine integral rather than a totalized value. *Window independence:* if χ_1, χ_2 are both admissible and $\psi = \chi_1 - \chi_2$ vanishes near all insertions, then

$$\int \partial_a \psi\, S(T_{ab}, \dots)\, \mathrm{d}y = -\int \psi\, \partial_{y^a} S(T_{ab}, \dots)\, \mathrm{d}y = 0$$

by integration by parts (no boundary terms since ψ is compactly supported), the divergence-free condition, and the vanishing of ψ at the insertions. *Coverage:* for every separated configuration there is an admissible window (its size may depend on the configuration, and no common shell exists as points collide); this is certified for every local label, every finite list of spectators of arbitrary degrees, and every character phase, by `admissible_shell_exists`, `tensor_generates_translations`, `shell_integrand_is_integrable`, `shell_charge_is_window_independent`. □

Remark 9.5. This check is sharper than stress-tensor conservation or the trace two-point function alone, but it does not prove that all possible stress-tensor improvements vanish (an improvement may leave the translation charge unchanged); it is a separated-insertion shell identity and does not construct an unrestricted contact-term extension, a dilation Ward identity, or a comparison with an independently chosen Minkowski sign convention.

9.2 A nonzero trace and its short-distance law

In the selected component theory indexed by a phase ξ^*, define

$$\Theta := \sum_{\mu=0}^{3} T_{\mu\mu} \;-\; S_{\xi^*}\Big(\sum_{\mu=0}^{3} T_{\mu\mu}\Big)\, \mathbf{1}, \tag{9.2}$$

that is, for any list of labels $v_1, \dots, v_n$ and points $x_0, \dots, x_n$, with $v_0 := \sum_\mu T_{\mu\mu}$,

$$\begin{aligned} S_{\xi^*}\big(\Theta, v_1, \dots, v_n\big)(x_0, x_1, \dots, x_n) &= S_{\xi^*}\big(v_0, v_1, \dots, v_n\big)(x_0, x_1, \dots, x_n) \\ &\quad - S_{\xi^*}\big(v_0\big)(x_0)\, S_{\xi^*}\big(v_1, \dots, v_n\big)(x_1, \dots, x_n). \end{aligned}$$

Thus $\Theta \in \mathcal{F}_4$ and $S_{\xi^*}(\Theta) = 0$ (Appendix A, `Trace.centered`).

Assumption 9.6 ((Y3) short-distance input)**.** The RG analysis of the construction gives

$$K_\Theta(r) = \Big(\frac{\beta(g_r)}{2g_r}\Big)^2 \big\langle Q(\tfrac{r}{2}e_0)\, Q(-\tfrac{r}{2}e_0)\big\rangle_{\mathrm{free},c}\,(1+o(1)), \qquad Q = F^a_{\mu\nu}F^a_{\mu\nu}, \tag{9.3}$$

as $r \downarrow 0$, where $g_r = g(1/r)$, $\beta(g) = \mu\, \mathrm{d}g/\mathrm{d}\mu$, and $\langle\cdot\rangle_{\mathrm{free},c}$ is the connected two-point function of the free Gaussian curvature; the marginal character of the curvature-squared operator transmits the one-loop running into a logarithmic enhancement of the singularity, while the residue is the free-field value (certificates `YM.Trace.Over.Scribe`, `YM.Trace.Verse.tagline`).

Lemma 9.7 (Free Wick contraction coefficient, complete verification)**.** *With*

$$\big\langle F^a_{\mu\nu}(x) F^b_{\rho\sigma}(0)\big\rangle_{\mathrm{free}} = \frac{\delta^{ab}}{\pi^2|x|^4}\big(I_{\mu\rho}(x)I_{\nu\sigma}(x) - I_{\mu\sigma}(x)I_{\nu\rho}(x)\big), \qquad I_{\mu\nu}(x) = \delta_{\mu\nu} - 2\frac{x_\mu x_\nu}{|x|^2}, \tag{9.4}$$

one has

$$\big\langle Q(x)Q(0)\big\rangle_{\mathrm{free},c} = \frac{48\,\dim_{\mathbb{R}}\mathfrak{g}}{\pi^4|x|^8}. \tag{9.5}$$

Proof. Write $P_{\mu\nu} = x_\mu x_\nu/|x|^2$; then $P^2 = P$ (rank-one projection), $I = \mathbf{1} - 2P$, so

$$I^2 = \mathbf{1} - 4P + 4P^2 = \mathbf{1}, \qquad \mathrm{Tr}(I^2) = 4, \qquad \mathrm{Tr}(I^4) = 4.$$

Let $A_{\mu\nu,\rho\sigma} = I_{\mu\rho}I_{\nu\sigma} - I_{\mu\sigma}I_{\nu\rho}$. The connected two-point function receives two (equal) Wick pairings and the color contraction $\sum_{a,b}\delta^{ab}\delta^{ab} = \dim_{\mathbb{R}}\mathfrak{g}$; the index contraction is

$$\sum_{\mu\nu\rho\sigma} A^2_{\mu\nu,\rho\sigma} = \underbrace{\Big(\sum I^2_{\mu\rho}\Big)\Big(\sum I^2_{\nu\sigma}\Big)}_{4\cdot 4 = 16} + \underbrace{16}_{\text{second term}} \underbrace{-2\,\mathrm{Tr}(I^4)}_{4} = 16 + 16 - 8 = 24.$$

Hence $\langle Q(x)Q(0)\rangle_{\mathrm{free},c} = 2\cdot\dim_{\mathbb{R}}\mathfrak{g}\cdot\frac{24}{\pi^4|x|^8} = \frac{48\,\dim_{\mathbb{R}}\mathfrak{g}}{\pi^4|x|^8}$. □

Remark 9.8 (Correction to the source document). The source document (§7.2, Step 2) states "using $I_{\mu\nu}I_{\mu\nu} = 4$ and $I_{\mu\rho}I_{\nu\rho} = I_{\mu\nu}$, the full tensor contraction evaluates to $24/(\pi^4|x|^8)$". The strict computation is $I_{\mu\rho}I_{\nu\rho} = \delta_{\mu\nu}$ (since $I^2 = \mathbf{1}$), and the cross term is $\mathrm{Tr}(I^4) = 4$; the value 24 is unaffected. We use the matrix computation of Lemma 9.7 and do not rely on the misprinted identity.

Lemma 9.9 (One-loop and anomaly coefficients, complete derivation)**.** *With* $\beta(g) = -b_0 g^3 + O(g^5)$, $b_0 = 11C_{\mathfrak{g}}/(48\pi^2)$, *and* $\ell = \log\frac{1}{\Lambda r}$,

$$g^2(1/r) = \frac{1}{2b_0\ell}\,(1+o(1)), \qquad \frac{\beta(g_r)}{2g_r} = -\frac{1}{4\ell}\,(1+o(1)).$$

Proof. From $\frac{\mathrm{d}(g^{-2})}{\mathrm{d}\log\mu} = -\frac{2\beta(g)}{g^3} = 2b_0 + O(g^2)$, integration gives $g^{-2}(\mu) = 2b_0\log(\mu/\Lambda) + O(\log\log\mu) = 2b_0\ell\,(1+o(1))$ at $\mu = 1/r$. Hence $g_r^2 = (2b_0\ell)^{-1}(1+o(1))$, and

$$\frac{\beta(g)}{2g} = \frac{-b_0 g^3 + O(g^5)}{2g} = -\frac{b_0}{2}g^2 + O(g^4) = -\frac{b_0}{2}\cdot\frac{1}{2b_0\ell}(1+o(1)) = -\frac{1}{4\ell}(1+o(1)).$$ □

Theorem 9.10 (Trace asymptotic)**.** *With the centered trace* Θ *of (9.2),*

$$\lim_{r\downarrow 0} r^8\log^2(1/r)\, S_{\xi^*}\big(\Theta,\Theta\big)(-re_0/2, re_0/2) = \frac{3\,\dim_{\mathbb{R}}\mathfrak{g}}{\pi^4} =: A_G > 0, \tag{9.6}$$

and Euclidean covariance upgrades this to the all-directions law: with $K_\Theta(x) = \mathrm{Re}\; S_{\xi^*}(\Theta,\Theta)(x,0)$,

$$|x|^8\log^2(1/|x|)\, K_\Theta(x) \longrightarrow \frac{3\,\dim_{\mathbb{R}}\mathfrak{g}}{\pi^4} \qquad (x\to 0,\ x\neq 0). \tag{9.7}$$

Proof. Algebraic assembly. By (Y3), Lemmas 9.7 and 9.9, writing $c(r) = r^8 \log^2(1/r)$ and $\log(1/r) = \ell + O(1)$,

$$c(r)\, K_\Theta(r) = r^8 \ell^2 \cdot \frac{1}{16\ell^2} \cdot \frac{48 \dim_{\mathbb{R}} \mathfrak{g}}{\pi^4 r^8} \,(1 + o(1)) = \frac{3 \dim_{\mathbb{R}} \mathfrak{g}}{\pi^4} \,(1 + o(1)).$$

Note that b_0^2 from the two anomaly insertions cancels against $g_r^4 \sim (2b_0\ell)^{-2}$ exactly, so the amplitude depends only on $\dim_{\mathbb{R}} \mathfrak{g}$; the Yang–Mills value of b_0 is fixed separately in Section 9.4. Inside the formalization the constant appears as `award` $= 12(d_2/b_0)^2 \dim_{\mathbb{R}} \mathfrak{g}/\pi^4$ and the identity `Trace.Over.force_b` ($d_2 = -b_0/2$) reduces it to $3 \dim_{\mathbb{R}} \mathfrak{g}/\pi^4$ (certificate `YMAFBridge.canonical_award`).

From lattice-approximated points to the full limit (squeeze). The formalized statement `TraceTest.shortDistance` first proves the limit at approximated positions: for s near 1,

$$c(r_n)\, G_\Theta\big(r_n \delta_n(s)\big) \longrightarrow \frac{A_G}{s^8}, \qquad \delta_n(s) \to s,$$

where $\delta_n(s)$ corrects for approximating physical points separated by $r_n s$ by nearby lattice points. By reflection positivity and the spectral representation, G_Θ is real-valued and nonincreasing, $G(r) = \int_0^\infty e^{-r\lambda}\, \mathrm{d}\rho(\lambda)$ with $\rho \geq 0$. Choosing $s_- < 1 < s_+$, for all sufficiently large n, $r_n \delta_n(s_+) \geq r_n \geq r_n \delta_n(s_-)$, so monotonicity gives

$$G_\Theta\big(r_n \delta_n(s_+)\big) \leq G_\Theta(r_n) \leq G_\Theta\big(r_n \delta_n(s_-)\big);$$

multiplying by $c(r_n) > 0$ and taking $\liminf/\limsup$,

$$\frac{A_G}{s_+^8} \leq \liminf_n \, c(r_n) G_\Theta(r_n) \leq \limsup_n \, c(r_n) G_\Theta(r_n) \leq \frac{A_G}{s_-^8},$$

and $s_\pm \to 1$ squeezes the limit to A_G. E1 (Euclidean covariance) then promotes the time-direction result to all directions, giving (9.7). □

Remark 9.11. Although a classical smooth connection has traceless stress tensor, Θ is in general nonzero as a result of renormalization—the trace anomaly. In particular the two-point function is nonzero for sufficiently small r, so some smeared $\Phi_\Theta(f)$ is nonzero. The choice $\Lambda = 1$ records the construction's reference units; it does not define an intrinsic physical scale. A position-space RG-improved version of the computation, in the same normalization, is given in [BM, eqs. (3.72)–(3.76)].

9.3 Excluded free-field descriptions

For $m \geq 0$, an integer $q \geq 0$, and a continuous real-valued function h on $\mathbb{R}^3$, homogeneous of nonnegative integral degree, consider

$$K_{m,q,h}(r) = \int_{\mathbb{R}^3} \frac{E_m(p)^q\, h(p)^2}{2E_m(p)}\, e^{-rE_m(p)}\, \mathrm{d}p, \qquad E_m(p) = \sqrt{m^2 + |p|^2}, \tag{9.8}$$

so that (9.8) and the factors K_{ab} of (9.9) carry the same normalization; the convention with denominator $E_m(p)$ instead of $2E_m(p)$ differs by the constant factor 2, which is absorbed into the arbitrary complex coefficients c_a of (9.9) and leaves the class $\mathcal{W}$ of finite sums of finite products literally unchanged. and the class of finite complex linear combinations of finite products

$$W(r) = \sum_{a=1}^{N} c_a \prod_{b=1}^{n_a} K_{ab}(r), \qquad K_{ab}(r) = \int_{\mathbb{R}^3} \frac{E_{ab}(p)^{q_{ab}}\, h_{ab}(p)^2}{2E_{ab}(p)}\, e^{-rE_{ab}(p)}\, \mathrm{d}^3p, \tag{9.9}$$

with $E_{ab}(p) = \sqrt{m_{ab}^2 + |p|^2}$, $N, n_a \in \mathbb{N}_0$, $c_a \in \mathbb{C}$, $m_{ab} \geq 0$, $q_{ab} \in \mathbb{N}_0$, and h_{ab} continuous, real-valued, and homogeneous of degree $j_{ab} \in \mathbb{N}_0$. This class includes the suitable polynomial observables and derivatives in free fields with finitely many masses.

Lemma 9.12 (Finite asymptotic expansion)**.** *Each $K_{m,q,h}$ has, as $r \downarrow 0$, a finite expansion $\sum_{\nu=1}^{M} a_\nu r^{\alpha_\nu} (\log r)^{k_\nu} + O(1)$ with $a_\nu \in \mathbb{R}$, $\alpha_\nu \in \mathbb{Z}$, $k_\nu \in \mathbb{N}_0$; the class is closed under finite products and finite linear combinations of finitely many masses.*

Proof. Homogeneity of h of degree j makes the radial part of the integrand $p^{2j} E^{q-1} \varphi(\hat{p})$ with φ smooth on the sphere; the radial integral is $\int_0^\infty p^{2j} E^{q-1} \varphi\, e^{-rE}\, \mathrm{d}p$. If $m = 0$, then $E = p$, the integrand is $\propto p^{2j+q+1} e^{-rp}$, and $\int_0^\infty p^s e^{-rp}\, \mathrm{d}p = \Gamma(s+1) r^{-(s+1)}$ gives integer powers. If $m > 0$, the large-p region $E = p + \frac{m^2}{2p} + \cdots$ gives integer powers $r^{-(2j+q+2)}, r^{-(2j+q)}, \ldots$, and the small-$p$ region $E = m + \frac{p^2}{2m} + \cdots$ with $e^{-rE} = e^{-rm}(1 + O(rp^2))$ gives $r^0, r^1, \ldots$. When powers from the two regions resonate, the critical radial integral of type $\int p^{-1}\, \mathrm{d}p$ produces $\log r$; each resonance raises the logarithmic power by one, and the number of resonances is finite, so k_ν is finite and ≥ 0. Products of two finite expansions are finite expansions of the same type ($r^{\alpha+\alpha'} (\log r)^{k+k'}$), and finite linear combinations are clear; hence every W of the form (9.9) has a finite expansion of this type. Certificates: `YM.TraceExtension.Wick.real_profile`, `AtOrder`, `Hierarchy.wickProfile`. □

Lemma 9.13 (Finite limit forces zero, at the correct weight)**.** *Let $c(r) = r^8 \log^2(1/r)$. If W belongs to the class of Lemma 9.12 and $\lim_{r \downarrow 0} c(r) \operatorname{Re} W(r)$ exists and is finite, then the limit is 0.*

Proof. Write $\operatorname{Re} W(r) = \sum_\nu a_\nu r^{\alpha_\nu} (\log r)^{k_\nu} + O(1)$ (taking real parts affects only the coefficients, and $(\log r)^k$ differs from $\log^k(1/r)$ by a sign, irrelevant to powers and divergence). If $\alpha_\nu < -8$, $r^{\alpha_\nu + 8} \log^{2+k_\nu}(1/r) \to \infty$, forcing $a_\nu = 0$; if $\alpha_\nu = -8$, the term contributes $a_\nu \log^{2+k_\nu}(1/r)$, which diverges unless $a_\nu = 0$ (note $k_\nu \geq 0$); if $\alpha_\nu > -8$, the term tends to 0; and the $O(1)$ term is killed by $c(r) \to 0$. Hence a finite limit forces every surviving coefficient to vanish, and the limit is 0. □

Proposition 9.14 (Free-field exclusion)**.** *There is no W of the form (9.9) and no $\delta > 0$ such that $K_\Theta(r) = W(r)$ on $(0, \delta)$.*

Proof. By Theorem 9.10, $\lim_{r \downarrow 0} c(r) K_\Theta(r) = A_G = 3 \dim_{\mathbb{R}} \mathfrak{g} / \pi^4 > 0$. If $K_\Theta = W$ on $(0, \delta)$, then $\lim_{r \downarrow 0} c(r) \operatorname{Re} W(r) = A_G > 0$, contradicting Lemma 9.13. □

Remark 9.15 (Erratum on the weight). The source document (§7.3, Steps 2–3, and Appendix A, `TraceTest.excludes`) phrases the obstruction with the weight $r^8 \log(1/r)$ and asserts "by (9.7), $\lim r^8 \log(1/r) K_\Theta = A_G > 0$"; however the weight of (9.6)/(9.7) is $c(r) = r^8 \log^2(1/r)$, and with the $\log^1$ weight one has $r^8 \log(1/r) K_\Theta = A_G / \log(1/r) \to 0$, which gives no positive limit. The two formulations are inconsistent. We adopt the weight $c(r) = r^8 \log^2(1/r)$ of Theorem 9.10 and prove the exclusion by Lemma 9.13; the argument is then more direct and does not require the source classification that "no term with $\alpha = -8$, $k \geq 1$ occurs" (under the $\log^2$ weight every term with $\alpha = -8$, $k \geq 0$ diverges). The conclusion is unchanged and the proof is shorter. Note that the Lean certificates `YM.TraceExtension.Wick.no_powerlog` and `euclidean_excludes` check the $\log^1$-weight version; what is proved here is a mathematically self-consistent replacement argument in the $\log^2$ weight of Theorem 9.10, and we do not claim that those certificates cover the $\log^2$ weight.

Remark 9.16. The exclusion covers only the finite class (9.9): finitely many masses, finitely many factors, finite sums. It does not address descriptions by infinite composites of Gaussian fields, and it does not use reflection positivity or the squeeze argument of Theorem 9.10—two-point positivity enters only through $A_G > 0$. The general exclusion of continuous-mass generalized free fields is outside the scope of the construction.

9.4 Asymptotic freedom and the running coupling

Definition 9.17 (g_W)**.** Starting from the bare couplings supplied by Theorem 5.17, perform RG blocking down to the first lattice spacing below r, and define $g_W(r)$ as the inverse square root of the resulting normalized coefficient of the marginal (curvature-squared) term in the effective action, after removal of the initial cutoff and extraction of a subsequential infinite-volume limit; the normalization requires the quadratic term of the blocked effective action to have identity coefficients (the kinetic normalization of Lemma 4.15). Independence of $g_W(r)$ from the subsequential limit is not proved.

Assumption 9.18 ((Y4) the one-loop value)**.** At one loop, the coefficient of the marginal operator receives the gauge-boson loop, the Faddeev–Popov ghost loop, and the measure (BRS) contribution, which for a simple algebra with adjoint Casimir $C_{\mathfrak{g}}$ sum to

$$b_0 = \frac{1}{(4\pi)^2} \cdot \frac{11}{3}\, C_{\mathfrak{g}} = \frac{11\, C_{\mathfrak{g}}}{48\pi^2}, \tag{9.10}$$

with no scalar or fermion loops. This is a pure integer/rational computation from the Casimir of $\mathfrak{g}$ (certificates `YM.Pert.Asymptotics.b0`, `YM.Gauge.casimir`).

Theorem 9.19 (Effective coupling at short distances)**.** *For small r and any selected subsequential infinite-volume limit, $g_W(r) > 0$, and*

$$\log(1/r)\, g_W(r)^2 \longrightarrow \frac{1}{2b_0} \qquad (r \downarrow 0). \tag{9.11}$$

Proof. Step 1 (recursion). The recursion of Lemma 5.12, $\frac{\mathrm{d}}{\mathrm{d}\log\mu} g^{-2} = 2b_0 + O(g^2)$, is equivalent to $\beta(g) = -b_0 g^3 + O(g^5)$; integrating (Lemma 9.9) gives $g^{-2}(\mu) = 2b_0 \log(\mu/\Lambda) + O(\log\log\mu)$. *Step 2 (value).* This is (Y4). *Step 3 (from the lattice trajectory to g_W).* By construction, g_W lies on the same RG trajectory as the bare couplings $\beta_{N,p}$, so $g_W^{-2}(r) = 2b_0 \log(1/r) + O(\log\log(1/r))$; taking reciprocals and multiplying by $\log(1/r)$ gives (9.11). □

Remark 9.20. Two convention dependencies are explicit: Λ is scheme-dependent (a change of scheme adds a constant to $\log(1/r)$ and does not affect (9.11)), and the normalization of g_W is fixed by the kinetic normalization. The argument is perturbative in the marginal coupling only: it controls $g_W(r)$ for small r, where $g_W(r) \to 0$, and does not address the infrared side of the flow, Landau-pole type questions, or uniqueness of g_W across subsequences. A second coupling defined through separated continuum correlations—normalizing a scalar two-point derivative and an integrated tensor–scalar–scalar correlation to their free values and reading the coupling from a derivative of a tensor–scalar correlation—has ratio to g_W tending to 1, so the running law is visible in continuum observables as well as in the effective action.

Remark 9.21 (Free-field agreement at small scales). Along one common sequence $r_k \downarrow 0$, the averaged n-point stress-tensor correlations, multiplied by r_k^{4n} and evaluated at $r_k x_1, \dots, r_k x_n$, converge to those of $\dim \mathfrak{g}$ independent free Maxwell fields (the Gaussian two-form curvature obtained as the exterior derivative of a $\mathfrak{g}$-valued one-form Gaussian free field), with every fixed number of derivatives, uniformly on compact sets avoiding collisions. This comparison holds for the mixed theory along the particular scaling sequence, whereas Theorems 9.10 and 9.19 hold as $r \downarrow 0$ arbitrarily, and Theorem 9.10 applies to the selected phase ξ^*.

9.5 Operator product expansion and contact extensions

Theorem 9.22 (Operator product expansion)**.** *Let v_i and u_j be local labels of degrees d_i and $w = \sum_i d_i$. There exist labels $o_\alpha \in \mathcal{F}_{d_\alpha}$, finitely many per degree, such that for every $D \geq 0$ and $\varepsilon > 0$,*

$$S^s\big(v, u(rx)\big) = \sum_{d_\alpha \leq D} c_\alpha(v; rx)\, S^s\big(o_\alpha, u(0)\big) + R^s_D, \qquad |R^s_D| \leq r^{D+1-w-\varepsilon}, \tag{9.12}$$

where s denotes the mixed theory or a component theory. Uniform formulation: *for input fields $A_i \in \mathcal{F}_{d_i}$ at positions x_i, a spectator product $B = \prod_\ell B_\ell(y_\ell)$, and an expansion point o, the same expansion reads*

$$S^s\Big(\prod_i A_i(x_i)\, B\Big) = \sum_{d_\alpha \leq D} c_\alpha(A; x, o)\, S^s\big(O_\alpha(o)\, B\big) + R_{D,s}, \qquad |R_{D,s}| \leq r^{D+1-w-\varepsilon}, \tag{9.13}$$

with the spectators held at a fixed separation from o and from each other (see (9.14) below); the short forms $u(rx)$, $u(0)$ of the source document, which suggest spectators scaled with r or placed at the origin on a coincidence diagonal, are replaced by this single convention. More precisely: for fixed constants $M \geq 0$, $\eta, c > 0$, the allowed configurations at scale $r > 0$ satisfy

$$\begin{aligned} \|x_i - o\| &\leq Mr \quad \text{for every } i, \\ \|x_i - x_j\| &\geq \eta r \quad \text{for } i \neq j, \\ \|y_\ell - o\| &\geq c \quad \text{for every } \ell, \\ \|y_\ell - y_m\| &\geq c \quad \text{for } \ell \neq m, \end{aligned} \tag{9.14}$$

the inputs forming a cluster near the expansion point o while the spectators u_j at y_ℓ stay separated; with $O_\alpha := \texttt{T.actor}(d, i)$ and $c_\alpha(A; x, o) := \texttt{T.coefficient}(A, d, i, x, o)$, define, for $s = \text{mix}$ or $s = \xi \in X$,

$$R_{D,s} = S^s\Big(\prod_i A_i(x_i)\, B\Big) - \sum_{d_\alpha \leq D} c_\alpha(A; x, o)\, S^s\big(O_\alpha(o)\, B\big). \tag{9.15}$$

The bounds are

$$|c_\alpha(A; x, o)| \leq r^{d_\alpha - w - \varepsilon}, \qquad |R_{D,s}| \leq r^{D+1-w-\varepsilon}, \tag{9.16}$$

holding for all sufficiently small $r > 0$ for every $\varepsilon > 0$; the coefficient bound requires only the first two conditions of (9.14), the remainder bound requires all four and holds uniformly over all states $s \in \{\text{mix}\} \sqcup X$. Setting $o = 0$ and input positions $x_i = r\hat{x}_i$ with fixed $\hat{x}_i$ (and the spectator product B held fixed) recovers the short form of (9.12); (9.13) is the same statement with the geometric conditions made explicit. The coefficients are independent of spectators, state, and truncation degree; their independence of the chosen subsequential infinite-volume limit is not proved, and common coefficients of two expansions do not by themselves imply equal correlations. The cutoff D counts local-observable label degrees, not loop orders; perturbative agreement to all orders remains unproved.

Proof: extraction of the coefficients and bounds. The coefficients $c_\alpha(A; x, o)$ are produced by a fixed algebraic prescription on the label spaces before any correlation is evaluated; their independence of spectators, phases, and truncation orders is automatic rather than dynamical. (1) *Degree decomposition:* for each d fix a splitting $\mathcal{F}_d = \mathcal{F}_{d-1} \oplus E_d$; the channels $\alpha = (d, i)$ index a conjugation-invariant basis of the new layer E_d (`YM.OPE.Orbits.Reader.channels`, `.basis`, `.actor`), with $n_d = \dim E_d$ fields per layer (`YM.OPE.Orbits.Reader.fresh`). (2) *Test-field formation:* form a degree-d field from finitely many correlations of the inputs against additional test insertions supported on a fixed "stratum" configuration (`YM.OPE.Orbits.Reader.cutTest`, `.stratum`); the separation hypotheses

(9.14) keep the stratum tests away from both the input cluster and the spectators. (3) *Symmetrization and projection:* average over $O(4)$ (`YM.OPE.Orbits.averageV`), symmetrize under label conjugation (`YM.OPE.Orbits.checkReal`), and strip the $\mathcal{F}_{d-1}$ components (`YM.OPE.shadow`); the coordinates in the basis $O_{(d,i)}$ define $c_\alpha(A; x, o)$ (`YM.OPE.Orbits.Reader.coef`, packaged in `YM.OPE.Prescription.ope`). (4) *Bounds:* because the prescription is fixed, the coefficient bound in (9.16) requires only the first two conditions of (9.14) (`YM.OPE.Orbits.coefficient_separated`); the remainder bound requires all four, uniformly over s (`YM.OPE.Orbits.quantitative_separated`). The analytic estimates entering Step (4) are those of the submitted proof; no perturbative computation or comparison of coefficient values is involved. □

Proposition 9.23 (Compact-set uniform bounds)**.** *Fix labels, truncation degree D, and $\varepsilon > 0$, and let $\mathcal{K}$ be a compact set of configurations satisfying (9.14) (input cluster and spectator sets in fixed compacts, uniformly separated), with $N_D = \#\{\alpha : d_\alpha \le D\} < \infty$. Suppose the pointwise bounds (9.16) hold on $\mathcal{K}$ for all s. Then any norm of the complete coefficient vector $c = (c_\alpha)_\alpha$ is uniformly bounded on $\mathcal{K}$, and the remainder bound holds uniformly in $(x, o, y) \in \mathcal{K}$ and s.*

Proof. From the pointwise bound, $\|c\|_\infty \le \max_\alpha r^{d_\alpha - w - \varepsilon} =: B_D$, a finite constant; in the finite-dimensional space $\mathbb{C}^{N_D}$ all norms are equivalent, so there is C_D with $\|c\| \le C_D \|c\|_\infty \le C_D B_D$, independent of the configuration. The remainder bound's right side contains no configuration or state, so $\sup_{\mathcal{K},s} |R_{D,s}| \le r^{D+1-w-\varepsilon}$. (Compactness is not needed for the uniform bound—it ensures the quantification over "every compact separated cluster and spectator set" is well formed.) □

Remark 9.24 (Contact extensions: the submission's predicate). The contact-term extension is isolated behind an explicit predicate in the submission's public interface (`ContactSpec.lean`). For a hierarchy T, a state option s : `Option T.Character`, and any number n of insertions with labels v : `Fin n` $\to$ `T.Tag`,

$$\texttt{Contacts} :\Longleftrightarrow \exists\, D : \text{ test functions on the } \textit{full} \text{ configuration space } (\mathbb{R}^4)^n \longrightarrow \mathbb{C},$$

continuous and linear, such that every test function f whose tubular support lies in the separated insertion region satisfies $D(f) = \int_{(\mathbb{R}^4)^n} f(x)\, T.\texttt{correlation}(s, v, x)\, \mathrm{d}x$. The distribution D lives on the whole space, including every collision diagonal, and restricts to the already constructed correlation on separated insertions. The submission's bridge theorem (`ContactBridge.lean`) proves `Contacts` for the constructed hierarchy, transporting the OPE-based multiple-contact development (harmonic extension around the collision locus, with the extension recorded so that evaluation on separated tests reproduces the original correlation). Three limitations remain: the predicate does not assert uniqueness of the extending distribution; it does not prove the full OS axioms (reflection positivity) for the extended correlations on unrestricted Euclidean test spaces; and lattice convergence of the extended functionals is not established. Relatedly, diagonal-flat Schwartz tests annihilate distributions supported on the diagonals, so the independent OS check of Appendix C does not touch contact-term uniqueness.

9.6 The field-content audit

The dictionary between the constructed operator spaces and classical gauge-invariant polynomials is checked in six steps (`YMReview.Audits.OperatorContent`, worked example `OperatorContent-20261028.lean`; check A2 of Table 1). The registered label collection is `Labels.Pol G` with its invariant subspace `Labels.invariantSpace G`; a label in `Labels.Label G d` carries a finite weight bound d, and no fixed upper degree is imposed on the full collection. The audit is stated in the identifiers of Appendix A, which we use verbatim; in the notation of the

main text, `S.Field` d is the quotient $\mathcal{F}_d = V_d/N_d$ of Definition 2.7, `S.null` d is the correlation null space N_d, `S.Union` is the label space V, `F.`α`.space.`X is the phase space X of Proposition 7.10, `F.`α`.ensemble`(i) is the phase hierarchy of character i, and `F.action.operators.field` is the field-operator map of Section 8.

1. *Exhaustiveness of the register.* Every invariant polynomial $p \in$ `Labels.invariantSpace G` appears at some register level: there exist d and $v \in$ `Labels.Label G d` with `v.val` $= p$ (`every_registered_polynomial_appears`, from `OperatorCompleteness.raw_register_exhaustive`). This is exhaustiveness relative to the submitted algebraic dictionary; human review must still identify that dictionary with the intended local gauge-invariant fields, and the theorem does not address nonlocal observables such as an independently specified class of line defects.

2. *Free on-shell comparison space.* At filtration d, the free on-shell polynomial comparison space `Maxwell.shellSector G d` is exactly the invariant part of the free lower-weight cap (`free_sector_dictionary`, from `OperatorCompleteness.free_sector_exhaustive`), and the separated free correlations leave no further residual null space: `Maxwell.residual G d` $= \perp$ (`free_residual_vanishes`, from `free_no_extra_nulls`). These identify the comparison space; they are not claims that interacting correlation functions are Gaussian.

3. *Interacting quotient dictionary.* Writing `S.null` d for the kernel of the map from raw labels to all separated spectator correlations and `S.Field` $d = V_d/$`S.null` d as in Section 2, the rank comparison of the submission, reached through `F.phase.transport.comparator`, yields a complex linear equivalence

$$\texttt{S.Field}\, d \;\simeq_{\mathbb{C}}\; \texttt{Maxwell.shellSector G d} \qquad \text{for every } d$$

   (`interacting_field_dictionary`, from `OperatorCompleteness.local_field_content`). This constrains the size of the interacting quotient at every filtration level. An equivalence of these vector spaces is not an equality of interacting and free correlators, an algebra isomorphism, or a unitary equivalence of quantum field theories; the source comparison also tracks the real signed-coordinate structure, which the display above suppresses.

4. *No field is lost in a phase.* For every character phase $i \in$ `F.`α`.space.`X—not merely almost every phase under the decomposition measure—the phase and mixed null spaces agree, and the preceding equivalence transports:

$$\texttt{F.}\alpha\texttt{.ensemble}(i)\texttt{.Field}\, d \;\simeq_{\mathbb{C}}\; \texttt{Maxwell.shellSector G d} \qquad \text{for every } d,\ i$$

   (`each_phase_retains_the_dictionary`, from `OperatorCompleteness.phase_field_content`, via `F.phase.phase.retains i d`). This is stronger than the statement that the mixed correlation function is the average of its phase values.

5. *Cyclicity uses the full collection.* `F.action.operators.vector` assigns a Hilbert-space vector to every registered field word applied to the vacuum, and the linear span of its range is dense (`full_word_domain_is_dense`, from `OperatorCompleteness.registered_words_dense`). The assertion uses the full field collection: the smaller observable family behind the scalar OS bridge of check A1 does not by itself prove this cyclicity statement. Cyclicity is partly a feature of reconstructing a Hilbert space from field words; its physical force comes from combining it with the field-content comparison and the faithful representation, not from relabeling the constructed space as an independently defined canonical one.

6. *Surviving fields act nontrivially.* Every nonzero local quotient field has some nonzero smeared operator: for $p \in$ `S.Union`, $p \neq 0$, there is a test f with `F.action.operators.field` $p\ f \neq 0$ (`surviving_field_has_nonzero_operator`, from `OperatorCompleteness.quotient_fields_not_lost`). Passing to the reconstructed operator representation therefore kills no surviving local field class; it does not assert that every smearing of a nonzero field is nonzero.

The audit concerns the full registered invariant collection, not only the charge-conjugation even stress trace; it does not supply an independently checked charge-conjugation-equivariant dictionary or a classification of charge-conjugation odd particles, which would need their own precise statements. In the notation of Appendix A, the quotient and injectivity steps correspond to `Hierarchy.Field`, `toLabel`, and `Properties.injective`, `covers`, `represents`, `acts`, `faithful`. The transitive axiom dependencies of all six theorems are printed in the file and use only the permitted foundations.

10 Open problems: precise formulations and a logical boundary

The ten open problems recorded in the source documents are restated here as statements of the form "which additional input is missing", together with two proved logical boundaries: the sixteen registered inputs do not imply a mass gap, and a mass gap is equivalent to a uniform exponential rate in E4 (Proposition 10.22). Throughout, (MG1)–(MG4) and (O1)–(O10) are *proposed* conditional inputs, of the form "if (MGk) holds, then ..."; none of them is claimed as proved, none is registered in Appendix B, and none carries a repository certificate. The four candidates (MG1)–(MG4) are projections of a single uniform infrared estimate (MG), registered in Proposition 10.34.

10.1 A spectral representation toolkit

Lemma 10.1 (Exponential decay versus spectral support)**.** *Let ρ be a finite positive measure on $[0,\infty)$ and $G(\tau) = \int_{[0,\infty)} e^{-\tau\lambda}\,\mathrm{d}\rho(\lambda)$. For any $m > 0$, the following are equivalent:*

(i) $\rho\big((0,m)\big) = 0$*;*

(ii) *there exists $K < \infty$ such that $\tilde{G}(\tau) := \int_{(0,\infty)} e^{-\tau\lambda}\,\mathrm{d}\rho(\lambda) \le K e^{-m\tau}$ for all $\tau \ge 0$.*

Proof. Write $\tilde{\rho} = \rho - \rho(\{0\})\delta_0$. (i)⇒(ii): if $\operatorname{supp}\tilde{\rho} \subset [m,\infty)$, then $\tilde{G}(\tau) = e^{-m\tau}\int_{[m,\infty)} e^{-\tau(\lambda-m)}\,\mathrm{d}\tilde{\rho}(\lambda) \le e^{-m\tau}\|\rho\|$, so $K = \|\rho\|$ works. (ii)⇒(i): suppose $a := \tilde{\rho}\big((0,m-\varepsilon)\big) > 0$ for some $\varepsilon > 0$; then $\tilde{G}(\tau) \ge \int_{(0,m-\varepsilon)} e^{-\tau\lambda}\,\mathrm{d}\tilde{\rho} \ge a\,e^{-(m-\varepsilon)\tau}$, so $a\,e^{\varepsilon\tau} \le K$ for all τ, which is impossible. Hence $\tilde{\rho}\big((0,m-\varepsilon)\big) = 0$ for all $\varepsilon > 0$, i.e. $\rho((0,m)) = 0$. □

Lemma 10.2 (Gap implies uniform exponential clustering)**.** *Under (PH4), if $\sigma(H) \subset \{0\} \cup [\Delta,\infty)$, then for all $\psi, \chi \in \mathcal{H}$,*

$$\big|\langle\psi, E(s)\chi\rangle - \langle\psi,\Omega\rangle\langle\Omega,\chi\rangle\big| \le e^{-\Delta s}\,\big\|\psi - \langle\Omega,\psi\rangle\Omega\big\|\,\big\|\chi - \langle\Omega,\chi\rangle\Omega\big\|.$$

Proof. Let $\chi_\perp = \chi - P_0\chi$ and $\psi_\perp = \psi - P_0\psi$; by (PH4) $P_0 = |\Omega\rangle\langle\Omega|$, so $\langle\psi, E(s)\chi\rangle - \langle\psi, P_0\chi\rangle = \langle\psi_\perp, E(s)\chi_\perp\rangle$. By the spectral theorem this equals $\int_{[\Delta,\infty)} e^{-s\lambda}\,\mathrm{d}\mu_{\psi_\perp,\chi_\perp}(\lambda)$, and $\|\mu_{\psi_\perp,\chi_\perp}\| \le \|\psi_\perp\|\,\|\chi_\perp\|$ by polarization and Cauchy–Schwarz; the claim follows. □

Proposition 10.3 (Dense channel coverage gives the full gap)**.** *Let $\mathcal{D}_0 \subset \Omega^\perp$ be dense, and suppose there exist $\mu > 0$ and $K(\psi) < \infty$ such that*

$$\big|\langle \psi, E(s)\psi\rangle\big| \le K(\psi)\, e^{-\mu s} \qquad (\psi \in \mathcal{D}_0,\ s \ge 0).$$

Then $\sigma(H) \subset \{0\} \cup [\mu, \infty)$.

Proof. For $\psi \in \mathcal{D}_0$ we have $\rho_\psi(\{0\}) = \|P_0\psi\|^2 = 0$, so Lemma 10.1 gives $\rho_\psi((0,\mu)) = 0$. Now suppose $(0,\mu) \cap \sigma(H) \neq \varnothing$ and pick $\lambda_0 \in (0,\mu) \cap \sigma(H)$ and $\varepsilon > 0$ with $I := [\lambda_0 - \varepsilon, \lambda_0 + \varepsilon] \subset (0,\mu)$. Then $E_H(I) \neq 0$; take a non-zero ϕ in its range, so that $\rho_\phi(I) = \|\phi\|^2 > 0$. Choose a continuous $h : [0,\infty) \to [0,1]$ with $h \equiv 1$ on I and $\operatorname{supp} h \subset (0,\mu)$ (Urysohn's lemma). By the functional calculus $h(H)$ is bounded with $\|h(H)\| \le 1$. Since $\mathcal{D}_0$ is dense in $\Omega^\perp$, choose $\psi_j \in \mathcal{D}_0$ with $\psi_j \to \phi$. Then

$$\int h \,\mathrm{d}\rho_{\psi_j} = \langle \psi_j, h(H)\psi_j\rangle \longrightarrow \langle \phi, h(H)\phi\rangle = \int h \,\mathrm{d}\rho_\phi \ge \rho_\phi(I) = \|\phi\|^2 > 0,$$

while $\rho_{\psi_j}((0,\mu)) = 0$ and $\operatorname{supp} h \subset (0,\mu)$ force $\int h \,\mathrm{d}\rho_{\psi_j} = 0$ for every j, a contradiction. Hence $(0,\mu) \cap \sigma(H) = \varnothing$. □

Remark 10.4 (The density hypothesis is the field-word density)*.* The dense subspace $\mathcal{D}_0$ is supplied by the field-word cyclicity (`registered_words_dense`, field-audit step 5); the smaller observable family behind the scalar OS bridge (check A1) does not suffice. An exponential rate in a single channel (e.g. the trace two-point function) therefore yields only a *channel gap*, not the full gap of Definition 1.2.

Moreover the rate must be common to all registered label words: (Y2) supplies no degree-uniformity (Proposition 10.26), and a model with per-tower exponential rates but no common rate exists (Proposition 10.28, calculus-law part).

10.2 A model within the registered inputs without a mass gap

Let

$$\rho := \delta_0 + \sum_{k \ge 1} 2^{-k} \delta_{2^{-k}} \quad \text{on } [0,\infty),$$

a finite positive measure with $\|\rho\| = 2$ and $\operatorname{supp}\rho \subset [0, \frac{1}{2}]$. Let $\mathcal{H} = L^2([0,\infty), \mathrm{d}\rho)$, let H be the multiplication operator $(H\psi)(\lambda) = \lambda\psi(\lambda)$ ($H \ge 0$), let $\Omega = \mathbf{1}/\sqrt{2}$, and let P_0 be the projection onto the one-dimensional atom at $\lambda = 0$. Take the generalized free field with Källen–Lehmann spectral measure ρ [GJ, OS1]: its two-point function is

$$G(\tau) = \int_{(0,\infty)} e^{-\tau\lambda} \,\mathrm{d}\tilde\rho(\lambda) = \sum_{k \ge 1} 2^{-k} e^{-\tau 2^{-k}}, \qquad \tilde\rho = \rho - \delta_0, \tag{10.1}$$

and the full hierarchy is given by the Wick formula.

Proposition 10.5 (Logical independence of the mass gap)**.** *There is an OS hierarchy satisfying every clause of (AN1) (E0, E0′, E1, E2, E3, E4 clustering, normalization, analyticity) and (PH1)–(PH4), whose reconstructed Hamiltonian satisfies*

$$\sigma(H) = \{0\} \cup \{2^{-k} : k \ge 1\},$$

so that 0 is an accumulation point of the spectrum and there is no mass gap; yet E4 clustering and vacuum uniqueness hold.

Proof. The OS axioms. E1 (Euclidean covariance) and E3 (permutation) are inherited from the construction of the generalized free field; E0 holds because the compactly supported ρ gives finite moments, holomorphy of the Laplace transform for $\operatorname{Re}\tau > 0$, and tempered kernels at separated points (diagonal singularities are handled by the $\mathcal{S}_0$ conditions of Appendix C); E0′ holds by Wick counting: the $2n$-point function is a sum of $(2n-1)!!$ products of two-point functions, and compact support gives uniform bounds on separated sets, yielding growth of type $\|\mathcal{S}_{2n}(f)\| \le \beta^{2n}(n!)^A \|f\|_{(ns,ns)}$; E2 holds because G is the Laplace transform of a positive measure, hence completely monotone, and reflection positivity of the two-point function is inherited by the Gaussian (Wick) hierarchy; analyticity follows from holomorphy of the Laplace transform on $\operatorname{Re}\tau > 0$. (PH1)–(PH4): the model is time-translation invariant with trivial ground-state algebra action, commutativity, faithfulness of the vacuum vector for the scalars, and P_0 one-dimensional by construction.

(a) Clustering. By dominated convergence, $\tilde{G}(\tau) \to 0$ as $\tau \to \infty$, i.e. the connected two-point function tends to 0; all connected functions of the Wick hierarchy are products of connected two-point functions, so E4 holds.

(b) The clustering rate is only power law. Let $k_0 = \lfloor \log_2 \tau \rfloor$ ($\tau \ge 2$). *Lower bound:* taking $k = k_0$, we have $2^{k_0} \le \tau < 2^{k_0+1}$, so $2^{-k_0} \ge 1/\tau$ and $1 \le \tau 2^{-k_0} < 2$, hence $G(\tau) \ge 2^{-k_0} e^{-\tau 2^{-k_0}} \ge e^{-2}/\tau$. *Upper bound:* for $k > k_0$, $\sum_{k>k_0} 2^{-k} = 2^{-k_0} \le 2/\tau$; for $k \le k_0$, writing $j = k_0 - k \ge 0$, we have $2^{-k} = 2^j 2^{-k_0} \le 2^j/\tau$ and $\tau 2^{-k} \ge 2^{j-1}$, so this part is at most $\frac{1}{\tau}\sum_{j\ge 0} 2^j e^{-2^{j-1}} = C/\tau$. Therefore

$$\frac{e^{-2}}{\tau} \le G(\tau) \le \frac{C+2}{\tau} \qquad (\tau \ge 2),$$

i.e. $G(\tau) \asymp \tau^{-1}$; in particular $e^{m\tau} G(\tau) \to \infty$ for every $m > 0$.

(c) No spectral gap. $\operatorname{supp}\tilde{\rho} = \{2^{-k}\}$ accumulates at 0, so $\inf\big(\sigma(H) \setminus \{0\}\big) = 0$ and $\rho((0,\varepsilon)) > 0$ for every $\varepsilon > 0$; by Lemma 10.1, no spectral gap holds.

(d) Logical conclusion. If "the theory has a mass gap" were derivable from (AN1), (PH1)–(PH4), and the standard OS theorems, it would hold in every model of those axioms; (c) exhibits a counterexample. Hence a uniform mass gap requires input beyond the sixteen registered assumptions. By Propositions 10.7 and 10.9 this remains true even when the ultraviolet trace law, the Ward identities, the field-content dictionary and the OPE data are imposed: those data are insensitive to the spectral weight near zero. □

Remark 10.6 (What the model does and does not say). The model does *not* claim that the constructed Yang–Mills theory has no mass gap; it is a generalized free field proving that the axioms together with the registered inputs do not force one. It lies precisely outside the exclusion of Proposition 9.14, which covers only the finite class of (9.9) (finitely many masses, factors, and sums); the general exclusion of continuous-mass generalized free fields remains open (problem O6 below). Consequently O6 is a *necessary but not sufficient* prerequisite for a mass gap: until such infrared descriptions are excluded, a gapless model compatible with all registered inputs stands; and even after exclusion, (MG) must supply a positive lower bound (Proposition 10.24). This also shows that ultraviolet data — the trace anomaly (a renormalization effect, Remark 9.11), asymptotic freedom, the OPE error bounds — constrain the infrared spectral weights of such a model not at all.

Proposition 10.7 (Ultraviolet data do not constrain the infrared spectrum)**.** *There exists a tempered generalized-free-field hierarchy satisfying (AN1) E0–E4, vacuum uniqueness, the field-content dictionary (B1), the stress-tensor candidate calculation (Y1), the calculus and flux laws (Y2), the exact short-distance constant $A_G = 3\dim_{\mathbb{R}}\mathfrak{g}/\pi^4$ of Theorem 9.10, and the OPE coefficients and remainder bound of Theorem 9.22, whose Hamiltonian has no spectral gap.*

Proof. The spectral weight. Work with a Källen–Lehmann *weight* ρ on $[0,\infty)$ which is locally finite and of polynomial growth (tempered); for every Schwartz test f the induced vector spectral measure $\rho_f = |\hat{f}|^2\rho$ is finite, so Lemma 10.1 and all spectral arguments apply to ρ_f. Set

$$\rho = \delta_0 + \rho_{\mathrm{UV}} + \rho_{\mathrm{IR}}, \qquad \rho_{\mathrm{IR}} = \sum_{k\ge 1} 2^{-k}\delta_{2^{-k}}, \qquad \rho_{\mathrm{UV}}(\mathrm{d}E) = c\,\frac{E^5}{\log^2(E+e)}\,\mathbf{1}_{E\ge 1}\,\mathrm{d}E.$$

The two nontrivial pieces have disjoint supports: $\operatorname{supp}\rho_{\mathrm{IR}} \subset (0,\frac{1}{2}]$, $\operatorname{supp}\rho_{\mathrm{UV}} \subset [1,\infty)$. Let $G(\tau) = \int e^{-\tau E}\,\mathrm{d}\rho(E)$ and let the full hierarchy be Gaussian/Wick with two-point kernel $G(x-y)\,\langle v,w\rangle_{\mathrm{label}}$, the label Gram matrix being the one of Section 2.

Short distance. The four-dimensional Euclidean propagator of mass E is $\Delta_E(r) = \frac{E}{4\pi^2 r}K_1(Er)$, with $\Delta_E(r) \asymp \frac{1}{4\pi^2 r^2}\mathbf{1}_{Er\lesssim 1}$ and exponential suppression for $Er \gg 1$. Hence

$$K(r) \;\asymp\; \frac{1}{4\pi^2 r^2}\int_0^{1/r} \mathrm{d}\rho(E) \;\sim\; \frac{c}{24\pi^2}\,r^{-8}\log^{-2}(1/r),$$

because $\int_1^{1/r} E^5\log^{-2}E\,\mathrm{d}E \sim \frac{1}{6}(1/r)^6\log^{-2}(1/r)$. The infrared piece has finite total mass and bounded support, so it contributes at most $\|\rho_{\mathrm{IR}}\|/(4\pi^2 r^2) = O(r^{-2})$, a strictly lower-order singularity. Choosing c fixes the leading coefficient to A_G; the assembly of Theorem 9.10 (Lemma 9.7 times Lemma 9.9) involves only the large-E Taylor coefficients and never touches the spectrum below $E=1$, so the constant $3\dim_{\mathbb{R}}\mathfrak{g}/\pi^4$ is unchanged.

Long distance. The ultraviolet piece gives $O(e^{-\tau})$ as $\tau\to\infty$; the infrared piece gives $\sum_k 2^{-k}e^{-\tau 2^{-k}} \asymp \tau^{-1}$ by the computation of Proposition 10.5(b). Hence E4 clustering holds with a power-law rate, $e^{m\tau}G(\tau)\to\infty$ for every $m>0$, and $\inf(\sigma(H)\setminus\{0\}) = 0$ because $\{2^{-k}\}$ accumulates at 0. By Lemma 10.1 there is no gap.

The axioms and the dictionary. $\rho\ge 0$ gives E2 (complete monotonicity of G); E0/E0′ follow from polynomial growth plus Wick counting exactly as in Proposition 10.5, with $\|\rho\|$ replaced by the polynomial-growth constant; E1/E3 and analyticity are inherited; ρ has no atom at 0 other than δ_0, so by Proposition 7.2 E4 is equivalent to vacuum uniqueness, which holds. Because the infrared modification multiplies the *scalar* two-point factor and leaves the label Gram matrix untouched, the correlation null spaces N_d are unchanged, hence (B1) (a rank comparison, Remark 2.13) is unchanged. (Y1) is the purely algebraic computation of Lemma 9.2 on the free shell and is mass-independent. For (Y2), the calculus law holds for any translation-invariant Gaussian covariance, since ∂_{x_0} differentiates the scalar factor; the flux law's shell integral is the translation generator and likewise depends only on the fact that *every* insertion, including T_{ab}, shares the same two-point factor. Finally, the OPE coefficients c_α are produced by the fixed algebraic prescription of Theorem 9.22 before any correlation is evaluated, from the short-distance Taylor coefficients; the $O(r^{-2})$ infrared contribution does not enter the $r^{-8}\log^{-2}$ layer, and the remainder $\int_{E>1/r}$ is controlled by ρ_{UV}. □

Remark 10.8. This upgrades Proposition 10.5: the counterexample already satisfies the ultraviolet law, the Ward identities, the field-content dictionary and the OPE data. It also shows that the trace anomaly — a renormalization effect (Remark 9.11) — is a *necessary constraint, not a dynamical explanation* of the infrared spectrum. The model remains a generalized free field, so (Y3)/(Y4) and the RG blocking mechanism have no counterpart in it; this is recorded in Remark 10.42.

Proposition 10.9 (The gap is not continuous in ultraviolet data)**.** *Let T be a generalized free field with a gap $\Delta>0$ (a single pole at $E=\Delta$ plus the ultraviolet tail of Proposition 10.7), and let T_ϵ be obtained by adjoining $\epsilon\rho_{\mathrm{IR}}$, $\epsilon>0$. Then T_ϵ and T have identical full formal short-distance asymptotics, identical OPE coefficients, identical Ward identities and identical field-content ranks; all*

separated-point correlations differ by at most $C\epsilon$; yet $\mathrm{gap}(T) = \Delta$ *while* $\mathrm{gap}(T_\epsilon) = 0$ *for every* $\epsilon > 0$. *Consequently no criterion built from finitely many — or from all formal orders of — short-distance data can decide the mass gap.*

Proof. Immediate from the support separation in Proposition 10.7: the modification is $O(r^{-2})$ at short distance and carries a positive spectral weight arbitrarily close to 0. The bound on separated correlations follows from $\|\epsilon\rho_{\mathrm{IR}}\| = \epsilon$; the gap statement from Lemma 10.1. □

Proposition 10.10 (Dimensional transmutation produces no physical mass in this construction; source-level reading)**.** *All registered scale-bearing objects are ultraviolet or conventional: the short-distance theorems (Theorems 9.10, 9.19, 9.22, Proposition 9.14) concern $r \to 0$, $a_N \to 0$; the Ward identity (Proposition 9.4) is a separated-insertion shell law, and Remark 9.5 records that no dilation Ward identity is constructed; the calibrations of Lemma 4.15 are conventions ($a(0) = 1$, $\beta(0) = \bar{\beta}$). Remark 9.11 states that $\Lambda = 1$ "records the construction's reference units; it does not define an intrinsic physical scale", and Remark 9.20 states that a scheme change only shifts* $\log(1/r)$ *by a constant. Hence no proved object links the ultraviolet scale to an infrared spectral lower bound: the inference "trace anomaly → broken scale invariance → Λ_{QCD} → $m \sim \Lambda_{\mathrm{QCD}}$" lacks its middle term in this construction. A proof would require a full-scale dilation Ward identity / Callan–Symanzik relation, which is absent.*

Proposition 10.11 (Boundary: what cannot yield a gap)**.** *No argument using only ultraviolet data (Theorems 9.10, 9.19, 9.22), rateless clustering (Proposition 7.15 and Appendix C), and pointwise finite-volume gaps can produce a mass gap.*

Proof. The model of Proposition 10.5 satisfies (AN1) and (PH1)–(PH4) with clustering (rateless, by (a) and (b)) and with ultraviolet behavior configurable independently of its infrared weights; every finite-volume truncation of any model has a pointwise gap. If a gap were derivable from such premises, it would hold in every model satisfying them; the counterexample shows otherwise. □

10.3 Spectral-gap bridges

Lemma 10.12 (Doeblin spectral gap)**.** *Let (X, μ) be a probability space and $P(x, \mathrm{d}y)$ a Markov kernel with stationary measure ν. If $P(x, \mathrm{d}y) \geq \varepsilon\, \nu(\mathrm{d}y)$ for some $\varepsilon > 0$ and all x, then for every f with $\int f \,\mathrm{d}\nu = 0$,*

$$\|Pf\|_{L^2(\nu)} \leq (1 - \varepsilon)\|f\|_{L^2(\nu)}.$$

If moreover $T \geq 0$ is the self-adjoint reflection-positive transfer operator with continuous kernel $c_1 \leq K \leq c_2$ and Perron root λ_1, then with $R = \log(c_2/c_1)$ the h-transform gives a Doeblin constant of order e^{-2R} and

$$\Delta \;\geq\; -a_t^{-1} \log\left(1 - e^{-2R}\right) \;>\; 0.$$

Proof. Write $P = \varepsilon\nu + (1 - \varepsilon)Q$ with Q a Markov kernel with the same stationary measure ν. For $\int f \,\mathrm{d}\nu = 0$, $Pf = (1-\varepsilon)Qf$. By Jensen $|Qf|^2 \leq Q(f^2)$, and stationarity gives $\int Q(f^2)\,\mathrm{d}\nu = \int f^2 \,\mathrm{d}\nu$, hence the L^2 bound.

For the second part, let $Th = \lambda_1 h$, $h > 0$, and set $P(x, \mathrm{d}y) = K(x, y)h(y)(\lambda_1 h(x))^{-1}\mathrm{d}\mu(y)$; then $\int P(x, \mathrm{d}y) = 1$ and $\nu = h^2\mu / \int h^2 \,\mathrm{d}\mu$ is stationary by symmetry of K. From $c_1 \int h \,\mathrm{d}\mu \leq \lambda_1 h(x) \leq c_2 \int h \,\mathrm{d}\mu$ one gets $h(y)/h(x) \geq c_1/c_2$ and $\lambda_1 \leq c_2$, hence $P(x, \mathrm{d}y) \geq (c_1/c_2)^2 \mathrm{d}\mu(y) = e^{-2R}\mathrm{d}\mu(y)$; since h is bounded above and below, $\mathrm{d}\mu$ and $\mathrm{d}\nu$ are equivalent with a normalization-dependent constant. The spectral radius on the zero-mean subspace is then $\leq 1 - \varepsilon$ and $\Delta = -a_t^{-1} \log(\lambda_2/\lambda_1) \geq -a_t^{-1} \log(1 - \varepsilon)$. □

Remark 10.13. The numerical constant depends on the normalization of h; the invariant statement is the *form* of the bound: a uniform positive lower bound on the kernel gives a positive gap. The degeneration of this bound with the volume is the content of Proposition 10.16.

Proposition 10.14 (Change of layer: the gap is scale-invariant)**.** *Assume exact blocking, i.e. the spatial transfer operator commutes with the one-step transfer matrix $\mathcal{T}_m$. Then blocking one step $a^{(m)} \mapsto b_m a^{(m)}$ replaces $\mathcal{T}_m$ by $\mathcal{T}_m^{b_m}$ and leaves the physical gap unchanged:*

$$\Delta_{m+1} = -\frac{1}{b_m a^{(m)}} \log \lambda_2\big(\mathcal{T}_m^{b_m}\big) = -\frac{b_m}{b_m a^{(m)}} \log \lambda_2(\mathcal{T}_m) = -\frac{1}{a^{(m)}} \log \lambda_2(\mathcal{T}_m) = \Delta_m.$$

Consequently a uniform gap may equivalently be read at the reference layer $a(0) = 1$.

Proof. Under exact blocking the coarse-grained transfer matrix is the b_m-fold time composition of the fine one, so its eigenvalues are $\lambda_j(\mathcal{T}_m)^{b_m}$ (the Perron root stays 1 after normalization). Substituting into $\Delta_m = -a^{(m)\,-1} \log \lambda_2$ gives the displayed identity. At the reference layer Lemma 4.15(3) gives $\beta_{N,p}(0) = \bar{\beta}$, a *fixed* constant; the weak-coupling running (4.8) is a statement about the bare layer only. Hence a strong-coupling / Dobrushin-type analysis at the reference layer is not excluded by asymptotic freedom, and the three conditions "$\lambda < 1/3$", "J_* lies in a spectral redundancy direction", and "a strong-coupling seed Δ_0 exceeds the accumulated Kato losses" — which arise only when one perturbs the transfer matrix layer by layer — are *not necessary* for (MG2). □

Remark 10.15. Actual (RG2) truncations break exact blocking; the resulting remainder is controlled at the schedule level by Lemma 5.9, not by a Kato perturbation series. This removes an apparent conflict between asymptotic freedom and strong-coupling estimates that was noted in earlier drafts of the supplement.

Proposition 10.16 (The global positivity bound degenerates with the volume)**.** *At the reference layer the one-step kernel is $K_0 = \exp[-L_E]$ with $L_E \geq 0$ the Wilson action on a whole spatial slice. Then the Doeblin lower bound of Lemma 10.12 degenerates as $p \to \infty$.*

Proof. Take $U = \mathbf{1}$, so every plaquette term $1 - \operatorname{Re}\operatorname{Tr}(U_\square)/m$ vanishes and $L_E(U) = 0$. Take U_* obtained by setting a fixed fraction of the spatial links to $-\mathbf{1}$ (possible for odd side length, `Patch.Large.odd`); at least cL_p^3 plaquettes then have $1 - \operatorname{Re}\operatorname{Tr}(U_\square)/m \geq c' > 0$, so $L_E(U_*) \geq c_0 \bar{\beta} L_p^3$. Hence

$$R_0(p) = \sup L_E - \inf L_E \;\geq\; c_0 \bar{\beta} L_p^3 \to \infty.$$

With $K_0 \in [e^{-R_0(p)}, 1]$, Lemma 10.12 gives $\varepsilon_p \lesssim e^{-2R_0(p)} \to 0$, so the lower bound $\Delta_M(p) \approx -a_t^{-1} \log(1 - \varepsilon_p) \approx e^{-2c_0 \bar{\beta} L_p^3}$ degenerates. □

Remark 10.17. This is a degeneration of a *bound*, not of the gap. Lüscher-type finite-volume corrections $\Delta m(L) \sim C L^{-1} e^{-\alpha m_0 L}$ show that on the massive side the true gap is volume-uniform. The corollary is that (MG2) cannot be proved by a single global Doeblin estimate; a *local* mixing estimate whose rate is independent of L_p is required.

Proposition 10.18 (The kernel family fixes the finite-volume spectral measure)**.** *Fix a volume p. Suppose two Wilson families satisfying (RG1)–(RG5), (UV1), and (AN1) E2 give the* same *fixed-volume separated-point kernels $S_p(v; x)$ for all label lists and all $x \in \Omega_n$. Then for every field-word vector ψ the OS scalar spectral measure $\mu_p^\psi(B) = \langle \psi, E_{H_p}(B)\psi \rangle$ is the same, hence $\sigma(H_p)$ is the same.*

Proof. Let $\mathcal{D}_p$ be the dense span of positive-time field words (`registered_words_dense`, field-audit step 5). For $\psi \in \mathcal{D}_p$, ψ is a finite linear combination of words, so by reflection positivity $\langle \psi, E_{H_p}(t)\psi\rangle$ is a finite linear combination of separated-point correlations $S_p(A^*A_t)$, i.e. of pointwise values of the kernels. These are given for *every* $t > 0$: (RG5) with an exhausting sequence of compact sets K_m and Theorem 5.17's smoothness on Ω_n supply each fixed configuration, however large the time separation. The Laplace transform $\widehat{\mu_p^{\psi}}(t) = \int_{[0,\infty)} e^{-t\lambda}\,\mathrm{d}\mu_p^{\psi}(\lambda)$ is therefore determined on $(0,\infty)$; by Lerch's uniqueness theorem for finite measures it determines μ_p^{ψ}. Polarization determines all $\mu_p^{\psi,\chi}$, and density of $\mathcal{D}_p$ determines the projection-valued measure E_{H_p}, hence the spectrum. □

Remark 10.19. Lerch's theorem needs only *pointwise* values; no t-uniform bound is required. What is not fixed by this argument is the *effective readability* of $\mu_p((0,\varepsilon))$ for large t; that is the remaining gap (Proposition 10.34).

Proposition 10.20 (Pointwise Laplace convergence suffices; strengthens Proposition 10.32)**.** *Assume*

$$\text{(MG2")} \qquad \exists\, \varepsilon_0 > 0 : \quad \mu_p^{\psi}\big((0,\varepsilon_0)\big) = 0 \quad \textit{for all } p \textit{ and all } \psi \perp \Omega,$$

equivalently $\|\mathcal{T}_{N,p}|_{\Omega_{N,p}^{\perp}}\| \le e^{-a_N\varepsilon_0}$ *for all* N, p*. Let* μ_{∞}^{ψ} *be the spectral measure of* any *infinite-volume subsequential limit of Theorem 6.4. Then* $\sigma(H) \subset \{0\} \cup [\varepsilon_0, \infty)$*, with the* same ε_0 *for every subsequence.*

Proof. Fix ψ in the dense field-word domain with $\psi \perp \Omega$. By (MG2"), μ_p^{ψ} is supported on $\{0\} \cup [\varepsilon_0, \infty)$, and $\rho_p(\{0\}) = 0$ for $\psi \perp \Omega$, so

$$\widehat{\mu_p^{\psi}}(t) = \int_{[\varepsilon_0,\infty)} e^{-t\lambda}\,\mathrm{d}\mu_p^{\psi}(\lambda) \le e^{-\varepsilon_0 t}\,\|\mu_p^{\psi}\| \le C_{\psi} e^{-\varepsilon_0 t},$$

with $C_{\psi} = \|\psi\|^2$ independent of p. Theorem 6.4 gives pointwise convergence of separated-point kernels, hence $\widehat{\mu_{p_k}^{\psi}}(t) \to \widehat{\mu_{\infty}^{\psi}}(t)$ for every $t > 0$. Suppose $\mu_{\infty}^{\psi}((0,\varepsilon_0 - \delta)) > 0$ for some $\delta > 0$. Then $\widehat{\mu_{\infty}^{\psi}}(t) \ge a\, e^{-(\varepsilon_0-\delta)t}$ with $a > 0$, and passing to the limit in the displayed bound gives $a\, e^{\delta t} \le C_{\psi}$ for all $t > 0$, a contradiction as $t \to \infty$. Hence $\mu_{\infty}^{\psi}((0,\varepsilon_0)) = 0$; the Urysohn/density argument of Proposition 10.3 then gives $\sigma(H) \cap (0,\varepsilon_0) = \varnothing$. Because ε_0 was fixed before the subsequence was chosen, the conclusion is subsequence-independent. □

Remark 10.21. This proof uses *only* pointwise convergence of Laplace transforms and Lemma 10.1. It needs neither strong convergence of transfer matrices, nor tightness of $\{\mu_p^{\psi}\}$, nor the restriction of t to integer multiples of a_N. It therefore supersedes the proof of Proposition 10.32 and adds the statement that the gap does *not* require uniqueness of the infinite-volume limit (O1/O9), only volume-uniformity — an asymmetry with the repository README, which lists the two open problems side by side.

Proposition 10.22 (A rate in E4 is equivalent to the mass gap)**.** *Under (AN1) E2 and* `registered_words_dense`*, let* $\mathcal{D}$ *be the dense field-word domain. The following are equivalent for* $\varepsilon_0 > 0$*:*

(1) $\sigma(H) \subset \{0\} \cup [\varepsilon_0, \infty)$*;*

(2) *there is* $C_{\psi,\chi} < \infty$ *for all* $\psi, \chi \in \mathcal{D}$ *with*

$$\big|\langle \psi, E(t)\chi\rangle - \langle \psi, \Omega\rangle\langle \Omega, \chi\rangle\big| \;\le\; C_{\psi,\chi}\, e^{-\varepsilon_0 t} \qquad (t \ge 0).$$

Proof. (1)⇒(2) is Lemma 10.2. (2)⇒(1): take $\chi = \psi$, put $\psi_\perp = \psi - P_0\psi$; then $\langle\psi, E(t)\psi\rangle - |\langle\Omega,\psi\rangle|^2 = \widehat{\rho_{\psi_\perp}}(t) \le C_\psi e^{-\varepsilon_0 t}$ and $\rho_{\psi_\perp}(\{0\}) = 0$, so Lemma 10.1 gives $\rho_{\psi_\perp}((0,\varepsilon_0)) = 0$. The Urysohn/density argument of Proposition 10.3 upgrades this to $\sigma(H) \cap (0,\varepsilon_0) = \varnothing$. □

Remark 10.23. Combined with Proposition 7.2, the three properties *clustering / vacuum uniqueness / (PH4)* are the same $o(1)$ statement, and the mass gap is exactly its *uniform-rate* version. Appendix C's sentence "No rate is asserted anywhere in the chain" must therefore be read as: the rate is absent at the *entrance* (`RemainingObligations.lean` keeps E4 as the single explicit premise of `completeOS`), not lost inside the six transport modules. The smaller scalar Θ-hierarchy of check A1 gives only a channel rate; the rate in (2) must be common to *all registered label words*, with constants allowed to depend on the word.

Proposition 10.24 (Every subexponential clustering rate is compatible with no gap)**.** *Let $\varphi : [0,\infty) \to (0,\infty)$ be non-increasing, $\varphi(t) \to 0$, satisfy the doubling condition $\varphi(t/4) \le C_4\varphi(t)$ and a polynomial lower bound $\varphi(t) \ge t^{-A}$ for some $A > 0$. Then there is a finite positive measure $\tilde\rho$ on $(0,1]$ with no atom at 0, accumulating at 0, such that $G(\tau) = \int e^{-\tau\lambda}\,\mathrm{d}\tilde\rho(\lambda) \le C\varphi(\tau)$ for all large τ, while $\tilde\rho((0,\varepsilon)) > 0$ for every $\varepsilon > 0$ — hence, by Lemma 10.1, no spectral gap.*

Proof. Put $\lambda_k = 2^{-k}$ and $w_k = \varphi(2^k) - \varphi(2^{k+1}) \ge 0$. Since $\varphi \to 0$, telescoping gives $\sum_{k\ge1} w_k \le \varphi(2) < \infty$ and infinitely many $w_k > 0$, so $\tilde\rho = \sum_k w_k \delta_{2^{-k}}$ is finite, has no atom at 0, and accumulates at 0. For $\tau \ge 2$ set $k_0 = \lfloor \log_2 \tau \rfloor$, so $2^{k_0} \le \tau < 2^{k_0+1}$ and $1/\tau < 2^{-k_0} \le 2/\tau$.

Tail ($k \ge k_0$): $\sum_{k\ge k_0} w_k e^{-\tau 2^{-k}} \le \sum_{k\ge k_0} w_k = \varphi(2^{k_0}) \le \varphi(\tau/2)$.

Head ($k < k_0$): write $j = k_0 - k \ge 1$. Then $2^{-k} = 2^j 2^{-k_0} \ge 2^j/\tau$, hence $e^{-\tau 2^{-k}} \le e^{-2^j}$; also $w_k \le \varphi(2^{k_0-j})$ and $2^{k_0-j} \ge \tau/2^{j+1}$, so $w_k \le \varphi(\tau/2^{j+1}) \le C_4^{(j+1)/2+1}\varphi(\tau)$ by iterating the doubling condition. Therefore the head is at most $\varphi(\tau)\sum_{j\ge1} C_4^{(j+1)/2+1} e^{-2^j} = C'\varphi(\tau)$, the series converging because e^{-2^j} dominates any exponential in j.

Combining, $G(\tau) \le (\varphi(\tau/2) + C'\varphi(\tau)) \le C''\varphi(\tau)$. Since $\tilde\rho((0,\varepsilon)) > 0$ for all $\varepsilon > 0$, Lemma 10.1 (contrapositive) gives no gap. □

Remark 10.25. The class covers τ^{-p}, $\log^{-q}\tau$, $\exp(-(\log\tau)^\beta)$ and the $1/m$-type rates supplied by the schedule estimates (`Fits.stepSlow` gives $J(m) \le J_0 + m$, hence $J(m)^{-1} \gtrsim m^{-1}$). Propositions 10.5, 10.7 and the tower model of Proposition 10.28 are instances. Consequently "each channel clusters exponentially" and "clustering holds with every subexponential rate" are both *insufficient*: only a common exponential rate $\varepsilon_0 > 0$ produces a gap. This also corrects the claim in (O6) that excluding continuous-mass generalized free fields is *equivalent* to an (MG3)-type estimate: the implication goes one way only.

Proposition 10.26 (Ward identities do not supply degree-uniformity)**.** *Assume (Y2). Then:*

(1) *if $v = D_b v'$, the spectral support of v is contained in that of v'; in particular the spectral threshold cannot decrease along a D_b-tower;*

(2) *D_b carries a free vector index b, so it raises tensor rank by one;* no scalar label lies in the image of any *D_b;*

(3) *consequently (Y2) relates spectral thresholds only* within *each tower of iterated derivatives; the thresholds of the tower bases are free parameters of (Y2), and no relation between distinct towers follows.*

Proof. (1) The calculus law $S(D_b v_0, v_1, \dots) = \partial^b_{x_0} S(v_0, v_1, \dots)$ multiplies the Wightman spectral measure by the momentum weight $|p_b|^2$ (a non-negative factor); a non-negative weight can only shrink the support. (2) In the flux law the index a is summed while b is free, so $D_b v_0$ has index set $\{b\} \cup \mathrm{idx}(v_0)$; a rank-zero scalar cannot arise. (3) By Section 9.5, $\mathcal{F}_d = \mathcal{F}_{d-1} \oplus E_d$ with $n_d = \dim E_d$ *fresh* fields per layer; (1)–(2) leave each E_d's base unconstrained, and Remark 9.5 states that no dilation Ward identity is constructed. □

Remark 10.27. This retracts an earlier claim that (Y2) gives the degree-uniformity needed for (MG3′) for free. The Ward identities do *not* provide a common rate across towers. The remedy is not to prove a degree-uniform estimate tower by tower, but to use the *full-spectrum* form (MG2"), which bypasses all channel/tower covering questions (Proposition 10.20).

Proposition 10.28 (Tower model; calculus-law part; partially constructed)**.** *Let the label filtration decompose into derivative towers indexed by $t \in \mathbb{N}$, each with a base label. There is a tempered Gaussian hierarchy satisfying (AN1), (B1), (Y1), the* calculus law *of (Y2), the ultraviolet law of Theorem 9.10 and the OPE coefficients of Theorem 9.22, in which every tower clusters exponentially with rate $m_t = 2^{-t} > 0$, yet $\inf_t m_t = 0$ and there is no mass gap.*

Proof. Take a block-diagonal KL weight: tower t carries $\rho_t \geq 0$ supported in $[m_t, \infty)$, has no atom at 0, and all towers share the *same* large-E ultraviolet tail of Proposition 10.7. Within a tower, a derivative $P(D)w_t$ multiplies ρ_t by a polynomial momentum weight, so its support stays in $[m_t, \infty)$ and the calculus law holds by construction. Block diagonality preserves the label null spaces, hence (B1); the shared ultraviolet tail preserves the short-distance constant and the OPE coefficients; $\rho \geq 0$ gives E2, and no 0-atom gives E4 together with vacuum uniqueness (Proposition 7.2). Each tower clusters at rate m_t, but $\rho((0,\varepsilon)) > 0$ for every $\varepsilon > 0$, so Lemma 10.1 gives no gap. □

Remark 10.29 (Open identification point: the flux law in tower models). In this tower model the *flux (shell) law* is not checked: T_{ab} is a single rank-two label, whereas Wick contractions with a base label from a different tower vanish, so the left-hand side of the flux law and $\partial^b_{x_0} S(v_0, \dots)$ need not match. By contrast, the *uniform* superposition of Proposition 10.7 satisfies both laws because all insertions share one two-point factor. Accordingly Proposition 10.28 is a *calculus-law* model, and the statement "each channel clusters exponentially yet there is no gap" is established under the calculus law alone. Whether the flux law admits a tower-wise construction is left open; the degree-uniformity conclusion of Proposition 10.26 does not depend on it.

Proposition 10.30 (No gap-shaped predicate in the source signatures; source-level reading)**.** *The records read directly from the submitted sources (`YM.Overlay.Schedule.Priced / Tabs / consumes`, `YM.RG.History.Fits`, `YM.RG.Calendar`, `YM.Bid.Fuels.Circle`, `YM.Next.MULTISOURCE.Mixed.mixed_circle`, `YM.RG.ActivityBounds`, `YM.Source.Insertion.TransferData.Bo`) contain no predicate about a large-time translation s, a spectral projection $E_H((0,\varepsilon))$, or a second transfer-matrix eigenvalue λ_2. Specifically:*

- *`Tabs` is produced by `Stock.consumes` from `stepsize` (a polylog doubling condition) and `afford` (a $z^{17} = o(s)$ cost bound) together with* Tendsto $(J(i+1)^{-1}) \to 0$*;*
- *`Fits.deep` asserts $J \to \infty$ with* no rate*, `Fits.stepSlow` gives $J(m+1) \leq J(m) + 1$, and `Calendar.hatBeta` $= \bar{\beta} + \mathrm{g} \cdot \mathrm{scale}(m)$, `coverSize`, `monitor`, `forcedAllowance` are logarithmic/polylogarithmic functions of $\hat{\beta}$;*
- *`Mix.Circle C d u` is an existence statement about `Story/Lane/Stubs.Holds/Integrable` witnesses with `width C r k` i-type envelopes; `Bid.installs` consumes it to produce*

`Input.Metered C`, a field-content/ultraviolet metering *record;* *`mixed_circle`* *derives* *`Circle`* *from* *`Responsive`*, *`Margins`*, *`Addresses`*, *`Tabs`*;

- *`Activities.KBound/EBound`* *are summability caps for polymer gas/regular occurrences;* *`TransferData.Based.front/back`* *are coefficient-norm bounds with factors* columns $\cdot\, b_m^d$ *(Remark 5.21: no uniform inverse bound).*

By the soundness observation of Proposition 10.5(d), no statement whose predicates are absent from these signatures can be derived from them. Hence (MG) is not a receipt waiting to be discharged inside the formalization*: it is a new estimate external to the sources.*

Remark 10.31. This is a structural signature check, *not* a formal independence theorem: the six `Cluster*` modules were not read verbatim in this study (Appendix C's module-by-module description was used). The README's actual wording — *"a positive mass gap ... outside the scope of the public challenge"* — together with the reached closure of 100,168 declarations, is consistent with this reading.

10.4 Conditional routes to a mass gap

We record four candidate infrared inputs. Each is a proposal; none is proved anywhere in this paper or in the repositories.

(MG1) Renormalized cluster expansion. At each blocking level m a string tension/mass lower bound $\sigma_m > 0$, with summable cross-level losses, such that $\inf_m \sigma_m \geq \sigma_* > 0$ in physical units as $a_N \to 0$. (Strong coupling seeds exist [Wilson, OS78, Seiler]; the bare couplings here run to weak coupling, (4.8), so a no-phase-transition bridge is additionally needed.)

(MG2) Transfer-matrix gap, uniformly. There is $\Delta > 0$ with $\|\mathcal{T}_{N,p}|_{\Omega_{N,p}^{\perp}}\| \leq e^{-a_N \Delta}$ for all N, p, preserved under blocking.

(MG3′) Infrared exponential clustering, uniform rate. There exists $\varepsilon_0 > 0$ such that for every finite label word A there is $C_A < \infty$ with $\sup_p \left|S_p(A^* A_s) - |S_p(A)|^2\right| \leq C_A e^{-\varepsilon_0 s}$ $(s \geq 0)$; the constants may depend on A, the rate ε_0 may not (Proposition 10.26 shows the Ward identities do not supply this uniformity).

(MG4) Wilson loops with area law. Non-local Wilson loop observables (outside the local dictionary of Proposition 2.12) together with a uniform area-law bound $\langle W_C \rangle \leq C\, e^{-\sigma\, \mathrm{area}(C)}$, $\sigma > 0$.

Proposition 10.32 ((MG2) implies a continuum gap)**.** *Under (MG2) and the limit hierarchy of Theorem 6.4,* $\sigma(H) \subset \{0\} \cup [\Delta, \infty)$.

Proof. On the finite lattice define the semigroup $E_{N,p}(t) = \mathcal{T}_{N,p}^{t/a_N}$ (with t an integer multiple of a_N). By (MG2), for $\psi \perp \Omega_{N,p}$, $\|E_{N,p}(t)\psi\| \leq e^{-\Delta t}\|\psi\|$. The uniform kernel convergence of Theorems 5.17 and 6.4 gives strong convergence $E_{N,p}(t) \to E(t)$ on the dense domain generated by field words (the restriction of t to integer multiples of a_N is removed by strong continuity of the limiting semigroup). Hence for ψ in that dense domain, $\psi \perp \Omega$,

$$\langle \psi, E(t)\psi \rangle = \lim \langle \psi_{N,p}, E_{N,p}(t)\psi_{N,p} \rangle \leq e^{-\Delta t}\|\psi\|^2.$$

Lemma 10.1 gives $\rho_\psi((0, \Delta)) = 0$, and the density argument of Proposition 10.3 gives $\sigma(H) \cap (0, \Delta) = \varnothing$. □

Remark 10.33 (Superseded proof). A stronger and shorter proof is given in Proposition 10.20: it uses only pointwise convergence of Laplace transforms, yields the same Δ for every infinite-volume subsequence, and shows the gap does not require uniqueness of the infinite-volume limit (O1/O9), only volume-uniformity.

Goal	Minimal additional input	Consequence
Full gap $\Delta > 0$	(MG) = (MG2), or (MG3′) + full channel coverage	Definition 1.2; subsequence-independent (Proposition 10.20)
Channel spectral threshold	(MG3′) in one channel	$\rho_\psi((0,m)) = 0$ for that channel only
Continuum gap from the lattice	(MG1)/(MG2)	Proposition 10.20
Exclusion of gapless models	O6 below	necessary-not-sufficient (Proposition 10.24)

Proposition 10.34 (The single remaining predicate)**.** *The four candidate inputs (MG1)–(MG4) and the volume-uniform transfer-matrix condition (MG2) are equivalent projections of one statement. Register*

> ***(MG)*** *There exists $\varepsilon_0 > 0$ such that for every finite label word A there is $C_A < \infty$ with*
>
> $$\sup_p \left| S_p(A^*A_s) - |S_p(A)|^2 \right| \;\le\; C_A\, e^{-\varepsilon_0 s} \qquad (s \ge 0),$$
>
> *where C_A may depend on A but the rate ε_0 may not, and the supremum is over all volumes.*

Then:

(1) *(MG) ⇒ (MG2") ⇒ (by Proposition 10.20) $\sigma(H) \subset \{0\} \cup [\varepsilon_0, \infty)$ for* every *infinite-volume subsequential limit;*

(2) *(MG) restricted to all registered label words* implies *the mass gap of the limit theory, and the two rates coincide (Propositions 10.20, 10.22); the converse implication is open (Remark 10.42(6));*

(3) *(MG1), (MG2), (MG3′) and (MG4) are respectively the cluster-expansion, spectral, rate, and Wilson-loop projections of the same uniform infrared lower bound; (MG4) and single-channel versions give only a* channel *threshold (Remark 10.4, Proposition 10.26);*

(4) *dropping either uniformity — over volumes p, or over label words with a common rate — invalidates the conclusion (Propositions 10.16, 10.24, 10.28).*

Proof. (1) Apply Lemma 10.1 to the diagonal matrix element of each field-word vector and then Proposition 10.3's Urysohn argument to obtain $\mu_p^\psi((0,\varepsilon_0)) = 0$ uniformly in p; Proposition 10.20 finishes. (2) is Proposition 10.22. (3) is the content of Propositions 10.12–10.20 and Remarks 10.4/10.27. (4) is Propositions 10.16, 10.24 and 10.28. □

Form	Rate statement
(MG2″)	$\|\mathcal{T}_{N,p}\|_{\Omega^\perp}\| \le e^{-a_N \varepsilon_0}$, all N, p.
(MG3′)	$\sup_p \|S_p(A^*A_s) - \|S_p(A)\|^2\| \le C_A e^{-\varepsilon_0 s}$, all words A.
(MG)	as (MG3′), the common rate over *all* registered label words.
Gap	$\sigma(H) \subset \{0\} \cup [\varepsilon_0, \infty)$, same ε_0, every subsequence.

Remark 10.35 (Final boundary). Everything else in Section 10 is now either proved or reduced to (MG): O6 is a necessary-not-sufficient sanity check (Proposition 10.24), O1/O9 are not needed for the gap (Proposition 10.20), O5 requires in addition an atom in the spectral measure (Proposition 10.38), O8 gates O2 (Lemmas 10.36–10.37), and O7/O10 are ultraviolet/perturbative questions independent of the spectral gap. **A positive mass gap is not proved.**

10.5 Sharpened statements of the remaining open problems

(O1) Convergence and uniqueness of the infinite-volume limit. Missing input: a volume-direction estimate — e.g. (RG2) upgraded from volume-uniform bounds to volume-uniform difference estimates — or a uniqueness statement for the RG fixed point. *Reduction proved here:* if the limits obtained from any two volume subsequences agree on the countable index set $\mathcal{L} \times \{K_m\} \times \{\alpha\}$ of Theorem 6.4, then they are equal as kernels on Ω_n (uniform continuity on compacts plus exhaustion), and hence define the same functionals.

(O2) A continuum probability measure. Missing input: positivity on the *unrestricted* joint Schwartz test space, uniqueness of contact extensions (O8), and Minlos-type growth and compatibility conditions.

(O4) Wilson loops and the area law. Missing input: incorporation of non-local line observables into the hierarchy; the local dictionary of Proposition 2.12 covers only $\bigcup_d \mathcal{F}_d$. Note the area law would give at most a channel threshold, not the full gap (Proposition 10.3 and Remark 10.4).

(O5) Single-particle states. A gap does not suffice: Proposition 10.38 below shows a reflection-positive generalized free field with spectral gap m whose Källen–Lehmann measure has no atoms, hence no isolated one-particle shell and no single-particle states [GJ].

(O6) General exclusion of continuous-mass generalized free fields. This is a necessary, not sufficient, prerequisite for the mass gap (Remark 10.6); (MG3′) implies it, but excluding free-field descriptions does not by itself supply the uniform rate, since interacting gapless models are not generalized free fields (Proposition 10.24). It should be read as a sanity check on (MG3′), not as an equivalent formulation.

(O7) Order-by-order OPE/perturbation matching. Missing input: identification of the algebraic prescription coefficients c_α of Theorem 9.22 with perturbative coefficients order by order, including operator-valued renormalization conditions and state-dependent parts.

(O8) Uniqueness and positivity of contact extensions. Two proved obstructions: diagonal-flat tests annihilate diagonal-supported distributions (Lemma 10.36), so the check A1 comparison cannot decide uniqueness; and extensions are never unique without further requirements (Lemma 10.37). Missing input: fixation of the diagonal part by reflection positivity on the unrestricted test space, plus lattice convergence of the extended functionals.

(O9) Subsequence independence and the mixed vacuum. Missing input: the same volume-direction control as (O1), plus clustering of the mixed theory (which would collapse all phases by Lemma 6.6).

(O10) Infrared flow and Euclidean invariance, part (iii). The running-coupling argument is perturbative in the marginal coupling and does not address the infrared side; the third part of Euclidean invariance (stress-energy flux decay) is recorded as a construction input in Section 8.4. An independent expansion of the latter would reduce the reliance on construction inputs without changing the conclusions.

Lemma 10.36 (Diagonal-flat tests annihilate diagonal-supported distributions)**.** *Let $T \in \mathcal{D}'\big((\mathbb{R}^4)^n\big)$ be supported on a coincidence diagonal Δ_{ij}, and let $f \in \mathcal{S}_0\big((\mathbb{R}^4)^n\big)$ (all derivatives vanish on Δ_{ij}). Then $\langle T, f\rangle = 0$.*

Proof. By the local structure theorem for distributions supported on a closed submanifold, $T = \sum_{|\beta|\le N} \partial^\beta T_\beta$ with T_β measure-type distributions on Δ_{ij} (normal δ functions and their derivatives). Hence $\langle T, f\rangle = \sum_\beta (-1)^{|\beta|} \int_{\Delta_{ij}} \partial^\beta f \, \mathrm{d}T_\beta = 0$, since all derivatives of f vanish on Δ_{ij}. □

Lemma 10.37 (Contact extensions are never unique)**.** *Let T_0 be a distribution on the full configuration space satisfying the* `Contacts` *predicate of Remark 9.24, and let $T_\gamma \neq 0$ be any distribution supported on the coincidence set $\Delta = \bigcup_{i<j} \Delta_{ij}$. Then $T_0 + \gamma T_\gamma$ satisfies* `Contacts` *for every $\gamma \in \mathbb{C}$.*

Proof. A separated test f has tubular support avoiding Δ, so $\langle T_\gamma, f\rangle = 0$; the new functional therefore agrees with T_0, and a fortiori with $\int f\, T.\texttt{correlation}$, on separated tests, and is continuous and linear. □

Proposition 10.38 (A spectral gap does not give single-particle states)**.** *There is a reflection-positive generalized free field whose spectrum begins at $m > 0$ (so the gap is $\Delta = m$) and whose two-point function decays on the exponential scale $e^{-m\tau}$ (up to powers), but whose Källen–Lehmann measure $\mathrm{d}\rho(\mu^2) = \mathbf{1}_{[m^2,\infty)}\, r(\mu^2)\, \mathrm{d}\mu^2$, with $r \ge 0$ locally integrable, polynomially growing, and atomless, has no isolated one-particle shell; consequently the theory has no single-particle states [GJ].*

Proof. Take any such r (e.g. $r \equiv 1$ on $[m^2, \infty)$); the corresponding measure ρ is positive with $\operatorname{supp}\rho \subset [m^2, \infty)$, so the Hamiltonian of the associated generalized free field satisfies $\sigma(H) \subset \{0\} \cup [m, \infty)$, and Lemma 10.1 gives exponential decay at rate m. Since ρ has no atoms, there is no mass shell $\delta(\mu^2 - M^2)$ in the Källen–Lehmann representation and hence no one-particle subspace. □

Remark 10.39 (Dependency summary). The structural edges are: O3 ⇔ E4-with-a-rate (Proposition 10.22); O6 → O3 (necessary-not-sufficient, Proposition 10.24); O8 → O2. The bridges of Lemmas 10.1–10.2 and Propositions 10.3, 10.18, 10.20 reduce the mass gap to the single volume-uniform infrared estimate (MG) (Proposition 10.34); uniqueness of the infinite-volume limit (O1/O9) is *not* needed (Proposition 10.20). Everything else in this section is a precise statement of what is missing.

Remark 10.40 (Retractions from the exploratory study). The following intermediate claims were superseded during the source-level analysis; they must not be carried into citations of this paper.

#	Retracted claim	Replaced by
R1	(PH4) is an independent seventeenth input	Proposition 7.2; sixteen inputs
R2	(MG2) requires a new infrared *predicate* absent from (RG1)–(RG5)	Proposition 10.18: the kernel family already fixes μ_p; what is missing is an *estimate* on it (Proposition 10.30)
R3	(Y2) gives degree-uniformity for free	Proposition 10.26
R4	A filtration-layer model with per-layer thresholds violates no Ward law; a per-tower model satisfies all of (Y2)	Proposition 10.28: the layer model violates the calculus law; the tower model satisfies the calculus law only, with the flux law left open
R5	`Mix.Circle` is the source-level registration of (MG2)	Proposition 10.30: `Circle` $\to$ `Input.Metered` is a UV/field-content chain
R6	The left half of the feasibility condition $K/(2b_0 \log b_{\min}) < \bar{\beta}$ is an open obstacle	Lemma 5.9 (`Rolls.exists_rolls` + weighted Cesàro) closes it; the right half *is* (MG)
R7	`registered_words_dense` might cover only the scalar Θ channel; the E4 chain includes `ArityGrowth`	The full field basis is used (Remark 10.4); the six E4 modules are `ClusterKernel`, `ClusterCore`, `ClusterCones`, `ClusterDomination`, `ClusterTarget`, `ClusterSchwartz`; `ArityGrowth` belongs to E0′
R8	The counterexample is Gaussian; a non-Gaussian interacting upgrade was claimed	The Gaussian model suffices (Proposition 10.7); a mixture attempt loses (PH4) or is conditionally Gaussian (Remark 10.41)
R9	A pointwise gap at each volume suffices for the rate	The family of Proposition E.2 has pointwise gap $m_p = 1/(2p+1)$ but no uniform one; (MG) with $\sup_p$ is not ruled out by it
R10	"Not a receipt dischargeable inside the formalization"	Tightened: reached closure of 100,168 declarations; no gap-shaped predicate in the signatures (Proposition 10.30)
R11	The autocorrelation criterion applies to non-Gaussian mixing	The Markov-chain reduction is unverified beyond the Gaussian case (Theorem E.5 is proved for Gaussian models)
R12	The Lean sources rule out the E.2 family	Not read; the family is consistent with the reached closure (Remark 10.31)

Remark 10.41 (Why the counterexample need not be non-Gaussian). Two attempts to upgrade Proposition 10.5 to a genuinely interacting model give a dichotomy: a *mixture* $\int T_\lambda \, \mathrm{d}\nu(\lambda)$ of Gaussian hierarchies with a non-atomic ν is non-Gaussian and preserves (B1)/(Y1)/(Y2)/UV data, because (B1) is a rank comparison and the Ward/OPE prescriptions are linear in correlations; but its Hilbert space is a direct integral with a *trivially acting* base, so $\ker H \supseteq L^2(\nu)$ and vacuum uniqueness is lost. Making the base *ergodic* under the spatial translation action restores vacuum uniqueness by covariance, but the model remains *conditionally Gaussian*. Since Proposition 10.7 already provides a gapless model satisfying the registered ultraviolet, Ward and field-content conditions, the question of a genuinely interacting gapless model is a question of *model classification*, not of the decidability of the gap. It does not affect (MG).

Remark 10.42 (Open identification points (not theorems)). 1. *Drift identity.* Lemma 5.9 identifies the source's `drift` (`Intent.Large.slow/kick`, `Bench.recover`'s step error) with the one-

step remainder ε_m. The identification is supported by the shared shape $|\beta_{m+1} - \beta_m + \mathrm{g} \log L| \leq$ drift, but the definitions of `Survey.melt`, `stage`, `profile` were not read; Lemma 5.9 is a theorem *about the schedule-controlled remainder*, and its application to the submitted recursion is a source-level reading.

2. *`Fits` discharge.* Whether `Nonperturbative.existence` discharges `Fits` or it enters through `WilsonRegularity` was not read. `Rolls.exists_rolls` is a closed theorem, so the downstream argument does not depend on this.

3. *Tightness of $\{\mu_p^{\psi}\}$ at $t \downarrow 0$.* Not used in Proposition 10.20, but still needed for a full probability-measure construction (O2).

4. *Spectral type of H_p.* At fixed (N, p) the transfer kernel is continuous on a compact group product, hence Hilbert–Schmidt, so $\mathcal{T}_{N,p}$ is compact and $\Delta_{N,p} > 0$. After the UV limit, compact spatial support alone does *not* imply a compact resolvent: the model of Proposition 10.5 has $0 \in \sigma_{\mathrm{ess}}(H_p)$. By Proposition 10.18 the spectral type is fixed by the kernel family; the open question is the positivity of $\inf_p \Delta_p$, i.e. (MG2").

5. *Six `Cluster*` modules.* Their verbatim sources were not read; Propositions 10.22/10.30 use Appendix C's module-by-module description.

6. *Reverse implication.* (MG) $\Rightarrow$ gap is proved; gap $\Rightarrow$ (MG) is not. Lüscher-type finite-volume analysis suggests the reverse in known massive model classes, but no general proof is given here.

7. *Flux law in tower models.* See Remark 10.29.

11 How the formal proof was checked

The committee distinguishes three activities. *Formal verification* checks that a proof establishes a stated mathematical claim from its definitions and permitted axioms. In Lean, proofs are represented by terms that a small kernel checks against the inference rules; automated tactics must produce terms passing this check. *Validation* asks the further question whether the formal claim correctly expresses the intended Yang–Mills construction. This is addressed through review of the definitions and through independent adapters—Lean proofs that the constructed objects satisfy requirements from previous formalizations. The fields and correlations were traced from the specified Wilson lattice theory through the ultraviolet and large-volume limits, the phase decomposition, and the relativistic reconstruction.

Lean's flexibility also permits unwanted changes: a submitted development can alter the problem statement or an underlying definition during verification and make an invalid proof appear to verify. For this reason a separate comparison against a fixed reference specification is needed; this is performed by the tool *Comparator* [Comparator], which checks a submitted solution against a separately frozen reference statement, compares target statements and their underlying definitions, audits the proof's axiom dependencies, and rechecks the exported proof with a Lean kernel. The precise targets and verification configurations are recorded in Section 14.

After the original submission, OpenAI informed the group that its original public Comparator target `YangMills.jointExistence` omitted some of the connections to the lattice construction established by its proof. This weakness does not undermine the review: the combined target `YMReview.Acceptance.combined_acceptance` requires those connections by retaining the native

construction certificates for the lattice observables, limits, phases, and reconstructed fields. The assessment of these connections therefore rests on the stronger target and the checks below, rather than on the public existence statement alone.

11.1 Clean rebuild, inventory, and kernel check

The recorded clean build used Lean 4.32.0 on Apple Silicon, with the pinned official Mathlib dependency cache and no pre-existing project build artifacts. It rebuilt 9,470 project modules, completed in about 184 minutes, and reported zero errors. A subsequent incremental confirmation also passed. All 9,477 supplied Lean source files matched the submission manifest. This is record M1.

The inventory inspected all 250,319 project declarations imported by the construction, including generated declarations. It found no project axiom declarations, no direct references to Lean's placeholder axiom `sorryAx`, and no declarations marked unsafe. The 302 generated declarations are partial recursion helpers; none occurs among the dependencies of the audited principal results. The sole source-level `sorry` belongs to the reference statement in `Challenge.lean`, which is loaded separately from the solution; the solution supplies a proof without importing that file.

The submission's verifier `VerifyJoint.lean` compared the target statement and its reachable shared definitions in the reference and solution environments. It then rechecked the 100,168 project declarations on which the 17 principal results depend, using a fresh kernel environment initialized with the pinned external libraries, and audited 14,364 external declarations referenced at the boundary of those project dependencies. The principal proofs use only the permitted foundational axioms `propext`, `Classical.choice`, and `Quot.sound`. The inventory covers all imported project declarations; the kernel check covers the dependencies of the 17 principal results. This submitted verifier is separate from the upstream Comparator used for the three runs reported in Section 11.2.

The clean build emitted 2,453 warnings in 1,063 files, most concerning proof-script cleanup or deprecated syntax; the audited solution proofs have no dependency on the placeholder warning.

Remark 11.1 (Code quality concerns). Of 250,319 project declarations (including auto-generated auxiliary declarations), only 100,168 (including auxiliary declarations) are used in the dependency tree of the main theorems; the difference is less drastic for source declarations but still significant. The submission also uses unconventional naming, e.g. `Farewell.dead` in `TicketsChart.lean` and `Cake` and `cookie` in `Kitchen.lean`. These are recorded as concerns; they do not affect the kernel acceptance of the proofs.

11.2 Comparator results

Comparator passed on October 30, 2026 for the submitted public theorem `YangMills.jointExistence`, and on November 3, 2026 for the independent combined theorem `YMReview.Acceptance.combined_acceptance`. The November 3 verification record preserves the checked target's identity, configuration, source-file hashes, and logs; the corresponding Lean files in the release are identified by the review source manifest. The run passed statement and definition comparison, the permitted-axiom check, and the built-in Lean kernel. Four controls behaved as expected: a valid proof was accepted, while a changed statement, a forbidden axiom, and a changed definition were rejected. These runs are record M2.

The combined run rebuilt its 40 required review modules in a new project and reused the previously built submission artifacts and dependency cache. The separate clean rebuild of the submission is documented in Section 11.1. All three runs used pinned, unmodified upstream Comparator and exporter revisions and Comparator's built-in Lean kernel. The project and exporter

used Lean 4.32.0; the Comparator kernel used Lean 4.34.0. On macOS the runs used the upstream development launcher `fake-landrun.sh`, which runs the checks without Linux Landrun's process isolation; the records therefore document statement, definition, axiom, and kernel checks, but do not test isolation of potentially hostile build code. An additional external kernel implementation was not used. Each verification record identifies the exact target, inputs, and configuration.

A further Comparator run on November 5, 2026 passed for the OPE statement-alignment theorem described in Section 13.5. These are two distinct records: the submission's own `VerifyJoint.lean` kernel replay (record M1) and the committee's three upstream Comparator runs reported here; the latter used `fake-landrun.sh` on macOS and therefore document statement, definition, axiom and kernel checks, but not process isolation.

11.3 Additional review after the kernel check

Lean kernel acceptance determines whether the proof follows from its formal definitions and permitted axioms; it does not settle whether those definitions express the intended natural language problem. For this reason the group carried out extensive (but not exhaustive) human review of the Lean definitions, checking them for alignment with the problem statement and with expectations based on prior discussion; the glossary of Appendix A documents these definitions.

The definitions were also compared with previous formalizations. In particular, the group proved that the Schwinger functions of the centered stress trace in the selected phase satisfy the OS axioms and the E0′ growth condition in the formulation of [DMY]; Section 12 describes the comparison of test spaces and axiom formulations.

The scalar OS, lattice, field-content, translation Ward, leading trace, and common-scale checks are assembled in `YMReview.Acceptance.combined_acceptance`, which compiled and passed a Comparator check. This theorem retains the stronger construction data supplied by `YangMills.constructed`, as well as the public existence conclusions. The retained data connect all the additional properties to the same construction. The OPE statement-alignment check of Section 13.5 is a separate endpoint with its own Comparator run.

11.4 Identifying the constructed theory as Yang–Mills

The combined target retains the construction certificates `C.Makes` and `Spec.Final`, connecting the Wilson lattice fields to the continuum correlations, phase decomposition, and reconstructed QFT described in Sections 4–8. These data fix the lattice fields' classical curvature-polynomial meaning and renormalization, identify continuum field words across both limits, and connect the selected phase's boundary distributions to reconstructed vacuum expectations.

For each group and positive construction parameter, one Wilson family is fixed before the extraction order is chosen; one refinement, correlation table, and final witness then satisfy all the combined checks. In particular, the Ward identity, the scalar OS axioms, and the leading trace law concern the same stress tensor and the same selected phase. Together with the field-content and short-distance results, these connections support identification with Yang–Mills theory. Universality, infrared uniqueness, and full perturbative agreement remain open (Section 10).

Remark 11.2 (Worked construction-chain audit). The commented `OctoberChain.lean` proves seven consequences of the retained construction data, without adding hypotheses or auditing the submitted proof's individual steps. Its boundary-uniqueness lemma concerns a fixed analytic continuation, not uniqueness of the infrared limit. The compilation record and reproduction instructions document this supplementary Lean check, separate from the three Comparator runs.

The seven consequences of the construction-chain audit

The audit fixes a family `C` (`C.Makes`), an extraction order, the unselected mixed hierarchy `S : Table C.book`, a final witness `F : Spec.Final` retaining fields, stress tensor, phase construction, analytic continuation and relativistic reconstruction, and the acceptance predicate `Accepted` adding the review requirements (finite Wilson law, full-UV convergence). Combining the proved existence of an accepted witness (`YMReview.Acceptance.combined_acceptance`) with the following seven lemmas gives consequences for any accepted construction; the proofs inspect definitions, not tactic history.

1. *Insertion formula* (`actual_insertion_formula`, proved by `rfl`). The lattice insertion is

$$I_{N,p}(v,x;U) = a_N^{-d}\,(\texttt{places}\,\widetilde{v})(x;U), \qquad \widetilde{v} = \big(\texttt{terminal}^{-1} \circ \mathrm{diag}(t^{-1})\big)(v),$$

   where `terminal` undoes the terminal renormalization and $t \neq 0$ is the nonzero RG threshold scale; the composite is a linear equivalence of the raw label space at each cutoff. This checks what `C.O` denotes; classical Taylor matching of local sections belongs to `C.Makes`, retained in `Accepted.lattice.makes`.

2. *No coefficient direction is removed* (`every_primary_is_represented`). Every placed primary section is represented by an actual renormalized insertion up to the common factor a_N^{-d}, at every site of the fixed lattice. Finite-cutoff statement: no uniform bound on the inverse coefficient map as the cutoff is removed is provided, so this does not by itself exclude a degenerate continuum limit (the field-content and short-distance results do).

3. *Nonconstant primaries survive renormalization* (`nonconstant_primary_survives`). An invertible coefficient change and a nonzero common scale cannot turn every nonconstant primary observable into a constant. It does not assert that every local section is nonconstant or that chosen fields have nonzero limiting correlations.

4. *Lattice origin of continuum field words* (`every_field_word_has_lattice_origin`). Every field word has fixed raw representatives, chosen once depending on the table and the word but not on cutoff or volume, and its entire mean correlation function is an iterated limit of lattice moments:

$$\mathrm{smear}(N,p,v,t) \to \mathrm{row}(p)\ (N\to\infty), \qquad \mathrm{row}(\texttt{o.per}\,k) \to \int t(x)\,\texttt{S.useAt}(x;a)\,\mathrm{d}x\ (k\to\infty).$$

   No interchange of limits, no rate, no uniqueness across volume subsequences; the right-hand side uses the *mixed* table `S`, and replacing it by a selected phase is not justified by this lemma.

5. *Phase readings are ground-state algebra products* (`phase_second_moments_are_ground_products`). For distinct words w, u in positive-time slabs,

$$\int_X \xi(w)\,\xi(u)\,\mathrm{d}\nu(\xi) = \big\langle \Omega_0,\, \texttt{slot}(w)\,\texttt{slot}(u)\,\Omega_0 \big\rangle,$$

   via multiplicativity of characters: products of phase readings are tied to products in the actual commutative ground-state algebra. This is stronger than the barycenter identity alone, but it is not a factorization of a phase's multipoint correlation into one-point functions, and it does not say that each phase is itself a lattice subsequential limit or that there is only one phase.

6. *Reconstruction uses the selected phase's boundary distribution* (`reconstructed_words_are_selected_boundary`). The vacuum expectation of every reconstructed relativistic field word equals the expectation of the *same* word in the selected phase's boundary distribution, through the realization certificate `qt.meets` belonging to the same `F`.

7. *A fixed continuation has a unique boundary value* (`same_continuation_same_boundary`). If two contours from the same selected phase have equal analytic continuations, their distributional boundary pairings against every joint Schwartz test agree, by uniqueness of limits in $\mathbb{C}$. The hypothesis `same` is essential: this does not establish uniqueness of analytic continuation from Euclidean data, of the infrared subsequential limit, or of the phase.

The transitive axiom dependencies of all seven lemmas are printed by `#print axioms` in the same file; they use only the permitted foundations `propext`, `Classical.choice`, and `Quot.sound`. The ordinary compilation of these lemmas is supplementary evidence, not a Comparator run or a fresh rebuild.

12 Independent verification of the OS axioms

The OS comparison uses the centered trace of the stress tensor Θ, defined by (9.2), in the selected component ξ^*. The adapter between the two formalizations changes coordinates to match the [DMY] choice of time axis and proves the corresponding identities for kernels and measures. It preserves the field, the state, and the complex correlations at every order.

12.1 The test space and the Schwinger functional

The domain of test functions for this check is `ZeroDiagonalSchwartz`, defined as

$$\mathcal{S}_0\big((\mathbb{R}^4)^n\big) = \big\{ f \in \mathcal{S}\big((\mathbb{R}^4)^n\big) : \partial^\alpha f\big|_{\Delta_{ij}} = 0 \text{ for all } \alpha \text{ and } i \neq j \big\}, \tag{12.1}$$

where $\Delta_{ij} = \{x : x_i = x_j\}$. These Schwartz tests vanish to every order on the coincidence diagonals, but need not vanish in a neighborhood of them or have compact support. They include the compactly supported smooth tests away from coincidences used in the lattice-to-continuum convergence statement. The Schwinger functional used in this independent verification is

$$\mathcal{S}_n(f) = \int_{(\mathbb{R}^4)^n} K^{\xi^*}_{\Theta,n}(x)\, f(x)\, \mathrm{d}x, \qquad K^{\xi^*}_{\Theta,n}(x) := S^{\xi^*}_{\Theta,\dots,\Theta}(x). \tag{12.2}$$

The adapter proves that (8.1) is absolutely convergent and defines a continuous linear functional in the inherited Schwartz topology, and proves the E0′ bound controlling growth with the number of insertions. This is check A1.

Remark 12.1 (Distributional scope). Several distinct distributional statements enter the construction. Lattice convergence first determines correlations away from coincidences. The independent OS comparison integrates the selected stress-trace kernels against joint Schwartz tests flat to all orders on every coincidence diagonal; these tests cannot detect contact terms supported on those diagonals. The submission's reconstruction separately supplies tempered Wightman distributions on unrestricted joint Schwartz tests. Euclidean contact extensions are additional constructions. The diagonal-flat OS comparison does not establish the uniqueness of these extensions or their satisfaction of OS positivity on unrestricted Euclidean test spaces.

12.2 Independence and scope of the check

The target definitions come from the formulation in the repository `xiyin137/OSreconstruction` [DMY], also in the Palomar registry. Its OS axioms and the corresponding Wightman axioms were formulated independently and before the group learned of the OpenAI Yang–Mills construction. The adapter proves that this selected scalar hierarchy satisfies those target OS axioms and the associated growth condition, using the submission's analytic estimates and reconstruction theorems. This is a check against independently specified axioms; its analytic input comes from the submitted proof.

The extracted target declarations were checked against the pinned upstream definitions, and the adapters compiled in the submission's Lean 4.32.0 environment. The comparison concerns the centered-trace hierarchy, rather than a separate translation of every tensor field into the upstream formalization.

13 Other independent checks and their implications

13.1 The finite lattice law

The independent lattice definitions come from `mrdouglasny/lgt` [LGT], at the pinned revision identified there. The checked adapter `LGTWilsonBridge.lean` reindexes sites, links, and plaquettes and proves equality of the normalized finite Wilson laws. In the conventions of that repository,

$$\beta_{\text{submission}} = m\,\beta_{\text{lgt}}, \tag{13.1}$$

where m is the matrix representation dimension. With this conversion, the action, partition function, probability measure, and integrals of integrable observables agree. The proof also matches the gauge actions and plaquette orientations. The comparison was done on Lean 4.32.0. This is check A3. The comparison is restricted to the finite-volume law; the pinned upstream repository also contains strong-coupling and mass-gap material (`LGT/MassGap/StrongCoupling.lean`, `LGT/YMMeasure.lean`), which the bridge does not invoke.

13.2 The Ward-identity adapter

The Ward-identity adapter derives equation (9.1) from the charge and differentiation laws in the OpenAI submission. It covers every local field label and every component, and proves existence of admissible support for the bump function χ, integrability, and independence of this support. This checks the formulation and translation normalization of the submission's Ward identity. This is check A6.

13.3 Common scale, normalization, and a physical probe

One needs to ensure that different lattice cutoffs and box sizes use the same reference units and field normalizations. The adapter extracts the following identities for the same family, establishing a common normalization convention across volumes (check A5):

1. the reference RG spacing is $a(0) = 1$, fixing a common unit of length;
2. for every cutoff N and volume index p, the RG trajectory reaches $\beta_{N,p}(0) = \beta_{\text{ref}} > 0$; i.e. the RG flow starting at smaller lattice spacings (with coupling $\beta_{N,p}(N)$) results in the same reference coupling at lattice spacing 1;

3. kinetic normalization: the quadratic term in the expansion of the effective action is normalized to have identity matrix coefficients;

4. insertion normalization: for a degree-d label v, its normalized coefficient at the fixed reference level $r(d)$ equals v, independently of cutoff and volume, whenever $N \geq r(d)$.

A second check connects this calibration to a nonzero observable in the mixed theory. One can choose a fixed separated trace two-point test function f such that

$$\lim_{k\to\infty} \lim_{N\to\infty} \left|C_{N,p(q(k))}(\Theta, \Theta; f)\right|^{1/8} = \left|S^{\mathrm{mix}}(\Theta, \Theta; f)\right|^{1/8} > 0. \tag{13.2}$$

Here the lattice insertion uses the renormalized lattice representative of Θ, which is centered by subtracting the selected component's expectation. Once chosen, the test and centering remain fixed at every cutoff and volume.

13.4 The short-distance adapter

The short-distance adapter expresses the submission's trace asymptotic (9.6) for approaches to the origin in all directions. It identifies the same centered trace used in the scalar OS check and the same stress tensor tested by the translation Ward identity. This is check A4.

13.5 OPE statement alignment

The group checks that the submitted OPE theorem implies the finite asymptotic expansion with explicit geometric conditions and bounds uniform over the phases. The purpose is to check the meaning and scope of the formal statement; it does not compute the coefficients independently or compare their values with Yang–Mills perturbation theory. The analytic estimates are supplied by OpenAI's proof.

The review target `YMReview.OPE.CommonOPE` states two bounds directly, without invoking the source's `OperatorProductExpansion` predicate. For nearby fields of total degree w, it requires

$$|c_\alpha| \leq r^{d_\alpha - w - \varepsilon}, \qquad |R_{D,s}| \leq r^{D+1-w-\varepsilon}, \tag{13.3}$$

where $R_{D,s}$ is the difference between the actual correlation and its finite OPE sum through degree D, evaluated in state s. The coefficients are independent of spectators, phases, and truncation orders. The check still uses the OpenAI submission's definitions of fields, coefficients, and correlations. Additional checks on the consistency of the field decomposition were performed; the two review modules compile without warnings, and all five audited declarations use only the permitted standard axioms. A separate Comparator run for the OPE statement passed on November 5, 2026. This is check A8.

14 Verification records and reproduction

The mathematical statements and their source definitions are collected in the Lean statement glossary of Appendix A. This section records which additional checks were proved, which versions were checked, and where the build and Comparator results can be inspected.

14.1 Versions and repository layout

The verification material is in `yang-mills-group/YM` [YMRepo]. The October package is under `submissions/openai-20261028/`. Within that package, the submitted source is under `submission/`, the group's additional Lean code under `checks/`, the commented examples under `docs/examples/`, and the dated records under `evidence/`. Source and evidence links use this package layout; the assessment PDF is at the repository-level path `docs/assessment-1.pdf`.

The submitted files are preserved byte for byte. The provenance directory records their hashes, the imported review sources, and the upstream definition snapshots. The review source manifest identifies the included copies of the review files and their hashes, linking the recorded Comparator runs to the sources in this release. For the declaration-by-declaration alignment with the upstream axiom systems, the files `ImportedAFGauge.lean`, `ImportedAFSpecification.lean`, `ImportedCenterWeights.lean`, `ImportedDiagonalFlatBounds.lean`, `ImportedLGTData.lean`, `ImportedReflectionApproximation.lean`, `ImportedSchwartzCutoffs.lean`, `ImportedSchwartzDomain.lean`, and `ImportedSchwartzProduct.lean` re-export the pinned upstream definitions used by the adapters; their hashes are part of the provenance record.

Four counts appear in the records; they count different objects, not conflicting measurements:

Count	Object
9,477	Lean source files in the submission manifest.
9,470	project modules rebuilt in the clean build.
9,458	prebuilt modules whose bytes matched the original build inputs (the remainder being covered by import and kernel-declaration matches).
9,564	Lean files in the full source archive (submission, checks, evidence, and provenance).

14.2 Checks and their records

The identifiers M1–M2 are mechanical verification records, and A1–A8 are the group's additional checks; their mathematical meaning is explained in Sections 11.4–13.5. Table 1 links each check to its specification, proof, and verification record. Each A-entry also links a commented Lean example explaining its definitions, hypotheses, and scope. Where a check uses a specific numerical coefficient or another result from the literature, the corresponding worked example gives the supporting references and explains the conventions used in the comparison. These references support particular checks; they are not a general bibliography of the construction. The examples distinguish informal literature comparisons from the statements verified in Lean.

The commented examples explain the existing checks without changing their targets. The initial compilation record covers `OctoberChain.lean` and `LeadingTrace-20261028.lean`; the additional record covers the other five examples. All compiled without errors or warnings using only the permitted standard axioms. These ordinary Lean compilations are supplementary to the historical Comparator runs.

14.3 Reproducing the checks

From the release repository root, run

```
cd submissions/openai-20261028
python3 scripts/verify.py integrity
python3 scripts/verify.py all
```

The runner checks the imported files and OS-definition correspondence, builds the submission and review modules, audits axioms, builds the pinned tools, runs the controls, and checks all three Comparator endpoints. It writes a new dated record under `runs/`. The source and review use Lean 4.32.0 and Mathlib revision `81a5d257c8e410db227a6665ed08f64fea08e997`, and the pinned Comparator kernel uses Lean 4.34.0. A fresh source rebuild requires a checkout without existing submission build artifacts; a run that reuses those artifacts is recorded as incremental.

Acknowledgements

The source assessment document was prepared by Michael R. Douglas, Arthur Jaffe, Ron Nissim, and Fred Rajasekaran, with contributions from other group members and consultants, and approved by the entire group. The group used OpenAI's Astra to help translate the Lean code to informal mathematics and to develop the additional tests of the validation report, and Anthropic's Fable to give feedback on the document.

A Lean statement glossary

The construction theorem `YangMills.constructed` supplies the data used to prove the public statement `YangMills.JointExistence`, displayed below. The October 30 Comparator run checked the proof of this public statement.

```
def JointExistence : Prop := forall g : GaugeGroup, g.IsSimple ->
  exists C : WilsonFamily g,
    WilsonRegularity C /\\
    forall p : N -> N, StrictMono p ->
    exists q : N -> N, StrictMono q /\\
      exists T : Hierarchy C, T.Properties (Function.comp p q)
```

See `YangMills.constructed`; `YangMills.jointExistence`. The definitions of `GaugeGroup` and `IsSimple` can be found in `PublicTheory.lean`; `IsSimple` requires the Lie algebra to be noncommutative and to have no proper nonzero ideals. The mathematical notation follows the main text; S^s denotes a statement applying to either the mixture ($s = \mathrm{mix}$) or a component ($s = \xi \in X$). The selected Wightman theory has Hilbert space $\mathcal{H}_{\xi^*}$, domain $\mathcal{D}_{\xi^*}$, vacuum Ω_{ξ^*}, and fields $\Phi_v^{\xi^*}$. Below, each entry records the identifier, its mathematical content, the construction's instantiation, and the principal Lean certificates.

`YangMills.WilsonFamily`

Fix a closed matrix gauge group $G \subseteq U(m)$. The construction supplies these fields through `YangMills.wilsonFamily`.

`YangMills.WilsonFamily.period.` Positive integers L_p giving the physical side lengths. The construction uses $L_p = 4(M + p + 1)$, where $M \in \mathbb{N}$ is fixed sufficiently large to satisfy minimum-size and initial-volume bounds. Certificates: `YM.RG.Supply.Program.book`; `YM.RG.Supply.Wait.engine`.

`YangMills.WilsonFamily.denominator.` Positive integers D_N giving lattice spacing $a_N = D_N^{-1}$. The construction sets $D_N = \prod_{k<N} b_k$, where $b_N \geq 3$ is odd and $b_{N+1} \geq b_N$, and ensures $\log b_N / \log D_N \to 0$. The ratios b_k are chosen in order to satisfy an explicit inequality in terms of constants of the RG transformation.

`periodStrict`, `denominatorStrict`. Require (L_p) and (D_N) to be strictly increasing; clear from the definitions. Certificates: `YM.Schedule`; `YM.Schedule.denominator`; `YM.Schedule.den_strict`; `YM.Schedule.Gentle`; `YM.RG.Calendar.Fits`.

`YangMills.WilsonFamily.beta`. The coupling in $e^{-\beta_{N,p}E(U)}$. For each cutoff N and volume p, the bare coupling $\beta_{N,p}$ is chosen so that after N RG steps the resulting coupling equals a fixed reference value $\bar{\beta}$. Certificates: `YangMills.wilsonLaw`; `YM.Family.beta`.

`haar`, `probability`, `invariantHaar`. A measure h on G with $h(G) = 1$, invariant under left translation; the construction uses normalized Haar measure. Certificates: `YM.Gauge.probHaar`; `YM.Gauge.probHaar_isProbability`; `YM.Gauge.probHaar_isHaar`.

`YangMills.WilsonFamily.Token`. An abstract common set of labels for lattice observables at all spacings and volumes: pairs (d, c), where c is the coordinate vector, in a chosen basis of V_d, of a gauge-invariant polynomial in the curvature and its covariant derivatives at a point. The pair records the degree bound and the polynomial before passing to the quotient $\mathcal{F}_d$. Certificates: `YM.Register.Token`; `YM.Labels.Frame.coords`.

`YangMills.WilsonFamily.insertion`. Assigns a complex value $I_{N,p}((d, c), x; U)$ to a token, a site $x \in \Lambda_{N,p} = (\mathbb{Z}/(L_p D_N)\mathbb{Z})^4$, and a configuration U ($U_{y,\mu} \in G$ on each positively oriented edge). For each degree bound d, the construction chooses a finite offset set $S_d \subset \mathbb{Z}^4$ shared by all labels in V_d and independent of N, p, and lattice polynomials $P_{d,1}, \ldots, P_{d,k}$ representing the chosen basis $v_1, \ldots, v_k$ of V_d, each depending only on the matrix entries (and conjugates) of links with endpoints in S_d. The polynomial $P_{d,j}$ reproduces the continuum label v_j in the sense that, on link matrices obtained by parallel transport of polynomial test gauge potentials, its Taylor expansion in a agrees through order d with the scaled continuum polynomial v_j. The construction also adjusts the coefficient vector c: for sufficiently fine spacings, RG transport gives an invertible linear map from coefficients at a_N to coefficients at a fixed reference spacing $a_{r(d)}$, $r(d) \le N$ depending only on d; applying its inverse to c and then coordinate-dependent rescaling factors gives coefficients $\tilde{c}_j^{N,p}$. The code defines this map and proves its invertibility but does not give explicit formulas for the inverse's matrix entries. With $P_{d,j}(x; U)$ the polynomial translated to x,

$$I_{N,p}((d, c), x; U) = a_N^{-d} \sum_{j=1}^{k} \tilde{c}_j^{N,p} \, P_{d,j}(x; U), \tag{A.1}$$

depending only on links with endpoints in $x + S_d$. For the token representing v, this is the insertion $I_{N,p}(v, x; U)$ of (4.7), with spacing and normalization factors included in $c_j(N, p; v)$. Certificates: `YM.Family.seen`; `YM.Family.O`; `YM.Source.SectionData`; `YM.Source.TransferData.cylinder`.

`YangMills.WilsonRegularity`

This predicate imposes four conditions on a Wilson family; `YangMills.regular` proves them for the constructed family when G is simple.

`running`. There is a positive slope b^* such that the normalized couplings converge to b^* uniformly over all volume indices:

$$\exists b^* > 0 \; \forall \varepsilon > 0 \; \exists N_0 \in \mathbb{N} \; \forall N \ge N_0 \; \forall p \in \mathbb{N}, \qquad \left| \frac{\beta_{N,p}}{\log(a_N^{-1})} - b^* \right| < \varepsilon. \tag{A.2}$$

The construction uses $b^* = 2b_0 = 11C_{\mathfrak{g}}/(24\pi^2)$, with $C_{\mathfrak{g}}$ the adjoint Casimir constant. Certificates: `YM.Pert.Asymptotics.b0`; `YM.Gauge.casimir`.

`continuous.` For every N, p, v, x, the function $U \mapsto I_{N,p}(v, x; U)$ is continuous in the product topology on configurations. Certificate: `YM.Family.continuous0`.

`invariant.` For every sitewise gauge transformation $t : \Lambda_{N,p} \to G$, with $U^t_{y,\mu} = t(y)\, U_{y,\mu}\, t(y + e_\mu)^{-1}$,

$$I_{N,p}(v, x; U^t) = I_{N,p}(v, x; U), \tag{A.3}$$

obtained from gauge-invariant local sections. Certificate: `YM.Family.invariant`.

`localSupport.` Each token v has a finite set $S_v \subset \mathbb{Z}^4$, independent of N, p, x, such that the insertion depends only on links with both endpoints in $x + S_v$ (periodically wrapped). For a constructed token (d, c), S_v is the fixed neighborhood chosen for the $P_{d,j}$. Certificate: `YM.Family.local0`.

`YangMills.Hierarchy`

Further requirements for some of these fields are in `YangMills.Hierarchy.Properties`. The construction fills the fields through `YangMills.hierarchy`. Below, S^{mix} denotes the continuum Euclidean correlations obtained from the lattice limit, before selecting a phase.

`YangMills.Hierarchy.H.` The complex Hilbert space $\mathcal{H}_{\xi^*}$ of the selected phase: finite linear combinations of words of smeared fields, with inner products from the selected phase's reconstructed Wightman functions, quotiented by null vectors and completed. Certificates: `YM.Boundary.harvest`; `YM.Kernel.host`.

`YangMills.Hierarchy.labels.` The complex vector space V labeling the fields of the Wightman theory. Given a token $t = (d, c)$, associate $R_t(n, b_1, \dots, b_n; x_0, \dots, x_n) = S^{\mathrm{mix}}(t, x_0; b_1, x_1; \dots; b_n, x_n)$; then $V = \mathrm{span}_{\mathbb{C}}\{R_{(d,c)} : d \in \mathbb{N},\ c \in V_d\}$. Certificates: `YangMills.labelsOf`; `YM.Table.Union`.

`YangMills.Hierarchy.theory.` A bundled structure containing the field operators $\Phi_v^{\xi^*}$, their common domain $\mathcal{D}_{\xi^*}$ and vacuum Ω_{ξ^*}, the adjoint operation on labels, translations, an energy–momentum measure, and the Poincaré action. Certificates: `YangMills.view`; `YM.Boundary.Soil.starter`.

`YangMills.Hierarchy.Field.` A finite-dimensional complex normed vector space $\mathcal{F}_d$ for each degree bound d. The construction sets $\mathcal{F}_d = V_d/N_d$; writing $[c]$ for the equivalence class of a coefficient vector, the norm is the quotient norm from the supremum norm $\|c\| = \max_i |c_i|$ on the chosen basis of V_d. Certificates: `YM.Table.null`; `YM.Table.Field`.

`toLabel.` A complex-linear map $\mathcal{F}_d \to V$ sending $[c]$ to $R_{(d,c)}$; injective because N_d is precisely the kernel of $c \mapsto R_{(d,c)}$. Certificates: `YM.Table.spectra`; `YM.Table.pin`.

`project.` Sends a Token to a Field by the quotient map $(d, c) \mapsto (d, [c])$. Certificate: `YM.Table.classTokens`.

`YangMills.Hierarchy.Character.` The phase space X: a compact metric space with its Borel σ-algebra. Let $\mathcal{H}_{\mathrm{OS}}$ be the Hilbert space reconstructed from S^{mix}, with time semigroup $E(t) = e^{-tH}$; let $\mathcal{H}_0 = \ker H$ with projection P_0. For each word A with distinct points in a time slab $0 < t < s$, the construction supplies a bounded operator $B_{s,A}$: on vectors represented

by positive-time words, it shifts existing insertions forward by s and prepends the insertions of A. Define

$$M_A = P_0 B_{s,A}\big|_{\mathcal{H}_0}, \qquad \mathcal{A} = C^*\big(M_A : s > 0,\ A \text{ contained in } 0 < t < s\big), \tag{A.4}$$

where $C^*(\cdot)$ is the norm-closed unital algebra generated by these operators and their adjoints. The construction proves $\mathcal{A}$ commutative and takes its Gelfand spectrum as X, so $\mathcal{A} \cong C(X)$; points are characters $\xi : \mathcal{A} \to \mathbb{C}$, determining phase correlations by $S^\xi(A) = \xi(M_A)$ after translating the word A into a positive-time slab, independently of the translation and slab width. Certificates: `YM.OS.Semigroup.ground`; `YM.OS.Spread`; `YM.OS.Spread.algebra`; `YM.OS.Ground.at_any`; `YM.OS.Spread.Resolution.receives`; `YM.OS.Ground.Theta`; `YM.OS.Ground.resolution`.

`phaseMeasure`**.** A probability measure ν on X representing the vacuum state of the mixed theory via Gelfand and Riesz–Markov–Kakutani:

$$\int_X \xi(a)\,\mathrm{d}\nu(\xi) = \langle \Omega_0, a\,\Omega_0 \rangle \qquad (a \in \mathcal{A}), \tag{A.5}$$

where Ω_0 is the vacuum vector of the ground-state subspace $\mathcal{H}_0$.

`selected`**.** A distinguished phase $\xi^* \in X$. Writing $\tau = \sum_{\nu=0}^{3} T_{\nu\nu}$ and $\tilde\tau = \tau - e\mathbf{1}$ ($e \in \mathbb{R}$ arbitrary), let G_ξ be the ground-state space of phase ξ. The construction takes a sequence $r_n > 0$, $r_n \to 0$, and shows there are strictly increasing indices n_k such that the set

$$E = \Big\{ \xi \in X : \dim G_\xi = 1,\ \text{and } \forall e \in \mathbb{R},\ \lim_{k\to\infty} r_{n_k}^{8}\, S^\xi\big(\tilde\tau(0)\,\tilde\tau(r_{n_k} e_0)\big) = 0 \Big\} \tag{A.6}$$

satisfies $\nu(E) = 1$, where $e_0 = (1,0,0,0)$. The same subsequence works for every $\xi \in E$ and every $e \in \mathbb{R}$; the construction chooses $\xi^* \in E$. Certificates: `YM.OS.Cluster.lifts`; `YM.Trace.submits`; `YM.Trace.Echo.retreats`; `YM.Trace.echo_tab`.

`correlation`**.** Euclidean n-point functions of fields (d_i, v_i) at positions $x_i \in \mathbb{R}^4$, with no phase specified (S^{mix}) or a character ξ (S^ξ). Certificates: `YM.OPE.spectator`; `YM.Table.useAt`; `YM.OS.Spread.Resolution.receives`.

`continuation`**.** A function taking n field labels, n complex times z_i, and n spatial positions $x_i \in \mathbb{R}^3$ to a complex number, written $F_n^{\xi^*}(v; z, x)$, using the analytic continuation of S^{ξ^*} so that $S^{\xi^*}(v;\tau,x) = F_n^{\xi^*}(v; i\tau, x)$ for $\tau_1 < \cdots < \tau_n$. Certificates: `YM.Boundary.Waves.Region`; `YangMills.hierarchy`.

`Channel`**,** `actor`**,** `coefficient`**.** The index set for degree-d fields on the right-hand side of the OPE: decomposing $\mathcal{F}_d = \mathcal{F}_{d-1} \oplus E_d$, the channels index a conjugation-invariant basis of E_d (with $n_d = \dim E_d$); $\alpha = (d, i)$ has $d_\alpha = d$, and `actor`(d, i) is the i-th basis element of $E_d \subseteq \mathcal{F}_d$. `coefficient` gives the complex OPE coefficient $c_\alpha(A; x, o)$ for input fields $A_k = (d_k, v_k)$ at positions x_k about an expansion point o: the construction forms a field in $\mathcal{F}_d$ from finitely many correlations of the inputs with additional test insertions, then averages over $O(4)$, symmetrizes under conjugation, removes lower-degree components, and takes coordinates in the basis $O_{(d,i)}$. Certificates: `YM.OPE.Orbits.Reader.cutTest`; `YM.OPE.Orbits.Reader.stratum`; `YM.OPE.Orbits.Reader.channels`; `YM.OPE.Orbits.Reader.basis`; `YM.OPE.Orbits.Reader.actor`; `YM.OPE.Prescription.ope`; `YM.OPE.shadow`; `YM.OPE.Orbits.averageV`; `YM.OPE.Orbits.checkReal`; `YM.OPE.Orbits.Reader.fresh`; `YM.OPE.Orbits.Reader.coef`.

`centeredTrace`. The element of $\mathcal{F}_4$ given by the trace of the Euclidean stress tensor minus its one-point expectation in the selected phase:

$$\Theta = \sum_{\mu=0}^{3} T_{\mu\mu} - S^{\xi^*}\Big(\sum_{\mu=0}^{3} T_{\mu\mu}\Big)\mathbf{1}, \tag{A.7}$$

where $\mathbf{1} \in \mathcal{F}_4$ is the identity field, so $S^{\xi^*}(\Theta) = 0$. Certificates: `YM.TraceExtension.center`; `YM.TraceExtension.stressTrace`; `YM.TraceExtension.theta`.

`YangMills.Hierarchy.Properties`

These structure fields concern the same hierarchy T and the selected volume indices $p(q(k))$ as in Theorem 6.4. The construction gives the proofs together in `YangMills.constructed`; the fields `reconstruction` and `nonfree` depend on further structures explained below.

`wilson`. Requires the full UV limit of the smeared Wilson correlations at every fixed volume, followed by the infinite-volume limit along $p(q(k))$:

$$\lim_{k\to\infty} \lim_{N\to\infty} C_{N,p(q(k))}(v; f) = \int_{(\mathbb{R}^4)^n} f(x)\, S^{\mathrm{mix}}(v; x)\, \mathrm{d}x, \tag{A.8}$$

where v_i is a lattice token representing a classical label and $\bar{v}_i = \texttt{T.project}(v_i)$ (the main text suppresses the bar). The limit holds for every number of insertions, token list v, and smooth test $f \in C_c^\infty$ away from coincidences; at fixed k, the support of f must lie strictly inside the finite-volume box (which holds for all sufficiently large k). In the notation of (4.9), Lean's smear is

$$C_{N,p}(v; f) = a_N^{4n} \sum_{x_1,\dots,x_n \in \Lambda_{N,p}} f(a_N x_1, \dots, a_N x_n) \int \prod_{i=1}^{n} I_{N,p}(v_i, x_i; U)\, \mathrm{d}\mu_{N,p}(U), \tag{A.9}$$

with $\Lambda_{N,p} = (\mathbb{Z}/(L_p D_N)\mathbb{Z})^4$, $a_N = D_N^{-1}$, the sum over ordered n-tuples, and $\mu_{N,p}$ the normalized Wilson measure. The construction obtains S^{mix} by these ordered limits and shows that tokens equivalent under `project` have equal limiting correlations. Certificates: `YangMills.WilsonFamily.smear`; `YangMills.WilsonFamily.finiteMoment`; `YangMills.WilsonFamily.place`; `YangMills.wilsonLaw`; `YangMills.WilsonFamily.Converges`; `YangMills.wilsonLimit`.

`ope`. Requires bounds on the OPE coefficients and the truncation error. For input fields $A_i \in \mathcal{F}_{d_i}$ at positions x_i, set $w = \sum_i d_i$; let B_ℓ be spectator fields at positions y_ℓ and $B = \prod_\ell B_\ell(y_\ell)$. For fixed constants $M \geq 0$ and $\eta, c > 0$, the allowed configurations at scale $r > 0$ satisfy

$$\begin{aligned} \|x_i - o\| &\leq Mr \quad \text{for every } i, \\ \|x_i - x_j\| &\geq \eta r \quad \text{for } i \neq j, \\ \|y_\ell - o\| &\geq c \quad \text{for every } \ell, \\ \|y_\ell - y_m\| &\geq c \quad \text{for } \ell \neq m, \end{aligned} \tag{A.10}$$

so the inputs form a cluster near the expansion point o while spectators stay separated; the theorem applies uniformly over all configurations satisfying (A.10). With $O_\alpha := \texttt{T.actor}(d, i) \in \mathcal{F}_d$ and $c_\alpha(A; x, o) := \texttt{T.coefficient}(A, d, i, x, o)$, define, for $s = \mathrm{mix}$ or $s = \xi \in X$,

$$R_{D,s} = S^s\Big(\prod_i A_i(x_i)\, B\Big) - \sum_{d_\alpha \leq D} c_\alpha(A; x, o)\, S^s\big(O_\alpha(o)\, B\big). \tag{A.11}$$

The bounds are

$$|c_\alpha(A;x,o)| \le r^{d_\alpha - w - \varepsilon}, \qquad |R_{D,s}| \le r^{D+1-w-\varepsilon}, \tag{A.12}$$

holding for all sufficiently small $r > 0$ for every $\varepsilon > 0$; the coefficient bound requires only the first two conditions of (A.10), while the remainder bound requires all four and holds uniformly over all states $s \in \{\mathrm{mix}\} \sqcup X$. The coefficients are independent of spectators, state, and truncation degree. Setting $o = 0$ and input positions rx_i recovers (9.12). Certificates: `YangMills.Hierarchy.Within`; `YangMills.Hierarchy.Clear`; `YangMills.Hierarchy.OperatorProductExpansion`; `YM.OPE.Orbits.coefficient_separated`; `YM.OPE.Orbits.quantitative_separated`; `YangMills.publicOPE`.

`nonfree`. Four conditions on the degree-four field Θ = `T.centeredTrace`. Abbreviate, with $e_0 = (1,0,0,0)$ and ξ^* = `T.selected`,

$$G_\Theta(r) = S^{\xi^*}\big(\Theta(-\tfrac{r}{2}e_0)\,\Theta(\tfrac{r}{2}e_0)\big). \tag{A.13}$$

The construction proves the conditions together in `YangMills.publicTrace`. Certificates: `YangMills.Hierarchy.TraceTest`; `YangMills.Hierarchy.traceTwo`.

`TraceTest.centered`. The vacuum expectation vanishes for every complex Schwartz test: $\langle \Omega_{\xi^*}, \Phi_\Theta^{\xi^*}(f)\Omega_{\xi^*}\rangle = 0$ for all $f \in \mathcal{S}(\mathbb{R}^4;\mathbb{C})$, by analytically continuing the vanishing Euclidean one-point function. Certificates: `YM.TraceExtension.theta`; `YM.TraceExtension.vev_boundary`; `YM.TraceExtension.pinned_centered`.

`TraceTest.amplitudePositive`. The short-distance coefficient is strictly positive:

$$A_G := \frac{3\,\dim_{\mathbb{R}}\mathfrak{g}}{\pi^4} > 0, \tag{A.14}$$

made explicitly available to downstream proofs. Certificates: `YM.Trace.Verse.award`; `YM.Trace.Over.force_b`; `YM.Trace.Verse.award_pos`; `YangMills.publicTrace`.

`TraceTest.shortDistance`. With $c(r) = r^8 \log^2(1/r)$,

$$\lim_{r \downarrow 0} c(r)\, G_\Theta(r) = A_G. \tag{A.15}$$

The construction first shows the limit at approximated positions using lattice grid points; the ultraviolet comparison, logarithmic normalization, and free two-point calculation give $c(r_n)\,G_\Theta(r_n\delta_n(s)) \to A_G/s^8$ with $\delta_n(s) \to s$, where $\delta_n(s)$ corrects for approximating physical points separated by $r_n s$ by nearby lattice points (s fixed near 1). The positive spectral representation makes G_Θ real-valued and nonincreasing; choosing $s_- < 1 < s_+$ gives eventually

$$G_\Theta(r_n\delta_n(s_+)) \le G_\Theta(r_n) \le G_\Theta(r_n\delta_n(s_-)),$$

and multiplying by $c(r_n)$, taking $n \to \infty$, and letting $s_\pm \to 1$ yields (A.15). Certificates: `YM.Trace.Verse.tagline`; `YM.Trace.Verse.monotone_axis`; `YM.Trace.Verse.border`; `YM.Trace.Verse.outside`; `YangMills.publicTrace`.

`TraceTest.excludes`. G_Θ disagrees near zero with every finite function of the form

$$W(r) = \sum_{a=1}^{N} c_a \prod_{b=1}^{n_a} K_{ab}(r), \qquad K_{ab}(r) = \int_{\mathbb{R}^3} \frac{E_{ab}(p)^{q_{ab}}\, h_{ab}(p)^2}{2E_{ab}(p)}\, e^{-rE_{ab}(p)}\, \mathrm{d}^3p, \tag{A.16}$$

with $E_{ab}(p) = \sqrt{m_{ab}^2 + \|p\|^2}$, $N, n_a \in \mathbb{N}_0$, $c_a \in \mathbb{C}$, $m_{ab} \geq 0$, $q_{ab} \in \mathbb{N}_0$, and $h_{ab} : \mathbb{R}^3 \to \mathbb{R}$ continuous and homogeneous of degree $j_{ab} \in \mathbb{N}_0$. Precisely: $\forall W$ of this form, $\forall \delta > 0$, $\exists r \in (0, \delta) : G_\Theta(r) \neq W(r)$. The construction proves that $\operatorname{Re} W$ admits an expansion around 0, $\operatorname{Re} W(r) = \sum_{\nu=1}^{M} a_\nu r^{\alpha_\nu} (\log r)^{k_\nu} + O(1)$ with $\alpha_\nu \in \mathbb{Z}$, $k_\nu \in \mathbb{N}_0$, and that any finite limit of $r^8 \log(1/r) \operatorname{Re} W(r)$ must be zero; by (A.15), agreement near zero would instead give the strictly positive limit A_G (see Remark 9.15 for the weight correction), a contradiction. Certificates: `YM.TraceExtension.Wick.real_profile`; `YM.TraceExtension.Wick.AtOrder`; `YM.TraceExtension.Wick.no_powerlog`; `YM.TraceExtension.Wick.euclidean_excludes`; `YangMills.Hierarchy.Mode`; `YangMills.Hierarchy.Track`; `YangMills.Hierarchy.wickProfile`.

`reconstruction`. Fifteen conditions on the hierarchy T, proved together in `YangMills.publicReconstruction`; see `YangMills.Hierarchy.Reconstructed`.

`axioms`. The Wightman axioms W0–W3 for the smeared fields $\Phi_v^{\xi^*}(f) : \mathcal{D}_{\xi^*} \to \mathcal{D}_{\xi^*}$, complex-linear in $v \in V$ and $f \in \mathcal{S}(\mathbb{R}^4; \mathbb{C})$:

W0. *Translation spectrum and uniqueness of the vacuum.* Spacetime translations act through a strongly continuous unitary representation $U(a)$ with projection-valued energy–momentum measure E, $\langle \psi, U(a)\psi \rangle = \int_{\mathbb{R}^4} e^{ip\cdot a} \, \mathrm{d}\mu_\psi(p)$, $\mu_\psi(B) = \langle \psi, E(B)\psi \rangle$; E is supported in the closed forward light cone $V_+ = \{p : p^{\mathrm{time}} \geq \|p^{\mathrm{space}}\|\}$; and every translation-invariant vector is a complex multiple of Ω_{ξ^*}.

W1. *Common domain, regularity, adjoints, cyclicity.* $\mathcal{D}_{\xi^*}$ is dense, contains the normalized vacuum, and is preserved by every smeared field; for fixed v and ψ, $f \mapsto \Phi_v^{\xi^*}(f)\psi$ is continuous from the Schwartz to the Hilbert norm topology; with the conjugate-linear adjoint map $v \mapsto v^*$, $\langle \Phi_v^{\xi^*}(f)\psi, \chi \rangle = \langle \psi, \Phi_{v^*}^{\xi^*}(\bar{f})\chi \rangle$; and $\mathcal{D}_{\xi^*}$ is exactly the span of field words applied to the vacuum, the empty word being Ω_{ξ^*}.

W2. *Poincaré covariance.* The proper orthochronous Poincaré group acts through a strongly continuous unitary representation extending the translations, and on labels through complex-linear bijections $\rho(g)$ with $\rho(1) = \mathrm{id}$, $\rho(gh) = \rho(g)\rho(h)$; it preserves $\mathcal{D}_{\xi^*}$, fixes the vacuum, and $U(g)\Phi_v^{\xi^*}(f)\psi = \Phi_{\rho(g)v}^{\xi^*}(f_g)\, U(g)\psi$ with $f_g(x) = f(\Lambda^{-1}(x - a))$ for $gx = \Lambda x + a$.

W3. *Local commutativity.* If every point of $\operatorname{supp} f$ is spacelike separated from every point of $\operatorname{supp} h$, then $\Phi_v^{\xi^*}(f)\Phi_w^{\xi^*}(h)\psi = \Phi_w^{\xi^*}(h)\Phi_v^{\xi^*}(f)\psi$.

The construction proves these in the selected phase: Wightman boundary values define the inner product on finite words; the isometric map from $\mathcal{H}_{\xi^*}$ into the selected phase's OS Hilbert space intertwines translations and sends the Wightman vacuum to the OS vacuum; nonnegativity of energy and Lorentz covariance give the forward-cone spectrum; and translation-invariant vectors on the Wightman side map into the vacuum subspace on the OS side, so by injectivity they lie in the Wightman vacuum subspace. Certificates: `YangMills.FieldSystem`; `YangMills.WightmanTheory.Axioms`; `YM.Kernel.host`; `YM.Boundary.Soil.starter`; `YM.Boundary.Soil.sharp`; `YM.Boundary.Soil.usable`; `YM.Boundary.Soil.regional`; `YM.Boundary.Soil.mobilize`; `YM.Boundary.Soil.cultivates`; `YangMills.view`; `YM.Boundary.Outfield.daylight`; `YM.Boundary.harvest`; `YM.Boundary.Action.spectrum`; `YM.Boundary.Action.only`; `YangMills.view_axioms`.

`faithful`. Every nonzero label v has a test f with $\Phi_v^{\xi^*}(f) \neq 0$: vanishing of all smeared Wightman correlations containing a label forces vanishing of all Euclidean correlations con-

taining a representative of that label, with arbitrary spectators at pairwise distinct points, and these are exactly the labels quotiented out. Certificates: `YangMills.FieldSystem.Faithful`; `YM.Work.C26.answer`; `YM.Boundary.Soil.sees`.

`injective.` For each d, distinct fields in $\mathcal{F}_d$ have distinct images under `toLabel`. Certificates: `YM.Table.faithfully`; `YM.Table.pin_inj`.

`acts.` Some field in some $\mathcal{F}_d$ has a nonzero smeared operator (via the identity field in $\mathcal{F}_0$), excluding vacuous faithfulness. Certificate: `YangMills.publicReconstruction`.

`covers.` Every label in V is $\texttt{T.toLabel}(d, v)$ for some d and $v \in \mathcal{F}_d$. Certificate: `YM.DiffPlan.covered`.

`represents.` Every pair (d, v), $v \in \mathcal{F}_d$, is the image of a lattice token under `project` (a quotient map). Certificates: `YM.Table.representative`; `YangMills.publicReconstruction`.

`normalized.` ν has total mass one: $\|\Omega_0\|^2 = 1$. Certificate: `YM.OS.Ground.resolution`.

`fullSupport.` Every nonempty open subset of X has positive ν-measure, since no nonzero element of the ground-state algebra annihilates Ω_0. Certificates: `YM.OS.Ground.attains`; `YM.OS.Ground.supports`.

`permutation.` Correlations (mixed and component) are invariant under simultaneously permuting fields and insertion points, inherited from lattice integrals and passed to the continuum limit and each phase. Certificates: `YM.Table.class_perm`; `YangMills.publicReconstruction`.

`smooth.` Correlations (mixed and component) are C^∞ away from coincidences. Certificates: `YM.Table.class_smooth`; `YangMills.publicReconstruction`.

`characterContinuity.` For fixed fields at distinct points, correlations are continuous in $\xi \in X$, via $S^\xi(A) = \xi(M_A)$ and continuity of character evaluation. Certificates: `YM.OS.Spread.Resolution.Out`; `YangMills.publicReconstruction`.

`mixture.` Averaging phases recovers mixed correlations: $\int_X S^\xi(v; x)\,\mathrm{d}\nu(\xi) = S^{\mathrm{mix}}(v; x)$, via $\int_X \xi(M_A)\,\mathrm{d}\nu(\xi) = \langle \Omega_0, M_A \Omega_0 \rangle = \langle \Omega_0, P_0 B_{s,A} \Omega_0 \rangle = \langle \Omega_0, B_{s,A} \Omega_0 \rangle = S^{\mathrm{mix}}(A)$. Certificates: `YM.OS.Spread.Resolution.barycenter`; `YangMills.publicReconstruction`.

`analytic.` For each fixed list of fields and spatial positions, $F_n^{\xi^*}(v; z, x)$ is holomorphic in z on $\operatorname{Im} z_1 < \cdots < \operatorname{Im} z_n$. Certificates: `YM.Work.C03.defined_exists`; `YM.Work.C03.defined_region`.

`root.` At purely imaginary times the continuation equals the selected-phase Euclidean correlation: for $\tau_1 < \cdots < \tau_n$, $F_n^{\xi^*}(v; i\tau, x) = S_n^{\xi^*}(v; (\tau_i, x_i)_{i=1}^n)$. Certificates: `YM.Work.C02.answer`; `YangMills.publicReconstruction`.

`boundary.` Approaching real times from any fixed strictly ordered imaginary direction recovers the smeared Wightman vacuum expectation: for $h_1 < \cdots < h_n$, Schwartz tests f_i, and quotient fields $v_i \in \mathcal{F}_{d_i}$ with images $\bar{v}_i = \texttt{T.toLabel}(d_i, v_i)$,

$$\lim_{\varepsilon \downarrow 0} \int F_n^{\xi^*}(v; t + i\varepsilon h, x) \prod_{i=1}^n f_i(t_i, x_i)\,\mathrm{d}^{4n}(t, x) = \left\langle \Omega_{\xi^*}, \Phi_{\bar{v}_1}^{\xi^*}(f_1) \cdots \Phi_{\bar{v}_n}^{\xi^*}(f_n) \Omega_{\xi^*} \right\rangle = W_n^{\xi^*}(\bar{v}_1, f_1; \ldots; \bar{v}_n, f_n). \quad \text{(A.17)}$$

Certificates: `YM.Boundary.Waves.Opens`; `YM.Boundary.Soil.relishes`; `YangMills.publicMoment`; `YangMills.publicReconstruction`.

`YangMills.PublicWilson` and `YangMills.ConcreteSU2` (submission interfaces)

Two files of the submission's public layer complete the picture.

`YangMills.PublicWilson`. The four-dimensional periodic Wilson lattice and its ordered continuum limit. `Site` is the discrete torus $(\mathbb{Z}/\ell\mathbb{Z})^4$; `Configuration` assigns group elements to positively oriented links; `plaquette` is the ordered plaquette product. The energy `wilsonEnergy` sums $1 - \operatorname{Re}\operatorname{Tr}(U_\square)/(g.n + 1)$ over sites and directions, so the trace normalization by the matrix dimension is built into the law; `wilsonLaw` is the product Haar prior with the normalized density $e^{-\beta E}$. The structure `WilsonFamily` packages `period`, `denominator` (mesh $= 1/$`denominator`), `beta`, Haar data, `Token`, and `insertion`; `finiteMoment` is the lattice correlation (4.9) before smearing, and `separated` is the open configuration domain.

`YangMills.ConcreteSU2`. The defining complex two-dimensional representation `su2`: `GaugeGroup`, with `su2Simple`: `su2.IsSimple` (the Mills definition) and `su2Theory`, the concrete specialization of `jointExistence` to $SU(2)$: there is a Wilson family with the regularity and hierarchy properties of the main theorem.

B Explicit assumptions and certificates

Every analytic input used in the main text is listed here with its certificate in the formalization or the validation document (denoted VD). None is used beyond its stated content, and none is re-proved in this paper. The register contains sixteen external entries. (RG4) is derived from (RG2)+(RG3) in Lemma 5.14 and is listed for completeness only; (PH4) is derived in Proposition 7.2 from (AN1) E2 + E4 and is likewise listed as a proved spectral conclusion rather than an input. Conditions (5.2) and the bootstrap of Lemma 5.11 are sufficient but not necessary (Lemma 5.9).

Assumption	Content	Certificates
(RG1)	Gauge fixing and blocking kernel preserve the finite lattice symmetry group.	VD §4.3 Rem. 4.2, §6.4; `YM.RG.Program`.
(RG2)	One-step estimates: one fixed contraction rate $\lambda \in (0,1)$, geometric decay of irrelevant terms, one-loop recursion for the marginal term, lower bound $\beta_m \geq \beta_* > 0$ with the dominance condition $2b_0 \log 3 > K/\beta_*$ of (5.2); uniform in volume. The source-level form is `Large.slow/Fits.driftSmall` (Lemma 5.9); the dominance condition is sufficient but not necessary.	`YM.RG.Supply.Program.book`, `YM.RG.Calendar.Fits`, `YM.Pert.Asymptotics.b0`, `YM.Price.Intent.Rolls.exists_rolls`.
(RG3)	Reference-level insertion readout is stable; RG transport invertible with controlled inverse.	`CommonScale.reference_insertion`, `TransferData.cylinder`.
(RG4*)	Kernel-level difference estimate (5.3): derived from (RG2)+(RG3) as in Remark 5.15.	Same as (RG4); `YM.Family.continuous0`, `YM.Family.local0`.
(RG5)	Uniform smoothness estimates for the lattice kernels on compact separated sets.	VD §4.3, §5.2 Rem. 5.2.
(UV1)	Local boundedness and equicontinuity of all derivatives, uniform in volume.	VD §5.2 Rem. 5.2; reflection positivity, chessboard, transfer operator, Cauchy.
(PH1)	Time-translation conjugation structure of word operators.	`YM.OS.Spread`, `YM.OS.Semigroup.ground`.
(PH2)	The ground-state algebra $\mathcal{A}$ is commutative.	`YM.OS.Spread.algebra`.
(PH3)	Ω_0 is separating for $\mathcal{A}$ (faithfulness).	`YM.OS.Ground.attains`, `.supports`.
(AN1)	OS axioms E0/E0′/E1–E4, normalization, and analytic continuation for the selected phase.	Appendix C; `YMOSBridge.verifiedOS`, `complete_OS_on_this_hierarchy`.
(Y1)	The stress-tensor candidate space is one-dimensional, with normalization fixed by the Ward identity.	VD §7.1; `WardFormulationAudit.lean`.
(Y2)	Flux (shell) law and calculus law for the abstract field derivative.	`phase.flux`, `phase.calculus`; `TranslationWard-20261028.lean`.
(Y3)	Reduction of the trace two-point function to the free-field residue times the squared anomaly coefficient.	`YM.Trace.Over.Scribe`, `YM.Trace.Verse.tagline`.
(Y4)	One-loop value $b_0 = 11C_{\mathfrak{g}}/(48\pi^2)$ from the loop sum.	`YM.Pert.Asymptotics.b0`, `YM.Gauge.casimir`.
(B1)	Field-content dictionary: $N_d = \ker \psi_d$, rank comparison and phase transport (Proposition 2.12).	`OperatorCompleteness.local_field_content`, `phase_field_content`, `F.phase.transport.comparator`, `F.phase.phase.retains`.
(C1)	Finite-law comparison with the independent lattice formalization (Proposition 4.16). Record only: it does not enter any subsequent derivation.	`LGTWilsonBridge.lean`, `same_plaquette`, `coupling_conversion`, `normalized_law_on_each_lattice`, `finite_moments_are_the_family_moments`.

Derived, not an input: (RG4) — $|C_{N+1,p} - C_{N,p}| \leq A(v,f)\lambda^N$ (Lemma 5.14); its kernel-level form (RG4*) by Remark 5.15. *Proved spectral conclusion, not an input:* (PH4) — $\mathcal{H}_0^{\xi^*} = \mathbb{C}\,\Omega_{\xi^*}$ (Proposition 7.2; certificates `YM.OS.Cluster.lifts`, `YM.Trace.submits`, `selected`, (A.6)). *The number seventeen in the repository records counts the headline results; it is unrelated to the sixteen external inputs above.*

C Comparison with the [DMY] OS axioms

The comparison of Section 12 is made against the independently developed OS formulation of [DMY] (repository `xiyin137/OSreconstruction`, also in the Palomar registry). The target definitions were extracted from a pinned upstream snapshot, checked against it declaration by declaration, and the adapters were compiled in Lean 4.32.0. The comparison concerns the centered-trace hierarchy in the selected phase ξ^* only; it is not a translation of every tensor field into the upstream formalization. The adapter changes coordinates to match the [DMY] choice of time axis (`configEquiv`); it preserves the field, the state, and the complex correlations at every order.

Bridge objects

Upstream object	Construction counterpart
`SpacetimeDim 3`	Euclidean spacetime $\mathbb{R}^4$ (dimension index $3+1$).
`NPointDomain 3 n`	Configuration domain of n-tuples in $\mathbb{R}^4$.
`scalarKernel F n`	The selected phase's n-point function $K^{\xi^*}_{\Theta,n} = S^{\xi^*}_{\Theta,\dots,\Theta}$ of the centered trace Θ (Section 9.2); reality is supplied by the source reality theorem.
`targetKernel F n`	Real kernel on `NPointDomain 3 n`, set to zero on coincidence diagonals *by convention* (a measurability device; not a claim that the physical distribution has no contact terms).
`schwinger F n`	The Schwinger functional $\mathcal{S}_n(f) = \int K^{\xi^*}_{\Theta,n}(x) f(x)\,\mathrm{d}x$ of Section 12.1.
`configEquiv`	Coordinate change matching the [DMY] time axis.

Axiom-by-axiom correspondence

The test space throughout is $\mathcal{S}_0((\mathbb{R}^4)^n)$ (`ZeroDiagonalSchwartz 3 n`), the joint Schwartz tests flat to all orders on every coincidence diagonal Δ_{ij}; it contains the compactly supported smooth tests away from coincidences used in the lattice-to-continuum convergence.

[DMY] item	Statement in upstream form	Construction counterpart and certificate
E0 (temperedness)	$\mathcal{S}_n$ is continuous on $\mathcal{S}_0$ (Schwartz topology); integrability is proved before the functional is identified with the integral (Lean's integral is totalized).	Kernels are smooth away from coincidences with a finite inverse-power bound in the minimum separation: `separated_bound` (`YMOSBridge.targetKernel_bound`); `pairing_is_integrable` (`schwinger_integrable`); `E0_continuity` (`E0_tempered`).
E0′ (arity growth)	$\exists$ slope s, base β, factorial exponent A *chosen before* n and f: $\|\mathcal{S}_n(f)\| \le \beta^n (n!)^A \, \|f\|_{(ns,ns)}$ (Schwartz seminorm at index (ns, ns)).	`YM.Work.B11.phase_cap` supplies a common bound $K^n (n/r)^{An}$ at separation r; `ArityGrowth.lean` transports it through the diagonal-flat integral estimate, absorbing n^{An} into an exponential times $(n!)^A$: `E0_prime_uniform_arity` (`YMOSBridge.E0_prime_growth`).
E1 (Euclidean invariance)	Rotations (`os_7` ← `Trace.rotation`, full rotation) and translation covariance (`os_8` ← `Trace.translated`).	Section 8.4 (gauge fixing preserving lattice symmetries, divergence-free stress tensor, no spontaneous breaking); packaged in `YMOSBridge.verifiedOS`.
E2 (reflection positivity)	Re OSInnerProduct$(B, B) \ge 0$ for every Borchers sequence B supported in the ordered positive-time region.	Section 7.1: the finite compact-test reflection form is identified entry by entry with the source positive form; support-preserving cutoffs and E0 continuity extend it to the full domain: `E2_full_target_form` (`YMOSBridge.E2_reflection_positive`).
E3 (permutation)	Invariance under simultaneous permutation of fields and insertions (`os_2` ← `Moments.Perm`).	`Properties.permutation` (inherited from lattice integrals, passed to the limit and each phase); packaged in `verifiedOS`.
E4 (clustering)	Separated groups factorize: for $f \in \mathcal{S}_0^n$, $g \in \mathcal{S}_0^m$, $\|\mathcal{S}_{n+m}(f\, g_a) - \mathcal{S}_n(f)\mathcal{S}_m(g)\| < \varepsilon$ for $\|a\|$ large.	The selected phase's vacuum factorization (Proposition 7.15), extended to joint Schwartz pairings; `CompleteOS.E4_cluster` supplies the precise spatial-translation clause: `E4_separated_groups` (`YMOSBridge.E4_cluster_norm`).
Normalization	Probability/normalization (`os_1` ← `Moments.probability`); `verifiedOS` also packages the required linearity.	$\nu(X) = 1$, $\|\Omega_0\|^2 = 1$ (`phaseMeasure`, `Properties.normalized`).
Analyticity	Analyticity of the Schwinger functions (`os_3` ← `Moments.analytic`).	Section 8: analytic continuation $F_n^{\xi^*}(v; z, x)$, $\operatorname{Im} z_1 < \cdots < \operatorname{Im} z_n$ (`Properties.analytic`, `root`).

Complete record	`verifiedOS` packages E1–E4, normalization, and linearity into `OsterwalderSchraderAxioms 3` with `S` *equal* to the Schwinger pairing above (no second choice of theory); `complete_OS_on_this_hierarchy` appends `OSArityLinearGrowthCondition 3` (E0′), compiled against the imported upstream definition, with the source-correspondence audit recording the snapshot match.

Scope of the comparison

1. Diagonal-flat test functions annihilate distributions supported on coincidence diagonals. The comparison therefore does not establish uniqueness of Euclidean contact extensions, nor OS positivity of such an extension on unrestricted Euclidean test spaces; these are different questions from the axioms proved here.

2. Neither clustering statement specifies an exponential rate, proves a mass gap, or makes the infrared extraction unique.

3. The file `OSAxioms-20261028.lean` verifies this instance of the imported OS definition; it does not prove or invoke the entire Douglas–Mei–Yin reconstruction theorem. The submission's own Wightman reconstruction is retained elsewhere in the combined target.

Proof-chain details: clustering and the packaged record

The E4 clustering chain runs through six review modules; we record the role of each:

- `ClusterKernel.lean`: geometric preliminaries—a large signed coordinate can be produced by a coordinate swap (`exists_signed_large_coordinate`), the configuration coordinate change preserves norms, and separated configurations stay separated under translation-joined insertion of a far group (`joinConfig_injective`, $\|x\| + \|y\| < \|a\|$).

- `ClusterCore.lean`: the dynamical core—for a chamber V with unique heat, ground-state projection of heat-evolved vectors, the mixed heat limit, and clustering of arbitrary registered words along the chamber axis; specialized to every character phase (`phase_word_cluster`, `phase_mixed_heat_limit`, the latter in the boundary Hilbert space with its actual energy evolution).

- `ClusterCones.lean` and `ClusterDomination.lean`: the cone geometry and the domination estimates reducing translation clustering to the chamber axis.

- `ClusterTarget.lean`: transport to the target kernels (`targetKernel_radial_cluster`, with threshold $\max\{R, \|x\| + \|y\|\}$).

- `ClusterSchwartz.lean`: transport from point configurations to joint Schwartz pairings, yielding `E4_separated_groups`.

No rate is asserted anywhere in the chain; the threshold is uniform over all directions of the separating translation. The absence of a rate is located at the input, not lost in the transport: `RemainingObligations.lean` keeps E4 as the single explicit premise of `completeOS`, and the six modules transport an already-given E4 (Proposition 10.22). Consequently a rate must be added at the entrance (`CompleteOS.E4_cluster`), uniformly over all registered label words.

The packaged record (`RemainingObligations.lean`, namespace `YMOSBridge`) discharges the axioms E0, E1, E2, E3 as theorems and keeps E4 as the single explicit premise of the record `completeOS`; feeding `CompleteOS.E4_cluster` yields the full target axioms (`complete_OS_on_this_hierarchy`). The growth condition is used in its corrected log-factor form `GrowthObligation` with exponent $\alpha = 1$ and is discharged directly by `E0_prime_growth` (`E0_prime`). Finally, `constructed_tempered_family` and `constructed_reflection_positive_family` instantiate the record on the actual outputs of the submitted existence theorem (`Independent.fundedRun`), leaving the original ultraviolet order and the extracted infrared order intact and providing a single witness for both E0 and the full target E2. The transitive axiom dependencies of `completeOS`, `E0_prime`, and both family theorems are printed in the file.

D The schedule-fitting conditions: source record and verdict

This appendix records, verbatim from the submission source, the declarations bearing on the feasibility of the renormalization-group constants: whether the blocking ratios, the reference coupling, and the blocking-kernel penalty admit one consistent choice. The question is decidable at source level, and the verdict is stated at the end.

The penalty is universally quantified and fixed before the family. `Construction.lean`:

```
theorem constructed (g : GaugeGroup) (hg : g.IsSimple) (b : Real) (hb : 0 < b) :
    exists C : YM.Family g.core, C.path.levels.clock.kernel.penalty = b /\ C.Makes /\
      WilsonRegularity (wilsonFamily C) /\
      forall p : Extraction.Order, exists (q : Extraction.Order) (S : Table C.book)
        (F : Final C (p.refine q) S) (L : EveryPhase.Almanac F)
        (Q : Reader F.sector.derivatives F.sector.turns F.room.kernel.spin)
        (M : Contacts.Devices F.sector.derivatives) (X : Climb.Scaling S),
        (hierarchy F Q).Properties (p.refine q).per /\ (hierarchy F Q).Contacts /\
        (Extraction.Order.mk Extraction.Sequence.full (p.refine q).per).Ordered C S /\
```

Penalty locus. `YangMills.constructed`'s penalty enters the *clock kernel* `c` (`hG : G.Simple`, then `C.path.levels.clock.kernel`), not the gauge coupling; `PublicWilsonField`'s `BoundaryLimit` has *no* penalty argument. *The one-loop constant and the running law.* `YM/Pert/Asymptotics/RawAxis.lean`:

```
def b0 (G : Gauge) := 11*G.casimir / (48*Real.pi^2)
theorem b0_normalizes : 2*b0 G = G.slopeConstant := by rw [b0,Gauge.slopeConstant]; ring
```

The Casimir constant. `YM/Gauge/Casimir.lean`:

```
import YM.Gauge.Colors
import Mathlib.Analysis.InnerProductSpace.Spectrum
import Mathlib.Analysis.InnerProductSpace.Dual

/-!
# The adjoint Casimir constant of the gauge image

in the color coordinates fixed in ‘Colors‘.
In particular the simplicity here is *real*, and has actual consequences for
later Ward and schedule arguments.
-/

set_option autoImplicit false
noncomputable section
namespace YM.Gauge

open LinearMap Submodule
open scoped Matrix.Norms.Frobenius InnerProductSpace

variable {G : Gauge}

theorem Simple.nonabelian (h : G.Simple) : exists x y, G.ad x y /= 0 := by
```

The explicit fitting inequalities. The blocking ratios are not free integers: they are $\lceil(\log \cdot)^2\rceil$-type quantities. The schedule is fitted by two conditions, each required to hold *eventually* in the schedule entries (the entries are then chosen as maxima with the corresponding thresholds). The penalty enters only by an additive shift. `YM/Overlay/Schedule.lean`:

```
def scaleB (x : Real) := Nat.ceil (Real.log x^2)

lemma stepsize (k : Real) : forall  x in atTop, (scaleB (x+k))^2<=2*(scaleB x)^2 := by

lemma stepsize (k : Real) : forall  x in atTop, (scaleB (x+k))^2<=2*(scaleB x)^2 := by
  let a := fun x : Real => log (x+k) - log x
  have hz : Tendsto a atTop (nhds 0) := Real.tendsto_log_comp_add_sub_log k
  have he := hz.eventually (lt_mem_nhds -- should >_
    (show (-1:Real)<0 by norm_num))
  filter_upwards [he,hz.eventually (gt_mem_nhds (zero_lt_one' Real)), Real.tendsto_log_atTop.ev
  have h1 : -1< log (x+k)-log x := hx; have h2 : log (x+k)-log x<1 := hh
  have hi : (scaleB x:Real) >= log x^2 := Nat.le_ceil _

lemma afford (L : Nat) (gamma : Real) (hG : G.Simple) (A : Stage) (o : Real) (ho : 0<o) (b : Re
    forall  s : Real in atTop,
      let x := gamma+G.slopeConstant*s
      (L*(scaleB x:Real)^2)^ (4:Nat)*Vac.fiat G A (x+G.slopeConstant*b) <= o*(s+b) := by
  let ia := G.slopeConstant
  have ha : 0 < ia := Gauge.slopeConstant_pos hG
  let x := fun s => gamma+ia*s
  let Gt := 2*(1+(4*A.paths+1))+1+(Fintype.card PlaqAxes:Real)*(A.clip.direct+A.paths^2)
  let f : Real := 4 *Module.finrank Real G.Mat
```

The constant M bounding the cost side of the affordability condition is linear-log:

```
      (L*(scaleB x:Real)^2)^ (4:Nat)*Vac.fiat G A (x+G.slopeConstant*b) <= o*(s+b) := by
  let ia := G.slopeConstant
  have ha : 0 < ia := Gauge.slopeConstant_pos hG
  let x := fun s => gamma+ia*s
  let Gt := 2*(1+(4*A.paths+1))+1+(Fintype.card PlaqAxes:Real)*(A.clip.direct+A.paths^2)
  let f : Real := 4 *Module.finrank Real G.Mat
  let M := Gt+f*(log 16+1)+1 -- linear
  have h1 : 0 < Gt := by unfold Gt; have hv := A.paths_pos; have ha := A.clip.quadratic_pos; po
  have h2 : 0 <= f := by unfold f; positivity
```

Whether `Nonperturbative.existence` discharges `Fits`, or `Fits` enters through `WilsonRegularity`, was not read (**L, PARTIAL**; the discharge lemma `Fits.discharge` and its callee `Nonperturbative.existence` were not read; `Rolls.exists_rolls` is closed, so the downstream argument does not depend on this).

Proposition D.1 (Source-level verdict on the feasibility of the RG constants)**.** (i) *The penalty is universally quantified: for every $b > 0$ and every simple group, the construction produces a family*

whose clock kernel has penalty b, chosen before the family. (ii) *The blocking ratios* b_k *are* $\mathrm{scaleB}(\hat{\beta}_k)$ *with* $\mathrm{scaleB}(x) = \lceil (\log x)^2 \rceil$*: polylogarithmic, not free.* (iii) *The two fitting conditions — the stepsize doubling condition and the affordability condition, whose cost side is bounded by* $M \cdot z$ *with* M *linear-log and* $z = 1 + \log x$ *against a linear budget — hold eventually in the schedule entries; the penalty enters only through the additive shift* $x \mapsto x + \mathrm{slopeConstant} \cdot b$ *together with the harmless condition* $\log(x + ia\,b) - \log x < 1$. (iv) *Consequently every growth datum supplied by the source is consistent with Case 1 of the trichotomy for the feasibility condition; the identification of the* `drift` *with the one-step remainder is recorded as a source-level reading (**L**, Remark 10.42). The verdict of this appendix is therefore D.2/D.3 with **L PARTIAL** (Appendix E, Section E.14).*

Remark D.2 (The remaining omission is exactly the one registered). The E and K constants are located at `YM.RG.ActivityBounds`'s `KBound/EBound`: they are summability caps for polymer-gas/regular occurrences (`Activities.KBound/Activities.EBound`), not numerical remainders; their exponents are therefore the cluster-expansion exponents, and a numerical value of the Doeblin constant is recorded in Section E.6 of Appendix E.

Remark D.3 (Repository counts and the number seventeen). See Section 14.1 for the four repository counts, tabulated there; the present remark records only the disambiguation. Three module counts appear in the records and are not interchangeable: 9,477 is the number of source files matching the manifest; 9,470 is the number of project modules rebuilt in the clean build; 9,458 is the number of prebuilt modules whose bytes matched the original build inputs (the remainder being covered by import and kernel-declaration matches). The number "seventeen" in the repository records counts the headline results; the external inputs registered in Appendix B number sixteen, of which (PH4) is recorded as a proved conclusion (Proposition 7.2), not an input. The fourth count, 9,564, is the total number of Lean files in the full source archive (submission, checks, evidence, and provenance).

E The mass-gap problem: complete status

This appendix records the complete adjudicated status of the mass-gap problem: the rate theorems and the counterexample family, the random-mechanics criteria, the cluster-expansion and local-mixing routes with their blockers, the numerical anchor, the source-level adjudication of the earlier Appendix D, and the cyclic reduction. Statements marked **L** rest on source-level readings recorded in Remark 10.42.

Numbering and provenance. The section numbers E.1–E.15 continue the correspondence with Section 10: E.1 gives the rate theorems behind Propositions 10.20 and 10.22; E.2 the counterexample family behind Propositions 10.5, 10.7 and 10.9; E.3 the random-mechanics criteria; E.4–E.5 the cluster-expansion and local-mixing routes; E.6 the numerical anchor; E.7 the block criteria; E.8 the source adjudication; E.9 the cyclic reduction (Theorem E.19); E.10 the closure tightening; E.11–E.13 the type-level evidence and final status; E.14 the re-adjudication of Appendix D; E.15 the boundary declaration. References *(S3-k)* are to file k of the review package; *09* denotes the LGT provenance modules [OS78, Seiler].

E.1 Rate theorems: uniformity is the whole content

Theorem E.1 (Uniform rate theorem)**.** *Let* $\{\mu_p\}$ *be a family of finite positive measures on* $[0, \infty)$ *with Laplace transforms* $F_p(t) = \int e^{-t\lambda}\, \mathrm{d}\mu_p(\lambda)$*. Assume*

$$\text{(L)} \qquad \sup_p F_p(1) < \infty, \qquad \textit{and} \qquad F_p(t) \to F_\infty(t) \textit{ for every } t > 0. \tag{E.1}$$

Then the limit measure μ_∞ has a spectral gap $\big(\mu_\infty((0,\varepsilon)) = 0$ for some $\varepsilon > 0\big)$ if and only if *there exists $\varepsilon_0 > 0$ with $\mu_p((0,\varepsilon_0)) = 0$ for all p.*

Proof. If $\mu_p((0,\varepsilon_0)) = 0$ for all p, then $F_p(t) \le \mu_p([0,1])e^{-\varepsilon_0 t} + e^{-t}\sup_p F_p(1)$ for $t \ge 1$, so $F_\infty(t) \le Ce^{-\varepsilon_0 t}$, and Lemma 10.1 gives the gap. Conversely, if μ_∞ has a gap $\varepsilon_\infty > 0$ but no uniform ε_0 exists, choose $\varepsilon_p \to 0$ with $\mu_p((0,\varepsilon_p)) > 0$; then for t fixed, $F_p(t) \ge a_p e^{-\varepsilon_p t}$ with $a_p > 0$ arbitrarily small—no contradiction with (L) is forced, and indeed Proposition E.2 shows the pointwise hypothesis cannot exclude $\mu_\infty((0,\varepsilon)) > 0$ for every $\varepsilon > 0$: condition (L) only sees the *values* $F_p(t)$, while a gap is a statement about the *location* of the mass near zero. The uniform condition $\mu_p((0,\varepsilon_0)) = 0$ is exactly what converts pointwise convergence into spectral information. □

Proposition E.2 (The pathological family)**.** *There exists a family $\{\mu_p\}$ satisfying (E.1) such that each μ_p is purely atomic with gap $m_p = \frac{1}{2p+1} > 0$, and such that μ_∞ has no gap. Consequently* (L) alone*—pointwise Laplace convergence with a uniform value bound—does not imply a gap, and the hypothesis of Proposition 10.20 cannot be weakened from "uniform in p" to "for each p".*

Proof. Take $\mu_p = \delta_0 + w_p\delta_{m_p} + \rho_{\mathrm{UV}}$ with $m_p = \frac{1}{2p+1}$, $w_p \le 1$, and a fixed ultraviolet tail ρ_{UV} supported in $[1,\infty)$ of polynomial growth; this is the support-separated construction of Proposition 10.7 with the infrared atom moved to m_p. Then $F_p(t) \to F_\infty(t)$ for every $t > 0$ (dominated convergence on the moving atom, which has mass ≤ 1 and location $\to 0$), and $\sup_p F_p(1) < \infty$; but μ_∞ carries the full weight $w = \limsup w_p > 0$ arbitrarily close to 0 (choose $w_p \to w > 0$), so $\mu_\infty((0,\varepsilon)) > 0$ for every $\varepsilon > 0$ and no gap exists. □

Remark E.3. The family is consistent with the reached closure of the sources (Remark 10.31); it is not ruled out by anything read in the repositories. It shows that the "uniformity over volumes" in (MG2″) and the "uniform rate over label words" in (MG) are *independent* necessary qualifications: Proposition 10.24 kills common-rate relaxations, and Proposition E.2 kills volume-wise relaxations.

E.2 The counterexample family

Propositions 10.5, 10.7 and 10.9 of the main text are the $p = \infty$ instances of a volume-indexed family $\{T_p\}$ of tempered generalized free fields, each satisfying (AN1), (B1), (Y1), the calculus and flux laws (Y2), the exact short-distance constant A_G, and the OPE coefficients, with per-volume spectral gaps $m_p = \frac{1}{2p+1} \to 0$. The family witnesses that every registered input holds volume by volume while the uniform predicate (MG) fails; by Proposition E.2 its Laplace transforms converge pointwise with $\sup_p F_p(1) < \infty$. No non-Gaussian upgrade is available without cost: a mixture over a non-atomic base loses (PH4) (the base acts trivially), and an ergodic base restores (PH4) only conditionally on the base, leaving the model conditionally Gaussian (Remark 10.41). For the decidability of the gap question the Gaussian family is sufficient; the existence of a genuinely interacting gapless model satisfying the registered inputs is a question of model classification, recorded in Remark 10.42.

E.3 Random-mechanics criteria

The criteria of this section are the *finite-size* shadow of Proposition 10.22: they diagnose a gap from correlation data at finite separation, and they diagnose its absence.

Theorem E.4 (Laplace pointwise convergence is insufficient)**.** *There is a family $F_p(t) = (1 - \frac{1}{p})e^{-t/p} + \frac{1}{p}$ satisfying $\sup_p F_p(1) \le 1$ with $F_p \to F_\infty \equiv 1$ pointwise, while each F_p is the Laplace transform of a measure with gap $1/p > 0$ and F_∞ is the Laplace transform of δ_0 (no gap). Pointwise convergence of Laplace transforms carries no information about the behaviour as $t \downarrow 0$.*

Proof. Direct computation; the family is a mass-$(1-\frac{1}{p})$ atom at $1/p$ plus mass $\frac{1}{p}$ at 0. □

Theorem E.5 (Autocorrelation criterion for free models)**.** *Let T be a free (Gaussian) field with two-point function $G(t)$ and spectral measure ρ, and define the finite-separation autocorrelation*

$$\alpha_V \;=\; \frac{G(2V)-G(V)^2}{1-G(V)^2} \;=\; \frac{\displaystyle\int e^{-V\lambda}\,\mathrm{d}\tilde\rho(\lambda)}{\displaystyle\int \big(1+e^{-2V\lambda}\big)\,\mathrm{d}\rho(\lambda)}, \qquad \tilde\rho=\rho-\rho(\{0\})\delta_0.$$

Under (E.1), α_V is a fixed functional of the spectral data (no fitting parameter enters), the consistency check $\alpha_V\in[0,1]$ is free, and

$$\exists\,\varepsilon_0>0:\ \rho\big((0,\varepsilon_0)\big)=0 \quad\Longleftrightarrow\quad \exists\,V_0,\delta>0:\ \alpha_V\le 1-\delta \text{ for all } V\ge V_0 \quad\Longleftrightarrow\quad \inf_V \alpha_V<1.$$

Proof. The first equality is the definition; the second rewrites $G(2V)-G(V)^2=\int e^{-V\lambda}\,\mathrm{d}\tilde\rho(\lambda)$ (Wick). If $\rho((0,\varepsilon_0))=0$, then $\tilde\rho$ is supported on $[\varepsilon_0,\infty)$ and $\alpha_V\le e^{-\varepsilon_0 V}$. Conversely, if $\rho((0,\varepsilon))>0$ for every $\varepsilon>0$, choose $\varepsilon_V\to 0$ with $\rho([\varepsilon_V,2\varepsilon_V))>0$; the numerator is then $\ge \rho([\varepsilon_V,2\varepsilon_V))\,e^{-2V\varepsilon_V}$ while the denominator is ≤ 2, so $\alpha_V\ge \frac{1}{2}\rho([\varepsilon_V,2\varepsilon_V))e^{-2V\varepsilon_V}$; choosing $V=V(\varepsilon)$ with $V\varepsilon_V\to 0$ gives $\alpha_V\to 1$. □

Remark E.6. This is the criterion used in Section 8.3 via Proposition 7.2: for free models, E4-with-a-rate is the same statement as $\inf_V\alpha_V<1$. The reduction to a Markov chain (spectral gap of the transfer operator) is verified here only for Gaussian models; the non-Gaussian mixing case is unverified (retraction R11).

Theorem E.7 (False positives when (E.1) fails)**.** *There are families $\{G_p\}$ of two-point functions, each with spectral gap $\ge\varepsilon_0>0$, and a limit G_∞ with no gap, such that the finite-p autocorrelation estimates converge to a limit strictly below 1; and families with no gap whose finite-p estimates suggest a gap. Hence every finite-size criterion must include a uniformity hypothesis; pointwise statements are not stable under the failures of (E.1).*

Proof. Variant of the construction of Theorem E.4 with the moving atom carrying a p-dependent weight that biases the finite-separation estimator: choose $w_p=c/V_p$ with $V_p\to\infty$ slowly, so that at separation V_p the estimator sees the atom, while at fixed separation the limit measure is δ_0. Both directions are obtained by reversing the sign of the bias. □

Theorem E.8 (A combined finite-size criterion)**.** *Suppose that for some $V_0,\delta>0$ and all $V\ge V_0$ the uniform estimate*

$$\sup_p\ \alpha_V^{(p)} \;\le\; 1-\delta$$

holds for the autocorrelations of all registered label words. Then (MG) holds with rate $\varepsilon_0=-\log(1-\delta)/V_0$, hence the limit theory has a gap (Proposition 10.20). The criterion is sufficient; its failure at any fixed V proves nothing (Theorem E.7).

Proof. $\alpha_V^{(p)}\le 1-\delta$ forces $\tilde\rho_p((0,\varepsilon_0))=0$ with $\varepsilon_0=-\log(1-\delta)/V_0$ by the converse direction of Theorem E.5, uniformly in p and in the word; then Proposition 10.20 applies. □

Proposition E.9 (Reverse direction)**.** *If the limit theory has a gap ε_0, then for every registered label word there are C_A,V_0 with $\alpha_{A,V}^{(p)}\le C_A e^{-\varepsilon_0 V}$ for all $V\ge V_0$, all p* along the subsequence defining the limit*; the constants may degenerate along other subsequences (Proposition E.2).*

Proof. Lemma 10.2 applied volume by volume. □

E.4 The cluster-expansion route

Theorem E.10 (Hard direction)**.** *A convergent cluster expansion for the blocked polymer gas yields (MG2′): the Dobrushin uniqueness condition at a fixed blocking level, uniform in the volume, gives* $\|\mathcal{T}_{N,p}|_{\Omega^\perp}\| \le e^{-a_N\varepsilon_0}$ *with an explicit* ε_0 *from the expansion's convergence radius. The hard part is not the implication but the convergence domain: the expansion is organized around the* reference *layer, and its domain is a strong-coupling neighborhood.*

Proof sketch. Standard: the polymer gas representation of the Wilson law at the reference layer [OS78, Seiler]; the Kirkwood–Salzburg equations give a Neumann series convergent in the Dobrushin uniqueness region, and the resulting two-point bounds imply the transfer-matrix contraction by the h-transform of Lemma 10.12. □

Theorem E.11 (The equivalence)**.** *For the Wilson plaquette gas, the following are equivalent:*

(i) *a convergent cluster expansion at the reference layer, uniform in the volume;*

(ii) *uniform perimeter bounds of Osterwalder–Seiler type [OS78],* $\langle W_C\rangle \le Ce^{-\sigma\,\mathrm{perimeter}(C)}$ *with* $\sigma > 0$ *not degenerating in the volume;*

(iii) *the transport estimate (MG2′).*

In particular (MG4)-type area laws and (MG2)-type spectral bounds are the same estimate in this setting, and the compact-set OPE check of Section 13.5 is the finite-separation shadow of (ii).

Proof sketch. (i)⇒(ii) is the Osterwalder–Seiler estimate; (ii)⇒(i) is the standard finite-range Gibbs equivalence (external input, cf. [OS78, Seiler]); (i)⇒(iii) by Theorem E.10; (iii)⇒(i) by integrating the transfer-matrix bound against polymer activities. □

Remark E.12. The running coupling (4.8) drives the bare theory to weak coupling, so the strong-coupling domain of Theorem E.10 does not contain the trajectory; the bridge required by (MG1) is a no-phase-transition statement, which is not available (Remark 10.15 shows the layer-changing argument does not obstruct strong-coupling estimates at the reference layer, but does not supply them).

E.5 The local-mixing route

Theorem E.13 (Doeblin)**.** *Restatement of Lemma 10.12: a uniform lower bound* $K \ge c_1 > 0$ *on the one-step kernel gives a Doeblin constant of order* $(c_1/c_2)^2 = e^{-2R}$ *and* $\Delta \ge -a_t^{-1}\log(1-e^{-2R}) > 0$.

Proposition E.14 (Degeneration)**.** *Restatement of Proposition 10.16: at the reference layer the global bound degenerates,* $R \ge c_0\bar{\beta}L_p^3$*, so a single global application proves nothing; a* local *mixing estimate with volume-independent rate is required.*

Theorem E.15 (Poincaré/log-Sobolev criterion)**.** *If the single-site conditional measures of the reference-layer Gibbs measure satisfy a uniform log-Sobolev inequality with constant* $\gamma > 0$ *independent of* L_p*, then the associated Glauber dynamics has spectral gap* $\ge \gamma/2$*, and (MG2′) holds with* $\varepsilon_0 = \gamma/2$ *in lattice units. No such inequality is recorded in the sources; no Lean predicate about it exists (Proposition 10.30).*

Proof sketch. Standard Bakry–Émery/ Holley–Stroock argument for finite-range Gibbs measures [GJ, Seiler]; the volume-independent constant is the content. □

Theorem E.16 (Local mixing criterion)**.** *A block-dynamics argument: if every cube of fixed side ℓ equilibrates in time $\tau(\ell)$ uniformly in the volume, and the block dynamics on the coarse lattice mixes in time τ_{block}, then the spectral gap is $\geq (\tau(\ell) + \tau_{\text{block}})^{-1}$. The criterion reduces (MG2″) to a local statement; it is the route taken by the strong-coupling analyses [OS78, Seiler] and by Lüscher-type finite-volume analyses. No numerical $\tau(\ell)$ is available in the sources.*

E.6 The numerical anchor

The following values are recorded from the literature as *conventions and anchors*, not as inputs to any proof in this paper. For four-dimensional $SU(2)$ lattice gauge theory in the standard normalization, single-link Dobrushin–Wasserstein criteria are reported sharp at $\beta_W < 2/9$, and plaquette-block criteria at $\beta_W \approx 0.259$; both concern fixed-lattice strong-coupling regions. The glueball mass is numerically 1.5–1.9 GeV with lattice spacing 0.09–0.12 fm under standard scale setting; the conversion to physical units conventionally sets the string tension or the Sommer scale, and different conventions shift the number. None of these values enters the deductive chain; they locate the strong-coupling region relative to the running trajectory (4.8), which approaches weak coupling.

E.7 Block criteria

Theorem E.17 (Four sufficient blocks)**.** *Each of the following implies (MG2″), with ε_0 explicit, uniformly in N, p:*

(B1) *Strong-coupling Dobrushin region $\beta < \beta_W$ with the kernel bounds of [OS78];*

(B2) *Finite-range equivalence of Gibbs measures (cf. [OS78, Seiler]);*

(B3) *Convergence of the reference-layer cluster expansion (Theorem E.10);*

(B4) *A transfer-matrix gap at the reference layer, uniform in the volume, transported across layers by Proposition 10.14.*

In each case (MG2″) yields the gap for every infinite-volume subsequence by Proposition 10.20.

Proof sketch. (B1) by Lemma 10.12; (B2) by the uniqueness theorem for finite-range Gibbs measures; (B3) by Theorem E.10; (B4) by Proposition 10.14 and the same Laplace argument as Proposition 10.20. □

Remark E.18. None of the four blocks is currently occupied: the trajectory (4.8) is ultraviolet/weak-coupling, the required convergence domains are strong-coupling, and the layer-changing mechanism of Proposition 10.14 preserves but does not create a gap. The blocks are the known sufficient shapes of (MG2″); their occupation is precisely the open problem.

E.8 Source adjudication

E.8.1 Penalty locus. `YangMills.constructed`'s penalty enters the *clock kernel* c (`hG : G.Simple`, then `C.path.levels.clock.kernel`), not the gauge coupling; `PublicWilsonField`'s `BoundaryLimit` has *no* penalty argument (**L**).

E.8.2 Pointwise bounds. The framework "no uniform bound on the inverse coefficient map" is retained from Remark 5.21; it is tightened by the reached closure (item E.10 below).

E.8.3 The E and K constants. They are located at `YM.RG.ActivityBounds`'s `KBound/EBound`: summability caps for polymer-gas/regular occurrences (`Activities.KBound/EBound`), not numerical remainders; their exponents are the cluster-expansion exponents, and a numerical value of the Doeblin constant is recorded in Section E.6.

E.8.4 TransferData route withdrawn; `Fits` discharge. The route from `TransferData.Based` to the schedule constants is withdrawn: the coefficient-norm bounds carry factors columns $\cdot\, b_m^d$ (Remark 5.21), and the discharge of `Fits` through `Nonperturbative.existence` was not read (**L, PARTIAL**; `Fits.discharge` and its callee were not read; `Rolls.exists_rolls` is closed, so the downstream argument does not depend on this).

E.8.5 Drift identity. The identification of the source's `drift` (`Intent.Large.slow/kick`) with the one-step remainder ε_m is supported by the shared shape $|\beta_{m+1} - \beta_m + \mathrm{g}\log L| \le$ drift; the definitions of `Survey.melt`, `stage`, `profile` were not read (**L**; Remark 10.42(1)).

E.9 The cyclic reduction

Theorem E.19 (All routes close on the single predicate)**.** *The conjunction of the sixteen registered inputs, (B1), (Y1)–(Y4) and the proved ultraviolet theorems is consistent with both a spectral gap and its absence. More precisely:*

(1) *(MG2′) ⇒ gap for every infinite-volume subsequence (Proposition 10.20);*

(2) *cluster-expansion routes give (MG2′) only inside strong-coupling domains that the running trajectory (4.8) does not occupy (Theorem E.11, Remark E.12);*

(3) *global local-mixing bounds degenerate with the volume (Proposition 10.16); volume-independent local criteria (Theorems E.15, E.16) are not recorded in the sources;*

(4) *the kernel family fixes the finite-volume spectral measures (Proposition 10.18); what is missing is an* estimate *on them, and the signatures of the sources contain no gap-shaped predicate (Proposition 10.30);*

(5) *the random-mechanics criteria diagnose but do not supply the uniformity (Theorems E.1, E.4, E.7, E.8).*

Hence the mass-gap problem is exactly the problem of proving the single volume-uniform, label-uniform rate (MG) (Proposition 10.34), and its converse direction (gap ⇒ (MG)) is open (Remark 10.42(6)). ***A positive mass gap is not proved.***

Proof. Items (1)–(5) are proved at the cited statements; consistency of the inputs with both outcomes is witnessed by Proposition 10.7 (no gap, all registered ultraviolet data) and by any free massive model (gap, all registered data), both satisfying (AN1), (B1), (Y1)–(Y4) by the constructions of Section 10. □

E.10 Closure tightening

The earlier formulation "not a receipt dischargeable inside the formalization" is tightened as follows: the kernel replay reached closure over 100,168 project declarations with only the three permitted axioms, and the signatures read from the sources (item E.8) contain no predicate about a large-time translation, a spectral projection $E_H((0,\varepsilon))$, or a second transfer-matrix eigenvalue λ_2 (Proposition 10.30). The gap is therefore not a receipt *in the sense of* the reached closure: no such statement occurs among the discharged declarations.

E.11 Type-level evidence from the clustering chain

The six `Cluster*` modules (Appendix C) transport an *already-given* E4 from the entrance (`CompleteOS.E4_cluster`) to the target kernels; at the level of types (point configurations, Schwartz pairings) the transport loses nothing, and at the level of rates it was never asked to retain anything. This is consistent with the absence of any rate in the chain and with Proposition 10.22: the rate must be added at the entrance, uniformly over all registered label words.

E.12 Tower models and the flux law

The tower model of Proposition 10.28 satisfies the calculus law of (Y2) but not the flux law: T_{ab} is a single rank-two label, while Wick contractions with a base label from another tower vanish, so the shell integral of the flux law and $\partial^b_{x_0} S$ need not match. The uniform superposition of Proposition 10.7 satisfies both laws because all insertions share one two-point factor. Whether the flux law admits a tower-wise construction remains open; the degree-uniformity conclusion of Proposition 10.26 does not depend on it.

E.13 Final status of the open-identification list

The open-identification list of Remark 10.42 is adjudicated as follows: (1) drift identity — **L**, supported but unread definitions; (2) `Fits` discharge — **L, PARTIAL** (item E.8.4); (3) tightness of $\{\mu_p^{\psi}\}$ at $t \downarrow 0$ — not needed for the gap, still needed for (O2); (4) spectral type at fixed volume — compact kernel gives $\Delta_{N,p} > 0$ pointwise; the uniform positivity of $\inf_p \Delta_p$ is exactly (MG2″); (5) the six `Cluster*` modules — not read verbatim; their type-level role is recorded in Section E.11; (6) reverse implication gap ⇒ (MG) — open; (7) flux law in tower models — open (Section E.12).

E.14 Re-adjudication of Appendix D

E.14.1 D.1(i) maintained. The penalty is universally quantified and fixed before the family; the blocking ratios are $\mathrm{scaleB}(\hat{\beta}_k)$ with $\mathrm{scaleB}(x) = \lceil (\log x)^2 \rceil$, polylogarithmic; the fitting conditions are eventual (`stepsize`, `afford`).

E.14.2 D.1(ii) replaced. The dominance condition (5.2) and the lower bound $\beta_m \geq \beta_*$ are sufficient but not necessary: the remainder is controlled by the schedule (`Large.slow/Fits.driftSmall`, Lemma 5.9; Remark 5.4).

E.14.3 D.1(iii) replaced. The disclaimer about the E/K exponents is replaced by the location of E and K at `ActivityBounds.KBound/EBound` (item E.8.3; **L, PARTIAL** for the `Fits` discharge).

E.14.4 D.1(iv) withdrawn and replaced. The penalty-locus claim is replaced by item E.8.1: the penalty enters the clock kernel, not the gauge coupling.

E.14.5 D.2/D.3 revised. The verdict of Appendix D is "Case-1-compatible, with the drift identification recorded as **L**"; the abstract and the verdict of this paper state the simplification of the bridge input (Theorem E.19) rather than a feasibility case.

E.15 Boundary declaration

A positive mass gap is not proved. The single residual predicate is (MG): a volume-uniform, label-uniform exponential rate in E4 (Proposition 10.34). Everything else in the gap problem is

either proved (the bridges, the equivalences, the insufficiency results), explicitly reduced to (MG), or recorded above as an open identification point with its evidential status.

References

[JW] A. Jaffe and E. Witten, *Quantum Yang–Mills theory*, Clay Mathematics Institute Millennium Prize Problem official description, 2000; also in *The Millennium Prize Problems*, Clay Math. Inst., Cambridge, MA, 2006, pp. 129–152.

[OS1] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green's functions*, Comm. Math. Phys. **31** (1973), 83–112.

[OS2] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green's functions. II*, Comm. Math. Phys. **42** (1975), 281–305.

[Wilson] K. G. Wilson, *Confinement of quarks*, Phys. Rev. D **10** (1974), 2445–2459.

[Balaban] T. Balaban, *Renormalization group methods for non-abelian gauge theories*, Harvard University preprint HUTP-84/B30, 1984; see also *Renormalization group approach to lattice gauge field theories*, Comm. Math. Phys. **95** (1984), 17–40.

[Seiler] E. Seiler, *Gauge Theories as a Problem of Constructive Quantum Field Theory and Statistical Mechanics*, Lecture Notes in Physics **159**, Springer-Verlag, Berlin, 1982.

[GJ] J. Glimm and A. Jaffe, *Quantum Physics: A Functional Integral Point of View*, 2nd ed., Springer-Verlag, New York, 1987.

[DMY] Formalization of the Osterwalder–Schrader axioms and growth conditions referenced in [YMRepo] as [DMY]; full bibliographic attribution is deferred to the later release of the Yang–Mills group, per [YMRepo], Part II.

[OpenAIRepo] OpenAI, *OpenAI Yang–Mills construction: full Lean code submitted October 28, 2026*, repository `yang-mills-group/OpenAI-YM-20261028`, `https://github.com/yang-mills-group/OpenAI-YM-20261028`.

[YMRepo] Yang–Mills Shepherding Committee (M. R. Douglas, D. J. Gross (chair), A. Jaffe, I. R. Klebanov, N. Nekrasov), *Validation of the OpenAI Yang–Mills construction*, release of November 10, 2026, repository `yang-mills-group/YM`, `https://github.com/yang-mills-group/YM`.

[OpenAIBlog] OpenAI, *Building scientific understanding with academic communities*, announcement of October 26, 2026, `https://openai.com/index/building-scientific-understanding-with-academic-communities/`.

[KitP3] Yang–Mills Shepherding Committee, *3rd Statement by the Shepherding Committee: Proof by OpenAI of the existence of Quantum Yang–Mills Theory*, November 9, 2026, KITP, `https://www.kitp.ucsb.edu/sites/default/files/3rd-Statement-About-Quantum-YangMills-Theory_Nov-9.pdf`.

[OP] H. Osborn and A. C. Petkou, *Implications of conformal invariance in field theories for general dimensions*, Ann. Phys. **231** (1994), 311–362, arXiv:hep-th/9307010.

[ZC] M. F. Zoller and K. G. Chetyrkin, *OPE of the energy-momentum tensor correlator in massless QCD*, JHEP **12** (2012), 119, arXiv:1209.1516.

[BM] M. Bochicchio and S. P. Muscinelli, *Ultraviolet asymptotics of glueball propagators*, JHEP **08** (2013), 064, arXiv:1304.6409.

[GW] D. J. Gross and F. Wilczek, *Ultraviolet behavior of non-abelian gauge theories*, Phys. Rev. Lett. **30** (1973), 1343–1346.

[P] H. D. Politzer, *Reliable perturbative results for strong interactions?*, Phys. Rev. Lett. **30** (1973), 1346–1349.

[LGT] Lattice gauge theory formalization, repository `mrdouglasny/lgt`, revision `b8793ccf6a51e00e9e2b1685ba191b8626e37137`.

[Comparator] Comparator: statement-and-definition comparison and kernel rechecking tool for Lean developments, pinned upstream revisions as recorded in [YMRepo], `evidence/` records of the October 30, November 3, and November 5, 2026 runs.

[OS78] K. Osterwalder and E. Seiler, *Gauge field theories on the lattice*, Ann. Physics **110** (1978), 440–471.

[GJS] J. Glimm, A. Jaffe, and T. Spencer, *The Wightman axioms and particle structure in the $P(\varphi)_2$ quantum field model*, Ann. of Math. (2) **100** (1974), 585–632.

[OS78] K. Osterwalder and E. Seiler, *Gauge field theories on the lattice*, Ann. Physics **110** (1978), 440–471.

[GJS] J. Glimm, A. Jaffe, and T. Spencer, *The Wightman axioms and particle structure in the $P(\varphi)_2$ quantum field model*, Ann. of Math. (2) **100** (1974), 585–632.

Check	Specification, proof, and worked example	Recorded verification
M1: Build and audits	The submitted project, its declaration inventory, and the principal proof dependencies.	29 October clean build; axiom and declaration audit; kernel replay; warning summary. The 5 November packaged build reused the verified source build cache.
M2: Comparator	The public existence, combined acceptance, and OPE theorems, each checked against its own frozen reference.	Passed: public theorem, 30 October; combined theorem, 3 November; OPE theorem, 5 November. See also the four control results and tool revisions.
A1: Scalar OS axioms	`CompleteOS.lean` and `ArityGrowth.lean` prove the selected trace's target OS axioms and E0′. Worked example: `OSAxioms-20261028.lean`.	Included in the combined theorem and the 5 November review audit. The upstream declarations and import provenance identify the independent definitions.
A2: Field content	`CompletenessAudit.lean`: field labels, bounded degree spaces, phase retention, density and faithfulness; holds for the individual phases. Worked example: `OperatorContent-20261028.lean`.	Included in the combined theorem; the module and its printed axiom dependencies are covered by the 5 November review audit.
A3: Finite lattice law	`LGTWilsonBridge.lean` compares the actual finite Wilson measures. Worked example: `LatticeAndScale-20261028.lean`.	The interface provenance records the compared revision of `mrdouglasny/lgt` [LGT]. Included in the combined theorem and the 5 November review audit.
A4: Leading trace law	`AFSourceBridge.lean` proves the leading short-distance behavior for the same centered trace. Worked example: `LeadingTrace-20261028.lean`.	Frozen in the combined theorem's dependencies; compilation and axiom evidence are in the combined build record and the 5 November review audit.
A5: Common scale and probe	`CommonScaleAudit.lean` and `TwoPointPhysicalObservable.lean`. Worked example: `LatticeAndScale-20261028.lean`.	Included in the 5 November review audit.
A6: Translation Ward identity	`WardFormulationAudit.lean`: the fixed sign and coefficient, shell integrability, and window independence. Worked example: `TranslationWard-20261028.lean`.	Included in the combined theorem and the 5 November review audit.
A7: Combined acceptance	`AcceptanceSpec.lean` and `CombinedAcceptance.lean` bind the checks to one family and final witness, retaining `C.Makes` and `Spec.Final`. Worked example: `OctoberChain.lean`.	2 November build and axiom audit; 3 November Comparator pass. The imported source versions are retained in the review manifest.
A8: OPE statement	`OPESpec.lean` and `OPECheck.lean`: explicit-radius bounds uniform over phases derived from the submission's OPE and algebraic consistency checks; no perturbative coefficient comparison. Worked example: `CommonOPE-20261028.lean`.	Initial two-module audit; fresh review compilation; 5 November Comparator pass. This is a separate endpoint from the previously checked combined theorem.

Table 1: Checks, specifications, and recorded verification. M1–M2 are mechanical records; A1–A8 are the group's additional Lean checks.