Deriving Standard Model Gauge Structure from Observer Overlap Consistency

B. Müller Alexander Osika Mario Poneder Kai Xue Peter Nguyen
Maarten Antonie Visser David Matscheko

July 26, 2026

Paper release: `r1583` **Released:** July 26, 2026

Abstract

Observer-Patch Holography models an observer as a finite self-reading patch with local state, boundary ports, readback records, and repair moves. This paper gives two gauge-reconstruction routes with distinct proof boundaries. The categorical route starts from overlap transport and reconstructs a compact group through a receipt-certified bosonic sector category and its fiber functor. The finite carrier route derives an inverse-port response from icosahedral incidence and target-blind readback, then constructs the current algebra $\mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{u}(1)$. Under the declared fermionic rank-15 matter category, anomaly balance fixes the Standard Model charges up to conjugation and the realized tensors have common kernel $\mathbb{Z}_6$. Thus their maximal faithful matter image is $(\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1))/\mathbb{Z}_6$. The local tensors also admit the cover and the $\mathbb{Z}_2$ and $\mathbb{Z}_3$ quotients. Physical global-form selection, laboratory-current attachment, and identification of the finite current with the categorical group require separate evidence.

Two reconstruction routes

The first route is structural:

$$\begin{aligned}\text{overlap gluing and transport receipts} &\Longrightarrow \text{rigid symmetric sector category}\\ &\Longrightarrow \mathrm{Aut}_{\otimes}(\mathcal{F}).\end{aligned}$$

It reconstructs a compact group on the declared bosonic branch. It does not select the Standard Model without a realized sector witness.

The second route is finite and carrier-specific:

$$\begin{gathered}\text{icosahedral incidence} \Longrightarrow 10J = A^3 - 4A^2 - 5A + 10I \Longrightarrow R = -J\\ \Longrightarrow \mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{u}(1) \Longrightarrow \text{conditional Standard Model matter image.}\end{gathered}$$

The response signs, equivariant lift, anomaly balance, charge lattice, and common tensor kernel are independently recomputed by released certificates. No Minimal Admissible Realization premise enters this gauge implication. That economy rule is used only in separately marked generation, scalar, and extra-sector statements.

Contents

1 Recovered Core: Foundations and Structural Branches

Observer-Patch Holography is a reconstruction program for fundamental physics. On the declared screen-first global-support branch, physical data are charted by a horizon screen S^2, and the effective world is encoded by the mutual consistency of overlapping patch descriptions. This support screen is distinct from a local carrier boundary and a federation overlap nerve. On the producer branch, S^2 enters only after the carrier-to-support map satisfies the spherical incidence, mesh, cross-ratio, normalization, and refinement receipts. The legacy finite reference fixture does not supply that map because it has no composable seam triangle and its higher-overlap condition is vacuous. The source-derived producer instead supplies a twelve-chart, thirty-seam, twenty-face incidence nerve, nonvacuous seam-triangle cocycles, confluent phase repair, and a controlled oriented-S^2 support limit on one declared tower. Event, modular, scale, and laboratory attachments remain separate gates. This recovered-core section develops fixed-cutoff overlap repair, transportable sector reconstruction, the source-derived finite current algebra, and the conditional Standard Model matter image.

1.1 Overview

Overlap consistency first turns local patch states into a constrained gluing problem. The fixed-cutoff branch gives collar decompositions, edge centers, and finite repair dynamics on the declared physical quotient. The transportable edge-sector branch classifies fixed-stage zero-obstruction sectors and, on a cofinal tail carrying the compact-gauge refinement receipt, reconstructs a compact internal gauge group in the bosonic branch.

The finite-current route is separate. Icosahedral incidence and inverse-port pairing determine J, with $10J = A^3 - 4A^2 - 5A + 10I$. A target-blind impulse/readback protocol solves the common farthest-shell filter and produces $R = -J$. Its relative signs on $\mathbf{1} \oplus \mathbf{3} \oplus \mathbf{3'} \oplus \mathbf{5}$ are fixed up to common reversal, the conventional overall charge sign. Given the declared fermionic Spin category, the matter certificate derives $P_{\text{even}} - P_{\text{vac}}$ and $P_{\text{odd}} - P_{\text{top}}$ as the unique charge-conjugate pair of rank-15 chiral projectors. Anomaly balance and tensor descent give exact hypercharge, $N_c = 3$, the common $\mathbb{Z}_6$ kernel, and the maximal faithful matter image. This implication uses no Minimal Admissible Realization premise. The cover and its $\mathbb{Z}_2$, $\mathbb{Z}_3$, and $\mathbb{Z}_6$ quotients admit the same local tensors, so physical global-form selection requires independent line or bundle data. Physical matter typing, equality of the two reconstruction routes, laboratory-current attachment, scalar attachment and multiplicity, and physical family attachment are work in progress.

1.2 Model and Axioms

1.2.1 Observers and access model

The support-local algebra-state-record reduct of an observer patch is $O_{\text{red}} = (P_O, \mathcal{A}(P_O), \rho_O, R_O)$, where:

- $P_O \subset S^2$ is a connected screen patch (the observer's access region).
- $\mathcal{A}(P_O)$ is the von Neumann algebra associated to P_O.
- ρ_O is the local state, obtained by restricting the global state to $\mathcal{A}(P_O)$.
- R_O is a finite record algebra generated by stable internal record projectors within P_O.

An operational observer additionally carries overlap interface algebras and restriction maps, allowed update and repair instruments, and checkpoint data used for continuation. Observers are internal self-reading patterns in the global state. A patch and its compatible marginal supply a geometric or algebra-state chart; operational observer status begins when the interface, record, repair, and checkpoint tests pass.

1.2.2 Screen, patches, and algebra net

We work in a single static patch with a horizon screen S^2. Each connected subregion $P \subset S^2$ is assigned a von Neumann algebra $\mathcal{A}(P)$. The net satisfies isotony:

$$P \subset Q \implies \mathcal{A}(P) \subset \mathcal{A}(Q).$$

A global state ω is a positive linear functional on the inductive-limit algebra. Overlap consistency is imposed algebraically: for overlaps $P_1 \cap P_2$, ω restricted to $\mathcal{A}(P_1 \cap P_2)$ is the same from either side.

1.2.3 Five OPH axioms

This paper states the five OPH axioms because both specialist derivations share their observer-overlap foundation. Branch-local hypotheses are written at the theorem that consumes them. Transportability and the fixed-cutoff bosonic sector category are theorem-level results. Refinement/fiber descent and the cofinal compact-gauge witness require the explicit compact-gauge refinement receipt. The finite Standard Model current implication does not use Minimal Admissible Realization. Its fermionic matter conclusion uses the declared Spin category. Physical global-form selection and scalar multiplicity require separate premises.

The following screen axiom defines the screen-first support branch. The producer branch reaches the same axiom only after its carrier-to-support receipts establish the global support screen.

1. Screen Net. A horizon screen S^2 carries a net of algebras $P \mapsto \mathcal{A}(P)$.

2. Overlap Consistency. Local states agree on shared observables for any overlap.

3. Local MaxEnt and Refinement Stability. At the UV scale the realized state maximizes entropy subject to a finite family of gauge-invariant local constraints, and the resulting family of states is stable under coarse-graining so symmetry-allowed relevant operators are not held at zero by unexplained fine tuning.

4. Recoverable Generalized Entropy. A generalized entropy functional exists on caps, obeys quantum focusing on null generators, and comes with the recoverability structure used throughout the manuscript: collar tripartitions have small conditional mutual information (CMI) with controlled recovery maps, with stronger collar/null-strip hypotheses stated explicitly where needed.

5. Minimal Admissible Realization (MAR). On the admissible low-energy branch, the realized sector package is the lexicographically minimal one under the complexity vector $C(\mathfrak{S})$. MAR is an explicit structural-economy axiom on admissible realized branches, not a theorem derived from the preceding four axioms.

1.2.4 Closure values and theorem-local technical data

Transportability, the fixed-cutoff bosonic category, regulator bookkeeping, quasi-local propagation, endpoint-Lipschitz control, and collar decomposition are cited from their preceding theorems. Refinement/fiber descent and the cofinal compact-gauge witness additionally require the compact-gauge refinement receipt.

Local MaxEnt at the regulator scale. At the regulator scale ℓ_{UV}, the global state ω maximizes von Neumann entropy subject to:

1. A finite set $\{O_a\}$ of gauge-invariant local operator densities, each supported on a ball of radius $\leq r_0 = O(\ell_{\mathrm{UV}})$.
2. Constraint equations $\langle \sum_x O_a(x) \rangle = C_a$, one homogeneous global sum per density label a, where x runs over the cells of the UV lattice. The constraints are the N_{con} global sums, not one independent constraint per cell, so the number of independent constraints and of Lagrange multipliers is N_{con} (plus the optional global constraints below) on every lattice, independent of the cell count.
3. Optionally, a finite number of global constraints (total energy, charge).

This is the minimal specification needed to derive the local Gibbs form of Lemma 2.6 from Axiom 3 with a cutoff-independent multiplier count.

Clarification (MaxEnt $\neq$ thermal equilibrium). MaxEnt here is **local state selection** given constraints, not "the universe is in thermal equilibrium." Non-equilibrium physics is modeled in the enlarged *local-multiplier* family obtained by replacing the homogeneous multipliers λ_a with slowly varying fields $\lambda_a(x)$, equivalently by imposing per-cell constraints $\langle O_a(x) \rangle = c_a(x)$. That enlarged family has dimension proportional to the number of UV cells, so it is a function-space approximation regime with gradient error terms, not the finite-dimensional family of the refinement argument: the refinement-stable branch and the induced maps $R_{\ell \to L}$ below are always statements about the homogeneous N_{con}-parameter branch. Controlled violations of exact Markov additivity appear under the explicit mixing hypotheses stated later in Section 2.3. Equilibrium is an approximation regime with explicit error terms.

Rotationally invariant constraint branch. Constraint sets are SO(3)-invariant on S^2.

Gauge as overlap redundancy at fixed cutoff. Overlap identifications are not unique; the freedom that leaves overlap observables invariant forms a local groupoid.

Central-defect subbranch. Assume the overlap centers are identified with one fixed abelian unitary coefficient group Z_Σ on which overlap transport acts trivially. Choose unitary overlap implementers with $U_{ji} = U_{ij}^{-1}$. On triple overlaps, when the failure of strict coherence is central, define

$$\Omega_{ijk} := U_{ij} U_{jk} U_{ki} = z_{ijk}, \qquad z_{ijk} \in Z_\Sigma,$$

equivalently $U_{ij} U_{jk} = z_{ijk} U_{ik}$. The defect is data of the implementer lift; writing only $\mathrm{Ad}(z_{ijk})$ would erase it because conjugation by a central element is the identity.

Collar double-scaling hypothesis. There is one refinement family with $\delta \to 0$, $\ell_{\mathrm{UV}} \to 0$, and uniform finite-range and strong conditional Gibbs mixing constants such that

$$I(A_\delta : D_\delta \mid B_\delta)_\omega \leq c |\partial C|_{\mathrm{UV}} e^{-\delta/\xi}, \qquad \frac{\delta}{\xi} - \log |\partial C|_{\mathrm{UV}} \longrightarrow +\infty.$$

The second condition controls the complete boundary-prefactored envelope. The weaker ratio $\delta/\ell_{\mathrm{UV}} \to \infty$ does not imply it. Section 2.3 gives the finite-range theorem, explicit constants, a sufficient logarithmic schedule, and the finite-receipt fields. On the exact central-interface branch the CMI is zero without this mixing premise.

Derived quasi-local propagation and modular-locality control. At scale ℓ_{UV}, the automorphism group generated by the MaxEnt generator of Lemma 2.6, or more generally by any branch generator lying in the same bounded-support algebraic closure, has a finite Lieb–Robinson velocity v_{LR}: for local operators A, B supported on regions separated by distance d,

$$\|[A(t), B]\| \leq c\|A\|\|B\| \min(|R_A|, |R_B|)e^{-(d-v_{\mathrm{LR}}|t|)/\xi}$$

for $d > v_{\mathrm{LR}}|t|$, where $\xi = O(\ell_{\mathrm{UV}})$.

This is the branch-internal control statement that turns the quasi-local structure of the local-Gibbs branch into explicit support control for time-evolved operators. It is part of the third-axiom branch stated at regulator scale.

Approximate modular covariance on the fixed-cutoff realized branch. Under the fixed-cutoff realized presentation, Lemma 2.6 (local Gibbs), the derived quasi-local propagation bound above, and the collar double-scaling plus mixing hypotheses of Section 2.3, the modular flow $\sigma_t^{\omega,C}$ maps $\mathcal{A}(R)$ into a slightly thickened region algebra:

$$\sigma_t^{\omega,C}(\mathcal{A}(R)) \subseteq \mathcal{A}(R^{+v_{\mathrm{mod}}|t|}) \quad \text{up to error } \eta(d - v_{\mathrm{mod}}|t|),$$

where R^{+s} denotes the s-neighborhood thickening, v_{mod} is a "modular propagation velocity" controlled by v_{LR} and local norm bounds, and d is the distance from R to ∂C.

Heuristic motivation. On the reduced-modular-locality branch of Lemma 4.1a-b, modular additivity localizes the nonadditive defect of $K_C = -\log \rho_C$ to the collar. The associated propagation receipt then bounds support spreading under $e^{iK_C t}$. In the rate-controlled collar limit of Section 2.3, the thickening vanishes in macroscopic units. This remains heuristic support for the tangent-limit package; a Lieb–Robinson bound for the global Gibbs generator alone does not prove locality of the reduced modular Hamiltonian.

Continuum-limit heuristic. Define the induced region flow

$$f_t^C(R) := \lim_{\ell_{\mathrm{UV}} \to 0} R^{+v_{\mathrm{mod}}|t|}.$$

Then $\sigma_t^{\omega,C}(\mathcal{A}(R)) = \mathcal{A}(f_t^C(R))$ becomes exact in the continuum limit, with error controlled by

$$\eta(\delta) \lesssim 2\sqrt{c|\partial C|_{\mathrm{UV}}}\, e^{-\delta/(2\xi)}$$

when CMI is measured in nats, with the full rate condition imposed on the prefactor.

The discussion above is heuristic support for the tangent-limit package used in Theorem 4.2.

Refinement-stable MaxEnt branch. Choose any family of coarse-graining channels $\Phi_{\ell\to L}$ between UV scale ℓ and IR scale L that is compatible with the third OPH axiom, including its refinement-closure clause. Then on the realized branch one may write

$$\Phi_{\ell\to L}(\omega_\ell(\lambda)) = \omega_L(R_{\ell\to L}(\lambda)),$$

for an induced map $R_{\ell\to L}$ on the common finite-dimensional homogeneous multiplier family. The existence of $R_{\ell\to L}$ as a self-map of that same finite-dimensional family is exactly the closure clause, not a bookkeeping consequence of reusing the labels a: coarse-graining a finite-range Gibbs family generically generates interactions outside the retained density list, and the axiom asserts that on the realized branch those generated directions vanish or are absorbed into the retained family; the closure defect, the moment-matching I-projection realizing $R_{\ell\to L}$, and its trace-norm residual bound are constructed in Definition 2.6a and Lemma 2.6b, and what remains assumed rather than proved is that the defect vanishes along the realized branch. Granting closure, the regulator states lie in that shared multiplier family rather than in unrelated state spaces at different cutoffs, and the realized branch is the persistent trajectory or invariant subset selected inside it under the induced refinement maps. This is a state-side persistence statement; it does not by itself upgrade fixed-cutoff edge labels to a transportable refinement-persistent sector colimit.

Fixed-cutoff realized presentation. At a UV scale ℓ_{UV}, local patch algebras are type-I with finite-dimensional Hilbert spaces, and gauge-as-gluing is realized as a boundary action on a lifted finite-dimensional presentation. Section 2.3 proves the resulting fixed-cutoff collar package directly from overlap consistency on this realized presentation, while distinguishing the lifted fixed-point algebra from the sector-preserving collar algebra used in the edge-center (EC) theorem.

External mathematical inputs include SSA and recovery theorems (Petz 1986, 1988; Fawzi and Renner 2015). The categorical gauge route also uses Doplicher-Roberts reconstruction once localized zero-obstruction sectors are assembled in the small-region limit. Full citations appear in the References.

1.2.5 Notation

- ρ_C: reduced state on cap C.
- $K_C := -\log \rho_C^{\omega}$: modular Hamiltonian of the reference state.
- B_C: geometric generator of the cap-preserving conformal dilation.
- $S_{\mathrm{gen}}(C)$: generalized entropy on a cap.
- ℓ_{UV}: UV length scale of the refined screen net.
- δ: collar width around a cap boundary.

1.3 Information-Theoretic Tools

1.3.1 Strong subadditivity and Markov states

For any tripartite state ρ_{ABC},

$$I(A:C \mid B) := S(AB) + S(BC) - S(B) - S(ABC) \geq 0.$$

Exact Markov states satisfy $I(A:C \mid B) = 0$ and admit a recovery map:

$$\rho_{ABC} = (\mathrm{id}_A \otimes \mathcal{R}_{B\to BC})(\rho_{AB}).$$

1.3.2 Approximate recovery

If $I(A:C \mid B) \leq \varepsilon$ in nats, there exists a CPTP recovery map $\mathcal{R}$ with

$$\|\rho_{ABC} - (\mathrm{id}_A \otimes \mathcal{R})(\rho_{AB})\|_1 \leq 2\sqrt{1-e^{-\varepsilon}} \leq 2\sqrt{\varepsilon}.$$

1.3.3 Collar refinement and sufficient mechanisms

Fix a cap $C \subset S^2$ with boundary circle. For collar width δ define

$$B_\delta := \{x \in S^2 : \mathrm{dist}(x, \partial C) \leq \delta\}, \qquad A_\delta := C \setminus B_\delta, \qquad D_\delta := (S^2 \setminus C) \setminus B_\delta.$$

Then $S^2 = A_\delta \cup B_\delta \cup D_\delta$ with A_δ and D_δ interacting only through B_δ. The collar double-scaling hypothesis is the requirement that $I(A_\delta : D_\delta \mid B_\delta) \to 0$ in the refinement limit. At fixed regulator scale, the same finite-dimensional realized presentation supplies the patch-net and overlap-gluing data used below. We therefore isolate that realized presentation first, then separate the exact-Markov and quantitative routes.

Finite quotient ensemble theorem surface. Fix a finite regulator r. The physical presentation space is Σ_r, the presentation redundancy groupoid is Γ_r, and the finite physical quotient is

$$Q_r = \Sigma_r/\Gamma_r, \qquad \pi_r : \Sigma_r \to Q_r.$$

The quotient removes nonphysical presentation data: gauge representatives, port relabelings, mesh labels, shard or worker identifiers, queue order, repair schedule identifiers, retry counters, timestamps unless declared semantic, hidden carrier coordinates, and inert ancillary labels. If only settled configurations carry probability, the probability space is the normal-form subset

$$N_r = n_r(Q_r).$$

The map n_r is a normal-form map, not a probability law. Any promoted physical branch inherits this firewall: it must declare the quotient-intrinsic source law or action before a normal form can be read as a selection or prediction claim.

Observable algebras and reference states. In the finite classical case the quotient observable algebra is

$$\mathcal{O}_r = \ell^\infty(Q_r).$$

In the finite quantum case the physical algebra is a declared quotient algebra $\mathcal{A}_r^{\text{phys}}$ with state

$$\omega_r(A) = \mathrm{Tr}(\rho_r A).$$

When the reference object is obtained from a finite lifted carrier, the load-bearing data are not an abstract groupoid cardinality alone. They are a tracially pointed quotient

$$\left(\mathcal{A}_{r,b}^{\text{phys}}, \tau_{r,b}^0\right), \qquad \mathcal{A}_{r,b}^{\text{phys}} = z_{r,b} B(\widetilde{\mathcal{H}}_r)^{G_r} z_{r,b},$$

where $U_r : G_r \to U(\widetilde{\mathcal{H}}_r)$ is the compact gauge action and $z_{r,b}$ is the central projection for the declared boundary or superselection sector. The reference trace is

$$\tau_{r,b}^0(A) = \frac{\mathrm{Tr}_{\widetilde{\mathcal{H}}_r}(A)}{\mathrm{Tr}_{\widetilde{\mathcal{H}}_r}(z_{r,b})}.$$

If

$$z_{r,b}\widetilde{\mathcal{H}}_r \cong \bigoplus_\alpha V_\alpha \otimes M_\alpha,$$

with $d_\alpha = \dim V_\alpha$ and $m_\alpha = \dim M_\alpha$, then

$$\mathcal{A}_{r,b}^{\text{phys}} \cong \bigoplus_\alpha I_{V_\alpha} \otimes B(M_\alpha), \qquad p_{r,\alpha} = \frac{d_\alpha m_\alpha}{\sum_\beta d_\beta m_\beta}.$$

These are the induced central-sector weights only after the carrier representation and boundary sector have been fixed.

OPH quotient ensemble. An OPH quotient ensemble is specified by a quotient-intrinsic base weight and action

$$m_r : Q_r \to \mathbb{R}_{>0}, \qquad S_r : Q_r \to \mathbb{R} \cup \{+\infty\},$$

and

$$w_r(q) = m_r(q)e^{-S_r(q)}, \qquad Z_r = \sum_{q \in Q_r} w_r(q), \qquad \mu_r(q) = Z_r^{-1} w_r(q).$$

Equivalently, one may state an intrinsic projective prior ν_r on Q_r and set $\mu_r = (n_r)_\# \nu_r$. Uniform quotient counting, uniform representative counting pushed to the quotient, groupoid weights, and tracial central-sector weights are different physical claims. The paper must declare which one is being used.

Normal-form projector non-selection. For any retraction $N : Q \to Q_{\mathrm{nf}}$ onto a subset $Q_{\mathrm{nf}} \subseteq Q$, that is, any map whose restriction to Q_{nf} is the identity, the induced map on laws

$$\mathcal{C}_Q(\mu) = N_\# \mu$$

is idempotent:

$$\mathcal{C}_Q^2 = \mathcal{C}_Q.$$

Every law supported on Q_{nf} is fixed; both statements use the retraction property. Therefore settlement or canonicalization never selects a unique physical probability law by itself.

Selection-gap corollary. Let $X \subseteq Q_{\mathrm{nf}}$ be a finite set of quotient-normal candidates distinguished by visible invariants. Normal-form data determine X and its quotient-visible invariants, but they do not choose a member of X. If two laws μ, ν are supported on X and concentrate on different candidates, both are fixed by $\mathcal{C}_Q$. Unique sector selection therefore requires source data: an intrinsic action with a unique minimizer, a declared physical ensemble, or a refinement-stable gap certificate. A defect or holonomy classification can classify possible sectors without choosing the physical sector, and a contraction or repair generator can certify convergence toward a declared target without creating the target law.

Finite MaxEnt quotient ensemble. For finite Q, positive m, and quotient observables $F_1, \ldots, F_k$, maximizing

$$\mathcal{H}_m(\nu) = -\sum_q \nu(q) \log \frac{\nu(q)}{m(q)}$$

subject to

$$\sum_q \nu(q) = 1, \qquad \sum_q \nu(q) F_a(q) = c_a$$

has the full-support solution, when the feasible full-support surface is nonempty,

$$\mu(q) = \frac{m(q) \exp[-\sum_a \theta_a F_a(q)]}{Z(\theta)}.$$

Boundary optima obey the same formula after restricting to their support. On a finite noncommutative quotient algebra with faithful reference state σ_r,

$$\rho_r = \frac{\exp(\log \sigma_r - \sum_a \theta_a F_{r,a})}{\operatorname{Tr} \exp(\log \sigma_r - \sum_a \theta_a F_{r,a})}.$$

The finite constraint ledger must name every $F_{r,a}$, its units and support, the target expectation and source, sector or zero-mode treatment, refinement transformation, and proof that no run output or observational output entered the source definition.

Refinement compatibility and RG closure. For $s \succeq r$, let $c_{sr} : Q_s \to Q_r$ be the physical coarse map. Exact compatibility of weighted ensembles is equivalent to the fiber-sum identity

$$\sum_{q' : c_{sr}(q')=q} m_s(q')e^{-S_s(q')} = \alpha_{sr} m_r(q) e^{-S_r(q)}$$

for a constant $\alpha_{sr} > 0$ independent of q. Then

$$(c_{sr})_{\#}\mu_s = \mu_r.$$

If the one-step defects are

$$\delta_{k+1,k} = \|(c_{k+1,k})_{\#}\mu_{k+1} - \mu_k\|_{\mathrm{TV}},$$

then

$$\|(c_{nr})_{\#}\mu_n - \mu_r\|_{\mathrm{TV}} \le \sum_{k=r}^{n-1} \delta_{k+1,k}.$$

For exponential-family refinement, exact closure requires the fine conditional free energy

$$G_{sr,\theta}(q_r) = -\log \mathbb{E}_{m_s^0}\left[\exp[-\theta \cdot F_s(Q_s)] \mid c_{sr}(Q_s) = q_r\right]$$

to equal $\kappa_{sr}(\theta)+R_{sr}(\theta)\cdot F_r(q_r)$. If the residual is uniformly bounded by ε, the induced total-variation defect is bounded by $\tanh \varepsilon$.

Implementation invariance and representative lifting. If implementations A, B have quotient bijections $h_r : Q_r^A \to Q_r^B$ satisfying

$$m_r^B(h_r q) = m_r^A(q), \qquad S_r^B(h_r q) = S_r^A(q), \qquad h_r \circ c_{sr}^A = c_{sr}^B \circ h_s,$$

then

$$(h_r)_{\#}\mu_r^A = \mu_r^B.$$

For tracially pointed quantum quotients the corresponding equivalence is a trace-preserving quotient equivalence. It is invariant under unitary intertwiners preserving the gauge action and sector, and under inert trivial ancillas $A \mapsto A \otimes I_{\mathrm{anc}}$. It is not invariant under arbitrary changes of gauge-representation multiplicities.

If an implementation stores representatives, a representative-level law must be a conditional lift

$$\widetilde{\mu}_r(x) = \mu_r(\pi_r x)\kappa_r(x \mid \pi_r x), \qquad \sum_{x:\pi_r(x)=q} \kappa_r(x \mid q) = 1.$$

Then $(\pi_r)_{\#}\widetilde{\mu}_r = \mu_r$. Uniform representative sampling yields orbit-size weights and is physical only if representative counting is the declared base measure.

Quotient-lumpable kernels and sampler correctness. A representative kernel $\widetilde{P}(x, y)$ descends to Q_r only when

$$P_Q(q, q') = \sum_{y:\pi(y)=q'} \widetilde{P}(x, y)$$

is independent of the chosen representative $x \in \pi^{-1}(q)$. For $w(q) = m(q)e^{-S(q)}$ and proposal $R(q, q')$ with reciprocal support, the Metropolis–Hastings acceptance rule

$$a(q, q') = \min\left\{1, \frac{w(q')R(q', q)}{w(q)R(q, q')}\right\}$$

gives detailed balance

$$\mu(q)R(q,q')a(q,q') = \mu(q')R(q',q)a(q',q).$$

Repair-informed proposals must include the Hastings asymmetry term; otherwise the stationary law is generically changed.

Repair generators are not selectors. A repair generator of the form

$$L_{\mathrm{rep}} = \sum_C c_C(I - E_C)$$

is a relaxation or sampling object after a law has been selected. Conditional expectations E_C are defined on $L^2(X_r, \pi_r)$, so the reference law π_r is input. On overlapping collars the expectations need not commute. The correct finite gap certificate is the Poincare constant

$$\kappa_r = \inf_{f \perp 1} \frac{\sum_C \|(I - E_C)f\|^2}{\|f\|^2}.$$

If local fiber rates have a positive lower bound γ_*, then

$$L_{\mathrm{rep}} \geq \gamma_* \kappa_r (I - P_0).$$

Finite repair completeness gives $\kappa_r > 0$ at fixed regulator. A uniform refinement lower bound $\inf_r \kappa_r > 0$ is a separate theorem or receipt.

Finite evidence accuracy. For bounded coarse observables O, if

$$\|\widehat{\mu}_s - \mu_s\|_{\mathrm{TV}} \leq \epsilon_{\mathrm{samp}}$$

and the refinement defects sum to ϵ_{ref}, then

$$\left|\mathbb{E}_{\widehat{\mu}_s}[O \circ c_{sr}] - \mathbb{E}_{\mu_r}[O]\right| \leq 2\|O\|_\infty(\epsilon_{\mathrm{samp}} + \epsilon_{\mathrm{ref}}).$$

Continuum-facing observables require a realization map and correlation Cauchy bound in addition to a finite histogram.

Vacuum promotion gate. A stationary sampler is not a physical vacuum. For any faithful target law one can build a positive transfer operator with that law as ground state, so positivity alone is not a selector. Vacuum promotion requires source Euclidean slab data

$$\mathfrak{S}_r^E = (Q_r, m_r^0, J_r, V_r, a_{t,r})$$

whose conductance $J_r(q,q') = J_r(q',q) \geq 0$, local potential V_r, and slab thickness $a_{t,r}$ are derived without using the target law or sampler output. With connected event graph,

$$(H_r^E f)(q) = \frac{1}{m_r^0(q)} \sum_{q'} J_r(q,q')\big(f(q) - f(q')\big) + V_r(q)f(q)$$

is self-adjoint and bounded below on $L^2(Q_r, m_r^0)$; its finite Feynman–Kac semigroup is positivity improving. Perron–Frobenius gives a unique positive normalized ground state Ω_r, and the finite vacuum law is

$$\mu_r^{\mathrm{vac}}(q) = |\Omega_r(q)|^2 m_r^0(q).$$

For $T_r = e^{-a_{t,r}(H_r^E - E_{0,r})}$, the Doob kernel is stochastic and detailed-balanced with μ_r^{vac}. Continuum promotion additionally requires reflection positivity or equivalent reconstruction plus refinement compatibility of the transfer family.

Primordial and cosmological prediction firewall. A screen covariance contains incomplete radial information. The complete one-shell map

$$C_\ell = 4\pi \int_0^\infty \frac{dk}{k} \Delta_\zeta^2(k) j_\ell^2(k\chi_\star)$$

has an infinite-dimensional kernel that persists under positivity. OPH primordial promotion requires the source-only stress, single-clock, entropy-repair, curvature-evolution, adiabatic-mode, phase-coherence, physical mode, radial-null-space, and forward-projection receipts together with a scale-natural physical dilation intertwiner or complete radial cross-covariance tomography. A finite radial prior produces a conditional continuation. Observable CMB comparison also requires declared source, solver, dataset, covariance, nuisance, data-use, and pooled-reducer provenance.

Claim tiers and required receipts. Every ensemble-facing run records its ensemble id, claim tier, regulator, representative schema, gauge action, canonicalizer, base measure, action coefficients, coarse maps, zero-mode projector, amplitude convention, sampler, smoothing policy, source provenance, and explicit nonclaims. The seed belongs to the run receipt rather than the ensemble definition. The claim tiers are

$$\begin{aligned}
E0 &: \text{ seed noise, proposal noise, repair jitter,}\\
E1 &: \text{ conventional reference ensemble,}\\
E2 &: \text{ OPH-native quotient ensemble,}\\
E3 &: \text{ OPH vacuum,}\\
E4 &: \text{ OPH primordial field,}\\
E5 &: \text{ observable cosmological prediction.}
\end{aligned}$$

The evidence bundle must keep separate receipts for stationary-law schedule invariance, detailed balance of the aggregate kernel, and pathwise partition invariance. Deterministic replay of semantic random streams or a canonical serial chain is useful, but it is not pathwise partition invariance. Smoothing must preserve raw coefficients, raw spectra, smoothing kernels, smoothed coefficients, smoothed spectra, and hashes of each stage; it is not part of S_r unless explicitly declared.

Derived regulator realization. Choose a finite UV cellulation at scale ℓ_{UV} and let R be a finite union of cells. Hilbertize each cell's finite local data to a finite-dimensional space $\tilde{\mathcal{H}}_i \cong \mathbb{C}^{n_i}$. The extended algebra before quotienting by overlap redundancy is

$$\widetilde{\mathcal{A}}(R) = \mathcal{B}(\tilde{\mathcal{H}}_R), \qquad \tilde{\mathcal{H}}_R = \bigotimes_{i \subset R} \tilde{\mathcal{H}}_i.$$

Overlap-preserving changes of local trivialization act only on cut data. On each finite-dimensional chart, their unitary image has compact closure, so one may represent the boundary gluing redundancy by a compact group $G_{\partial R} \subset U(\tilde{\mathcal{H}}_R)$. The lifted boundary-invariant endomorphism algebra is

$$A_{\mathrm{inv}}(R) = \widetilde{\mathcal{A}}(R)^{G_{\partial R}} = \mathcal{B}(\tilde{\mathcal{H}}_R)^{G_{\partial R}}.$$

This is the fixed-cutoff realized presentation used throughout the collar analysis. The EC theorem works on the invariant-state realization and uses the sector-preserving induced collar algebra there; the entire fixed-point algebra on the unreduced tensor product is outside that theorem surface. When the consensus paper speaks of quotient repair, the physical repair law is defined on this fixed-point / overlap-invariant data and any representative-level map is only a lift of that quotient update. Descent to the gauge quotient is therefore built into the physical-algebra formulation as quotient-level repair: the local map is $\mathrm{locRep}_\lambda : Q \to Q$, and the global map is the finite quotient normal-form operator

$$\mathrm{Rep}_\lambda = \overline{\mathrm{nf}}_\lambda : Q \to Q.$$

On the declared fixed-cutoff collar branch, the local repair step itself is read from exact Markov splice or a declared Petz/Fawzi–Renner recovery channel. The declared fixed-cutoff branch adds repair completeness, together with control on the declared Petz domain where that branch is used. The consensus paper proves those clauses on a nontrivial rooted-tree packet-net export: finite packet labels live on a rooted tree, hidden labels are quotient-only, weighted parent-copy repair strictly lowers Φ, normal forms are exactly consistent packet assignments, and the classical full-support Petz channel is CPTP and trace-norm contractive on its positive-support-gap domain. Gauge-invariant observables factor through the quotient normal form on that verified branch. The declared touched-overlap acceptance contract makes Φ the finite-patch Lyapunov functional for primitive repair proposals on the scalar branch, but a proposal becomes a physical step only through the transactional acceptance layer: snapshot-current read sets, unchanged-outside-write validation, preserved boundary/sector/holonomy data, and exact well-founded descent. The read set is validation-complete: it contains every support of every measure term that can change under the write. Overlapping primitive proposals are aggregated by connected conflict component; the restriction-compatible union-collar glue supplies the canonical aggregate payload, while primitive members of that component do not commit independently. On the same fixed-cutoff quantum lift, the terminal expectation functional on each declared physical observable algebra is therefore unique even when microscopic representative lifts differ by gauge labels globally or by sector labels inside one quotient-local glued state. The unqualified coding statement at this stage is only a finite constraint-code statement: the overlap-net codewords are the globally consistent states $C = \Phi^{-1}(0)$. A bare graph does not determine code distance: the same graph can carry constant readouts with distance 1 or repetition constraints with distance $|V|$. Topological-code distance/min-cut, Knill–Laflamme resilience, spectral-gap convergence, BFT wall-clock liveness, and hardware search-work reduction therefore require their own certificates and are not imported by the core overlap graph. Exact descent gives termination. Confluence follows from an independently checked conflict-component diamond on the physical quotient, protected-support and protected-conflict completeness, and repair completeness. The theorem therefore gives one schedule-independent terminal quotient normal form from a fixed initial quotient state. A stronger same-boundary conclusion requires a preserved boundary/sector map and at most one consistent quotient extension in that boundary fiber. The layered functional carrier in the consensus paper proves the finite multi-edge, multi-step witness for $H_B \wedge H_{\mathrm{fib}}$; the functional selected-fiber theorem gives the rooted obstruction-check form by constructing the single candidate from the boundary/root data and then checking all non-tree, sector, and holonomy obstructions. On that branch, multiple same-boundary candidate interiors can be present, inconsistent candidates are eliminated, and all surviving candidates share one quotient normal form. Normal-form hashes are public equality receipts after quotienting, not selectors among physically distinct minimizing endpoints. If accepted schedules terminate at different observer-facing quotient normal forms without a declared holonomy or higher-gauge obstruction, that repair law is outside the OPH consensus theorem. The neutral result of Ref. [1] packages this quantifier split exactly: under observation preservation and completeness, agreement modulo a silent equivalence of normal endpoints reached from all same-boundary sources is equivalent to injectivity of the induced boundary map on the consistent quotient. This cross-source criterion, same-source confluence, weak normalization, and all-schedule liveness are separate obligations. The named layered and functional carriers discharge the cross-source premise only on their stated domains. For separated cofinal refinement systems, the consensus paper adds an inverse-limit bridge: if finite-stage restriction maps commute with the quotient normal-form maps and holonomy maps, and if compatible families are visibly separated on cofinal finite stages, then finite normal forms and holonomy obstructions assemble into unique refinement-limit classes with finite-stage witnesses for nonzero holonomy. It also adds the RG-facing comparison: when a chosen coarse-graining channel shadows the finite-stage normal forms and obstruction maps with declared defects ε^n and ε^h, reconciling first and then coarse-graining gives the same macroscopic law readout as coarse-graining first and then reconciling, up to $\max\{\varepsilon^n, \varepsilon^h\}$; exact naturality is the zero-defect case. Exact normal-form naturality is discharged by a repair-morphism witness, and holonomy naturality by the corresponding cochain map. For distributed implementations, the same paper proves that a worker run presents one finite universe only when it starts from one global carrier and each event projects to a legal monolithic repair path, a physical stutter, or a certified rollback to an earlier committed projection. The required certificate names the global graph, initial state, partition, cut interfaces, observer registry, run/config/code hashes, linearized commits, restart roots, final monolithic normal form, and recomputed readout. Executor order is separate from observer history. Observer-clock naturality additionally requires a history-augmented quotient with semantic event keys independent of worker ID, repair iteration, queue position, timestamp, and retry metadata; a global observer registry descended from local registry groupoids with disjoint patch/cap/future namespaces and explicit lineage arrows; a state-preserving observer-algebra extraction functor from terminal augmented normal forms; a support-visible cap chart map natural with that extraction; and an operational clock instrument whose readout is calibrated up to the declared affine reparameterization and residual bound. The listed objects or their proof-producing finite certificates are prerequisites. In their absence, modular-order/readback naturality is a named gate rather than a consequence of implementation bookkeeping. The same quotient boundary applies to neutral-bulk readout. For each shard s, raw record rows Σ_s must first be quotient by the declared presentation groupoid Γ_s and then sent through the terminal normal-form map n_s, producing

$$X_s = n_s(\Sigma_s/\Gamma_s).$$

Interface maps $\tau_{ts} : U_{st} \to U_{ts}$ between these terminal charts must satisfy the identity, inverse, cocycle, and zero-cycle-holonomy checks before rows from different shards are compared. A global neutral readout is the quotient

$$Q_{\text{vis}} = \left(\bigsqcup_s X_s\right) / \sim_\tau,$$

not the concatenation of shard-local tables. Channel features $F_{s,c}$ descend to Q_{vis} only when their domains and values are transported by the same atlas. For a complete declared channel set $\mathcal{C}_\star$, the neutral distance is

$$d_{\text{neu}}(x, y) = \left(\sum_{c \in \mathcal{C}_\star} w_c\, d_c(F_c(x), F_c(y))^p\right)^{1/p}.$$

Here $p \geq 1$, and $\mathcal{C}_\star$ is finite. An infinite declared channel set is allowed only when the displayed weighted p-sum is finite for every compared pair. Under that finite-distance premise this is stated first as a quotient-visible pseudometric. It becomes a metric only after quotienting zero-distance feature-collision classes or proving that the channels jointly separate points. Pairwise available-channel comparison is explicitly excluded: missingness is handled only by complete cases, a fixed quotient-visible missing symbol, or a train-only imputation map labelled as an imputed-representation metric. The certificate records a finite channel list or a pairwise weighted-ℓ^p summability witness. Presentation changes (gauge representatives, port relabelings, observer order, schedules, and shard partitions) preserve the metric only when they induce a bijection of Q_{vis} together with channel isometries; refinement claims additionally require a cofinally vanishing distance tail modulus. Finite Euclidean claims require the double-centered Gram certificate

$$B = -\frac{1}{2}H(D^{\circ 2})H \succeq 0$$

and noisy runs must report negative spectral mass, rank/effective-rank data, held-out stress, and planted/shuffled controls. Statistical certificates split independent generative shard batches before preprocessing; chart alignment, scaling, imputation, weights, graph construction, dimension choice, and thresholds are train/validation objects, and a test batch, seed, boundary condition, trajectory family, duplicate, or descendant appearing in more than one split blocks the claim. These theorems establish at most a common finite quotient metric or pseudometric. Identifying it with a physical Riemannian bulk or with Lorentzian/Einstein spacetime (via the conditional event-manifold packet of Theorem 4.3e, under its population, chart, cone, and reachability receipts) requires the separate modular/geometric branch below. The same consensus surface is classified explicitly as a finite-state exact theorem package plus this controlled inverse-limit bridge: decidable normal-form computation with the Lyapunov step bound, automatic approximate stability only through the collar-local splice and record controls, a conditional fair-block contraction branch for long-run noisy approximate consensus, and a verified rooted-tree packet-net subdomain. Computational universality for growing patch-net families belongs to the consensus paper's expressive-power boundary and is not a dependency for the recovered-core branches.

Fixed-cutoff packet closure map and invariant simplex. On any declared fixed-cutoff packet quotient Q whose repair relation is terminating and confluent, the normal-form map

$$N : Q \to Q_{\mathrm{nf}}$$

is well-defined and schedule-independent. It induces an affine OPH closure map on the probability simplex over packet states,

$$\mathcal{C}_Q : \Delta(Q) \to \Delta(Q), \qquad \mathcal{C}_Q\left(\sum_{q\in Q} p_q \delta_q\right) = \sum_{q\in Q} p_q \delta_{N(q)}.$$

Since Q is finite, $\Delta(Q)$ is a nonempty compact convex set and $\mathcal{C}_Q$ is a continuous affine self-map. Its image $\Delta(Q_{\mathrm{nf}})$ is invariant, and $\mathcal{C}_Q^2 = \mathcal{C}_Q$. This is an actual closure map on the fixed-cutoff packet-closed quotient branch. It is not the full Appendix/habitat closure-map theorem for arbitrary OPH state-and-law data; the first operation not internalized at that level is the general habitat lift from this finite packet quotient to the full compact-convex observer-supporting sector.

Finite quotient ensemble theorem surface. The closure map above is an idempotent normal-form projector. It sends any input law to terminal quotient normal forms, and every law supported on those normal forms is a fixed point. Hence the physical probability law at regulator r must be supplied as an additional quotient-intrinsic object. Let

$$Q_r = \Sigma_r/\Gamma_r$$

be the finite physical quotient, or let $N_r = n_r(Q_r)$ be the quotient-normal-form subset when only settled configurations are intended to carry probability. The declared OPH quotient ensemble is specified by an intrinsic base weight and action

$$m_r : Q_r \to \mathbb{R}_{>0}, \qquad S_r : Q_r \to \mathbb{R} \cup \{+\infty\},$$

and

$$w_r(q) = m_r(q)e^{-S_r(q)}, \qquad Z_r = \sum_{q\in Q_r} w_r(q), \qquad \mu_r(q) = Z_r^{-1} w_r(q).$$

Equivalently, one may state an intrinsic projective prior ν_r on Q_r and set $\mu_r = (n_r)_\# \nu_r$. The choice of m_r, S_r, or ν_r is load-bearing: uniform quotient counting, uniform representative counting pushed to the quotient, finite-groupoid weights, and trace weights of a declared quotient algebra are generally different.

Finite MaxEnt form. For a finite Q, positive m, and quotient observables $F_1, \dots, F_k$, maximizing

$$\mathcal{H}_m(\nu) = -\sum_q \nu(q) \log \frac{\nu(q)}{m(q)}$$

over the affine constraint surface

$$\sum_q \nu(q) = 1, \qquad \sum_q \nu(q) F_a(q) = c_a$$

has a unique full-support solution, when the feasible set has one, of the form

$$\mu(q) = \frac{m(q) \exp[-\sum_a \theta_a F_a(q)]}{Z(\theta)}.$$

This is the finite-dimensional Lagrange-multiplier argument with strict concavity of $\mathcal{H}_m$; boundary optima obey the same formula after restricting to their support. On a finite noncommutative quotient algebra the corresponding density operator with faithful reference state σ_r is

$$\rho_r = \frac{\exp(\log \sigma_r - \sum_a \theta_a F_{r,a})}{\operatorname{Tr} \exp(\log \sigma_r - \sum_a \theta_a F_{r,a})}.$$

Refinement and projective compatibility. For $s \succeq r$, let $c_{sr} : Q_s \to Q_r$ be the physical coarse-graining map. Exact compatibility of the weighted ensembles is equivalent to the fiber-sum identity

$$\sum_{q' :\, c_{sr}(q') = q} m_s(q') e^{-S_s(q')} = \alpha_{sr} m_r(q) e^{-S_r(q)}$$

for a constant $\alpha_{sr} > 0$ independent of q; necessarily $\alpha_{sr} = Z_s/Z_r$. The proof is just the equality

$$(c_{sr})_\# \mu_s(q) = Z_s^{-1} \sum_{q' :\, c_{sr}(q') = q} m_s(q') e^{-S_s(q')}$$

compared with $Z_r^{-1} m_r(q) e^{-S_r(q)}$. For the selected-prior route, if

$$c_{sr} \circ n_s = n_r \circ c_{sr}, \qquad (c_{sr})_\# \nu_s = \nu_r,$$

then functoriality gives $(c_{sr})_\# (n_s)_\# \nu_s = (n_r)_\# \nu_r$. Along a chain, finite-stage compatibility defines a unique projective-limit probability measure on the cylinder σ-algebra. If the one-step defects are

$$\delta_{k+1,k} = \|(c_{k+1,k})_\# \mu_{k+1} - \mu_k\|_{\mathrm{TV}},$$

then total-variation contraction gives

$$\|(c_{nr})_\# \mu_n - \mu_r\|_{\mathrm{TV}} \le \sum_{k=r}^{n-1} \delta_{k+1,k}.$$

Implementation invariance and representative lifting. If two implementations A, B carry quotient bijections $h_r : Q_r^A \to Q_r^B$ with

$$m_r^B(h_r q) = m_r^A(q), \qquad S_r^B(h_r q) = S_r^A(q), \qquad h_r \circ c_{sr}^A = c_{sr}^B \circ h_s,$$

then $(h_r)_\# \mu_r^A = \mu_r^B$. If an implementation stores representatives $\pi_r : \Sigma_r \to Q_r$, a representative-level law must be a conditional lift

$$\widetilde{\mu}_r(x) = \mu_r(\pi_r x)\kappa_r(x \mid \pi_r x), \qquad \sum_{x:\pi_r(x)=q} \kappa_r(x \mid q) = 1.$$

Then $(\pi_r)_\# \widetilde{\mu}_r = \mu_r$. Uniform representative sampling weights quotient states by orbit size and is physical only if representative counting is the declared base measure. A representative Markov kernel must also be quotient-lumpable: the transition probability to an orbit cannot depend on which representative of the source orbit is used.

Sampler correctness. A production sampler is a separate interface from a deterministic settler:

$$\texttt{settle}(q) \to n_r(q), \qquad \texttt{sample}(\mu_r) \to q.$$

For $w(q) = m(q)e^{-S(q)}$ and a proposal kernel $R(q, q')$ with reciprocal support, the Metropolis–Hastings acceptance rule

$$a(q, q') = \min\left\{1, \frac{w(q')R(q', q)}{w(q)R(q, q')}\right\}$$

gives detailed balance

$$\mu(q)R(q, q')a(q, q') = \mu(q')R(q', q)a(q', q),$$

and hence $\mu P = \mu$. Irreducibility and aperiodicity give uniqueness and convergence on the finite state space. If a repair-informed proposal is $q' = R_{\mathrm{repair}}(q) + \eta$ with density g, the Hastings ratio must include

$$\frac{g(q - R_{\mathrm{repair}}(q'))}{g(q' - R_{\mathrm{repair}}(q))};$$

dropping this term generally changes the stationary law. For a lazy reversible finite chain with spectral gap γ,

$$\|P^n(x, \cdot) - \mu\|_{\mathrm{TV}} \le \frac{1}{2}\sqrt{\mu(x)^{-1} - 1}\,(1 - \gamma)^n,$$

and the usual autocorrelation bound follows from the nontrivial spectral radius. Exact small regulators should therefore enumerate Q_r, w_r, P_r, detailed-balance error, stationarity error, refinement residuals, lumpability/orbit-size checks, and the exact spectral gap.

Gaussian screen continuation. For a real reduced screen field space V_r after background and dipole modes are removed, let K_r be positive definite and impose fixed expected quadratic release energy,

$$\mathbb{E}[q] = 0, \qquad \mathbb{E}[q^\mathsf{T} K_r q] = d_r A_r, \qquad d_r = \dim V_r.$$

The unique maximum-entropy density relative to the declared Euclidean volume is

$$p_r(q) = \frac{\sqrt{\det K_r}}{(2\pi A_r)^{d_r/2}} \exp\left[-\frac{1}{2A_r} q^\mathsf{T} K_r q\right], \qquad \mathrm{Cov}(q) = A_r K_r^{-1}.$$

This is the relative-entropy proof against the displayed Gaussian. If instead every sample obeys $q^\mathsf{T} K_r q = d_r A_r$, the MaxEnt law is uniform on the ellipsoid, not Gaussian. Under a linear coarse field map $C_{sr} : V_s \to V_r$, exact Gaussian refinement is the covariance criterion

$$A_r K_r^{-1} = C_{sr}(A_s K_s^{-1})C_{sr}^\mathsf{T}.$$

For harmonic truncation, choose real modes $Y_{\ell m}$, $2 \le \ell \le L$, set

$$K_L Y_{\ell m} = \kappa_\ell Y_{\ell m}, \qquad \kappa_\ell = [\ell(\ell+1)]^{1+\theta/2},$$

and sample independent coefficients

$$a_{\ell m} \sim \mathcal{N}(0, A/\kappa_\ell).$$

Projection from L' to L then discards high modes and is exactly refinement-compatible. For irregular meshes, covariance should be primary under coarse-graining:

$$\Sigma_r = C_{sr}\Sigma_s C_{sr}^{\mathsf{T}}, \qquad K_r = A_r \Sigma_r^{-1}.$$

Equivalently, when eliminating fine variables from a precision matrix

$$K_s = \begin{pmatrix} K_{cc} & K_{cf} \\ K_{fc} & K_{ff} \end{pmatrix},$$

the coarse precision is the Schur complement

$$K_r = K_{cc} - K_{cf}K_{ff}^{-1}K_{fc},$$

not the principal block K_{cc}.

Vacuum and primordial promotion gates. A stationary sampler is not a physical vacuum: the identity transition kernel leaves every measure stationary. A finite vacuum theorem needs a physical transfer operator. If T_r is real symmetric, positive definite, and strictly positive, Perron–Frobenius gives a simple largest eigenvalue $\lambda_{0,r}$ with positive eigenvector Ω_r. Then

$$H_r = -a_t^{-1}\log(T_r/\lambda_{0,r})$$

is positive semidefinite, Ω_r is the unique zero-energy state, and the Doob kernel

$$P_r(x,y) = \frac{T_r(x,y)\Omega_r(y)}{\lambda_{0,r}\Omega_r(x)}$$

is stochastic and detailed-balanced with

$$\mu_r(x) = \frac{\Omega_r(x)^2}{\sum_z \Omega_r(z)^2}.$$

Continuum promotion further requires reflection positivity or an equivalent reconstruction certificate plus refinement compatibility of the transfer family. A screen covariance alone does not determine an unrestricted primordial curvature spectrum. Even the complete noiseless one-shell sequence for the thin-shell map

$$C_\ell = 4\pi \int_0^\infty \frac{dk}{k}\Delta_\zeta^2(k) j_\ell^2(k\chi_\star)$$

determines the three-dimensional correlation only on $0 \le s \le 2\chi_\star$ and has an infinite-dimensional kernel that persists under positivity. Source-derived uniqueness requires a physical dilation-intertwiner theorem or complete radial cross-covariance tomography. A finite prior gives a conditional continuation. Conventional free-scalar and lattice-gauge baselines may therefore be recorded, but those baselines carry closed OPH-native-vacuum and primordial-promotion receipts unless these gates are supplied.

2 Finite Screen Spectrum Theorem Package

This fragment owns the screen-level spectrum theorem used by the staged cosmology branch. It does not own TT, TE, EE, lensing, likelihoods, or physical CMB promotion. Those remain downstream Boltzmann and data-contract gates.

Definition 2.1 (finite screen regulator). *For regulator r, a finite screen regulator is*

$$\mathfrak{S}_r = (\mathcal{T}_r, M_r, \Gamma_r, \mathcal{Q}_r, C_{sr}, J_{rs}),$$

where $\mathcal{T}_r$ is a finite cellulation of the screen, $M_r \succ 0$ is the area/quadrature mass matrix, Γ_r is the hidden-presentation groupoid, $\mathcal{Q}_r = \Sigma_r/\Gamma_r$ is the physical quotient, and C_{sr}, J_{rs} are the coarse-graining and interpolation maps. Shape regularity and quadrature convergence require

$$f^\top M_r g \to \int_{S^2} fg\, d\Omega$$

with an explicit band-limited error bound

$$\left| f^\top M_r g - \int fg\, d\Omega \right| \le \varepsilon_{M,r}(L) \|f\|_{H^s} \|g\|_{H^s}.$$

Patch count alone is therefore not an angular-resolution certificate.

Definition 2.2 (geometric screen scalar). *Let the quotient-visible collar-volume readout be*

$$J_{X,r}(x) = \lambda_r(x)\sqrt{\det \sigma_{AB,r}(x)} > 0,$$

and let $\bar{J}_{X,r}$ be emitted by an independently defined homogeneous or frozen-background branch, not by a CMB fit. Define

$$q_{0,r} = \frac{1}{3}\log\frac{J_{X,r}}{\bar{J}_{X,r}}.$$

On a certified spherical chart, with $B_r = [1, n_x, n_y, n_z]$,

$$\Pi_{<2,r} = B_r(B_r^\top M_r B_r)^{-1} B_r^\top M_r, \qquad q_r = (I - \Pi_{<2,r})q_{0,r}.$$

On an irregular screen, B_r is replaced by the certified generalized eigenprojector onto the finite $\ell = 0, 1$ subspace. Hard-coded feature z-scores or observer-summary weights are diagnostic proxies, not q_r.

Proposition 2.3 (quotient invariance). *If a recharting U preserves the screen inner product and geometric readout,*

$$U^\top M_r' U = M_r, \qquad q_{0,r}' = U q_{0,r}, \qquad B_r' = U B_r,$$

then $\Pi'_{\ge 2,r} U = U \Pi_{\ge 2,r}$, and $q_r' = U q_r$.

Proof. Substitution gives

$$\begin{aligned} \Pi'_{<2,r} U &= U B_r (B_r^\top U^\top M_r' U B_r)^{-1} B_r^\top U^\top M_r' U \\ &= U B_r (B_r^\top M_r B_r)^{-1} B_r^\top M_r = U \Pi_{<2,r}. \end{aligned}$$

Subtracting from U gives the high-mode identity. □

Definition 2.4 (physical scalar precision and action)**.** *The scalar precision K_r must have a physical origin on $V_r = \operatorname{im}\Pi_{\geq 2,r}$. Allowed branches are:*

$$K_r^{\mathrm{Hess}} = \nabla^2 \Phi_r \Big|_{q=0},$$

for a quotient-visible mismatch or release free energy; a repair Dirichlet form

$$K_r^{\mathrm{diss}} = H_r - T_r^{\mathsf{T}} H_r T_r, \qquad T_r^{\mathsf{T}} H_r T_r \preceq H_r;$$

or the reversible generator form

$$K_r^{\mathrm{Dir}} = -\frac{1}{2}(L_r + L_r^{\dagger_H}).$$

A raw nonsymmetric repair matrix or caller-supplied κ is not a precision operator. The absolute normalization must be fixed independently, for example by

$$K_0 Y_{\ell m} = \ell(\ell+1) Y_{\ell m}$$

on the $\theta = 0$ branch. Once normalized,

$$S_{\mathrm{scr},r}[q] = \frac{1}{2A_{q,r}} \langle q, K_r q\rangle_r = \frac{1}{2A_{q,r}} q^{\mathsf{T}} M_r K_r q.$$

Theorem 2.5 (finite MaxEnt screen covariance)**.** *Let $K_r e_i = \kappa_i e_i$ with an M_r-orthonormal basis of V_r, $\kappa_i > 0$, and $d_r = \dim V_r$. Among continuous densities in coefficients $q = \sum_i a_i e_i$ with $\mathbb{E}[a_i] = 0$ and*

$$\frac{1}{2}\mathbb{E}\langle q, K_r q\rangle_r = E_{q,r}^{\mathrm{src}},$$

the unique entropy maximizer is

$$d\mu_r(q) = Z_r^{-1} \exp\left[-\frac{1}{2A_{q,r}}\langle q, K_r q\rangle_r\right] dq, \qquad A_{q,r} = \frac{2E_{q,r}^{\mathrm{src}}}{d_r},$$

and

$$\mathbb{E}[a_i a_j] = \delta_{ij}\frac{A_{q,r}}{\kappa_i}.$$

Proof. The Euler–Lagrange equation for entropy with normalization and quadratic-energy constraints gives a Gaussian density with inverse temperature β_r. The expected energy is

$$E_{q,r}^{\mathrm{src}} = \frac{1}{2}\sum_i \kappa_i \frac{1}{\beta_r \kappa_i} = \frac{d_r}{2\beta_r}.$$

Thus $A_{q,r} = \beta_r^{-1} = 2E_{q,r}^{\mathrm{src}}/d_r$. Strict concavity of entropy gives uniqueness. □

Remark 2.6 (discrete quotient branch)**.** *For a discrete finite quotient with intrinsic base weight $m_r(q)$, the exact Gibbs law is*

$$\mu_r(q) = Z_r^{-1} m_r(q) e^{-\beta_r E_r(q)}.$$

A Gaussian statement on that branch requires a separate Laplace or central-limit theorem with controlled higher-order terms.

Definition 2.7 (source release energy)**.** *Let m_r^0 be the tracially pointed base weight on the finite physical collar quotient. Let $F_{r,a}$ be a finite ledger of primitive collar observables: scalar occupancy, protected-center presence, released volume, conserved source charges, and the release clock. The ledger excludes screen energy, sky spectra, CMB residuals, likelihoods, fitted amplitudes, and measurement-calibrated proxies. For source-emitted constraint values $c_{r,a}^{\mathrm{col}}$, the selected release law is*

$$\nu_r^{\mathrm{rel}}(x) = \frac{m_r^0(x)\exp[-\sum_a \lambda_{r,a}F_{r,a}(x)]}{Z_r}, \qquad \mathbb{E}_{\nu_r^{\mathrm{rel}}}F_{r,a} = c_{r,a}^{\mathrm{col}}.$$

Strict concavity of relative entropy gives uniqueness on full support. A boundary optimum uses the same form on its certified support. The constraint ledger, base weight, source DAG, multipliers, support, and refinement maps belong to the receipt.

For this independently selected quotient ensemble on released collar normal forms, define

$$E_{q,r}^{\mathrm{src}} = \frac{1}{2}\int \langle q_r(x), K_r q_r(x)\rangle_r \, d\nu_r^{\mathrm{rel}}(x).$$

The receipt must expose ν_r^{rel}, the base measure, K_r, d_r, $E_{q,r}^{\mathrm{src}}$, $A_{q,r}$, no-observation ancestry, and the same operator normalization used in the action. MaxEnt alone does not determine amplitude.

Proposition 2.8 (microscopic-law necessity)**.** *For fixed positive K_r, the Gaussian family $Z(A)^{-1}\exp[-\langle q, K_r q\rangle/(2A)]$ exists for every $A > 0$. The quadratic MaxEnt form therefore leaves one positive scale free. The source release law in Definition 2.7, or an equivalent source-only vacuum law, supplies the energy that fixes $A_{q,r} = 2E_{q,r}^{\mathrm{src}}/d_r$.*

Theorem 2.9 (rotational screen spectrum)**.** *If the continuum precision K_θ commutes with the $SO(3)$ action on scalar functions, then by Schur's lemma each $\mathcal{H}_\ell$ is an eigenspace. For the exact conformal-shell family,*

$$\Lambda_\ell(\theta) = \frac{\Gamma(\ell + 2 + \theta/2)}{\Gamma(\ell - \theta/2)}, \qquad \ell \geq 2.$$

Then $\Lambda_\ell(0) = \ell(\ell+1)$, and the MaxEnt covariance gives

$$C_\ell^q = A_q \, \frac{\Gamma(\ell - \theta/2)}{\Gamma(\ell + 2 + \theta/2)}.$$

Consequently $\mathcal{D}_\ell^q = \ell(\ell+1)C_\ell^q/(2\pi) \sim (A_q/2\pi)\ell^{-\theta}$.

Definition 2.10 (source-derived conformal precision)**.** *Let $L_r \succeq 0$ be the normalized detailed-balanced scalar collar Dirichlet operator on V_r, with continuum limit $-\Delta_{S^2}$. Set*

$$B_r = (L_r + \tfrac{1}{4}I)^{1/2}, \qquad K_{\theta,r} = \frac{\Gamma(B_r + \frac{3}{2} + \frac{\theta}{2})}{\Gamma(B_r - \frac{1}{2} - \frac{\theta}{2})}.$$

For $-2 < \theta < 4$ on the retained $\ell \geq 2$ branch, this operator is positive. The gamma recurrence gives $K_{0,r} = L_r$ exactly. The finite receipt carries detailed balance, positivity, the $\ell(\ell+1)$ normalization, anisotropy splitting, and refinement residuals.

Theorem 2.11 (finite refinement error)**.** *Suppose that for $2 \leq \ell \leq L$*

$$\left|\frac{\lambda_{\ell m,r}}{\Lambda_\ell(\theta)} - 1\right| \leq \varepsilon_{K,r}(L), \qquad \left|\frac{A_{q,r}}{A_q} - 1\right| \leq \varepsilon_{A,r}, \qquad \varepsilon_{K,r} < 1.$$

Then $C_{\ell m,r}^q = A_{q,r}/\lambda_{\ell m,r}$ satisfies

$$\left|\frac{C_{\ell m,r}^q}{A_q/\Lambda_\ell(\theta)} - 1\right| \leq \frac{\varepsilon_{A,r} + \varepsilon_{K,r}}{1 - \varepsilon_{K,r}}.$$

Proof. Write $A_{q,r} = A_q(1 + a_r)$ and $\lambda_{\ell m,r} = \Lambda_\ell(1 + b_{\ell m,r})$. Then

$$\frac{C^q_{\ell m,r}}{A_q/\Lambda_\ell} = \frac{1 + a_r}{1 + b_{\ell m,r}},$$

and the displayed bound follows from $|a_r| \le \varepsilon_A$, $|b_{\ell m,r}| \le \varepsilon_K$. $\square$

Theorem 2.12 (refinement-semigroup tilt)**.** *Let R_b be the scalar refinement map for scale ratio $b > 1$. If one isolated scalar covariance mode has positive eigenvalue $\lambda(b)$, $\lambda(b_1 b_2) = \lambda(b_1)\lambda(b_2)$, and λ is continuous, then there is a unique real θ with*

$$\lambda(b) = b^{-\theta}, \qquad \theta = -\frac{\log \lambda(b)}{\log b}.$$

Proof. Set $g(t) = -\log \lambda(e^t)$. The semigroup law gives $g(t + s) = g(t) + g(s)$. Continuity makes $g(t) = \theta t$, hence the result. $\square$

Definition 2.13 (edge-center reserve generator receipt)**.** *Write $s = \log b$ for logarithmic refinement thickness. Let $u_{\mathrm{full}}(s)$ be the scalar-conditioned covariance-survival cocycle across a full oriented collar, and let $u_q(s)$ be its source-facing half-collar restriction. The receipt requires*

$$u(s + t) = u(s)u(t), \qquad u(0) = 1,$$

strong continuity at zero, the source-derived full-collar density

$$-u'_{\mathrm{full}}(0) = \frac{P_\star}{24},$$

and the orientation-reversal coarea identity

$$-u'_q(0) = \frac{1}{2}[-u'_{\mathrm{full}}(0)] = \frac{P_\star}{48}.$$

These are infinitesimal generator statements. A one-step survival probability has exponent $-\log u_q(\log b)/\log b$.

Theorem 2.14 (edge-center tilt and repair-clock reconciliation)**.** *Under Definition 2.13,*

$$u_q(s) = e^{-\theta s}, \qquad \theta = \frac{P_\star}{48}, \qquad n_s = 1 - \frac{P_\star}{48}.$$

In the coordinate $\theta = \kappa_{\mathrm{rep}}(P_\star - \varphi)$, the same branch has

$$\kappa^{\mathrm{edge}}_{\mathrm{rep}} = \frac{P_\star}{48(P_\star - \varphi)}.$$

The value e is a separate diagnostic hypothesis unless an additional identity equates it with this source-derived coordinate.

Proof. For $g(s) = -\log u_q(s)$, the cocycle law gives the continuous Cauchy equation. Hence $g(s) = \theta s$. Differentiation at zero and the half-collar identity give $\theta = P_\star/48$. The formula for $\kappa^{\mathrm{edge}}_{\mathrm{rep}}$ is algebraic. $\square$

Theorem 2.15 (homogeneous radial source family)**.** *Assume source-stress closure, a single clock, freezeout, multicenter consistency, and translation/rotation covariance. Let* $C_\zeta = M_{\Delta^2_\zeta}$ *on the common physical* $d\ln k$ *mode basis. If a scale-natural source embedding transports refinement to physical dilation and its finite operator residual converges to*

$$D_s^{-1} C_\zeta D_s = e^{-\theta s} C_\zeta, \qquad (D_s f)(k, \hat{k}) = f(e^{-s}k, \hat{k}),$$

then every positive measurable representative satisfies

$$\Delta^2_\zeta(k) = A_\zeta (k/k_\star)^{-\theta}$$

almost everywhere. Continuity gives the pointwise identity. A single shell has an infinite-dimensional radial kernel, including positive ambiguities; complete radial cross-covariances give the independent spherical-Hankel tomography route.

Proof. Conjugation of the multiplication operator gives $\Delta^2_\zeta(e^s k) = e^{-\theta s}\Delta^2_\zeta(k)$ almost everywhere. In logarithmic coordinates, rational-translation invariance on a common full-measure set forces $e^{\theta t}\Delta^2_\zeta(e^t)$ to be constant almost everywhere. The one-shell and tomography statements follow from the correlation-restriction and spherical-Hankel theorems in the primordial bridge packet. □

Theorem 2.16 (thin-shell power-law lift)**.** *Assume* $q(\hat{n}) = Z_\star \Pi_{\ell\geq 2}\zeta_\star(R_\star \hat{n})$ *and*

$$\Delta^2_\zeta(k) = A_\zeta (k/k_\star)^{-\theta}.$$

Then

$$C^q_\ell = 4\pi Z_\star^2 \int d\ln k\, \Delta^2_\zeta(k) j_\ell^2(kR_\star),$$

and, for $-2 < \theta < 4$,

$$A_q = \pi^{3/2} Z_\star^2 A_\zeta (k_\star R_\star)^\theta \frac{\Gamma(1+\theta/2)}{\Gamma(3/2+\theta/2)}.$$

Thus

$$A_\zeta = \frac{A_q}{\pi^{3/2} Z_\star^2 (k_\star R_\star)^\theta} \frac{\Gamma(3/2+\theta/2)}{\Gamma(1+\theta/2)}.$$

Proposition 2.17 (finite-window bound)**.** *Let* $\Psi_\ell(k) = \int dr\, W(r) j_\ell(kr)$, $W \geq 0$, $\int W dr = 1$, *with mean* $R_\star$ *and variance* σ_R^2. *If*

$$\delta_\ell(k) = \frac{k^2\sigma_R^2}{2} \sup_{r\in \operatorname{supp} W} |j_\ell''(kr)|,$$

then

$$|C^q_{\ell,W} - C^q_{\ell,\mathrm{shell}}| \leq 4\pi Z_\star^2 \int d\ln k\, \Delta^2_\zeta(k) \left[2\delta_\ell(k) + \delta_\ell(k)^2\right].$$

Proposition 2.18 (radial non-identifiability)**.** *For a finite radial basis,* $C = Ap$. *If* $A \in \mathbb{R}^{N_\ell \times N_k}$ *has rank* r, *then* $\dim\ker A = N_k - r$, *and every* $p + v$ *with* $v \in \ker A$ *gives the same screen spectrum. Radial promotion therefore requires either a source theorem restricting* p, *or a declared prior with a published null basis, resolution kernels, positivity checks, and prior-sensitivity report.*

Theorem 2.19 (conditional source-spectrum promotion)**.** *One hash-locked source DAG may emit*

$$(q_r, S_{\mathrm{scr},r}, \theta, A_{q,r}, C^q_{\ell,r}, A_{\zeta,r}, \Delta^2_{\zeta,r})$$

as a conditional primordial packet only when the geometric-scalar, collar-law, release-energy, reserve-generator, conformal-precision, refinement, source-stress, single-clock, freezeout, adiabaticity, isocurvature, phase-coherence, thin-shell, radial-null, finite-window, and forward- residual receipts pass on that DAG. Any measurement, likelihood, fitted parameter, or measurement-calibrated ancestor fails the source packet. Physical temperature and polarization spectra require the independent Boltzmann, recombination, nuisance, covariance, and likelihood gates.

Target 2.20 (screen-spectrum receipt set)**.** *A concrete source-derived spectrum requires the following receipt families before primordial promotion:*

1. *geometry, scalar quotient, low-mode projector, scalar precision, and operator normalization;*
2. *quotient ensemble selection, scalar release energy, and MaxEnt screen covariance;*
3. *the infinitesimal reserve generator, scalar refinement tilt, and angular screen spectrum with a finite error budget;*
4. *source-stress, clock, freezeout, radial-window, radial-null, forward-residual, and transfer firewall receipts.*

The theorem and algorithm packet is complete. Construction of one finite source DAG that passes this receipt set is work in progress.

2.0.1 Source-only primordial bridge theorem

The finite screen covariance theorem is a screen theorem. To read it as a source-only primordial curvature spectrum one must add the following bridge objects and receipts. The dependency chain is

$$\boxed{\begin{array}{c}\mathrm{nf}(\mathcal{Q}_r) \to (\chi_r, u_r, h_r) \to \mathsf{BackgroundCurvatureStatus}_r \to (g, T_I) \to q_\zeta = \zeta_\rho \\ \to \text{single clock} \to \text{freeze-out} \to \text{coherent growing mode} \to \Delta_\zeta^2(k).\end{array}}$$

No successful fit of a screen C_ℓ or a TT spectrum may replace one of these hypotheses.

Definition (background curvature branch label). The primordial bridge carries a finite record

$$\mathsf{BackgroundCurvatureStatus}_r = (\mathsf{BranchType}, \kappa, I_K, I_{\Omega_K}, \mathsf{Basis}, \mathsf{Top})$$

where

$$\mathsf{BranchType} \in \{\textsc{FlatExact}, \textsc{FlatAssumed}, \textsc{OpenCurved}, \textsc{ClosedCurved}, \textsc{Unresolved}\}.$$

FlatExact is allowed only after a direct theorem or conditional CMH receipt. FlatAssumed may run flat kernels but its outputs remain assumption-conditioned. Unresolved blocks physical promotion. The H^3 observer-facing atlas is not this record; this record comes from the clocked FLRW slice and its spatial Levi–Civita holonomy or declared branch input.

Definition (source stress readout). On a reconstructed branch with relational metric g_{ab}, let

$$S_{\mathrm{src}} = \sum_I S_I[g, \Psi_I], \qquad T^I_{ab} = -\frac{2}{\sqrt{-g}}\frac{\delta S_I}{\delta g^{ab}}, \qquad T^{\mathrm{tot}}_{ab} = \sum_I T^I_{ab},$$

where I ranges over every active sector, including ordinary matter, radiation, neutrinos, OPH anomaly or repair sectors, and any auxiliary source kept in the run. At finite regulator this is a quotient-visible map

$$\mathsf{StressRead}_r : \mathrm{nf}(\mathcal{Q}_r) \longrightarrow \{T^I_{r,ab}\}_I.$$

All source stress components must be read from the same global normal form as the collar scalar.

Source-stress closure theorem. Assume each S_I is relationally diffeomorphism invariant, the source equations hold, sector exchanges satisfy

$$\nabla_a T_I^{ab} = Q_I^b, \qquad \sum_I Q_I^b = 0,$$

and $\mathsf{StressRead}_r$ descends to the physical quotient and converges. Then

$$\nabla_a T_{\mathrm{tot}}^{ab} = 0.$$

On the recovered Einstein branch,

$$G_{ab} + \Lambda g_{ab} = 8\pi G\, T_{ab}^{\mathrm{tot}}.$$

The evidence bundle must also report the non-circular comparison

$$E_T^{ab} = T_{\mathrm{src}}^{ab} - \frac{G^{ab} + \Lambda g^{ab}}{8\pi G};$$

using the geometric right-hand side alone is a geometry diagnostic, not a source-only primordial receipt.

Total-energy frame and curvature covector. When $T^a{}_b$ has a unique future timelike unit eigenvector u^a, define

$$T^a{}_b u^b = -\rho u^a, \qquad h_{ab} = g_{ab} + u_a u_b,$$

and decompose

$$T^{ab} = \rho u^a u^b + p h^{ab} + \pi^{ab}, \qquad u_a \pi^{ab} = 0, \qquad \pi^a{}_a = 0.$$

Let $\Theta = \nabla_a u^a$ and let σ_{ab} be the shear. Total energy conservation gives

$$\dot\rho + \Theta(\rho + p) + \sigma_{ab}\pi^{ab} = 0, \qquad \dot f := u^a \nabla_a f.$$

On an expanding branch, define

$$p_{\mathrm{eff}} = p + \frac{\sigma_{ab}\pi^{ab}}{\Theta};$$

then $\dot\rho + \Theta(\rho + p_{\mathrm{eff}}) = 0$.

Theorem (exact total-curvature evolution). Let

$$\dot\alpha = \Theta/3, \qquad \zeta_a = \nabla_a \alpha - \frac{\dot\alpha}{\dot\rho}\nabla_a \rho,$$

and

$$\Gamma_a^{\mathrm{eff}} = \nabla_a p_{\mathrm{eff}} - \frac{\dot p_{\mathrm{eff}}}{\dot\rho}\nabla_a \rho.$$

Then, exactly and without a long-wavelength approximation,

$$\mathcal{L}_u \zeta_a = -\frac{\Theta}{3(\rho + p_{\mathrm{eff}})}\Gamma_a^{\mathrm{eff}}.$$

Hence a barotropic effective source, $p_{\mathrm{eff}} = p_{\mathrm{eff}}(\rho)$, conserves ζ_a.

Collar scalar corollary. For relational rods Lie-dragged by the total-energy flow, the collar-volume Jacobian satisfies

$$J_X = \lambda\sqrt{\det \sigma_{AB}}, \qquad \frac{d}{d\tau}\frac{1}{3}\log J_X = \frac{\Theta}{3}.$$

Relative to a frozen background,

$$q_\zeta = \frac{1}{3}\log\frac{J_X}{\bar{J}_X}.$$

On the uniform-total-density cut,

$$q_\zeta = \zeta_\rho + \text{constant}.$$

After removing the monopole and dipole,

$$q = \Pi_{\ell\geq 2} q_\zeta = \Pi_{\ell\geq 2}\zeta_\rho, \qquad Z_q = 1.$$

This is the source-only quotient scalar on the OPH screen; by itself it has type ScreenCurvature, not PrimordialCurvature.

Definition (rank-one single-clock branch). Let

$$\mathcal{F}(x) = (\rho, p_{\text{eff}}, \{\rho_I, p_I, Q_I\}_{I=1}^N)$$

be the scalar source map on the source moduli space. A branch is rank-one single-clock when $\operatorname{rank} d\mathcal{F} = 1$, $d\rho \neq 0$, every level set of ρ is connected, all sector velocities agree with u^a, and orthogonal momentum transfer vanishes or is interval-bounded. At finite cutoff this requires a quotient-visible clock χ_r and a factorization

$$\mathcal{F}_r = \widehat{\mathcal{F}}_r \circ \chi_r.$$

Theorem (single-clock adiabaticity). On a rank-one single-clock branch,

$$p_{\text{eff}} = p_{\text{eff}}(\rho), \qquad S_a^{IJ} = 3(\zeta_a^I - \zeta_a^J) = 0,$$

and the intrinsic nonadiabatic pressure and transfer covectors obey

$$\Gamma_a^I = 0, \qquad \Xi_a^I = 0.$$

Therefore $\mathcal{L}_u\zeta_a = 0$. The proof is the chain rule through the single clock: every source scalar is a function of the same χ, and connected ρ-fibers remove residual branch labels.

Theorem (entropy repair gap). Let the linearized scalar repair/evolution block be

$$\delta Y_{n+1} = L_r \delta Y_n + \eta_n,$$

and let $P_\perp$ project away gauge directions, constrained zero modes, and the adiabatic growing direction. If there is a positive metric H_r such that

$$\|P_\perp L_r P_\perp\|_{H_r} \leq e^{-\gamma_r} < 1,$$

then

$$\|P_\perp \delta Y_n\|_{H_r} \leq e^{-n\gamma_r}\|P_\perp \delta Y_0\|_{H_r} + \sum_{j=0}^{n-1} e^{-(n-1-j)\gamma_r}\|P_\perp \eta_j\|_{H_r}.$$

In the unforced case the orthogonal entropy, decaying, and repair scalar modes vanish. Eigenvalues alone do not certify this theorem for nonnormal repair matrices; a norm bound or a discrete Lyapunov inequality

$$L_\perp^\top H L_\perp \preceq e^{-2\gamma} H$$

is required.

Theorem (growing-mode coherence). Let $X(k)$ collect the scalar initial-condition variables for curvature, radiation, baryons, CDM/anomaly, neutrinos, and repair-sector scalars, and let $v_+(k)$ be the normalized adiabatic growing vector. If the single-clock and repair-gap theorems hold through readout and no independent stochastic source acts afterward, then

$$X(k) = v_+(k)\zeta(k), \qquad \mathcal{P}_X(k) = \mathcal{P}_\zeta(k)v_+(k)v_+(k)^\dagger.$$

Thus rank $\mathcal{P}_X(k) = 1$, source-side isocurvature fractions vanish, and all species share a deterministic growing-mode phase. Finite runs must report λ_2/λ_1, $\beta_{\rm iso}(k)$, decaying-mode fraction, and phase-convention defects as source diagnostics.

Theorem (screen-to-primordial bridge). Assume the geometric record and causal-area metric receipts, the collar identity $q_\zeta = \zeta_\rho$, source-stress closure, rank-one normal form, positive repair gap, freeze-out, adiabatic growing-mode, isocurvature, and phase-coherence receipts all pass. Also assume the physical scale-bridge theorem: a source embedding with `source_embedding_hash`, a physical mode basis with `physical_mode_basis_id`, a calibrated comoving coordinate χ, mode-normalization and density-of-states receipts, and a no-post-hoc-calibration receipt, and a $\mathsf{BackgroundCurvatureStatus}_r$ that is not UNRESOLVED. Then the low-mode-removed OPH screen scalar is the source-only primordial curvature field restricted to the declared shell or radial window:

$$q(\hat{n}) = \Pi_{\ell_{\rm src}\geq 2} \int d\chi\, W(\chi)\zeta_\star(\chi\hat{n}).$$

For a thin shell, $q(\hat{n}) = \Pi_{\ell_{\rm src}\geq 2}\zeta_\star(R_\star\hat{n})$. The forward spectrum is exactly

$$C_\ell^q = 4\pi \int_0^\infty d\ln k\, \Delta_\zeta^2(k)|\Psi_\ell^{\mathsf{BranchType}}(k)|^2,$$

$$\Psi_\ell^{\mathsf{BranchType}}(k) = \begin{cases} \int dr\, W(r) j_\ell(kr), & \mathsf{BranchType} = \textsc{FlatExact} \text{ or } \textsc{FlatAssumed}, \\ \int dr\, W(r)\Phi_{\ell k}^{(-)}(r), & \mathsf{BranchType} = \textsc{OpenCurved}, \\ \int dr\, W(r)\Phi_{\ell k}^{(+)}(r), & \mathsf{BranchType} = \textsc{ClosedCurved}, \end{cases}$$

and for a thin shell

$$C_\ell^q = 4\pi \int_0^\infty d\ln k\, \Delta_\zeta^2(k) \left|\Psi_\ell^{\mathsf{BranchType}}(k; R_\star)\right|^2.$$

At finite regulator the continuum integral is represented by the finite spectral measure

$$d\nu_r(k) = \sum_j w_{rj}\delta(k - k_{rj}),$$

so the actual finite evidence receipt is

$$C_{\ell_{\rm src},r}^q = 4\pi \sum_j w_{rj}^{\log k}\, \Delta_{\zeta,r}^2(k_{rj})|\Psi_{\ell_{\rm src},r}^{\mathsf{BranchType}}(k_{rj})|^2.$$

Every radial node, window, and mode bin is bound to the geometry hash, source-embedding hash, scale-certificate hash, boundary-condition hash, and mode-basis identifier, and the evidence bundle must emit radial kernel geometry compatibility receipt. The observed multipole $L_{\rm CMB}$ is not $\ell_{\rm src}$; it is the output index of the line-of-sight transfer solver. Cap eigenmodes, inverse cap-opening angles, and relabeled unit strings remain diagnostic by the dimensional, angular, cap-mode, and unit-string no-go propositions.

Physical scale dependency. The k_{rj} grid in the preceding equation is not a free producer field. It is consumed only after

physical spatial k receipt

recomputes the comoving k-intervals from the physical geometry, scale certificate, scale factor, finite eigenvalues, normalization, and refinement residuals. The radial kernel and source-screen projection additionally require

screen to physical k association receipt.

The source covariance is a finite positive operator on the common physical mode projectors:

$$C_{\zeta,r} \succeq 0.$$

Homogeneity and isotropy on the safe band require

$$P_I C_{\zeta,r} P_I = \mathcal{P}_\zeta(k_I) P_I + E_I$$

with a declared residual bound, together with an off-diagonal leakage report

$$\sum_{I \neq J} |P_I C_{\zeta,r} P_J|^2.$$

The same physical projector family is used by primordial covariance and anomaly response:

common primordial anomaly mode basis receipt.

Mode freezeout and common initial surface are separate gates:

physical mode freezeout map receipt, *physical freezeout surface receipt.*

The latter contains initial data and normal derivatives on one common spacelike surface.

Radial prior, null space, and residual receipt. The map above is a forward projection, not a free inversion. With finite bins p_j for Δ_ζ^2,

$$A_{\ell j} = 4\pi \int_{\text{bin } j} d\ln k \, b_j(k) |\Psi_\ell(k)|^2, \qquad C^q = Ap.$$

The evidence bundle must declare the basis, support, prior p_0, Q, positivity conditions, singular spectrum of $AQ^{-1/2}$, effective rank, null basis, resolution kernels, prior sensitivity, and every forward residual

$$r_\ell = C_\ell^q - \widehat{C}_\ell.$$

A finite prior, regularizer, positivity constraint, or square discretization selects one representative from a nonunique shell equivalence class. Its typed output is ConditionalRadialContinuation. A source-derived primordial packet requires the common source, scale, mode, null-space, positivity, ancestry, and non-fitting forward-residual receipts together with one exact uniqueness branch from the radial theorem packet below. If the geometry packet is imported, the strongest possible claim tier is *conditional physical*; *source-native physical* requires the native $\mathsf{CosmoGeomRead}_r$ theorem.

Complete radial-lift theorem packet. The screen-to-radial map is a forward restriction with a nonunique unrestricted inverse. For a flat source window W and geometric normalization Z_q,

$$C_{\ell,W}^q = 4\pi Z_q^2 \int_0^\infty d\ln k \, \Delta_\zeta^2(k) |\Psi_{\ell,W}(k)|^2, \qquad \Psi_{\ell,W}(k) = \int W(dr) j_\ell(kr).$$

The physical radial packet must select exactly one of the two uniqueness branches below; a finite prior continuation has the type ConditionalRadialContinuation.

Theorem (one-shell correlation restriction and exact no-go). For a thin shell of radius R, the complete angular covariance satisfies

$$K_R(\mu) = Z_q^2 \xi_\zeta\left(R\sqrt{2-2\mu}\right) = \sum_{\ell \geq 0} \frac{2\ell+1}{4\pi} C^q_{\ell,R} P_\ell(\mu).$$

Hence even the noiseless sequence $\{C^q_{\ell,R}\}_{\ell\geq 0}$ determines the three-dimensional correlation only on $0 \leq s \leq 2R$. The full one-shell operator has an infinite-dimensional kernel, including positivity-preserving ambiguities. Explicitly, for any nonzero $g \in C_c^\infty((2R,\infty))$,

$$h_g(k) = \frac{2k^3}{\pi} \int_0^\infty r^2 g(r) j_0(kr)\, dr$$

has correlation perturbation $\xi_{h_g} = g$ and therefore changes no shell multipole. Choosing $B = |h_g| + w$ with any strictly positive integrable w makes B and $B + h_g$ two distinct positive spectra with the same complete shell covariance.

Proof. The spherical-Bessel addition theorem gives the displayed restriction identity. The map $\mu \mapsto R\sqrt{2-2\mu}$ covers exactly $[0, 2R]$. Radial Fourier inversion gives $\xi_{h_g} = g$ and is injective, so the kernel contains an infinite-dimensional copy of $C_c^\infty((2R,\infty))$. The positive pair differs by the same null element. $\square$

Corollary (exact low-mode-removed kernel). For a continuous signed correlation perturbation ξ_h, invisibility to every shell multipole $\ell \geq 0$ is equivalent to $\xi_h(s) = 0$ on $0 \leq s \leq 2R$. Invisibility to every retained OPH multipole $\ell \geq 2$ is equivalent to

$$\xi_h(s) = a + b\left(1 - \frac{s^2}{2R^2}\right), \qquad 0 \leq s \leq 2R,$$

for constants a, b. Legendre completeness gives the first statement. Removing $\ell = 0, 1$ leaves the span of P_0 and P_1 invisible, which gives the second.

Producer theorem (source-refinement orbit to physical dilation). Let $\mathcal{H}_r^{\mathrm{src}}$ be the scale-labelled scalar mode space generated by the cofinal refinement orbit of the released screen, including its scale label. A single angular cut is insufficient. Let U_r map the orbit unitarily onto the physical rank-one adiabatic source subspace. If

$$D_{s,r} U_r = U_r R_{s,r}, \qquad C_{\zeta,r} = U_r C_{\mathrm{src},r} U_r^*,$$

then

$$D_{s,r}^{-1} C_{\zeta,r} D_{s,r} - u_r(s) C_{\zeta,r} = U_r \left(R_{s,r}^{-1} C_{\mathrm{src},r} R_{s,r} - u_r(s) C_{\mathrm{src},r}\right) U_r^*.$$

Thus the source and physical operator residuals have equal norm. If the finite covariances converge strongly with a uniform norm bound, the finite dilation maps and their inverses converge strongly, $u_r(s) \to u(s)$, and this residual tends to zero, then $D_s^{-1} C_\zeta D_s = u(s) C_\zeta$ in the continuum. The commuting square states that source refinement followed by the physical embedding agrees with physical embedding followed by wavelength dilation, and it transfers the source covariance error exactly to the physical covariance.

Proof. Conjugate the source covariance by U_r and use the commuting square; unitary invariance gives the residual identity. For the limit, add and subtract the finite conjugated expression. Strong convergence, uniform boundedness, scalar convergence, and the norm-residual hypothesis send the three resulting terms to zero. $\square$

Theorem (source-dilation uniqueness). Let $C_\zeta = M_{\Delta_\zeta^2}$ on the common physical $L^2(d\ln k\, d\Omega_k)$ mode basis. Define

$$(D_s f)(k,\hat{k}) = f(e^{-s}k,\hat{k}).$$

If the source refinement packet proves

$$D_s^{-1} C_\zeta D_s = u_q(s) C_\zeta, \qquad u_q(s+t) = u_q(s)u_q(t), \qquad -\frac{d}{ds}\ln u_q(s)\Big|_{s=0} = \theta,$$

then $u_q(s) = e^{-\theta s}$ and there is $A > 0$ such that

$$\Delta_\zeta^2(k) = Ak^{-\theta} \quad \text{almost everywhere.}$$

For a continuous source representative and pivot $k_\star$,

$$\boxed{\Delta_\zeta^2(k) = A_\zeta (k/k_\star)^{-\theta}}, \qquad A_\zeta = \Delta_\zeta^2(k_\star).$$

Proof. Strong continuity solves the multiplicative Cauchy equation for u_q. Conjugation of the multiplication operator gives $\Delta_\zeta^2(e^s k) = e^{-\theta s}\Delta_\zeta^2(k)$ almost everywhere for each s. In logarithmic coordinates, $e^{\theta t}\Delta_\zeta^2(e^t)$ is invariant under all rational translations on one common full-measure set; bounded-transform convolution makes it constant almost everywhere. Continuity upgrades the identity to every k. □

Theorem (exact flat thin-shell lift). On the source-dilation branch, for $-2 < \theta < 4$ and $\ell \geq 2$,

$$\int_0^\infty \frac{dx}{x} x^{-\theta} j_\ell^2(x) = \frac{\sqrt{\pi}}{4} \frac{\Gamma(1+\theta/2)}{\Gamma(3/2+\theta/2)} \frac{\Gamma(\ell-\theta/2)}{\Gamma(\ell+2+\theta/2)},$$

so

$$C_{\ell,R_\star}^q = A_q \frac{\Gamma(\ell-\theta/2)}{\Gamma(\ell+2+\theta/2)},$$

with the complete amplitude conversion

$$\boxed{A_\zeta = \frac{A_q}{\pi^{3/2} Z_q^2 (k_\star R_\star)^\theta} \frac{\Gamma(3/2+\theta/2)}{\Gamma(1+\theta/2)}.}$$

Proof. Insert the source power law into the exact forward map, set $x = kR_\star$, and apply the Weber–Schafheitlin square integral. The convergence strip follows from $j_\ell(x) \sim x^\ell/(2\ell+1)!!$ at zero and $j_\ell(x) = O(x^{-1})$ at infinity. □

At the source-native pivot $k_\star R_\star = 1$ with $Z_q = 1$ and the source-law tilt θ, the conversion factor evaluates to approximately 0.16080676. It is an arithmetic consequence of the symbolic relation and supplies no source amplitude by itself.

Corollary (source-family injectivity and multipole consistency). Fix $\theta, Z_q, k_\star$, and a source window for which the forward coefficient is positive and finite. The forward map is injective in A_ζ on the family $\Delta_\zeta^2 = A_\zeta(k/k_\star)^{-\theta}$. For each retained multipole,

$$A_\zeta^{(\ell,W)} = \frac{C_{\ell,W}^q}{4\pi Z_q^2 k_\star^\theta \int d\ln k\, k^{-\theta} |\Psi_{\ell,W}(k)|^2}.$$

Every $A_\zeta^{(\ell,W)}$ must agree. A disagreement falsifies the declared window/dilation packet and cannot be repaired with an ℓ-dependent amplitude. Once the source law fixes θ, the screen has one amplitude left to determine, and the unused multipoles become consistency checks.

Theorem (finite-window bound). For $0 < \theta < 2\ell$, define

$$I_\ell(\theta) = \int d\ln x\, x^{-\theta} j_\ell^2(x), \quad J_\ell(\theta) = I_\ell(\theta-2) - \left[\ell(\ell+1) - \frac{\theta(\theta+1)}{2}\right] I_\ell(\theta),$$

$$\eta_{\ell,W} = \frac{2\sqrt{J_\ell(\theta)}}{\theta} \int W(dr)|r^{\theta/2} - R_\star^{\theta/2}|.$$

Then

$$|C^q_{\ell,W} - C^q_{\ell,R_\star}| \le 4\pi Z_q^2 A_\zeta k_\star^\theta \eta_{\ell,W} \left(2R_\star^{\theta/2}\sqrt{I_\ell(\theta)} + \eta_{\ell,W}\right).$$

Proof. In the Hilbert norm $\|f\|_\theta^2 = \int d\ln k\, k^{-\theta}|f(k)|^2$, $\|j_\ell(kr)\|_\theta = r^{\theta/2}\sqrt{I_\ell}$ and $\|\partial_r j_\ell(kr)\|_\theta = r^{\theta/2-1}\sqrt{J_\ell}$. The fundamental theorem of calculus, Minkowski, and $|\|f\|^2 - \|g\|^2| \le \|f-g\|(\|f\|+\|g\|)$ give the bound. ☐

Theorem (finite singular-value and prior continuation). For a declared finite radial basis, write $C = Ap$. If $A \in \mathbb{R}^{N_\ell \times N_k}$ has rank r, then $\dim\ker A = N_k - r$, and every exact solution is $p = p_{\text{part}} + V_0 a$ for a right-null basis V_0. Given a prior center p_0 and $Q \succ 0$, the unique minimizer of $\frac{1}{2}(p-p_0)^\mathsf{T} Q(p-p_0)$ subject to $Ap = C$ is

$$\boxed{p_* = p_0 + Q^{-1}A^\mathsf{T}(AQ^{-1}A^\mathsf{T})^+(C - Ap_0).}$$

The run publishes the singular values, rank threshold, right-null basis, resolution operator, unresolved projector, positivity active set, and prior sensitivity. The output type is ConditionalRadialContinuation.

Theorem (unrestricted tomographic uniqueness). For

$$C_\ell(r, r') = \frac{2}{\pi}\int_0^\infty k^2 P_\zeta(k) j_\ell(kr) j_\ell(kr')\, dk,$$

the covariance operator on $L^2(r^2 dr)$ obeys

$$Q_\ell = \mathcal{H}_\ell^{-1} M_{P_\zeta} \mathcal{H}_\ell.$$

Thus a complete radial cross-covariance operator determines P_ζ uniquely almost everywhere. Fixing one radius r_0, the family $C_\ell(r, r_0)$ known for all r suffices: its spherical-Hankel transform is $\sqrt{2/\pi}\, P_\zeta(k) j_\ell(kr_0)$, and the Bessel zeros are discrete. A finite collection of auto-spectra does not supply this tomography theorem.

Proof. Apply spherical-Hankel inversion in the first radial variable. Unitary conjugation exposes the multiplication operator; equal multiplication operators have equal multipliers almost everywhere. ☐

Theorem (branch-typed radial boundary). The ordinary spherical-Bessel and gamma-function formulas apply to FlatExact and FlatAssumed branches. Open and closed spatial branches use their declared hyperspherical eigenfunctions and spectral measures. An unresolved curvature branch blocks physical radial promotion, and the flat thin-shell amplitude formula cannot certify a curved branch.

Promotion rule. A source-derived primordial packet requires all source-stress, single-clock, freezeout, growing-mode, phase, physical-scale, source-amplitude, null-space, positivity, and non-fitting forward-residual receipts, together with either a radial-dilation intertwiner certificate or a radial-tomography certificate. A finite inverse selected before evaluation is a conditional radial continuation. It does not promote temperature or polarization spectra to physical observables.

Corollary (transfer firewall). Radial promotion changes no temperature, polarization, lensing, recombination, foreground, nuisance, covariance, or likelihood status. Those outputs require the independent transfer and likelihood system.

total stress closure receipt	$T_I^{ab}, Q_I^a, E_T, \nabla_a T^{ab}$
single clock normal form receipt	$\chi_r, \operatorname{rank} d\mathcal{F},$ connected fibers
entropy repair gap receipt	$L_\perp, H, \gamma,$ forcing/refinement bounds
curvature evolution receipt	$\zeta_a, \Gamma_a^{\mathrm{eff}},$ evolution residual, freeze bound
adiabatic mode receipt	$S_{IJ}, \Gamma_I, \Xi_I,$ growing-mode projection
isocurvature bound report	$\beta_{\mathrm{iso}}(k)$ and source covariance
primordial phase coherence receipt	$\lambda_2/\lambda_1,$ phase and decaying-mode defects
screen to radial lift receipt	$W(r), R_\star, \Psi_\ell(k),$ source hashes
radial null space report receipt	singular values, null basis, resolution kernels
forward projection residual receipt	$r_\ell,$ aggregate norms, quadrature and solver errors.

Source-provenance and pooled-reducer firewall. Let the source artifact DAG have parent-to-child edges and let M be the set of measurement, residual, likelihood, posterior, fitted-parameter, or measurement-calibrated nodes. Define

$$\tau(v) = \mathbf{1}_{v\in M} \vee \bigvee_{p\in \mathrm{Parents}(v)} \tau(p).$$

Then $\tau(v) = 1$ exactly when v has a forbidden measurement ancestor. The proof is induction over a topological ordering of the DAG. Every promoted source node for $\eta_R, \Gamma_{\mathrm{rec}}, A_\zeta, q_{\mathrm{IR}}, \ell_{\mathrm{IR}}$, $B_A(k,a)$, $\rho_A(a)$, and N_{CRC} must have $\tau = 0$. Unknown source metadata fails closed.

For nonlinear estimates, shards may not average nonlinear quantities. Each row must emit additive sufficient statistics in a commutative monoid $(\mathcal{M}, \oplus)$, for example

$$\sum w_i, \qquad \sum w_i x_i, \qquad \sum w_i x_i x_i^{\mathsf{T}}$$

or, for weighted response regression,

$$X^{\mathsf{T}} W X, \qquad X^{\mathsf{T}} W y, \qquad y^{\mathsf{T}} W y.$$

Associativity, commutativity, complete coverage, and duplicate removal make the pooled statistic partition-invariant; the deterministic estimator is then evaluated only after pooling.

Three-dimensional homogeneous lift. A rotationally invariant screen covariance about one observer is not a statistically homogeneous three-dimensional primordial field by itself. A frozen lift

$$\mathcal{L}_r : V_r^{\mathrm{screen}} \to \mathcal{H}_r^{(3)}$$

must define

$$\Sigma_r^{(3)} = \mathcal{L}_r \Sigma_r^{\mathrm{screen}} \mathcal{L}_r^*$$

and certify positivity, refinement convergence, rotation covariance, translation covariance or an equivalent multicenter consistency theorem, and explicit null-space treatment. If the limiting covariance commutes with translations and rotations, Fourier diagonalization gives

$$\langle \zeta_a(\mathbf{k}) \zeta_b^*(\mathbf{k}') \rangle = (2\pi)^3 \delta^{(3)}(\mathbf{k} - \mathbf{k}') P_{ab}(|\mathbf{k}|),$$

with $P_{ab}(k)$ positive semidefinite.

Transfer-ready initial data. The source-to-transfer map must emit a complete initial state

$$Y(k, \eta_i) = \mathcal{M}_{\rm init}(k; \mathfrak{B}_{\rm src})\xi(k),$$

including metric, photon, polarization, baryon, neutrino, anomaly, recipient, and auxiliary response variables. Near a singular early-time point, regular Frobenius modes solve

$$Y' = \left(\frac{M_{-1}}{\eta} + M_0 + O(\eta)\right)Y, \qquad Y = \eta^s(v_0 + v_1\eta + \cdots), \quad M_{-1}v_0 = sv_0.$$

The branch must classify regular adiabatic growing modes, allowed isocurvature modes, decaying modes, gauge modes, and singular interaction modes. On the adiabatic branch,

$$S_{IJ} = 3(\zeta_I - \zeta_J) = 0$$

for all active sectors. The finite evidence receipt must bound the initial Hamiltonian and momentum constraints and the propagated residual

$$\epsilon_{\rm constraint} = \sup_{k,\eta} \|\mathcal{C}(k,\eta)Y(k,\eta)\|.$$

Receipt contract. Every ensemble run must record its ensemble id, claim tier, regulator, representative schema, gauge action, quotient canonicalizer, base measure definition, action coefficients, coarse maps, zero-mode projector, amplitude convention, sampler kernel, partition-invariant random-event schema, smoothing policy, source provenance, and explicit nonclaims. The seed belongs to the run receipt rather than the ensemble definition. Claim tiers are:

$$\begin{aligned}
E0 &: \text{ seed noise, proposal noise, repair jitter,}\\
E1 &: \text{ conventional reference ensemble,}\\
E2 &: \text{ OPH-native quotient ensemble,}\\
E3 &: \text{ OPH vacuum,}\\
E4 &: \text{ OPH primordial field,}\\
E5 &: \text{ observable cosmological prediction.}
\end{aligned}$$

Smoothing must preserve raw coefficients, raw spectra, smoothing kernels, smoothed coefficients, smoothed spectra, and hashes of each stage; it is not part of S_r unless explicitly declared. For bounded coarse observables O, if the implemented law satisfies $\|\widehat{\mu}_s - \mu_s\|_{\rm TV} \leq \epsilon_{\rm samp}$ and the refinement defects sum to $\epsilon_{\rm ref}$, then

$$\left|\mathbb{E}_{\widehat{\mu}_s}[O \circ c_{sr}] - \mathbb{E}_{\mu_r}[O]\right| \leq 2\|O\|_\infty(\epsilon_{\rm samp} + \epsilon_{\rm ref}),$$

which is the finite evidence accuracy statement attached to this theorem surface.

Theorem 2.3 (EC from regulated overlap gluing). For a collar B_δ around a cap boundary Σ, there is a canonical decomposition

$$H_{B_\delta} = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{G_\Sigma} = \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

and the sector-preserving collar algebra

$$A_{\rm EC}(B_\delta) := \bigoplus_\alpha (B(H_{b_L^\alpha}) \otimes B(H_{b_R^\alpha}))$$

with

$$Z(A_{\mathrm{EC}}(B_\delta)) = \bigoplus_\alpha \mathbb{C}\,\mathbf{1}_\alpha,$$

such that $\mathcal{A}(A_\delta B_\delta)$ acts only on $H_{b_L^\alpha}$ and $\mathcal{A}(B_\delta D_\delta)$ acts only on $H_{b_R^\alpha}$ within each block.

Proof. Split the collar into half-collars B_L and B_R meeting on $\Sigma = \partial C$. By the realized regulator presentation above, the physical collar Hilbert space is the diagonal invariant subspace $(\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{G_\Sigma}$. Decompose each side into irreps:

$$\tilde{\mathcal{H}}_{B_L} = \bigoplus_\alpha (V_\alpha \otimes H_{b_L^\alpha}), \qquad \tilde{\mathcal{H}}_{B_R} = \bigoplus_\beta (V_\beta^* \otimes H_{b_R^\beta}).$$

Then

$$\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R} = \bigoplus_{\alpha,\beta} (V_\alpha \otimes V_\beta^*) \otimes (H_{b_L^\alpha} \otimes H_{b_R^\beta}).$$

By Schur's lemma,

$$(V_\alpha \otimes V_\beta^*)^{G_\Sigma} \cong \begin{cases} \mathbb{C}, & \alpha = \beta, \\ 0, & \alpha \neq \beta. \end{cases}$$

Therefore the invariant subspace is

$$H_{B_\delta} = \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

as claimed. The induced sector-preserving collar algebra is

$$A_{\mathrm{EC}}(B_\delta) = \bigoplus_\alpha (B(H_{b_L^\alpha}) \otimes B(H_{b_R^\alpha})),$$

so the center is generated by the block projectors. Adjacent region algebras act on the left or right factor only because the gluing action is supported on Σ. QED.

Remark. On the central-defect subbranch, replace G_Σ by its central extension. The sector label α then ranges over irreps of the extension; the decomposition is unchanged.

We refer to the decomposition in Theorem 2.3 as **edge-center completion (EC)**.

Exact Markov route. If $I(A_\delta : D_\delta \mid B_\delta) = 0$, the Hayden–Jozsa–Petz–Winter (HJPW) structure theorem yields a blockwise factorization of the state over *some* direct-sum/tensor decomposition of $\mathcal{H}_{B_\delta}$; that decomposition is state-dependent, and HJPW does not by itself identify it with the preselected EC factors b_L^α, b_R^α of Theorem 2.3. The exact identity used for literal Markov-modular equalities below is the **EC-aligned** normal form

$$\rho_{A_\delta B_\delta D_\delta} = \bigoplus_\alpha p_\alpha (\rho_{A_\delta b_L^\alpha} \otimes \rho_{b_R^\alpha D_\delta}), \qquad I_\omega(A_\delta : D_\delta \mid B_\delta) = 0,$$

and the identification of the HJPW factors with the EC factors is a genuine additional state condition, the **Markov-split alignment hypothesis (MSA)**. It is not implied by exact Markovity: with four qubits A, B_L, B_R, D and $\rho = \Phi_{AB_R} \otimes \Phi_{B_L D}$ (Bell pairs across the cut, one center block), one has $I(A : D \mid B) = 0$ but $I(A : B_R) = 2\ln 2$, so the state is exactly Markov yet not of the displayed EC-aligned form, which would force $I(A : B_R) = 0$. All literal Markov-modular equalities below therefore assume exact Markovity *plus* MSA, or an explicitly stated idealized limit that supplies both; The spacetime and Einstein paper states MSA formally and records this counterexample. MSA is checkable state by state: The spacetime and Einstein paper proves it equivalent to a splitting of $\log\rho$ into commuting one-sided terms supported on Ab_L^α and $b_R^\alpha D$, to the existence of a ρ-preserving conditional expectation onto either one-sided edge algebra (a Takesaki commuting square over the collar center [35]), and to blockwise vanishing one-sided mutual information $I(Ab_L^\alpha : b_R^\alpha D) = 0$. On the MaxEnt branch this reduces the derivation of MSA to a sharp condition on the effective Hamiltonian: every cross-cut coupling must act through the central sector label, as it does in lattice-gauge-type regulators where the interface energy is a function of the conserved flux. On that declared **central-interface branch** The spacetime and Einstein paper derives both MSA and exact collar Markovianity as theorems (descent of one-sided invariants and boundary charges through the edge-center quotient), and it adopts the central-interface normal form as an explicit axiom-level branch clause stated with its MaxEnt axiom, so the literal Markov-modular identities hold on the declared branch without a separate state hypothesis. The clause is independent of the repair/consensus axioms: The spacetime and Einstein paper exhibits a finite regulator package that satisfies overlap consistency, schedule-independent transactional repair, and an exactly closed refinement-channel family while its retained constraint family carries a boundary-invariant noncentral cross-cut coupling, so overlap-consistent repair does not force the clause and it stays a permanent declared branch input. Off the declared branch MSA is carried as an explicit hypothesis.

Interpreting collar refinement along the rate-controlled family stated in Axiom 3, Theorem 2.3 supplies the kinematic edge-center decomposition. Exact Markovity is one idealized route. The following lemma and theorem provide the quantitative decay route used when the manuscript keeps the approximation explicit.

Lemma 2.6 (MaxEnt with local constraints implies local Gibbs form). Under the local MaxEnt branch of Axiom 3, with its homogeneous global-sum constraints $\langle\sum_x O_a(x)\rangle = C_a$, and the finite-dimensional regulator realization above, the MaxEnt state has the Gibbs form

$$\omega = \frac{e^{-H_{\mathrm{eff}}}}{\operatorname{Tr} e^{-H_{\mathrm{eff}}}}, \qquad H_{\mathrm{eff}} = \sum_x \sum_a \lambda_a O_a(x) + H_{\mathrm{sec}},$$

where the sum runs over UV cells x and density labels a, with one cell-independent multiplier λ_a per label. For the finite-range collar theorem, H_{sec} is either a sum of uniformly bounded finite-range charge densities or is central and fixed on the chosen superselection sector. A genuinely nonlocal noncentral term falls outside the theorem. Under this condition the effective Hamiltonian is finite-range in UV graph units, and the multiplier count matches the constraint count N_{con} plus the number of retained sector terms on every lattice.

Proof. On a finite-dimensional algebra, the unique state maximizing $S(\rho) = -\operatorname{Tr}(\rho\log\rho)$ subject to linear constraints $\operatorname{Tr}(\rho O_i) = c_i$ is given by Lagrange multipliers:

$$\rho = \frac{e^{-\sum_i \lambda_i O_i}}{\operatorname{Tr} e^{-\sum_i \lambda_i O_i}}.$$

Strict concavity of von Neumann entropy ensures uniqueness. Here the constrained operators are the N_{con} global sums $O_a := \sum_x O_a(x)$ plus the retained sector terms, so the exponent has one shared multiplier per density label. MaxEnt gives the Gibbs form; the displayed locality restriction on H_{sec} is separate and load-bearing. Had the constraints instead been imposed per cell, $\langle O_a(x)\rangle = c_a(x)$, the multipliers would be cell-dependent fields $\lambda_a(x)$ and the family would grow with the lattice; that enlarged family belongs to the non-equilibrium approximation regime of the clarification above and is not the branch used in the refinement argument. QED.

Definition 2.6a (closure defect and induced refinement map). Fix successive regulator scales ℓ (finer) and L (coarser), let $\{\omega_L(\lambda')\}$ be the homogeneous global-sum exponential family at scale L, and let $\Phi_{\ell\to L}$ be an admissible coarse-graining channel. For fine-scale multipliers λ set $\sigma := \Phi_{\ell\to L}(\omega_\ell(\lambda))$ and define the **closure defect**

$$\varepsilon_{\ell\to L}(\lambda) := \inf_{\lambda'} D(\sigma \,\|\, \omega_L(\lambda')),$$

with D the quantum relative entropy. When the infimum is attained at a unique λ^*, set $R_{\ell\to L}(\lambda) := \lambda^*$.

Lemma 2.6b (I-projection residual bound). If the constrained operators at scale L together with the identity are linearly independent, then $\lambda' \mapsto D(\sigma\|\omega_L(\lambda'))$ is smooth and strictly convex; for faithful σ the minimizer $\lambda^* = R_{\ell\to L}(\lambda)$ exists, is unique, and is characterized by moment matching $\langle S_c\rangle_{\omega_L(\lambda^*)} = \langle S_c\rangle_\sigma$ for every constrained operator S_c; and

$$\left\|\sigma - \omega_L(R_{\ell\to L}(\lambda))\right\|_1 \le \sqrt{2\,\varepsilon_{\ell\to L}(\lambda)},$$

so the family reproduces every bounded expectation up to $\|B\|\sqrt{2\varepsilon_{\ell\to L}(\lambda)}$. The refinement-closure clause of Axiom 3 is exactly the statement $\varepsilon_{\ell\to L}(\lambda) = 0$ along the realized branch, in which case $\sigma = \omega_L(R_{\ell\to L}(\lambda))$ and the induced refinement map is this moment-matching I-projection.

Proof. Writing $\log\omega_L(\lambda') = -\sum_c \lambda'_c S_c - \log Z_L(\lambda')$, the objective is $-S(\sigma) + \sum_c \lambda'_c\langle S_c\rangle_\sigma + \log Z_L(\lambda')$, and the Hessian of $\log Z_L$ is the Duhamel (Kubo–Mori) covariance matrix of the constrained operators, positive definite under the stated independence; this gives strict convexity, uniqueness, and the moment-matching first-order condition, while attainment for faithful σ is standard finite-dimensional exponential-family duality [33]. The displayed residual bound is the quantum Pinsker inequality $D(\rho\|\tau) \ge \frac{1}{2}\|\rho-\tau\|_1^2$ [34] evaluated at λ^*, and $\varepsilon_{\ell\to L}(\lambda) = 0$ forces $\sigma = \omega_L(\lambda^*)$ by strict positivity of relative entropy off the diagonal. QED.

Acceptance check. A finite two-lattice test implementing this package (the constraint/multiplier count against the displayed N_{con}-dimensional family on successive lattices, the moment-matching projection, the residual bound, a generic decimation channel with strictly positive closure defect, and the exactly closed transverse-field subfamily where $R_{\ell\to L}$ acts as the identity) is included with the repository sources.

Strong conditional mixing premise. Let G_ℓ be the UV cell graph. Assume local dimension at most q, interaction degree at most Δ_0, interaction diameter at most r_0 graph layers, and

$$\sup_x \sum_{X\ni x} \|\beta\Phi_\ell(X)\| \le J_0$$

with constants independent of the cut and refinement stage. Let m_ℓ be the resolved collar width in graph layers and set $\delta = m_\ell \ell$. For the faithful Gibbs state, define

$$\mathbf{J}_\rho(A : D \mid B) := \log \rho_{ABD} + \log \rho_B - \log \rho_{AB} - \log \rho_{BD}.$$

The strong mixing premise is a boundary-anchor expansion

$$\mathbf{J}_\rho(A_\delta : D_\delta \mid B_\delta) = \sum_{z \in \partial^{\mathrm{UV}}_{r_0} C} E_{z,\ell,\delta}, \qquad \|E_{z,\ell,\delta}\|_\infty \le \kappa e^{-(m_\ell - r_0)/\zeta},$$

with κ, ζ uniform in the cut, boundary condition, and stage. This conditional or matrix mixing condition is stronger than ordinary two-point exponential clustering. Local Gibbs form, a Lieb–Robinson bound, or a Hamiltonian spectral gap does not imply it for a general noncommuting Gibbs state.

Theorem 2.5 (finite-range conditional Gibbs mixing gives collar recovery). Under Lemma 2.6 and the strong conditional mixing premise,

$$\boxed{I_\omega(A_\delta : D_\delta \mid B_\delta) \le \kappa |\partial C|_{\mathrm{UV}} e^{-(m_\ell - r_0)/\zeta} \le c |\partial C|_{\mathrm{UV}} e^{-\delta/\xi_\ell}}$$

with

$$\xi_\ell = \zeta \ell, \qquad c = \kappa e^{r_0/\zeta}.$$

Proof. Embedding the marginal logarithms in the tripartite algebra gives

$$I(A : D \mid B)_\rho = \mathrm{Tr}[\rho_{ABD} \mathbf{J}_\rho(A : D \mid B)] \,.$$

The boundary expansion and trace-norm duality bound this expectation by the sum of the boundary term norms. The second form follows from $\delta = m_\ell \ell$. QED.

For one uniform family, the sharp sufficient limit condition is

$$\frac{\delta}{\xi_\ell} - \log |\partial C|_{\mathrm{UV}} \longrightarrow +\infty.$$

If

$$|\partial C|_{\mathrm{UV}} \le a(\ell_0/\ell)^p, \qquad \delta(\ell) \ge \zeta(p + \eta)\ell \log(\ell_0/\ell), \quad \eta > 0,$$

then $\delta \to 0$, $\delta/\ell \to \infty$, and

$$I(A_\delta : D_\delta \mid B_\delta) \le ca(\ell/\ell_0)^\eta \longrightarrow 0.$$

The ratio $\delta/\ell \to \infty$ alone is insufficient. For $\ell_n = \ell_0 e^{-n^2}$, $\delta_n = \zeta n \ell_n$, and $|\partial C_n|_{\mathrm{UV}} = e^{n^2}$, that ratio diverges while the upper envelope grows as $e^{n^2 - n}$.

A finite receipt records the interaction range and norm bound, treatment of global sector terms, boundary-cell count, graph separation, regional CMI in nats, matrix-defect norm, predeclared κ, ζ, log-envelope slack, Fawzi–Renner error, held-out cuts, and the rate margin $\delta/(\zeta\ell) - \log |\partial C|_{\mathrm{UV}}$. Fitted constants and local random-triplet CMI remain finite proxies. They do not certify a cofinal limit.

This bound is the quantitative hinge for constructive gluing.

2.0.2 Concrete UV realization: quantum link models

Section 2.3 internalizes the regulator package: the fixed-cutoff type-I algebra and boundary fixed-point structure are the realized form of the screen net plus overlap gluing. Quantum link models are included here only as an explicit microscopic example of that structure, not as a separate axiom layer.

UV regulator. Triangulate S^2 at scale $\ell_{\rm UV}$, giving vertices v, oriented links ℓ, and plaquettes p. Refinement corresponds to $\ell_{\rm UV} \to 0$ with increasing lattice size.

Degrees of freedom. Attach to every oriented link ℓ a **finite-dimensional** Hilbert space $\mathcal{H}_\ell$. In ordinary Wilson lattice gauge theory, the continuum/refinement-limit edge description is modeled by $L^2(G)$ (infinite-dimensional for continuous G), but the microscopic OPH regulator is instead a **quantum link model** with finite-dimensional link Hilbert spaces that preserve gauge symmetry in operator form [52]. Optionally attach matter Hilbert spaces $\mathcal{H}_v$ at vertices. Then:

$$\tilde{\mathcal{H}}_{\rm total} = \bigotimes_\ell \mathcal{H}_\ell \otimes \bigotimes_v \mathcal{H}_v,$$

finite-dimensional on any finite lattice. This is a concrete realization of the extended type-I presentation used in Section 2.3.

Boundary gluing as Gauss-law invariants. Define a local gauge transformation group G_v at each vertex v acting on incident links (and matter at v). Physical states satisfy:

$$|\psi\rangle \in \mathcal{H}_{\rm phys} \quad \Longleftrightarrow \quad U(g_v)|\psi\rangle = |\psi\rangle \;\; \forall\, v,\, g_v \in G_v.$$

Equivalently: $\mathcal{H}_{\rm phys} = \tilde{\mathcal{H}}_{\rm total}^{\prod_v G_v}$.

Region algebras. For any region $R \subset S^2$, define an extended Hilbert space $\tilde{\mathcal{H}}_R$ from the links/vertices in R. The **boundary gauge group** $G_{\partial R}$ acts on the cut degrees of freedom (the "half-links" ending on ∂R). Define:

$$\mathcal{A}_{\rm inv}(R) = \mathcal{B}(\tilde{\mathcal{H}}_R)^{G_{\partial R}}.$$

This is the concrete quantum-link instance of the lifted fixed-point presentation used in Section 2.3. The same gauge constraints then induce the sector-preserving EC algebra on collars.

Why EC follows immediately. Take a cap C and a collar B_δ around ∂C. Because the *only* coupling between inside and outside is through the boundary gauge constraint, the collar Hilbert space decomposes into superselection blocks labeled by boundary irreps:

$$\mathcal{H}_{B_\delta} \cong \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

with center generated by the projectors P_α. This is precisely the Schur-lemma mechanism of Theorem 2.3. The labels α are the familiar edge-mode / electric-flux labels appearing whenever one factorizes gauge theories across an entangling cut [53]. Exact Markovity depends on the state and is not forced by the decomposition alone; nor does exact Markovity align the HJPW factors with these edge-mode labels, which is the separate Markov-split alignment hypothesis of Section 2.3.

Dynamics and MaxEnt. The natural Hamiltonian is a 2+1D lattice gauge Hamiltonian on the screen worldvolume: plaquette ("magnetic") terms, electric terms on links, vertex Gauss terms as constraints, plus local matter couplings. In quantum link form this is finite-dimensional per link while behaving like gauge theory in the continuum limit. Then the MaxEnt assumption becomes concrete: the MaxEnt state is a Gibbs state $\rho \propto e^{-\sum_i \lambda_i O_i}$ with quasi-local O_i, precisely the local-Gibbs regime.

Geometry and G. This microphysics naturally supplies the emergent geometric objects:

- **Edge entropy / area operator:** $L_C = \sum_\alpha (\log d_\alpha) P_\alpha$ becomes "log of boundary irrep dimension" in the gauge link model.
- **Newton constant G:** the conversion factor between edge entropy density per boundary UV cell and macroscopic geometric area.

Thus area is an operator living in the center of the boundary algebra, because in gauge systems the center is where the cut labels live.

Scope of this example. The quantum-link realization makes the fixed-cutoff matrix/fixed-point bookkeeping concrete. Modular flow on caps becomes geometric conformal dilation with the 2π KMS normalization through the support-visible BW scaling theorem on the geometric subnet. Viable architectures for this include holographic quantum error-correcting codes [54] and quantum double / string-net Hamiltonians [55], but any use of QECC distance, min-cut resilience, or Knill–Laflamme correction requires the corresponding code subspace, logical-operator, error-family, and recovery certificate.

2.0.3 Conformal-modular fixed-point microphysics

On the local finite-constraint MaxEnt branch of the third axiom, the logarithm of the selected state is a quasi-local UV generator, and the refinement-stable branch lies inside one common finite-dimensional multiplier family. At each regulator stage the patch and cap algebras are type-I matrix algebras. The finite regulator class is not refinement-closed, so the scaling-limit observer algebra may leave that class. The fixed-cutoff analysis proves a local-Gibbs/refinement-stable branch with propagation and endpoint-Lipschitz control. The geometric cap pair with geometric modular flow is obtained in the support-visible sense: the target algebra is the extracted geometric subnet, the support-visible modular matrix elements converge under regularization, and weak-$*$/GNS extraction plus support-readable modular covariance and ordered cut-pair rigidity gives the geometric cap automorphism.

Local finite-constraint MaxEnt branch. The constraint family $\mathcal{C}$ is generated by finitely many gauge-invariant local densities $\{O_a(x)\}$ of UV range $O(\ell_{\mathrm{UV}})$, constrained through their homogeneous global sums $\langle \sum_x O_a(x) \rangle = C_a$, with the same finite label set retained under refinement by the refinement-closure clause of the third OPH axiom. This clause restates that axiom at regulator level and adds no postulate beyond it; the closure clause itself is the substantive assumption.

Theorem 2.6 (Local constraints imply a local-Gibbs form). If the MaxEnt constraints are expectations of finitely many quasi-local operators $\{O_a\}$ with bounded support size at scale ℓ_{UV}, then the entropy maximizer is

$$\omega \propto \exp\left(-\sum_a \lambda_a O_a \right),$$

so the MaxEnt generator $H_{\mathrm{MaxEnt}} = -\log \omega$ is a UV-range quasi-local sum. This is exactly the local-Gibbs form used later.

Proof. Standard exponential-family result: maximum entropy subject to linear constraints $\langle O_a \rangle = c_a$ yields the Gibbs state with Lagrange multipliers λ_a. QED.

Derived propagation control on the same branch. Because H_{MaxEnt} is a finite-range or quasi-local sum on a finite type-I regulator net, standard Lieb–Robinson estimates [15] apply to the automorphism group it generates, or to any branch generator lying in the same bounded-support algebraic closure. Thus there is a finite propagation velocity v_{LR} and constants C, ξ such that for local observables A_X, B_Y,

$$\|[\tau_t(A_X), B_Y]\| \le C\, \|A_X\|\, \|B_Y\|\, \min(|X|,|Y|)\, e^{-(d(X,Y)-v_{\mathrm{LR}}|t|)/\xi}.$$

Because changing a bounded interval endpoint only adds or removes an $O(|\Delta v|)$ collar of local terms once the central endpoint piece is separated off, the same local branch also gives bounded-interval endpoint-Lipschitz matrix elements,

$$|\langle \psi, (K[I'] - K[I])\phi\rangle| \le C_{\psi,\phi,I_{\max}}\, |I'\Delta I|,$$

used later in the null-modular bridge. So quasi-local propagation and endpoint-Lipschitz control are not external regularity selectors; they are the local-constraint MaxEnt branch written in dynamical form.

Refinement-stable multiplier branch. By the refinement-closure clause of the third OPH axiom (a substantive renormalization condition, since coarse-graining a finite-range Gibbs family generically generates interactions outside any fixed finite density list), the same finite constraint family is preserved under coarse-graining, and the regulator states lie in one common finite-dimensional multiplier family rather than in unrelated state spaces at different cutoffs. The “refinement-stable” language used later therefore means persistence along the stable or fixed branch of this multiplier family. This state-side notion is enough to compare realized states across cutoffs. FixedCutoffBosonicSectorCategory constructs the fixed-stage tensor-generated category and its forgetful fiber. RefinementFunctorAndFiberDescent constructs the refinement pullbacks and compatible fibers only on a cofinal tail carrying the compact-gauge refinement receipt. The cofinal compact-gauge witness uses the same receipt. The separate finite Standard Model route derives its response from incidence and target-blind readback. Under the declared fermionic Spin category, the finite certificate derives the charge-conjugate rank-15 projector pair, hypercharge lattice, common $\mathbb{Z}_6$ kernel, and maximal faithful matter image. MAR enters the generation and family economy, no-extra-sector, charged-lepton, and D10 branches. The third OPH axiom does not supply physical matter typing, global-form selection, the sector-colimit data, or laboratory current attachment.

Scaling limit and algebraic type. The regulator presentation gives a family of finite type-I algebras

$$\mathcal{A}_\ell(C) \cong \mathcal{B}(\mathcal{H}_{C,\ell}).$$

A scaling limit of this family need not be type I. In the continuum-QFT case of interest one expects the local limit algebra $\mathcal{A}_\infty(C)$ to be non-type-I, typically type III. Accordingly the fundamental modular datum in the limit is the automorphism group of the pair $(\mathcal{A}_\infty(C), \omega_\infty^C)$, not a density matrix inside $\mathcal{A}_\infty(C)$.

Corollary 2.6 (Geometric modular action on paired-certified caps). For any support-visible extracted scaling-limit geometric cap pair $(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C})$ emitted on a branch carrying both the finite cap-normal certificate and the independently complete same-branch MGNS-1 package, let $\alpha_{\lambda_{\widehat{C}}(s)}$ be the automorphism induced by the independently normalized BW-framed cap flow. Then

$$\sigma_t^{\omega_\infty^{\mathrm{geo},C}} = \alpha_{\lambda_{\widehat{C}}(2\pi t)}.$$

If the limit cap algebra happens to be type I, this may be written as

$$K_C = 2\pi B_C + Z_C, \qquad Z_C \in Z(\mathcal{A}(C))_{\mathrm{sa}}.$$

In the generic continuum case, the automorphism identity is the complete statement.

Proof. This is The spacetime and Einstein paper's finite cap-net BW theorem applied to the extracted geometric cap pair. The finite cap-normal input supplies BW framing, support-order faithfulness, held-out oriented cross-ratio rigidity, the geometric support flow, independently normalized geometric 2π-KMS comparison, and wrong-normalization controls. The separate MGNS-1 input supplies the complete algebra-state reference tower, cap-family-uniform mixed-GNS convergence, comparison maps, and modular support covariance. QED.

Alternative derivation via net regularity. The same modular-covariance property can also be read off from a scaling-limit support map when the net satisfies the outer-regularity / minimal-support condition used later. This clarifies how the geometric labeling of the support-visible extracted limit net is read.

(NR) Outer regularity / minimal support. For any operator O, the intersection of all connected regions P with $O \in \mathcal{A}(P)$ is again a connected region, denoted $\operatorname{supp}(O)$.

Proposition 2.7 (Modular covariance from net regularity). Under (NR), define for any region $R \subset C$

$$f_t^C(R) := \bigcup_{O \in \mathcal{A}(R)} \operatorname{supp}\Big(\sigma_t^{\omega,C}(O)\Big).$$

Then $\sigma_t^{\omega,C}(\mathcal{A}(R)) = \mathcal{A}(f_t^C(R))$, which is exactly the desired modular-covariance property.

Proof. Since $\sigma_t^{\omega,C}$ is an automorphism of $\mathcal{A}(C)$, and (NR) allows one to read support from the net labeling, the map $R \mapsto f_t^C(R)$ is well-defined and consistent. QED.

Null-surface modular structure. On the same support-visible scaling branch and extracted geometric-subnet branch, the null-surface modular machinery narrows as follows:

- **Fixed-cutoff null-strip bridge.** The null-strip package proves transferred cut-center data, the theorem-local inherited left/right strip-split condition needed for the spatial-collar-type tensor decomposition, exact-or-controlled four-term strip additivity on one inherited strip model, renormalized endpoint-Lipschitz control up to the weak tail generator, and the derived half-sided modular pair whose Borchers–Wiesbrock consequence is an explicit positive null-translation generator on its Stone domain with the affine half-line modular relation; the same half-line family then fixes the generator/charge identification internally.

- **Derived half-sided modular inclusion.** Nested null half-line algebras satisfy half-sided modular inclusion on the declared geometric scaling branch by Corollary 5.2e; Borchers–Wiesbrock then identifies the positive null-translation generator on its Stone domain together with the affine half-line modular relation [19].

- **Weak continuity and finite variation.** The bounded-interval and half-line endpoint-Lipschitz control follow from the local MaxEnt branch and define the weak tail generator at fixed cutoff. The support-visible scaling theorem together with the extracted geometric cap pair supplies the continuum null-generator setting in which that weak-tail data can be matched to the geometric null modular action of the relativistic phase.

Constraint set specification. On the local finite-constraint branch, the "correct fixed-cap constraint set" becomes explicit: constraints are the local conserved charges of the symmetries used in the derivation:

1. **Edge/cap label constraints:** fix the distribution of collar-sector labels, equivalently $\langle L_C \rangle$ for each cap size, giving the area term.

2. **Gauge charges:** fix boundary flux or charge operators.

3. **Geometric (conformal) charges:** fix the expectation of the conformal Killing charges that preserve the cap, i.e. the generator B_C or its microscopic lattice approximation.

MaxEnt therefore selects the unique finite-stage invariant state compatible with those conserved charges. The support-visible BW scaling theorem is the continuum statement: if the limit algebra is non-type-I, the geometric modular action on that branch is outer rather than an inner density-matrix identity.

Focusing status. QNEC has rigorous QFT proofs in broad settings [29, 31], with hypotheses (null-surface locality and analyticity, modular-flow/causality/defect-OPE structure) that have not been verified on the reconstructed net, and QFC is strictly stronger than QNEC [30]. The Recoverable Generalized Entropy axiom therefore carries the focusing content of this branch itself (§5.10):

- EC + MaxEnt derive $S_{\text{gen}} = S_{\text{bulk}} + \langle L_C \rangle$ (Section 5.4), the object the axiom constrains.
- The null-stress and Einstein theorem branch supplies the kinematic inputs a future internal QNEC/QFC proof must consume; the derivation itself is open.

Summary. The CMFP package therefore separates the dependency structure into two layers:

- **Internal branch consequences:** the local-Gibbs form, quasi-local propagation, bounded-interval endpoint-Lipschitz control, and the realized state-side refinement branch along which later transport questions are asked, all from the local finite-constraint MaxEnt branch.
- **Support-visible scaling statement:** the scaling limit emits the support-visible geometric cap pair and ordered cut-pair rigidity holds on it. On that branch the limit algebra may be non-type-I, the modular action may be outer, and the 2π normalization and later null half-sided-inclusion bridge follow.

The theorem route works on the support-visible quotient. The realized transported cap-local system packages the geometric cap-local test family, the projectively compatible transported marginal family, and the asymptotic transport-equivalence certificate. The regularized transport theorem below replaces the unavailable full-algebra lower spectral floor. Local weak-$*$ extraction and GNS gluing then emit the support-visible scaling-limit cap pair, support-readable modular covariance reads the modular group as a cap-local support map, and ordered cut-pair rigidity collapses the residual cap-preserving conformal freedom to the unique hyperbolic subgroup.

Proposition 2.6a (Recoverability is not modular geometry: common-floor collapse countermodel). Exact or asymptotically exact Markov recovery at finite cutoff does not by itself imply a full-algebra lower spectral floor along refinement. In a two-dimensional matrix algebra, let

$$\rho_n = \begin{pmatrix} e^{-n} & 0 \\ 0 & 1 - e^{-n} \end{pmatrix}.$$

Each ρ_n is faithful at finite n, and tensoring it with any fixed finite exact-Markov collar factor gives a full-rank exact-Markov collar family. Nevertheless $\lambda_{\min}(\rho_n) = e^{-n} \to 0$, so there is no positive lower bound along the refinement family. On the off-diagonal matrix unit E_{12}, the modular generator carries the logarithmic gap

$$[-\log \rho_n, E_{12}] = n\, E_{12} + O(1) E_{12},$$

and the unregularized modular transport has no finite common-floor limit on that direction. Thus collar Markovity and finite-stage faithfulness are not enough for the false stronger full-algebra BW lift. The observer-facing theorem uses the support-visible regularized replacement, with the exact-Markov comparison family checked only on fixed collars and made replacement-independent after weak-$*$/GNS extraction.

Theorem 2.6b (regularized support-visible modular transport). Fix a local collar model of finite dimension $d_{m,\delta}$. Let ρ_n be the transported physical collar marginal, $\widehat{\rho}_n$ the exact-Markov comparison marginal on the same collar model, and $\Delta_n = \|\rho_n - \widehat{\rho}_n\|_1$. For $a > 0$, set $K_a(\rho) = -\log(\rho + a\mathbf{1})$. For every bounded collar observable O in the support-visible algebra,

$$\left|\operatorname{Tr} \rho_n O\big(K_a(\rho_n) - K_a(\widehat{\rho}_n)\big)\right| \le \|O\| \left(\frac{4\Delta_n}{a} + d_{m,\delta} a + 4\Delta_n |\log a|\right).$$

Consequently, if $a_n \downarrow 0$ is chosen with

$$\Delta_n/a_n \to 0, \qquad d_{m,\delta} a_n \to 0, \qquad \Delta_n |\log a_n| \to 0,$$

then the regularized support-visible modular matrix elements converge on that fixed collar model.

Proof sketch. On the spectral interval $[a, \infty)$, the logarithm is operator-Lipschitz at the scale used above by its integral representation. Splitting the comparison into the support above a, the a-tail, and the trace-distance error gives the displayed bound. The three displayed conditions make the three error terms vanish. QED.

Definition 2.6c (explicit multiresolution reference branch). The finite-to-continuum bridge is evaluated on a declared reference branch, not on an unstructured sequence of finite plots. A regulator is

$$r = (m, L, b),$$

where m is the refinement depth, L is the finite support or volume scale, and b is a boundary/phase label. The screen branch uses nested icosahedral subdivisions of S^2; Euclidean branches use dyadic mesh $a_m = a_0 2^{-m}$, with an additional temporal spacing a_t when a transfer matrix is present. Each added cell, collar, edge-center, or support-visible channel j carries

$$D_j = \bigoplus_{\alpha \in S_j} M_{d_{j,\alpha}}(\mathbb{C})$$

and a faithful detail state τ_j. At regulator r,

$$\widetilde{M}_r = \bigotimes_{j \in \mathcal{J}_r} D_j, \qquad \widetilde{\omega}_r = \bigotimes_{j \in \mathcal{J}_r} \tau_j.$$

A finite-depth local presentation circuit $\Gamma_r : \widetilde{M}_r \to M_r$ changes from multiresolution coordinates to the OPH patch, port, collar, and gauge coordinates. For $r \preceq s$, write

$$\widetilde{M}_s = \widetilde{M}_r \otimes D_{s \setminus r}, \qquad \tau_{s \setminus r} = \bigotimes_{j \in \mathcal{J}_s \setminus \mathcal{J}_r} \tau_j.$$

The canonical refinement and coarse-graining maps are

$$\iota_{rs}(A) = \Gamma_s\Big[\Gamma_r^{-1}(A) \otimes \mathbf{1}\Big],$$

and

$$Q_{sr}(X) = \Gamma_r\Big[(\mathrm{id} \otimes \tau_{s\setminus r})\big(\Gamma_s^{-1}(X)\big)\Big].$$

The embedded conditional expectation is $E_{sr} = \iota_{rs} Q_{sr}$. The reference state is $\widehat{\omega}_r = \widetilde{\omega}_r \circ \Gamma_r^{-1}$.

Theorem 2.6d (exact modular-compatible reference tower). The data

$$(M_r, \widehat{\omega}_r, \iota_{rs}, E_{sr})_{r \preceq s}$$

form a directed tower. The maps ι_{rs} are unital injective $*$-homomorphisms with $\iota_{st}\iota_{rs} = \iota_{rt}$. The Q_{sr} are faithful state-preserving completely positive coarse-grainings with $Q_{tr} = Q_{sr}Q_{ts}$. The maps E_{sr} are faithful conditional expectations onto $\iota_{rs}(M_r)$, and

$$\widehat{\omega}_s \circ \iota_{rs} = \widehat{\omega}_r, \qquad \widehat{\omega}_r \circ Q_{sr} = \widehat{\omega}_s, \qquad \widehat{\omega}_s \circ E_{sr} = \widehat{\omega}_s.$$

The finite reference modular groups are also exactly compatible:

$$\sigma_t^{\widehat{\omega}_s}(\iota_{rs}A) = \iota_{rs}(\sigma_t^{\widehat{\omega}_r}(A)).$$

This theorem uses no full-algebra spectral floor uniform in r; in bare coordinates it is only the identity $A \mapsto A \otimes \mathbf{1}$, partial expectation against faithful detail states, and conjugation by Γ_r, Γ_s.

Theorem 2.6e (canonical renormalization and correlation Cauchy bound). Let $\mathfrak{A}$ be the C^*-inductive limit with embeddings $j_r : M_r \to \mathfrak{A}$. There is a unique state $\widehat{\omega}$ with $\widehat{\omega} \circ j_r = \widehat{\omega}_r$, and the maps Q_{sr} induce state-preserving conditional expectations

$$\mathbb{E}_r : \mathfrak{A} \to j_r(M_r)$$

such that $\|\mathbb{E}_r(A) - A\| \to 0$ for every $A \in \mathfrak{A}$. The finite-regulator realization of a continuum test observable is

$$\mathcal{R}_r(A) = j_r^{-1}(\mathbb{E}_r A).$$

For $\|A_k\| \leq M$,

$$\left| \widehat{\omega}_r \left(\prod_{k=1}^{n} \mathcal{R}_r(A_k) \right) - \widehat{\omega} \left(\prod_{k=1}^{n} A_k \right) \right| \leq M^{n-1} \sum_{k=1}^{n} \|A_k - \mathbb{E}_r A_k\|.$$

Thus the operator mixing matrix and additive subtraction at finite cutoff are outputs of $\mathcal{R}_r$, not free fits to the target continuum answer. Independent lattice-spacing, volume, and boundary sweeps must publish the corresponding martingale-tail or physical-error envelopes before finite regulator data are read as continuum evidence.

Theorem 2.6f (compact-time modular convergence on fixed local algebras). Let $\rho_{s\downarrow r}^{\mathrm{phys}}$ be the density matrix of the physical state at fine stage s, restricted to the embedded fixed local algebra M_r, and let

$$\epsilon_{s,r} := \|\rho_{s\downarrow r}^{\mathrm{phys}} - \widehat{\rho}_r\|_1.$$

For a regularizer $\eta_s > 0$, define

$$K_{s\downarrow r}^{(\eta_s)} = -\log(\rho_{s\downarrow r}^{\mathrm{phys}} + \eta_s \mathbf{1}).$$

If $\lambda_r = \lambda_{\min}(\widehat{\rho}_r) > 0$, then for every $A \in M_r$ and every $T < \infty$,

$$\sup_{|t| \leq T} \left\| e^{-itK_{s\downarrow r}^{(\eta_s)}} A e^{itK_{s\downarrow r}^{(\eta_s)}} - \sigma_t^{\widehat{\omega}_r}(A) \right\| \leq 2T\|A\| \left[\frac{\epsilon_{s,r}}{\eta_s} + \log\left(1 + \frac{\eta_s}{\lambda_r}\right) \right].$$

Hence $\eta_s \to 0$ and $\epsilon_{s,r}/\eta_s \to 0$ give uniform-on-compact-time convergence on each fixed support-visible local algebra. For a countable cofinal local test tower, pointwise $\epsilon_{s,r} \to 0$ admits a diagonal subsequence and one cutoff schedule. The phrase "finite-stage modular action" below means this transported, regularized local action; it does not mean the unregularized full finite-cap modular group.

Cap-family uniformity clause. For application to the finite cap-net BW theorem, the compact-time estimate is required uniformly over the declared finite nondegenerate cap family and the fixed separating support-visible test tower. The operator-norm estimate then implies the quadratic mixed-GNS Cauchy residual because $\|XD\|_{2,\omega} \leq \|X\|\|D\|_{2,\omega}$. The reference tower supplies the analytic engine; the finite cap-normal certificate adds the support-order, BW-frame, oriented cross-ratio, geometric 2π-KMS, and wrong-scale controls.

Theorem 2.6g (finite positive transfer and Lorentzian handoff gate). On Euclidean repair branches, let $P_{C,r}$ be weighted resampling inside the fibers of the complete repaired visible datum for active collar C. Ref. [1] proves that this finite operator is the orthogonal conditional expectation and gives the exact support, equal-fiber-row, and weighted detailed-balance matrix receipt that recognizes it. A concrete active collar repair kernel is identified with $P_{C,r}$ only when its independently derived transition matrix passes that receipt; constructing the matrix from the target formula is not a verification. Choose $c_{C,r} \geq 0$, and set

$$L_r^{\mathrm{rep}} = \sum_C c_{C,r}(I - P_{C,r}), \qquad T_r(a_t) = e^{-a_t L_r^{\mathrm{rep}}}.$$

Then $L_r^{\mathrm{rep}} = L_r^{\mathrm{rep}\,*} \geq 0$, $0 < T_r(a_t) \leq I$, $T_r(a_t)\mathbf{1} = \mathbf{1}$, and the stationary Euclidean path measure is reflection positive:

$$\mathbb{E}_r[(\Theta F)F] \geq 0$$

for positive-time cylinder functions F. A Lorentzian unitary statement requires this finite positivity plus either exact shell-additive direct-limit semigroup identities or a Mosco / strong-resolvent substitute, including vacuum-projection and intertwiner convergence. Weak-$*$ state convergence alone is not a transfer theorem.

BW-side closure boundary. Proposition 2.6a blocks the stronger unregularized full-algebra conclusion from the proved finite-cutoff package. Theorems 2.6c–2.6g supply the explicit reference-tower certificate needed for the finite-to-continuum reading: continuum correlations come from the conditional-expectation renormalization map, compact-time modular group convergence is only for transported regularized physical marginals on fixed support-visible local algebras, and Lorentzian transfer needs reflection positivity or an equivalent positive-transfer certificate. Theorem 4.2 then uses exactly the observer-facing support-visible content to close the BW/geometric cap-pair automorphism statement. Without the multiresolution regulator certificate, finite regulator data remain regulator evidence. A black-box AQFT/BW reconstruction route would require separately verifying conformal-net properties such as locality, additivity, duality, standardness, positive energy, and suitable modular inclusions.

2.1 Overlap Consistency and Gluing

The constructive part of overlap consistency is the tree-gluing theorem below. The structural part is the origin of the gluing redundancy itself. On the ordinary or central-defect branch, gauge-as-gluing is the finite-regulator overlap redundancy of local chart presentations: once the overlap algebras are realized in finite-dimensional charts, overlap-consistent rechartings of a cut form a compact unitary transition system. The collar theorem of Section 2.3 should therefore be read with its boundary group G_Σ understood as shorthand for this derived compact boundary action, not as a model-specific add-on. On the genuinely noncentral branch, the same weak gluing data are encoded by a compact crossed-module change system, so the fixed-cutoff collar theorem upgrades to a higher-gauge statement rather than failing. The fixed-cutoff topological package is therefore closed on all three branches. A second structural point is UV underdetermination. If each patch Hilbert space is tensored with an inert finite ancillary factor and the observable patch algebras are embedded as $\mathcal{A}(P) \otimes \mathbf{1}$, then the physical observables, collar conditional mutual information, carried Markov errors, and quotient normal forms are unchanged. So the physical UV branch is fixed only modulo gauge or implementation hiding together with such ancillary stabilization, not a unique microscopic presentation; the unique theorem-grade object is the induced family of terminal expectation functionals on the declared physical observable algebras, not a preferred microscopic representative.

2.1.1 Constructive gluing on tree covers

Theorem 3.1 (tree gluing). Let a rooted tree of patches satisfy a tree-ordered overlap structure and a tripartite factorization (A_k, B_k, C_k) at step k. If a target state ρ^* obeys $I(A_k : C_k \mid B_k) \leq \varepsilon_k$, then there exist recovery maps $\mathcal{R}_k$ such that

$$\|\rho^*_{A_k B_k C_k} - (\mathrm{id}_{A_k} \otimes \mathcal{R}_k)(\rho^*_{A_k B_k})\|_1 \leq \delta_k,$$

with

$$\delta_k = 2\sqrt{1 - e^{-\varepsilon_k}} \leq 2\sqrt{\varepsilon_k},$$

The iteratively glued state $\hat{\rho}$ satisfies

$$\|\hat{\rho} - \rho^*\|_1 \leq \min\left(2, \sum_{k=2}^{n} \delta_k\right).$$

Proof. Induct on k. The recovery error contracts under CPTP maps, so the errors add. QED.

2.1.2 Gauge-as-gluing and loops

At finite regulator scale, the fixed-cutoff gauge-as-gluing package identifies the overlap-consistency redundancy of local chart presentations. Choose finite-dimensional local presentations of the overlap algebras on a connected cut Σ. Because the overlap algebras are matrix algebras, any overlap-consistent change of chart is inner and is implemented by a unitary on the cut Hilbert space. Fixing a reference chart, the compact closure of the subgroup generated by all such recharting unitaries is the boundary gluing group K_Σ; before fixing the reference chart, the same data form a compact unitary groupoid.

Proposition 3.2a (Derived gauge-as-gluing at finite regulator). Let a finite regulator chart be chosen for the patches meeting along a connected interface Σ. Then overlap consistency determines a compact unitary transition system on the cut data. On the ordinary or central-defect branch, this transition system reduces to a compact boundary group K_Σ, and when the triple-overlap defect is central its projective composition law lifts to a genuine action of a compact central extension $\widehat{K}_\Sigma$. Gauge is therefore the overlap redundancy itself, not an extra primitive.

Proof. On a finite-dimensional matrix algebra every *-automorphism is inner, so each overlap-consistent recharting is conjugation by a unitary on the cut Hilbert space. The subgroup generated by those unitaries has compact closure inside a finite-dimensional unitary group. If triple-overlap defects are central, the resulting projective composition law lifts to a central extension. QED.

Lemma 3.2b (trees vs loops). If the patch adjacency graph is a tree, one can choose local charts h_i so that the overlap labels satisfy $g_{ij} = h_i^{-1} h_j$ on all edges. If loops exist, the loop holonomy

$$H(\gamma) = g_{i_1 i_2} g_{i_2 i_3} \cdots g_{i_n i_1}$$

is invariant under local frame changes. Nontrivial holonomy is the obstruction to global trivialization. QED.

Corollary 3.2c (Collar consequence on the ordinary or central-defect branch). Let $B_\delta = B_L \cup B_R$ be a collar around a cap boundary Σ, and set $\widehat{K_\Sigma} = K_\Sigma$ on the ordinary branch. Then the EC theorem of Section 2.3 has the derived interpretation

$$H_{B_\delta} = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{\widehat{K}_\Sigma} \cong \bigoplus_\alpha (H_{b_L^\alpha} \otimes H_{b_R^\alpha}),$$

with

$$Z(A_{\mathrm{EC}}(B_\delta)) = \bigoplus_\alpha \mathbb{C}\,\mathbf{1}_\alpha.$$

The right half-collar carries the contragredient action because it sees the inverse transport across the same cut. Exact Markov normal forms used later require the additional state hypothesis $I_\omega(A_\delta : D_\delta \mid B_\delta) = 0$ *together with* the Markov-split alignment hypothesis of Section 2.3, or the explicitly stated idealized recoverability limit supplying both; the block decomposition itself is kinematic and follows from the derived boundary action.

Proof sketch. Decompose the left boundary data into irreps $(V_\alpha \otimes H_{b_L^\alpha})$ of $\widehat{K_\Sigma}$ and the right boundary data into the dual modules $(V_\beta^* \otimes H_{b_R^\beta})$. Then Schur's lemma leaves a singlet only when $\alpha = \beta$, producing the displayed direct sum. QED.

2.1.3 Loop obstruction class (central defect)

On the central-defect subbranch, identify the overlap centers with one fixed abelian unitary coefficient group Z_Σ and assume overlap transport acts trivially on it. Choose unitary overlap implementers normalized by $U_{ji} = U_{ij}^{-1}$ and define central defects z_{ijk} at the implementer level by

$$\Omega_{ijk} := U_{ij} U_{jk} U_{ki} = z_{ijk} \in Z_\Sigma,$$

or equivalently $U_{ij} U_{jk} = z_{ijk} U_{ik}$. This implementer-level relation records the projective multiplier; the automorphism-level expression $\mathrm{Ad}(z_{ijk})$ alone would be the identity and would not determine z_{ijk}.

Then $\{z_{ijk}\}$ is an untwisted Čech 2-cocycle with coefficients in Z_Σ, and its cohomology class $[z]$ is gauge invariant. Central defects do not obstruct the collar block decomposition: they only replace K_Σ by its central extension $\widehat{K_\Sigma}$. The condition $[z] = 0$ is exactly the strictifiability of the projective triangle defect. It is not by itself global path independence: at least one resulting strict edge 1-cocycle must also have trivial represented holonomy around every closed overlap loop. With a nontrivial action on varying centers, the corresponding statement uses the twisted/local-system Čech differential instead. (A full proof for the fixed-coefficient subbranch appears in Section 6.4 below, in the algebra-net language.)

2.1.4 Non-central obstruction (2-group cocycle)

When defects are not central, the natural coefficient data is a crossed module $(H \to G)$ with an action of G on H by conjugation. Here G is the reconstructed gauge group, and H is the unitary group acting on edge multiplicity spaces, with boundary map $\partial : H \to G$.

A crossed module is a homomorphism $\partial : H \to G$ together with an action of G on H such that

$$\partial(g \triangleright h) = g\,\partial(h)\,g^{-1}, \qquad \partial(h) \triangleright h' = hh'h^{-1}.$$

On a good cover $\{P_i\}$, a weakly coherent gluing is encoded by:

$$g_{ij} : P_{ij} \to G, \qquad h_{ijk} : P_{ijk} \to H,$$

obeying the 2-cocycle conditions

$$g_{ij}g_{jk} = \partial(h_{ijk})g_{ik},$$

and on quadruple overlaps,

$$h_{jkl}h_{ijl} = (g_{ij} \triangleright h_{ikl})h_{ijk}.$$

Gauge changes act by 1- and 2-cochains in the standard way for crossed-module cohomology. Write

$$q_\Sigma = [(g,h)] \in \check{H}^2(N_\Sigma, H_\Sigma \to G_\Sigma)$$

for the full crossed-module gluing orbit. It classifies the equivalence orbit of the pair (g,h) but does not in general determine a unique ordinary G_Σ-valued 1-cocycle class after strictification, because the map below need not be injective. Let

$$\begin{aligned} \iota_\Sigma : \check{H}^1(N_\Sigma, G_\Sigma) &\longrightarrow \check{H}^2(N_\Sigma, H_\Sigma \to G_\Sigma), \\ [g] &\longmapsto [(g,1)] \end{aligned}$$

be the natural strict-locus map. Define the separate $\{0,1\}$-valued strictification obstruction by

$$\begin{aligned} o^{(2)}_\Sigma(g,h) = 0 \quad &\Longleftrightarrow \quad q_\Sigma \in \mathrm{im}(\iota_\Sigma) \\ &\Longleftrightarrow \quad [(g,h)] \text{ has a representative } (g^{\mathrm{str}},1) \\ &\qquad\qquad \text{with } g^{\mathrm{str}}_{ij}g^{\mathrm{str}}_{jk} = g^{\mathrm{str}}_{ik}. \end{aligned}$$

and set $o^{(2)}_\Sigma(g,h) = 1$ otherwise.

Theorem 3.4 (non-central higher-defect strictification). The higher associator defect of (g_{ij}, h_{ijk}) is removable iff $o^{(2)}_\Sigma(g,h) = 0$. In that case a higher-gauge change gives $h_{ijk} = 1$ and a genuine G-valued edge 1-cocycle. Strict endpoint-only transport additionally requires that at least one allowed strict representative have trivial represented holonomy. The full orbit q_Σ need not determine a unique ordinary H^1 class.

Proof sketch. Removing the higher associator corresponds to $h_{ijk} = 1$ and $g_{ij}g_{jk} = g_{ik}$. Crossed-module coboundaries preserve the full orbit, so this strictification exists exactly when that orbit meets the image of the ordinary G-valued 1-cocycle locus, which is the definition of $o^{(2)}_\Sigma = 0$. The map ι_Σ need not be injective: H-valued edge changes can relate strict representatives and alter their ordinary holonomy through ∂. Thus the invariant residual test is whether at least one allowed strict representative has trivial represented holonomy, as in Theorem 3.4b. QED.

The central-defect case is the abelian truncation with H central and trivial action, which reduces to Section 3.3. A genuinely noncentral class is the point where the fixed-cutoff collar theorem upgrades from the ordinary-group package to the higher-gauge one below.

Corollary 3.4a (Fixed-cutoff higher-gauge EC and transportability). For a finite regulator chart on a connected cut Σ, the genuinely noncentral branch admits a compact crossed-module change system

$$\mathcal{T}_\Sigma = C^1(N_\Sigma, H_\Sigma) \rtimes C^0(N_\Sigma, G_\Sigma).$$

The physical higher-gauge collar is

$$\mathcal{H}_B^{2g} = (\widetilde{\mathcal{H}}_{B_L} \otimes \widetilde{\mathcal{H}}_{B_R})^{\mathcal{T}_\Sigma} \cong \bigoplus_\lambda (\mathcal{H}_{b_L^\lambda} \otimes \mathcal{H}_{b_R^\lambda}),$$

with

$$Z(\mathcal{A}_{2g}(B)) = \bigoplus_\lambda \mathbb{C}\,\mathbf{1}_\lambda,$$

and the full orbit q_Σ is invariant under local rechartings and classifies the fixed-cutoff genuinely noncentral gluing datum. Its separate invariant property $o_\Sigma^{(2)} = 0$ holds iff the higher associator is removable. The resulting strict representatives can have different ordinary G_Σ-valued 1-holonomies, so transportability is the separate existential test of Theorem 3.4b rather than a uniquely defined holonomy attached to q_Σ.

Proof sketch. Pair-overlap rechartings are inner, while triple-overlap associators strictify to compact crossed-module data. Finite-dimensional unitary $\mathcal{T}_\Sigma$-modules decompose semisimply, and Schur matching leaves only diagonal left/right blocks. Crossed-module cohomology classifies the higher defect; Theorem 3.4b separately tests the ordinary loop holonomy that remains after strictification. QED.

Theorem 3.4b (TransportabilityFromOverlapGluing). Fix a connected cut Σ and a finite regulator charting nerve N_Σ. Transport of a collar charge is defined by the path composite in the overlap recharting groupoid: along $p = (i_0 \to \cdots \to i_m)$, the charge block is moved by $U_p = U_{i_{m-1}i_m} \cdots U_{i_0 i_1}$. This construction uses the overlap unitary transition system itself, not an added DHR transportability assumption.

After a central or crossed-module strictification, let U^{str} denote a resulting genuine edge 1-cocycle. Write $\mathcal{C}_\alpha$ for the full transported collar-charge block, including its boundary carrier and multiplicity/intertwiner data, and define

$$\mathrm{Hol}_{\alpha,U^{\mathrm{str}}} : \pi_1(|N_\Sigma|, i_*) \longrightarrow U(\mathcal{C}_\alpha), \qquad [\gamma] \longmapsto U_\gamma^{\mathrm{str}}\big|_{\mathcal{C}_\alpha}.$$

On the ordinary branch, transport is strictly path-independent iff every closed overlap loop has trivial holonomy on the collar-sector block. On the central branch, elementary triangle moves accumulate the central cocycle z_{ijk}. Strict path-independent transport exists iff $[z]_\Sigma = 0$ and there is a central 1-cochain strictification for which $\mathrm{Hol}_{\alpha,U^{\mathrm{str}}} = 1$. The first condition removes the projective triangle defect; the second removes the residual ordinary loop action. If $[z]_\Sigma = 0$ but the second condition fails, the lifts form a genuine 1-cocycle and transport is path dependent.

On the genuinely noncentral branch, path comparisons are H_Σ-valued 2-morphisms in the crossed module $H_\Sigma \to G_\Sigma$. Strict ordinary transport exists iff $o_\Sigma^{(2)} = 0$ and at least one crossed-module strictification gives a genuine G_Σ-valued 1-cocycle with trivial represented holonomy on $\mathcal{C}_\alpha$. If $o_\Sigma^{(2)} \neq 0$, the sector has higher transport rather than ordinary compact-group DR transport. If $o_\Sigma^{(2)} = 0$ but every allowed strict representative has nontrivial represented G_Σ-holonomy, the higher associator has strictified but ordinary path dependence remains. The full orbit $q_\Sigma = [(g,h)]$ need not determine a unique ordinary H^1 class and is not itself the strictification test. Thus “zero obstruction” below means the combined associator-strictification and trivial-holonomy criterion, not vanishing of triangle or higher-associator data alone. QED.

Triangle-free cycle check. Let $N_\Sigma = C_n$, $n \geq 4$, have no filled 2-simplices. On $\mathcal{C}_\alpha = \mathbb{C}^2$, put identity transport on every oriented cycle edge except one, where the sector action is the non-scalar unitary $V = \mathrm{diag}(-1, 1)$, and use inverse transports on reversed edges. All central 2-data and higher associator data are vacuously trivial, so $[z]_\Sigma = 0$ and $o^{(2)}_\Sigma = 0$. Its ordered cycle holonomy is $V \neq 1$, which no central rephasing removes. For a genuinely noncentral instance, take the crossed module

$$\mathrm{SU}(2) \hookrightarrow \mathrm{U}(2)$$

with conjugation action and the standard action on $\mathcal{C}_\alpha$. The band $\mathrm{U}(2)/\mathrm{SU}(2) \cong \mathrm{U}(1)$ is detected by the determinant, and the cycle has band holonomy $\det V = -1$. An SU(2)-valued edge change has determinant one and a vertex-frame change only conjugates the based loop, so no allowed strict representative has trivial holonomy. The two routes between the endpoints of the distinguished edge differ, and the combined criterion rejects both the central and noncentral assignments.

2.2 Gauge Reconstruction and Standard Model Structure

2.2.1 Edge sector category and gauge group reconstruction

At any fixed UV cutoff, edge-center completion provides finitely many visible seed charges and finite-dimensional intertwiners. Their rigid tensor closure is defined below and may have infinitely many simple classes. The local MaxEnt / collar-mixing package established above controls only fixed-cutoff recoverability, modular-support localization, and carried error terms on the realized branch. It is logically separate from the question whether zero-obstruction edge sectors survive refinement. The fixed-stage category is theorem-produced; the refinement/fiber ladder and cofinal compact-gauge witness require the explicit compact-gauge refinement receipt defined below.

On the ordinary or central-defect branch, path-independent movement of collar charges is supplied by overlap gluing. TransportabilityFromOverlapGluing constructs transport from overlap paths and proves the exact combined criterion: the central triangle class $[z]_\Sigma$ must vanish and at least one allowed strictification must have trivial represented sector holonomy. On the genuinely noncentral branch, $o^{(2)}_\Sigma = 0$ strictifies the higher associator, after which at least one allowed strict G_Σ-valued representative must have trivial represented holonomy. A nonzero $o^{(2)}_\Sigma$ remains a higher-gauge sector; an orbit for which every strict representative has residual loop action is also outside the ordinary path-independent DR sector. The full orbit q_Σ need not determine a unique ordinary H^1 class and is not itself the strictification test. Thus the overlap obstruction calculus classifies and routes sectors; it does not by itself select the Standard Model.

For the tensor construction at cutoff r, choose one common stagewise strict edge cocycle U^{str}_r from the allowed representatives (the ordinary branch uses its given strict system). Retain only seeds on which this same representative has trivial action around every closed overlap loop. Sectors requiring incompatible strictifications define alternative fixed-stage categories; they are not unioned before tensor closure.

Classification is not realization.

Proposition (obstruction neutrality of the finite Standard Model gauge implication). The ordinary trivial-holonomy condition, the combined central condition $[z]_\Sigma = 0$ plus a trivial-holonomy strictification, and the combined noncentral condition $o^{(2)}_\Sigma = 0$ plus an allowed strict representative with trivial represented G_Σ-holonomy are transportability conditions. After one common stagewise strict representative is chosen, they permit an ordinary transportable bosonic sector category generated only by sectors trivial under that choice; they do not select

$$\frac{\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)}{\mathbb{Z}_6}.$$

On a cofinal tail carrying the compact-gauge refinement receipt, DR/Tannaka reconstruction returns $G = \mathrm{Aut}_{\otimes}(\mathcal{F})$ for the constructed sector category and fiber functor. If that category is trivial, G is trivial; in general G is whatever compact group the tensor-fiber data reconstructs. The separate finite response producer and the conditional charge-conjugate pair of rank-15 chiral projectors give determinant balance, the Standard Model charge lattice, and a common $\mathbb{Z}_6$ tensor kernel. Quotienting by the full kernel gives the maximal faithful matter image displayed above. The cover and its intermediate quotients carry the same local tensors, so the physical global form requires additional global data. This implication uses no MAR clause and does not identify the finite current with the independently reconstructed Tannaka group. QED.

Definition (fixed-cutoff tensor-generated sector category). Let $S_r = \{W_{\alpha,r} : P_{\alpha,r}$ is visible and has trivial holonomy under the common $U_r^{\mathrm{str}}\}$. Define

$$\mathsf{Sect}_r^{\mathrm{bos}} = \langle \mathbf{1}, S_r, S_r^{*}\rangle_{\oplus,\otimes,\mathrm{subobjects,isomorphism}} \subset \mathrm{Rep}_{\mathrm{fd}}(\widehat{K_{\Sigma,r}}).$$

This is the replete full additive Karoubi rigid tensor subcategory generated by the finitely many visible carriers. The original projectors $P_{\alpha,r}$ identify only one-collar seeds. The simple objects are all irreducible summands of finite tensor words in seeds and conjugates; they need not have projectors in the original one-collar algebra, and there may be infinitely many of them even though every object and Hom space is finite-dimensional. A subobject $X \subset T$ is supported by its invariant idempotent in $\mathrm{End}_{\widehat{K}_{\Sigma,r}}(T)$, or by the corresponding block in a supplied finite concatenated-collar realization of that tensor word.

Theorem (FixedCutoffBosonicSectorCategory). At every fixed regulator cutoff, on the ordinary or central-defect zero-obstruction branch and in the bosonic internal-gauge sector of the $3+1$-dimensional EFT regime, the preceding construction is a semisimple rigid symmetric C^*-tensor category. Complete reducibility of finite-dimensional representations of the compact group gives semisimplicity and subobjects; the operator adjoint and norm give the C^*-structure; conjugate representations give duals; and the standard flip gives the bosonic symmetry. For a closed loop, U_r^{str} acts as the identity on every retained seed, hence as the identity on direct sums, tensor products, conjugates, and invariant subobjects. Thus the whole generated category has one monoidally coherent path-independent transport, rather than an objectwise choice of strictification. Collar concatenation realizes tensor words whenever the corresponding finite tensor-realization receipt is supplied. The canonical carrier-forgetful functor

$$F_r : \mathsf{Sect}_r^{\mathrm{bos}} \longrightarrow \mathsf{Hilb}_{\mathrm{fd}}$$

is faithful, *-preserving, and symmetric strong monoidal. Fermionic signs, spinorial matter, and chirality are not part of this bosonic internal-gauge category; they belong to the later super-Tannakian or matter-sector lift.

Definition (compact-gauge refinement receipt). On a cofinal tail, first require a surjectivity/finite-extendability certificate for each finite-state restriction $\rho_{sr} : Q_s \to Q_r$. Only then is the commutative pullback $C(Q_r) \to C(Q_s)$, $f \mapsto f \circ \rho_{sr}$, injective. This does not produce a map of noncommutative collar algebras. For

$$\mathcal{A}_{\mathrm{EC},r} = \bigoplus_{\alpha} M_{d_{\alpha,r}}(\mathbb{C}), \qquad \mathcal{A}_{\mathrm{EC},s} = \bigoplus_{\beta} M_{d_{\beta,s}}(\mathbb{C}),$$

the receipt separately supplies multiplicities $m^{rs}_{\beta\alpha}$ and unitaries V^{rs}_{β}, with $d_{\beta,s} = \sum_\alpha m^{rs}_{\beta\alpha} d_{\alpha,r}$, and defines

$$\big(\jmath_{rs}(a)\big)_\beta = V^{rs}_\beta \left[\bigoplus_\alpha (a_\alpha \otimes \mathbf{1}_{m^{rs}_{\beta\alpha}})\right] (V^{rs}_\beta)^*.$$

Injectivity requires every coarse α to occur. The whole center maps into the fine center exactly when each β has one coarse shadow; on that branch $\jmath_{rs}(P_{\alpha,r}) = \sum_{\beta:\mathrm{sh}_{sr}(\beta)=\alpha} P_{\beta,s}$. Composition requires $m^{rt}_{\gamma\alpha} = \sum_\beta m^{st}_{\gamma\beta} m^{rs}_{\beta\alpha}$ and coherent V's. The same receipt supplies continuous surjections

$$\pi^\Sigma_{sr} : \widehat{K}_{\Sigma,s} \twoheadrightarrow \widehat{K}_{\Sigma,r}, \qquad \pi^\Sigma_{tr} = \pi^\Sigma_{sr} \circ \pi^\Sigma_{ts},$$

and one common strict representative U^{str}_r at each stage. The pullbacks carry coarse generators into $\mathsf{Sect}^{\mathrm{bos}}_s$, intertwine U^{str}_r with U^{str}_s, preserve trivial action of every fine overlap loop without changing representative object by object, and agree with the displayed collar supports. The receipt also supplies compatible finite concatenated-collar realizations for every finite tensor word and invariant subobject projector. The finite-state consensus maps alone imply none of these matrix, center, compact-group, or tensor-realization data.

Theorem (RefinementFunctorAndFiberDescent). On a cofinal tail carrying that receipt, define

$$U_{rs} := (\pi^\Sigma_{sr})^* : \mathsf{Sect}^{\mathrm{bos}}_r \longrightarrow \mathsf{Sect}^{\mathrm{bos}}_s.$$

Then U_{rs} is fully faithful, symmetric strong monoidal, and *-preserving; it preserves the tensor unit, duals, irreducibles, and the full zero-obstruction transport criterion, with $U_{st}U_{rs} = U_{rt}$. For the canonical forgetful fibers, $F_s U_{rs} = F_r$ literally on carrier Hilbert spaces and linear maps.

Proof. For $X = (H_X, \varrho_X)$, set $U_{rs}X = (H_X, \varrho_X \circ \pi^\Sigma_{sr})$ and $U_{rs}(f) = f$. Surjectivity gives the exact identity

$$\mathrm{Hom}_{\widehat{K}_{\Sigma,s}}(U_{rs}X, U_{rs}Y) = \mathrm{Hom}_{\widehat{K}_{\Sigma,r}}(X, Y),$$

so pullback is fully faithful and a coarse simple cannot split. Pullback commutes with tensor products, unit, conjugates, duality maps, and the bosonic flip, while composition of the π's gives composition of the U's. The receipt intertwines the common stagewise strict representatives, so zero holonomy persists for the whole tensor-generated category without a sector-specific gauge change. Since pullback changes only the action, the comparison $F_s U_{rs} X = H_X = F_r X$ is the identity; naturality, monoidality, and refinement coherence are identity diagrams. The displayed block formula, not finite-set duality, separately proves the noncommutative algebra embedding and its center behavior. QED.

The $\mathrm{U}(1)$ charge-one test. If $W_1(z) = z$ is visible and retained by the common stagewise representative U^{str}_r, the chosen category contains $W_n = W_1^{\otimes n}$ for $n \geq 0$ and $W_{-n} = (W_1^*)^{\otimes n}$ for $n > 0$. Charge two is $W_2 = W_1 \otimes W_1$; its subobject idempotent, equivalently the unique nonzero central projection in its intertwiner algebra, is $\mathbf{1}_{W_2} \in \mathrm{End}_{\mathrm{U}(1)}(W_1^{\otimes 2}) \cong \mathbb{C}$, represented physically in the certified two-collar algebra, not necessarily in the original one-collar algebra. At the next certified refinement, $U_{rs}W_2 = ((\pi^\Sigma_{sr})^* W_1)^{\otimes 2}$ remains simple and $F_s(U_{rs}W_2) = \mathbb{C} = F_r(W_2)$ with identity comparison. Thus the fixed collar remains finite while the tensor-generated simple skeleton may be infinite. A finite fusion/truncation alternative would require a separately stated associative fusion quotient.

Consequently, on a cofinal tail carrying the compact-gauge refinement receipt, the EFT branch has a directed ladder $(\mathsf{Sect}^{\mathrm{bos}}_r, U_{rs}, F_r)$ of tensor-generated bosonic edge-sector categories satisfying that zero-obstruction criterion, with fully faithful monoidal pullback functors and compatible forgetful fibers. Write

$$\mathsf{Sect}_\infty := \varinjlim_r \mathsf{Sect}_r^{\mathrm{bos}},$$

for the directed colimit retaining the sectors and intertwiners that persist in that certified system. Without the receipt, only the fixed-stage categories and forgetful functors are constructed; Sect_∞ and $\mathrm{Aut}_\otimes(\mathcal{F})$ are conditional objects. The theorem below is neutral about whether Sect_∞ is trivial or nontrivial: if the realized branch furnishes only the tensor unit, the reconstructed compact group is the trivial group.

Persistence lemma. Write $U_{rs} : \mathsf{Sect}_r^{\mathrm{bos}} \to \mathsf{Sect}_s^{\mathrm{bos}}$ for the pullback functor on a receipt-certified refinement tail. If a realized zero-obstruction edge-sector class α_{r_0} appears on a sufficiently fine collar, has overlap-visible support, and has representatives $\alpha_s \in \mathsf{Sect}_s^{\mathrm{bos}}$ with $U_{st}\alpha_s \simeq \alpha_t$ for all later $t \succeq s$, then it determines a unique object $[\alpha] \in \mathsf{Sect}_\infty$. Fully faithful pullback along a surjective compact-group map cannot erase or split a coarse simple, and the receipt preserves its full zero-obstruction data. This is a persistence result conditional on the certified ladder, not a consequence of the finite-state refinement maps and not a proof that a nontrivial colimit exists.

This gives a refinement-limit bosonic tensor category Sect_∞ of edge charges:

- objects: zero-obstruction sector labels that persist in the assumed directed system,
- morphisms: intertwiners between sectors,
- tensor product: fusion by collar concatenation, $\alpha \otimes \beta = \bigoplus_\gamma N_{\alpha\beta}{}^{\gamma}\, \gamma$,
- duals: orientation reversal / charge conjugation $\alpha \leftrightarrow \bar{\alpha}$,
- symmetric braiding in the EFT regime (no anyonic statistics in 3+1D).

The monoidal refinement maps descend the tensor unit, associators, duals, and bosonic symmetry to the colimit, while the compatible stagewise forgetful carriers descend to a faithful bosonic fiber functor $\mathcal{F} : \mathsf{Sect}_\infty \to \mathsf{Hilb}_{\mathrm{fd}}$. The fibers are finite-dimensional objectwise even though Sect_∞ can have infinitely many simple objects.

Theorem 6.1 (Receipt-conditional bosonic sector category and Tannaka/DR reconstruction). On the ordinary or central zero-obstruction bosonic EFT branch, suppose a cofinal tail carries the compact-gauge refinement receipt used by RefinementFunctorAndFiberDescent. Then the colimit Sect_∞ is a rigid symmetric C^* tensor category with faithful bosonic fiber functor $\mathcal{F}$. Let $\mathrm{Aut}_\otimes(\mathcal{F})$ denote the group of unitary symmetric monoidal natural automorphisms of $\mathcal{F}$. Then

$$G := \mathrm{Aut}_\otimes(\mathcal{F})$$

is a compact group and $\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G)$. In particular G is unique up to isomorphism, and this group is a compact subgroup of a product of unitary groups.

Proof. The MaxEnt/local-Gibbs/collar-mixing package is not used here to prove nontriviality or to supply the refinement receipt. The fixed-cutoff categories are constructed by FixedCutoff-BosonicSectorCategory, and on the certified tail the monoidal pullback functors and compatible forgetful fibers are constructed by RefinementFunctorAndFiberDescent. Since those functors preserve tensor products, unit, duals, and *-structure, the directed colimit inherits a rigid symmetric C^*-tensor structure, and the compatible stagewise carrier spaces descend to a faithful bosonic fiber functor $\mathcal{F}$. Fix a small skeleton of Sect_∞. For every object X, a unitary symmetric monoidal natural automorphism $\eta \in \mathrm{Aut}_\otimes(\mathcal{F})$ has a unitary component $\eta_X \in U(\mathcal{F}(X))$, so

$$\mathrm{Aut}_\otimes(\mathcal{F}) \hookrightarrow \prod_X U(\mathcal{F}(X)), \qquad \eta \mapsto (\eta_X)_X.$$

For each intertwiner $f : X \to Y$, naturality gives the closed relation $\mathcal{F}(f)\eta_X = \eta_Y \mathcal{F}(f)$. For each pair X, Y, the monoidal structure isomorphism $J_{X,Y}$ of $\mathcal{F}$ gives the closed relation

$$\eta_{X\otimes Y} = J_{X,Y}\,(\eta_X \otimes \eta_Y)\,J_{X,Y}^{-1}, \qquad \eta_{\mathbf{1}} = \mathrm{id}_{\mathbb{C}}.$$

Hence $G = \mathrm{Aut}_{\otimes}(\mathcal{F})$ is a closed subgroup of the compact product $\prod_X U(\mathcal{F}(X))$, so G is compact. The DR/Tannaka hypotheses are therefore satisfied on the constructed pair $(\mathsf{Sect}_\infty, \mathcal{F})$, and the Doplicher–Roberts/Tannaka reconstruction theorem gives a symmetric C^*-tensor equivalence $\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G)$ [40, 41]. If another compact group G' gave the same category through a faithful bosonic fiber functor, the induced symmetric tensor equivalence $\mathrm{Rep}(G) \simeq \mathrm{Rep}(G')$ compatible with the forgetful functors would identify both groups as the automorphism group of the same fiber functor. Thus $G \cong G'$, so the reconstruction is unique up to isomorphism. QED.

Definition 6.1r (DHR/DR realization hypotheses). The field-algebra step below uses the following net-realization package, stated explicitly because it is logically additional to the Tannakian data of Theorem 6.1:

1. **(R1) Observable net.** An isotonous net $O \mapsto \mathcal{A}(O)$ of C^*-algebras over the directed set of small-region collar localizations of the EFT regime, with quasi-local algebra $\mathcal{A}$ and a faithful vacuum representation.
2. **(R2) Locality, duality, Property B.** Spacelike commutativity; Haag duality $\mathcal{A}(O')' = \mathcal{A}(O)''$ in the vacuum representation, or the explicitly weaker essential-duality variant if that is the version invoked; and Property B.
3. **(R3) Localized endomorphisms.** Each persistent class $[\alpha] \in \mathsf{Sect}_\infty$ is realized by an endomorphism ρ^α_O of $\mathcal{A}$ acting trivially on $\mathcal{A}(O_1)$ for every localization region O_1 disjoint from O.
4. **(R4) Transportability.** For every admissible pair of regions $O, \widetilde{O}$ there exists a unitary $u \in \mathcal{A}$ with $u\,\rho^\alpha_O(\cdot)\,u^* = \rho^\alpha_{\widetilde{O}}(\cdot)$. This is a net-level statement about unitary intertwiners. Vanishing collar/represented holonomy supplies path-independent collar transport at the categorical level and is a necessary input to this hypothesis, but it is not by itself equivalent to DHR transportability.
5. **(R5) Statistics and conjugates.** Each ρ^α has finite statistics with permutation-symmetric (bosonic) statistics operator matching the symmetry of Sect_∞, and a conjugate endomorphism realizing the categorical dual $\bar{\alpha}$.
6. **(R6) Completeness.** The assignment $[\alpha] \mapsto [\rho^\alpha]$ is an equivalence of symmetric tensor C^*-categories between Sect_∞ and the category of localized transportable finite-statistics endomorphisms of $\mathcal{A}$.

This paper constructs the categorical side: the fixed-cutoff categories, the refinement functors and fibers, and the colimit pair $(\mathsf{Sect}_\infty, \mathcal{F})$. It does not construct the net realization (R1)–(R6); in particular no observable net is specified here on which the persistent sectors act as localized transportable endomorphisms. (R1)–(R6) therefore enter as explicit hypotheses, and every fixed-point identity $\mathcal{A} = \mathcal{F}_{\mathrm{net}}^G$ in this paper is conditional on them.

Corollary 6.1 (two-level gauge reconstruction on zero-obstruction sectors). The reconstruction conclusion splits into two levels with strictly different hypotheses.

(i) Tannakian level (proved). If in the small-region limit the edge sectors satisfy the zero-obstruction transport criterion of Theorem 3.4b and lie on a cofinal tail carrying the compact-gauge refinement receipt, then Theorem 6.1 yields a compact group $G = \mathrm{Aut}_{\otimes}(\mathcal{F})$ and a symmetric tensor equivalence $\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G)$. This level uses only the constructed category and fiber functor and involves no observable net.

(ii) DHR/DR field-net level (conditional). If in addition the realization hypotheses (R1)–(R6) of Definition 6.1r hold, then the Doplicher–Roberts reconstruction theorem [40, 41] supplies a field algebra $\mathcal{F}_{\mathrm{net}}$ with normal (here bosonic) commutation relations, carrying an action of G with field operators implementing every persistent sector, such that

$$\mathcal{A} = \mathcal{F}_{\mathrm{net}}^{G}.$$

Proof. Level (i) is Theorem 6.1. For level (ii): under (R1)–(R6) the localized transportable finite-statistics endomorphisms of $\mathcal{A}$ form a rigid symmetric C^*-tensor category equivalent to Sect_∞ by (R6), hence to $\mathrm{Rep}(G)$ by level (i). The Doplicher–Roberts theorem [40, 41], whose locality, duality, Property B, transportability, statistics, conjugate, and completeness inputs are exactly (R2)–(R6) on the net (R1), then produces the crossed-product field net $\mathcal{F}_{\mathrm{net}}$, its field operators, and the identification $\mathcal{A} = \mathcal{F}_{\mathrm{net}}^{G}$ with G as gauge group. Without (R1)–(R6) the displayed fixed-point identity is not asserted; the unconditional content of this corollary is level (i). QED.

Corollary 6.1a (The compact-gauge setup and realized witness). On the central branch, the categorical collar-transport input to Corollary 6.1 is the combined theorem-level condition $[z]_\Sigma = 0$ plus trivial represented holonomy for an allowed strictification; this condition feeds the sector-category construction and hypothesis (R4) of Definition 6.1r, but it is not itself the net-level DHR-transportability statement, which additionally requires localized endomorphisms and unitary transporters on an observable net. On the genuinely noncentral branch, strict ordinary transport requires both $o_\Sigma^{(2)} = 0$ and at least one allowed strict representative with trivial represented G_Σ-holonomy; a nonzero $o_\Sigma^{(2)}$ remains a higher-gauge fixed-cutoff sector, while an orbit for which every strict representative has residual loop action remains nontransportable for the ordinary DR corollary. The full orbit q_Σ classifies the crossed-module representatives but need not determine a unique ordinary H^1 class. FixedCutoffBosonicSectorCategory constructs the fixed-stage category. RefinementFunctorAndFiberDescent constructs the fully faithful pullback functors and compatible forgetful fibers only from the explicit compact-gauge refinement receipt. Thus DR/Tannaka reconstruction yields a compact group on the receipt-certified branch at the Tannakian level; the field-algebra identity $\mathcal{A} = \mathcal{F}_{\mathrm{net}}^{G}$ additionally requires the DHR/DR realization hypotheses (R1)–(R6). The finite Standard Model current is a separate source-derived construction, and equality between the two groups requires an additional commuting-square identification.

Theorem (gauge-sector classification and finite-current factorization). On the OPH compact-gauge lane, conditional on a cofinal tail carrying the compact-gauge refinement receipt, the categorical reconstruction factors as

$$\begin{aligned}&\text{overlap/gluing data} \longrightarrow \text{associator strictification test} \longrightarrow \text{strict-representative holonomy test}\\ &\longrightarrow \mathsf{Sect}_\infty \longrightarrow G_{\mathrm{Tan}} = \mathrm{Aut}_\otimes(\mathcal{F}).\end{aligned}$$

The first arrow computes and, where possible, removes the central or crossed-module associator defect. The second computes represented holonomy across the allowed strict representatives. The third keeps only sectors for which at least one such representative has trivial loop action. The fourth reconstructs the compact group from the persistent tensor category and fiber functor. The finite-current lane is

$$\begin{aligned}\text{twelve-port incidence} &\longrightarrow 10J = A^3 - 4A^2 - 5A + 10I\\ &\xrightarrow{\text{target-blind impulse/readback}} R = -J\\ &\longrightarrow \mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{u}(1).\end{aligned}$$

The finite producer solves a common filter from adjacency histories through graph diameter. Its relative A_5-sector signs are exact, and common reversal is the conventional overall charge sign.

Conditional on this response, the finite matter certificate derives $P_{\text{even}} - P_{\text{vac}}$ and $P_{\text{odd}} - P_{\text{top}}$ as the unique charge-conjugate pair of rank-15 chiral projectors. Determinant balance and tensor descent give the exact hypercharge lattice, $N_c = 3$, and the conjugation-insensitive $\mathbb{Z}_6$ kernel. Quotienting by it gives the maximal faithful matter image. No MAR clause enters that implication. Physical matter typing, global-form selection, laboratory current identification, and the equality $G_{\text{Tan}} \cong G_{\text{packet}}$ are open. MAR enters generation and family economy, no-extra-sector, charged-lepton, and D10 uses. Scalar attachment and multiplicity remain separate premises. QED.

2.2.2 Finite Standard Model gauge packet and separate MAR economy step

Theorem 6.1 yields *some* compact G_{Tan}. Independently, the certified twelve-port incidence data and inverse-port pairing determine J, with $10J = A^3 - 4A^2 - 5A + 10I$. The equivariant linear-response commutant is four-dimensional. A target-blind impulse/readback protocol solves the common farthest-shell filter and produces $R = -J$. Its restrictions to $\mathbf{1} \oplus \mathbf{3} \oplus \mathbf{3}' \oplus \mathbf{5}$ have fixed relative signs, with simultaneous reversal interpreted as the conventional overall charge sign. Under this contract the finite current algebra has Standard Model Lie type, and the finite matter certificate derives the two rank-15 projectors $P_{\text{even}} - P_{\text{vac}}$ and $P_{\text{odd}} - P_{\text{top}}$ as a unique charge-conjugate pair, and the finite hypercharge and common-kernel results do not depend on which representative is chosen. Axiom 5 (MAR, Minimal Admissible Realization) is absent from the finite gauge implication. Its consumers include generation and family economy, no-extra-sector, charged-lepton, and D10 branches. Physical matter typing, global-form selection, and laboratory current identification are open.

Axiom 5 (MAR): Minimal Admissible Realization. Among all OPH-realizable sector packages $\mathfrak{S}$ consisting of the connected Lie gauge-sector image relevant in the low-energy EFT, its admissible light chiral matter content, and one Higgs doublet, and which are (i) loop-coherent / transportable (the central triangle class vanishes, or the noncentral higher associator is strictifiable, and at least one allowed strict 1-cocycle representative has trivial represented sector holonomy), (ii) anomaly-free, (iii) refinement-stable with light chiral matter, (iv) single-Higgs Yukawa-completable with one connected abelian charge factor acting nontrivially on the coupled carrier, (v) intrinsically quark-sector CP-capable, (vi) weak-sector UV-completable on that same one-Higgs branch, the realized low-energy package is the lexicographically minimal one under

$$C(\mathfrak{S}) = (\chi_{\text{cpl}},\ N_{\text{nonab}},\ N_c,\ N_g).$$

The complete formal statement, admissibility definitions, and proof route are given in the rest of Section 6.2.

Here χ_{cpl} is the **coupled edge capacity**: the dimension of the minimal unitary carrier containing a common irreducible block on which the admissible pseudoreal and complex nonabelian charge types both act nontrivially. This is intentionally stronger than the abstract minimal faithful representation dimension. For $S(U(3) \times U(2))$, the block-diagonal action on $\mathbb{C}^3 \oplus \mathbb{C}^2$ is faithful of dimension 5, but it is not coupled and therefore does not enter MAR. The object minimized by MAR is the sector package $\mathfrak{S}$, not the bare tensor category by itself. MAR is thus an explicit structural-economy selector on the admissible class, not a theorem derived from the preceding axioms.

Definition 6.1b (MAR realization space and physical equivalence). Let $\mathfrak{A}_{\mathrm{MAR}}$ be the set of isomorphism classes of finite low-energy sector packages

$$\mathfrak{S} = (G^0, \mathcal{R}_{\text{light}}, H, \mathcal{Y}, \mathcal{F})$$

on the ordinary or central zero-obstruction bosonic branch of Theorem 6.1, together with the explicit realized one-Higgs chiral matter package used below. Here G^0 is the connected Lie gauge-sector image, $\mathcal{R}_{\text{light}}$ is the light chiral matter representation family, H is one Higgs doublet, $\mathcal{Y}$ denotes the Yukawa-completable charge data, and $\mathcal{F}$ is the descended finite bosonic fiber functor on the retained transportable sector package. A package is MAR-admissible exactly when it satisfies the six clauses in Axiom 5: zero obstruction, anomaly cancellation, refinement stability with light chiral matter, one-Higgs Yukawa completability with one connected abelian charge factor acting nontrivially on the coupled carrier, intrinsic quark-sector CP capability, and weak-sector UV completability on the same one-Higgs branch.

Two packages are physically equivalent when they are related by a compact-group isomorphism and a fiber-compatible symmetric monoidal equivalence that preserve the observer-visible representations, Yukawa invariants, anomaly polynomial, hypercharge lattice up to the allowed U(1) normalization, and the one-Higgs branch, modulo generation relabeling, charge-conjugation convention, gauge-center quotienting, implementation hiding, and inert ancillary stabilization.

Definition 6.1c (MAR order). For $\mathfrak{S} \in \mathfrak{A}_{\mathrm{MAR}}$, define

$$C(\mathfrak{S}) = (\chi_{\text{cpl}}, N_{\text{nonab}}, N_c, N_g) \in \mathbb{N}^4,$$

where N_{nonab} is the number of connected nonabelian simple factors acting nontrivially on the coupled carrier, N_c is the dimension of the minimal intrinsically complex color-type role, and N_g is the generation count on the realized one-Higgs chiral branch. MAR orders packages by the lexicographic order on $C(\mathfrak{S})$, then quotients ties by the physical equivalence relation above.

Proposition 6.1d (well-founded MAR minima). Every nonempty MAR-admissible class has a nonempty set of MAR-minimal packages. More precisely, if $\mathcal{A} \subseteq \mathfrak{A}_{\mathrm{MAR}}$ is nonempty, then the set $C(\mathcal{A}) \subseteq \mathbb{N}^4$ has a lexicographically least element. The MAR-minimal packages in $\mathcal{A}$ are exactly the packages whose complexity vector is that least element.

Proof. The lexicographic order on $\mathbb{N}^4$ is well-founded: first minimize χ_{cpl}, then N_{nonab} on that fiber, then N_c, then N_g. Each step minimizes a nonempty subset of $\mathbb{N}$. QED.

Proposition 6.1e (meaning of MAR uniqueness). MAR uniqueness is uniqueness inside the declared low-energy economy class modulo the equivalence relation in Definition 6.1b. It is not uniqueness of a microscopic regulator representative. The source-derived response and conditional rank-15 matter packet fix the connected Standard Model gauge image, exact hypercharge lattice after normalization, and $N_c = 3$ before MAR is applied. The scalar scan fixes compatible scalar charges and Yukawa channels, not scalar attachment or multiplicity. MAR enters the economy minimum $N_g = 3$, family economy, no-extra-sector, charged-lepton, and D10 branches. Physical matter typing, global-form selection, family attachment, and scalar attachment remain separate receipts.

Note on transport obstruction. The gluing obstruction is tracked explicitly so the gauge lane states its branch conditions. TransportabilityFromOverlapGluing identifies the central criterion as $[z]_\Sigma = 0$ together with an allowed trivial-holonomy strictification, and the genuinely noncentral criterion as $o^{(2)}_\Sigma = 0$ together with at least one allowed strict representative having trivial represented G_Σ-holonomy. When $o^{(2)}_\Sigma \neq 0$, the sector is handled by crossed-module data rather than by the ordinary DR field-algebra corollary; when $o^{(2)}_\Sigma = 0$ but every strict representative has residual holonomy, it also does not enter that corollary. The full orbit q_Σ classifies the crossed-module representatives but need not determine one ordinary H^1 class. FixedCutoffBosonicSectorCategory constructs the fixed-cutoff category. On a cofinal tail carrying the compact-gauge refinement receipt, RefinementFunctorAndFiberDescent constructs the refinement/fiber descent and the realized witness theorem supplies nonempty ordinary/central compact-gauge witness data. These transport statements are independent of the source-derived finite current, and identification of that current with the Tannaka group is open.

What the finite packet derives. Incidence determines the inverse-port pairing J. The target-blind impulse/readback producer derives $R = -J$, and the finite certificate derives the unique charge-conjugate pair of rank-15 chiral projectors. On that contract-conditional branch, the finite packet gives:

- **Product structure**: the rank-15 packet contains the coupled carrier $\mathbb{C}^3 \otimes \mathbb{C}^2$, which enforces commuting color and weak actions.
- **Finite sector content**: the response algebra supplies the weak and color Lie types, while either representative of the derived conjugate projector pair supplies the doublet and triplet matter roles.
- **Charge and global form**: anomaly-forced determinant balance fixes the normalized charge lattice, and tensor descent fixes the conjugation-insensitive $\mathbb{Z}_6$ kernel on the realized tensors.
- **Scalar compatibility**: the finite scan fixes the compatible scalar charges and Yukawa channels; it does not fix scalar attachment, multiplicity, or a one-Higgs economy branch.

MAR is absent from the displayed finite gauge implication. It enters the generation and family economy, no-extra-sector, charged-lepton, and D10 branches. Physical matter typing, global-form selection, and laboratory current identification are open.

The finite packet implications use the following standard lemmas:

Lemma 6.2 (Product factorization implies product group). If $\mathsf{Sect} \simeq \mathrm{Rep}(G)$ and $\mathsf{Sect} \simeq \mathsf{Sect}_1 \boxtimes \mathsf{Sect}_2$, then

$$G \cong G_1 \times G_2, \qquad \mathsf{Sect}_i \simeq \mathrm{Rep}(G_i).$$

QED.

Lemma 6.3 (SU(2) from a pseudoreal doublet). Let H be a positive-dimensional compact connected Lie group with a faithful *irreducible* 2D pseudoreal unitary representation V. Then the connected derived subgroup of the image is $SU(2)$ acting as the fundamental doublet. Irreducibility is load-bearing: $U(1)$ acting by $\mathrm{diag}(z, z^{-1})$ is faithful and pseudoreal but reducible, and contains no $SU(2)$. Proof: a compact connected abelian group has only one-dimensional irreducibles, so the image is nonabelian; by Schur the center acts by scalars, so the derived subgroup acts irreducibly on $\mathbb{C}^2$; the only connected compact groups acting irreducibly on $\mathbb{C}^2$ have derived subgroup $SU(2)$. Finite or disconnected counterexamples are not part of the theorem package, which only applies the lemma to the identity component on the relevant nonabelian image. QED.

Lemma 6.4 (SU(3) from an intrinsically complex triplet). Let H be a positive-dimensional compact connected Lie group with a faithful irreducible 3D unitary representation W that is *intrinsically complex on its nonabelian factor*: the connected derived subgroup of the image acts by an irreducible representation not equivalent to its conjugate. Abelian twisting does not qualify: a character twist can make the full representation inequivalent to its conjugate while the nonabelian factor stays real or pseudoreal, and such twists are excluded by the hypothesis. Then the connected derived subgroup of the image is conjugate to $SU(3)$ acting as the fundamental triplet, up to finite central kernel. The hypothesis rejects the defining representation of $SO(3)$ on $\mathbb{C}^3$ (real on the nonabelian factor) and its twists $(z, R) \mapsto zR$ under $U(1) \times SO(3)$ (real on the nonabelian factor), and retains the fundamental of $SU(3)$. Proof: by Schur the center acts by scalars, so the derived subgroup acts irreducibly; two nontrivial simple factors would force a tensor decomposition of dimension at least 4, so exactly one simple factor acts; among compact simple Lie algebras only $\mathfrak{su}(2)$ and $\mathfrak{su}(3)$ carry nontrivial irreducibles of dimension at most 3; the 3D representation of $\mathfrak{su}(2)$ is real, so the intrinsically complex hypothesis leaves the fundamental of $\mathfrak{su}(3)$. This restates Lemma `lem:su3class` of The Standard Model gauge paper. The intrinsically-complex-nonabelian hypothesis is load-bearing: without it, $SO(3)$ on $\mathbb{C}^3$ satisfies every remaining hypothesis and refutes the conclusion. Finite or disconnected counterexamples are not part of the theorem package, which only applies the lemma to the identity component on the relevant nonabelian image. QED.

Lemma 6.5 (Connected abelian factor criterion). If the admissible sector package contains a connected abelian charge factor acting nontrivially on the coupled carrier, then the identity component of the abelianized reconstructed group contains a one-torus. Under the single connected abelian-factor admissibility condition, this factor is $U(1)$. QED.

Proposition 6.6 (conditional maximal faithful matter image). Under the finite source-model inverse-port response, the finite matter certificate derives $P_{\text{even}} - P_{\text{vac}}$ and $P_{\text{odd}} - P_{\text{top}}$ as the unique charge-conjugate pair of rank-15 chiral projectors. For either representative, the kernel acting trivially on all realized sectors is Z_6. The maximal faithful matter image is therefore

$$G_{\text{packet}} = \frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}.$$

Charge conjugation changes the representative and leaves this kernel unchanged. The cover and its Z_2 and Z_3 quotients also carry the local tensors. The proposition does not select the physical global form, source physical matter typing, attach laboratory currents, or identify G_{packet} with the separately reconstructed Tannaka group. QED.

Proposition 6.6a (MAR-free finite Standard Model matter-image implication). Assume the certified twelve-port incidence branch and its inverse-port pairing J. Incidence leaves a four-dimensional equivariant linear-response commutant. The target-blind impulse/readback protocol solves the common farthest-shell filter and produces $R = -J$. Its restrictions to $\mathbf{1} \oplus \mathbf{3} \oplus \mathbf{3'} \oplus \mathbf{5}$ have fixed relative signs; common reversal is the conventional overall charge-conjugation choice. Under this response, the finite current algebra is

$$\mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{u}(1).$$

The finite matter certificate then derives $P_{\text{even}} - P_{\text{vac}}$ and $P_{\text{odd}} - P_{\text{top}}$ as the unique charge-conjugate pair of rank-15 chiral projectors. Anomaly-forced determinant balance and tensor descent on every realized matter tensor then give:

- either projector representative supplies the one-generation chiral matter module with the weak-doublet and color-triplet roles on a common coupled carrier;
- determinant balance and anomaly freedom fix the exact Standard Model hypercharge lattice, up to the conventional simultaneous charge reversal;
- the color multiplicity is $N_c = 3$;
- exhaustive center action on the realized tensors fixes the conjugation-insensitive kernel to $\mathbb{Z}_6$;
- the scalar scan fixes the compatible scalar charges and Yukawa channels, without fixing scalar attachment or multiplicity.

Consequently, the maximal faithful matter image is

$$G_{\text{packet}} = \frac{SU(3) \times SU(2) \times U(1)}{Z_6}.$$

No MAR clause enters this implication. Choosing one response sign and one projector representative is the conventional charge-conjugation choice. The cover and its Z_2 and Z_3 quotients carry the same local tensors. Physical matter typing, global-form selection, laboratory current identification, rank-45 physical family attachment, scalar attachment and multiplicity, QFT realization, and identification of G_{packet} with the independently reconstructed Tannaka group are open. QED.

Theorem 6.6aa (witness landing under separate economy hypotheses). On the ordinary or central zero-obstruction bosonic branch, suppose a cofinal tail carries the compact-gauge refinement receipt. Take the source-derived finite response of Proposition 6.6a and declare the fermionic Spin category used by its matter certificate. The certificate then supplies the rank-15 chiral matter module through its derived conjugate projector pair. Under separate formal three-generation and one-Higgs economy declarations, the conditional compact-gauge construction admits a formal cofinal heat-kernel edge-sector label witness

$$\mathcal{W}_{\text{SM}} = \{Q_i, u_i^c, d_i^c, L_i, e_i^c, H\}_{i=1}^3$$

with

$$Q_i = (3,2)_{1/6}, \quad u_i^c = (\bar{3},1)_{-2/3}, \quad d_i^c = (\bar{3},1)_{1/3}, \quad L_i = (1,2)_{-1/2}, \quad e_i^c = (1,1)_1, \quad H = (1,2)_{1/2}.$$

The fixed-cutoff edge heat-kernel law gives $p_R(t) \propto d_R e^{-tC_2(R)}$, hence positive support for every finite-dimensional witness sector at finite t. This positivity concerns the bosonic internal-gauge sector labels only; positive heat-kernel weight on a boundary label does not supply a light chiral fermion multiplet or scalar carrying that label. The conditional matter certificate supplies the one-generation chiral matter content. The three-family copy and one-Higgs scalar label are separate formal economy declarations, not physical attachments. The pullback functors constructed from the receipt carry the zero-obstruction labels cofinally. Proposition 6.6a gives the maximal faithful matter image

$$\frac{SU(3) \times SU(2) \times U(1)}{Z_6}, \qquad N_c = 3,$$

with the Standard Model hypercharge lattice without MAR. The CP-capability and weak-sector UV clauses give $3 \le N_g \le 5$, and MAR selects $N_g = 3$ inside that window while excluding extra light sectors in the declared economy class. Physical matter typing, global-form selection, laboratory current attachment, scalar attachment and multiplicity, rank-45 family attachment, equality with the Tannaka group, and QFT realization are open. QED.

Theorem (Conditional Standard Model packet with MAR economy completion). Assume the certified incidence branch, its source-derived finite response, the declared fermionic Spin category, the resulting conjugate pair of rank-15 chiral projectors, anomaly-forced determinant balance, and tensor descent on every realized tensor. Assume separately the ordinary or central zero-obstruction bosonic branch; a cofinal tail carrying the compact-gauge refinement receipt; and the fixed-cutoff bosonic sector category, refinement/fiber descent, and compact-group reconstruction above. Finally assume a nonempty MAR-admissible economy class, a one-Higgs scalar attachment witnessed by $\mathcal{W}_{\text{SM}}$, and the CP-capable and weak-sector UV clauses of Axiom 5. Then the completed matter and separately assumed scalar/economy packet has maximal faithful image data

$$\left(\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \mathcal{R}_{\text{SM}}, H, \mathcal{Y}_{\text{SM}}\right), \qquad N_c = 3, \qquad N_g = 3,$$

with

$$\mathcal{W}_{\text{SM}} = \{Q_i, u_i^c, d_i^c, L_i, e_i^c, H\}_{i=1}^3,$$

$$Q_i = (3,2)_{1/6}, \quad u_i^c = (\bar{3},1)_{-2/3}, \quad d_i^c = (\bar{3},1)_{1/3}, \quad L_i = (1,2)_{-1/2}, \quad e_i^c = (1,1)_1, \quad H = (1,2)_{1/2}.$$

The gauge group, matter type, charge lattice, and $N_c = 3$ implication use no MAR clause. Within the separate economy class, the CP-capability and weak-sector UV conditions give $3 \leq N_g \leq 5$, and MAR supplies the displayed $N_g = 3$ value and exclusion of extra light sectors in this theorem. MAR also enters the family, charged-lepton, and D10 economy branches. The one-Higgs scalar is an assumption in this completed packet; the finite scan fixes only compatible scalar charges and Yukawa channels. The displayed quotient is the maximal faithful image of the realized matter tensors. The theorem does not select the physical global form, attach a laboratory current, or attach the canonical rank-three screen band to the displayed three-family matter module. That promotion requires the source-derived complex rank-45 attachment, Spin/locality, residue, excluded-band, and complement-complete refinement receipts. The gauge proof is Proposition 6.6a, the hypercharge theorem, three-color corollary, and $\mathbb{Z}_6$ kernel computation. Proposition 6.9 and MAR provide the displayed generation/economy completion. Equality with the independently reconstructed Tannaka group requires a separate commuting-square receipt. QED.

Corollary 6.6b (Three colors from the finite matter packet). On the one-generation matter package supplied by the derived conjugate projector pair in Proposition 6.6a, the color factor is the fundamental triplet of SU(3), so

$$N_c = 3.$$

This exact finite implication uses no MAR clause. Physical matter typing requires an independent construction. QED.

Corollary 6.6c (structural electroweak force and charges). Under the finite response of Proposition 6.6a, assume separately that one compatible scalar slot is physically attached as a single Higgs doublet. On that declared one-Higgs branch, the electroweak gauge factor has Lie algebra

$$\mathfrak{su}(2)_L \oplus \mathfrak{u}(1)_Y.$$

The Higgs doublet $H = (1,2)_{1/2}$ selects a nonzero neutral vacuum vector

$$\phi_0 = \frac{v}{\sqrt{2}} \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \qquad v \neq 0,$$

with

$$Q\phi_0 = (T_3 + Y)\phi_0 = 0,$$

so its vector stabilizer has unbroken generator

$$Q = T_3 + Y,$$

and the unbroken gauge factor is $\mathrm{U}(1)_Q$. The charged weak generators give

$$W^{\pm} = \frac{1}{\sqrt{2}}(W^1 \mp iW^2),$$

while the neutral $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ gauge fields span the Z/A basis on the D10 quantitative branch. Explicitly,

$$e^{i\alpha T_3} e^{i\beta Y} \phi_0 = e^{i(\beta - \alpha)/2} \phi_0,$$

so exact vector invariance ties the phases by $\beta = \alpha$, locally, and leaves $\mathrm{U}(1)_Q$. The projective ray $[\phi_0]$, however, is unchanged for independent α and β: its stabilizer is the two-torus $\mathrm{U}(1)_{T_3} \times \mathrm{U}(1)_Y$, modulo the inherited finite center. Projectivization therefore classifies the carrier ray but cannot by itself select the electromagnetic diagonal. Thus, under the finite current theorem and separate one-Higgs attachment premise, the structural package contains the weak charged-current carriers $W^{\pm}$, the neutral weak carrier Z, the electromagnetic carrier A, and the exact charge operator $Q = T_3 + Y$. The mixing angle, v, and numerical W/Z masses belong to the D10 running/matching surface rather than to this finite structural corollary. QED.

Corollary 6.6d (conditional Maxwell sector on the realized electromagnetic branch). On the conditional, separately declared one-Higgs branch of Corollary 6.6c, assume the finite current is realized by a support-visible compact-gauge connection, and let A_Q be its component on the unbroken electromagnetic factor $\mathrm{U}(1)_Q$. The abelian restriction of the curvature gives

$$F_Q = dA_Q.$$

Assume in addition the Lorentzian low-energy Maxwell action

$$S_{\mathrm{EM}}[A_Q, J_Q] = -\frac{1}{2g_Q^2} \int F_Q \wedge *F_Q + \int A_Q \wedge *J_Q,$$

the Euler-Lagrange equations are

$$dF_Q = 0, \qquad d * F_Q = g_Q^2 * J_Q.$$

After canonical electromagnetic normalization, these are Maxwell's equations. The numerical value of g_Q, equivalently α_{em}, belongs to the Ward-projected electromagnetic readout. The compact-group reconstruction and connection label do not themselves supply this action or its nonzero kinetic coefficient. QED.

2.2.3 Refinement stability and unprotected relevant operators

The state-side refinement stability used later in the gauge branch is contained in the third OPH axiom through its refinement-closure clause. By that clause the same finite local constraint family is preserved across cutoffs, so the realized UV states lie in one common finite-dimensional MaxEnt family rather than in a new coupling space at every refinement step. The selected stable branch of this family is therefore not an extra state axiom beyond the MaxEnt branch itself. The closure clause it relies on is a substantive renormalization condition: coarse-graining a finite-range Gibbs family generically generates interactions outside any fixed finite density list. The closure defect, the moment-matching I-projection realizing $R_{\ell\to L}$, and its trace-norm residual bound are quantified in Definition 2.6a and Lemma 2.6b; the clause is exactly the statement that the defect vanishes along the realized branch, and that vanishing remains assumed rather than proved. This state statement does not supply the compact-gauge refinement receipt.

Concretely, if

$$\omega_\ell(\lambda) = Z_\ell(\lambda)^{-1} \exp\left(-\sum_x \sum_a \lambda_a O_a(x) - \sum_b \mu_b Q_b \right),$$

then any refinement channel $\Phi_{\ell\to L}$ compatible with Axiom 3, including its refinement-closure clause, acts on the realized branch by an induced finite-dimensional map

$$\Phi_{\ell\to L}(\omega_\ell(\lambda)) = \omega_L(R_{\ell\to L}(\lambda)).$$

The resulting refinement-stable state branch is simply a trajectory or invariant subset of this finite-dimensional multiplier map. Accordingly, when sections speak of a refinement-stable directed colimit of sectors, persistence along this MaxEnt branch is only the state-side input. The additional compact-gauge receipt must supply finite extendability, center-compatible block embeddings, coherent surjective compact-group maps, preservation of the full transport obstruction, and finite tensor realizations. The MaxEnt/mixing package supplies none of those data and does not decide whether the resulting certified colimit is trivial or nontrivial.

The consensus paper makes the state-side coarse-graining connection explicit. Given quotient normal-form maps n_r, obstruction maps h_r, and coarse-graining maps (ρ_{sr}, χ_{sr}), reconciliation commutes with coarse-graining at the macroscopic readout scale whenever the two square defects

$$d_r^Q(\rho_{sr} n_s(x), n_r \rho_{sr}(x)), \qquad d_r^{\mathcal{H}}(\chi_{sr} h_s(x), h_r \rho_{sr}(x))$$

are controlled. Exact refinement naturality gives zero defect; approximate RG matching is theorem-grade only to the extent that these defects are bounded on the selected branch. This keeps the MaxEnt refinement map $R_{\ell\to L}$ connected to the reconciliation law without promoting arbitrary coarse-graining channels to OPH laws.

Relevant operators that are neither symmetry-forbidden nor retained in the selected constraint family are precisely the directions that try to push the flow off that stable branch.

Lemma 6.7 (refinement-stable MaxEnt branch forbids unprotected relevant operators). Assume the local finite-constraint MaxEnt/refinement branch of the third OPH axiom. Let $\mathcal{O}$ be a gauge-invariant Lorentz-scalar relevant deformation in the emergent EFT sense ($\Delta < 4$ in $3+1$D), allowed by symmetry and absent from the retained constraint family. Then the selected branch can keep the coupling of $\mathcal{O}$ at zero only if symmetry forbids that direction or the constraint family explicitly retains it. Otherwise generic refinement induces a nonzero component along that direction and the flow leaves the selected branch. This is a branch-persistence statement, not a universal entropy-ordering theorem for arbitrary off-branch phases.

Proof sketch. Linearize the induced refinement map $R_{\ell\to L}$ on the finite-dimensional multiplier space around the selected branch. A relevant operator gives an unstable direction with scaling exponent $y > 0$. If $\mathcal{O}$ is not fixed by symmetry and is not part of the retained constraint family, generic UV mismatch produces a nonzero component along that direction. Because $y > 0$, repeated coarse-graining amplifies the component and pushes the flow off the selected branch unless one fine-tunes it away at every scale or protects it by symmetry or by the declared constraint family. QED.

Corollary 6.8 (chirality selector). A gauge-invariant Dirac mass term is a relevant scalar. If both chiralities exist in conjugate representations, the mass term is allowed and will be generated under refinement unless symmetry-forbidden or explicitly retained as a protected constraint. Therefore keeping light fermions on the selected refinement-stable branch requires chiral matter content or an explicit protecting mechanism. QED.

2.2.4 Generation number from CKM CP capability and weak-sector completability (derived from MAR)

Anomaly cancellation is generation-by-generation, so it does not fix the number of generations. On the separately declared one-Higgs economy branch, the lower and upper bounds come from the CP-capability and weak-sector UV clauses declared in MAR, and MAR then selects the minimum.

Proposition 6.9 (Conditional MAR minimum N_g = 3). On a separately declared one-Higgs quark economy class built over the conditional matter packet of Proposition 6.6a, with the derived $N_c = 3$ from Corollary 6.6b and with the CP-capability and weak-sector UV clauses contained in Axiom 5, the admissible generation window is $3 \le N_g \le 5$, and the fourth MAR economy coordinate selects

$$N_g = 3.$$

Inputs.

1. **Intrinsic CKM CP capability** is part of the declared quark economy branch through clause (v) of Axiom 5.
2. **Weak-sector UV completability** is part of the same declared one-Higgs economy branch through clause (vi) of Axiom 5.
3. **MAR minimality** acts on the same realized branch once the first three complexity entries are fixed.
4. Use the derived $N_c = 3$ from Corollary 6.6b.

Step 1: CKM CP-capability lower bound. The number of physical CP-violating phases in an N_g × N_g CKM matrix is:

$$\#(\text{CP phases}) = \frac{(N_g - 1)(N_g - 2)}{2}.$$

- For N_g = 1, 2: this is 0 → **no intrinsic CKM CP capability**.
- For N_g = 3: this is 1 → **intrinsic CKM CP capability is available**.

So intrinsic CKM CP capability requires:

$$N_g \ge 3.$$

Step 2: SU(2) asymptotic freedom upper bound. The one-loop coefficient is:

$$b_{\text{SU(2)}} = \frac{22}{3} - \frac{1}{3}N_g(N_c + 1) - \frac{1}{6},$$

where the final $-1/6$ is the contribution of one complex Higgs doublet. Asymptotic freedom means $b_{\text{SU(2)}} > 0$, i.e.,

$$N_g(N_c + 1) < \frac{43}{2}.$$

With N_c = 3, we have N_c + 1 = 4, so:

$$4N_g < \frac{43}{2} \quad \Rightarrow \quad N_g \le 5.$$

Combining: 3 ≤ N_g ≤ 5.

Step 3: MAR selection. These are not extra post-MAR selectors: they are the numerical content of clauses (v) and (vi) of Axiom 5 on the declared one-Higgs economy branch. With MAR lexicographic minimality selecting the window minimum, the upper bound $N_g \le 5$ from the SU(2) clause is a consistency check and does no selection work. Given the allowed window {3, 4, 5}, MAR (fourth component of the complexity vector $C(\mathfrak{S})$) selects the smallest viable choice:

$$N_g = 3.$$

QED.

Why this is convincing.

- The clause set selects a **single integer** inside the admissible window $\{3, 4, 5\}$; clause (v) (CKM CP capability) is a declared input that sets the lower edge, counted in the closure ledger.
- It uses **two explicit MAR clauses** (intrinsic CKM CP capability and weak-sector UV completability) plus MAR's lexicographic minimality, rather than a separate post hoc admissibility menu.
- It is not a fit to a continuous number.
- Under the stated finite gauge-packet and economy hypotheses, this is a conditional selection result. It is not forced by anomaly cancellation, the A_5 graph, or the target-free source reduct.

Corollary 6.9a (finite gauge packet with MAR economy minimum). Assume the incidence branch, its source-derived finite response, and the fermionic Spin category of Proposition 6.6a. On the ordinary or central zero-obstruction bosonic branch, assume the MAR-admissible economy class is nonempty and contains the separately declared one-Higgs chiral matter packet used through Proposition 6.9. Then the finite packet fixes

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_c = 3,$$

with the hypercharge lattice and finite quotient of Proposition 6.6, without MAR. Every MAR-minimal representative in the declared economy class also has $N_g = 3$ and no admitted extra light sector. The displayed quotient is the maximal faithful matter image. This corollary does not select the physical global form or assert scalar or family attachment, equality with the Tannaka group, QFT realization, or source-law uniqueness.

Proof. Proposition 6.6a derives the finite current algebra and the conditional matter, charge, and faithful-image implications. Corollary 6.6b fixes $N_c = 3$, and the hypercharge and quotient propositions fix the visible charge lattice and $\mathbb{Z}_6$ kernel. Proposition 6.1d supplies the MAR minima, and Proposition 6.9 fixes only the economy value $N_g = 3$; the same economy rule removes admitted extra light sectors. The equivalence relation removes relabelings, the conventional simultaneous charge reversal, gauge-center conventions, and inert implementation data. QED.

2.2.5 Hilbert-space formulation of gluing data

Let {P_i} be a good cover of the screen. For each patch, fix a representation

$$\pi_i : \mathcal{A}_i \to \mathcal{B}(\mathcal{H}_i).$$

For each overlap, choose a unitary intertwiner

$$U_{ij} : \mathcal{H}_j \to \mathcal{H}_i$$

such that for all $O \in \mathcal{A}_{ij}$,

$$\pi_i(O) = U_{ij} \pi_j(O) U_{ij}^{\dagger}.$$

Normalize U_ii = 1 and U_ji = U_ij†.

Lemma 6.10 (centrality on triple overlaps). On a triple overlap define

$$\Omega_{ijk} := U_{ij} U_{jk} U_{ki}.$$

For all $O \in \mathcal{A}_{ijk}$,

$$\Omega_{ijk}\pi_i(O) = \pi_i(O)\Omega_{ijk}.$$

Proof. Conjugation by U_ki sends π_i(O) to π_k(O), by U_jk to π_j(O), by U_ij back to π_i(O). Thus conjugation by Ω_ijk fixes π_i(O), so Ω_ijk commutes with π_i(O). QED.

Lemma 6.11 (gauge behavior). If Ũ_ij = V_i U_ij V_j† with V_i acting trivially on overlap observables, then

$$\tilde{\Omega}_{ijk} = V_i\Omega_{ijk}V_i^{\dagger}.$$

In particular, if Ω_ijk is central, its class is gauge invariant. QED.

2.2.6 Loop obstruction class (central defect)

Assume the defect is central, the overlap centers are identified with one fixed abelian unitary coefficient group Z_Σ, and overlap transport acts trivially on Z_Σ. With $U_{ji} = U_{ij}^{-1}$, define the implementer-level multiplier by

$$\Omega_{ijk} = U_{ij}U_{jk}U_{ki} = z_{ijk} \in Z_\Sigma,$$

equivalently $U_{ij}U_{jk} = z_{ijk}U_{ik}$. This is the abelian truncation of the full 2-group obstruction in Section 3.4. The multiplier must be defined at the implementer level: $\mathrm{Ad}(z_{ijk})$ is the identity for central z_{ijk} and does not by itself record the defect.

Theorem 6.12 (central defect strictification and residual holonomy). The family $\{z_{ijk}\}$ is an untwisted Čech 2-cocycle with coefficients in Z_Σ, and its class $[z] \in \check{H}^2(N_\Sigma, Z_\Sigma)$ is gauge invariant. On any quadruple overlap P_{ijkl},

$$z_{jkl}z_{ikl}^{-1}z_{ijl}z_{ijk}^{-1} = 1.$$

A central edge rephasing by a cochain in $C^1(N_\Sigma, Z_\Sigma)$ removes every triangle defect iff $[z] = 0$. After such a rephasing, write U_{ij}^{str} for the resulting genuine edge 1-cocycle. Let $\mathcal{C}_\alpha$ be the full transported charge block, including its carrier and multiplicity/intertwiner data. Loop-coherent, endpoint-only transport exists on that block iff, in addition,

$$\mathrm{Hol}_{\alpha,U^{\mathrm{str}}}([\gamma]) = U_\gamma^{\mathrm{str}}\big|_{\mathcal{C}_\alpha} = \mathbf{1}_{\mathcal{C}_\alpha} \qquad \text{for every } [\gamma] \in \pi_1(|N_\Sigma|, i_*).$$

Thus strict path independence is equivalent to the combined conditions $[z] = 0$ and trivial represented holonomy for one allowed strictification.

Proof. The fixed-coefficient and trivial-action hypotheses place every multiplier in the same group Z_Σ. Associativity compares $(U_{ij}U_{jk})U_{kl}$ with $U_{ij}(U_{jk}U_{kl})$ on a quadruple overlap. Substituting $U_{ij}U_{jk} = z_{ijk}U_{ik}$ gives the displayed Čech cocycle identity. A vertex-frame change $U_{ij} \mapsto V_iU_{ij}V_j^{\dagger}$ conjugates z_{ijk}, as in Lemma 6.11, and therefore leaves a central multiplier unchanged. Separately, a central edge rephasing $U_{ij} \mapsto a_{ij}^{-1}U_{ij}$, with $a \in C^1(N_\Sigma, Z_\Sigma)$, replaces z by the corresponding Čech coboundary. Hence $[z] = 0$ is equivalent to removing all triangle multipliers by a central 1-cochain. This yields a strict 1-cocycle, not automatically endpoint-only transport. For two paths p, p' with the same endpoints, the remaining discrepancy is the holonomy of the closed loop $p'^{-1}p$; it vanishes for every pair exactly when the displayed holonomy representation on $\mathcal{C}_\alpha$ is trivial. Conversely, strict path independence makes every triangle comparison and every represented closed-loop action trivial. QED.

If overlap transport acts nontrivially on a local system of centers, the same comparison contains the transported factor $\varphi_{ij}(z_{jkl})$ and must be written with the corresponding twisted Čech differential; that branch is not the untwisted theorem stated here.

2.2.7 EFT reduction to anomaly cancellation

Assume ExtEFT: a low-energy 3+1D chiral gauge theory exists with group G. Then the obstruction class $[z]$ is structurally analogous to the EFT 't Hooft anomaly class and is expected to map to it after a separate anomaly-descent construction. That map lies outside this manuscript, so $[z] = 0$ is used here as the internal triangle-defect strictifiability condition rather than a proved equivalence to EFT anomaly cancellation; the separate trivial-holonomy condition of Theorem 6.12 is required for transportability.

2.2.8 Hypercharge from anomaly freedom and Yukawas

Theorem 6.13 (Hypercharge from anomaly freedom and Yukawas). Assume gauge group SU(N_c) × SU(2) × U(1)_Y and one generation of left-handed Weyl fermions (Q, u^c, d^c, L, e^c), with a compatible scalar charge and Yukawa channels

$$QHu^c, \qquad QH^\dagger d^c, \qquad LH^\dagger e^c.$$

Then anomaly freedom and Yukawa invariance fix the hypercharges up to an overall normalization, yielding the Standard Model pattern for N_c = 3. Simultaneous reversal of all charges is the conventional charge-conjugation choice. In the OPH application, the conditional finite matter certificate derives the conjugate rank-15 projector pair and scans the compatible scalar charges and Yukawa channels. It does not fix scalar attachment or multiplicity.

Proof. Yukawa invariance gives

$$Y_u = -(Y_Q + Y_H), \quad Y_d = -Y_Q + Y_H, \quad Y_e = -Y_L + Y_H.$$

Anomaly cancellation yields

$$\begin{aligned} SU(2)^2U(1): &\quad N_cY_Q + Y_L = 0, \\ \mathrm{grav}^2U(1): &\quad 2N_cY_Q + N_cY_u + N_cY_d + 2Y_L + Y_e = 0. \end{aligned}$$

Solving gives

$$Y_L = -N_cY_Q, \quad Y_H = N_cY_Q, \quad Y_u = -(N_c+1)Y_Q, \quad Y_d = (N_c-1)Y_Q, \quad Y_e = 2N_cY_Q.$$

With these relations, SU(N_c)^{2}U(1) and U(1)^3 anomalies vanish automatically. Fixing the normalization by Q = T_3 + Y and Q(ν_L)=0 gives

$$Y_Q = \frac{1}{2N_c}.$$

For N_c = 3,

$$Y_Q = \frac{1}{6}, \quad Y_L = -\frac{1}{2}, \quad Y_e = 1, \quad Y_u = -\frac{2}{3}, \quad Y_d = \frac{1}{3}, \quad Y_H = \frac{1}{2}.$$

Without the compatible Yukawa channels, the cubic anomaly leaves two discrete branches (Y_u, Y_d exchange). The channel scan selects the compatible scalar charge branch. It does not select scalar attachment, multiplicity, or a one-Higgs economy. QED.

Corollary 6.13a (Exact rational hypercharges). With the conditional $N_c = 3$ of Corollary 6.6b and one fixed overall charge convention, the hypercharge assignments are uniquely fixed to exact rational values:

$$Y_Q = \tfrac{1}{6}, \quad Y_L = -\tfrac{1}{2}, \quad Y_u = -\tfrac{2}{3}, \quad Y_d = \tfrac{1}{3}, \quad Y_e = 1, \quad Y_H = \tfrac{1}{2}.$$

Why this is convincing.

- These are **exact rationals**, not approximate numbers.
- Their ratios are fixed by anomaly freedom + Yukawa invariance, and the absolute lattice is fixed by the standard normalization, with no continuous parameters to adjust.
- The resulting lattice coincides exactly with the Standard Model assignments on the declared matter packet.

2.2.9 Witten anomaly on the realized color branch

Theorem 6.14 (Witten anomaly consistency on the realized color branch). On the conditional one-generation package derived in Proposition 6.6a and Corollary 6.6b, the global SU(2) anomaly is absent generation by generation.

Inputs.

1. The realized color factor is the SU(3) triplet from Corollary 6.6b.
2. The matter content per generation includes:
   - one left-handed quark doublet Q which is an SU(2) doublet and carries color,
   - one left-handed lepton doublet L which is an SU(2) doublet and color singlet.
3. **Witten's global SU(2) anomaly constraint** (Witten, 1982): the number of left-handed SU(2) doublets must be even.

Proof. Count SU(2) doublets per generation:

- Quark doublets: N_c copies (one per color),
- Lepton doublets: 1 copy.

Total doublets per generation:

$$N_c + 1.$$

Substituting the realized value $N_c = 3$ gives

$$N_c + 1 = 4.$$

Witten anomaly cancellation therefore requires an even number of doublets, and the realized branch satisfies it exactly:

$$4 \equiv 0 \pmod 2.$$

So the realized triplet-doublet package is globally consistent generation by generation. QED.

Theorem-stack role. Under the source-derived response and declared matter typing, the D8 finite packet fixes $N_c = 3$ and carries it into D9. Witten's anomaly is a consistency check on that derived triplet-doublet package, not a residual selector that upgrades oddness to 3.

Why this is convincing.

- It checks a **global anomaly constraint** on the conditional matter branch without introducing a new selector.

- The parity constraint is independent of continuous parameters, RG running, masses, or Yukawa values.
- It confirms that the realized SU(3) triplet plus lepton doublet package is globally consistent generation by generation.

2.2.10 Bond-dimension gatekeeping

In tensor-network or code realizations, gauge actions act on edge factors of size χ, so emergent compact gauge groups embed in $\mathrm{U}(\chi)$. This shows a capacity constraint: accommodating SU(3) color and SU(2) weak factors shows $\chi \geq 6$ in the minimal case, consistent with the finite gauge packet.

2.2.11 Classical carrier modes and the quantum-particle gate

An abstract compact group identifies a symmetry type, and a classical field equation identifies a classical dynamical branch. Neither datum alone constructs a propagating quantum particle. The following receipt separates those levels.

The upstream structural results remain separate receipts. Conditional on the compact-gauge refinement receipt, Theorem 6.1 constructs the refinement-limit bosonic edge-sector category and its compact group; a field-algebra realization with that group as a local gauge symmetry requires the DHR/DR hypotheses (R1)–(R6) of Definition 6.1r. On the realized electroweak branch, a chosen nonzero neutral Higgs vacuum *vector* has stabilizer $\mathrm{U}(1)_Q$, whereas its projective ray alone has the larger two-torus stabilizer recorded above. On the gravitational side, Theorems 4.3c–4.3e and 5.1 conditionally supply the observer-facing geometry, Lorentzian event-manifold, and classical Einstein branch. None of these structural outputs supplies the action Hessian or quantum data required below.

Definition 6.18 (carrier-mode and quantum-particle receipts). Fix a structural invariant speed $c_\star$. This is the common invariant null-cone conversion speed. Its representation as $299\,792\,458\ \mathrm{m\,s^{-1}}$ is exact by the SI definition of the metre and is not a prediction of that decimal magnitude. A *classical massless carrier-mode receipt* for a field q_X consists of: (C1) a stated background and phase; (C2) an explicit quadratic action with positive nonzero physical kinetic coefficient; (C3) a gauge fixing or constraint reduction with a finite physical projector Π_X; and (C4) a positive reduced Hamiltonian. Separately, the action Hessian restricted to the physical subspace must have the form

$$K_X^{\mathrm{phys}}(\omega,\mathbf{k}) = Z_X\,(\omega^2 - c_\star^2|\mathbf{k}|^2)\Pi_X, \qquad Z_X > 0.$$

Equivalently, the free reduced Green function exhibits the action-level pole

$$D_X^{\mathrm{phys}}(\omega,\mathbf{k}) = \frac{i\,\Pi_X/Z_X}{\omega^2 - c_\star^2|\mathbf{k}|^2 + i0}.$$

This receipt proves a classical massless carrier mode and the absence of an algebraic mass term in that declared quadratic action. It does not prove a particle.

A *quantum-particle receipt* additionally requires: (Q1) a positive-energy vacuum quantization and a physical Hilbert space obtained by constraint reduction or BRST cohomology; (Q2) a physical two-point function whose Källén–Lehmann measure contains a positive-residue massless contribution

$$d\rho_X(\mu^2) = Z_X^{\mathrm{pole}}\,\delta(\mu^2)d\mu^2 + d\rho_X^{\mathrm{cont}}(\mu^2), \qquad Z_X^{\mathrm{pole}} > 0;$$

or an equivalent joint energy–momentum spectrum containing a positive-residue $p^0 > 0,\ p^2 = 0$ mass shell (at fixed $\mathbf{k}$, $\omega = c_\star|\mathbf{k}|$); and (Q3) the stability/asymptotic-state or LSZ hypotheses appropriate to the phase, including a deconfined asymptotic sector for a colored carrier. Only a receipt passing (C1)–(C4) and (Q1)–(Q3) licenses a quantum-particle claim. A reconstructed group, a connection label, or a classical field equation by itself fails this gate.

Theorem 6.18 (action-level modes on the Maxwell, pure-Yang–Mills, and Einstein branches). The following statements hold on the stated additional action and phase branches.

i. **Electromagnetic branch.** Suppose the source-free ordinary electromagnetic vacuum $J_Q = 0$ has an unbroken $\mathrm{U}(1)_Q$ connection with the Lorentzian Maxwell action

$$S_Q = -\frac{1}{4g_Q^2}\int F_{Q,\mu\nu}F_Q^{\mu\nu}\, d^4x, \qquad 0 < g_Q^2 < \infty,$$

about the ordinary vacuum, with no Higgs, Stueckelberg, medium, or nonlocal quadratic mass operator. In radiation gauge the reduced variables are the two transverse components and

$$H_Q^{(2)} = \frac{1}{2g_Q^2}\int (|\mathbf{E}_T|^2 + |\mathbf{B}|^2) d^3x, \qquad K_{Q,T} = g_Q^{-2}(\omega^2 - c_\star^2|\mathbf{k}|^2)\Pi_T, \qquad \operatorname{rank}\Pi_T = 2.$$

Equivalently, covariant ξ-gauge gives

$$D_{\mu\nu}^Q(k) = \frac{-ig_Q^2}{k^2 + i0}\left(\eta_{\mu\nu} - (1-\xi)\frac{k_\mu k_\nu}{k^2 + i0}\right),$$

whose longitudinal part decouples from a conserved current. Thus this branch has two classical massless electromagnetic carrier modes. Calling them photons requires the quantum-particle receipt.

ii. **Pure Yang–Mills branch.** Suppose a compact group G is supplied with the Lorentzian two-derivative pure Yang–Mills action with $0 < g^2 < \infty$, expanded about the topologically trivial background in a gauge $\bar{A}_\mu = 0$, in a perturbative or deconfined phase with no Higgs condensate. Writing the perturbation as a_μ^a, its quadratic action is

$$S_{\mathrm{YM}}^{(2)} = -\frac{1}{4g^2}\sum_a \int (\partial_\mu a_\nu^a - \partial_\nu a_\mu^a)(\partial^\mu a^{a,\nu} - \partial^\nu a^{a,\mu}) d^4x.$$

Lorenz gauge gives the standard covariant inverse, while Gauss' law and radiation gauge leave two transverse components per generator, with

$$H_{\mathrm{YM}}^{(2)} = \frac{1}{2g^2}\sum_a \int (|\mathbf{E}_T^a|^2 + |\mathbf{B}^a|^2) d^3x \geq 0,$$

with strict positivity on nonzero reduced finite-energy modes of $\mathbf{k} \neq 0$ under the stated boundary conditions. Their quadratic transverse kernel and free Green function are

$$K_{\mathrm{YM},T}^{ab} = \delta^{ab} g^{-2}(\omega^2 - c_\star^2|\mathbf{k}|^2)\Pi_T, \qquad D_{\mathrm{YM},T}^{ab} = \frac{ig^2\delta^{ab}\Pi_T}{\omega^2 - c_\star^2|\mathbf{k}|^2 + i0}, \qquad \operatorname{rank}\Pi_T = 2$$

so there are $2\dim G$ perturbative classical transverse modes and no hard mass in that quadratic expansion. It does not produce a gauge-invariant asymptotic gluon. In the confining QCD phase, and on a gauge-invariant Yang–Mills branch with a positive physical gap, the colored free pole need not occur in the physical spectrum.

iii. **Pure Einstein branch.** Suppose, in addition to the derived classical Einstein equation, that the field-content branch is the two-derivative Einstein–Hilbert action with $G > 0$ about a Minkowski vacuum with $\Lambda = 0$, without higher-curvature kinetic terms, bimetric mixing, or extra scalar/vector fields. After de Donder gauge and residual-gauge reduction, the two transverse-traceless components obey

$$S^{(2)}_{\mathrm{TT}} = \frac{1}{64\pi G}\int \Big[(\partial_t h^{\mathrm{TT}}_{ij})^2 - c_\star^2(\nabla h^{\mathrm{TT}}_{ij})^2\Big]\, d^4x,$$

$$H^{(2)}_{\mathrm{TT}} = \frac{1}{64\pi G}\int \Big[(\partial_t h^{\mathrm{TT}}_{ij})^2 + c_\star^2(\nabla h^{\mathrm{TT}}_{ij})^2\Big]\, d^3x, \qquad D^{\mathrm{TT}}_{ij,kl}(k) \propto \frac{i\Pi^{\mathrm{TT}}_{ij,kl}}{k^2 + i0}, \qquad \operatorname{rank}\Pi^{\mathrm{TT}} = 2.$$

Thus the pure Einstein linearization has two classical massless spin-two wave modes on this background. On a curved Einstein background the relevant statement is instead about the Lichnerowicz operator and its boundary conditions; a Poincaré particle pole is not automatic. The present OPH derivation does not construct the metric-field quantization, physical graviton Hilbert space, or positive-residue two-point pole required by (Q1)–(Q3).

Proof. For (i), expand the Maxwell action to second order. Gauss' law removes the longitudinal canonical pair, leaving two transverse polarizations with the displayed positive Hamiltonian and dispersion $\omega^2 = c_\star^2|\mathbf{k}|^2$. Inverting the gauge-fixed Hessian gives the displayed covariant Green function. For (ii), write $F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + O(A^2)$; the quadratic Hessian is one Maxwell Hessian per Lie-algebra generator, while the nonabelian interactions begin at cubic order. For (iii), the second variation of the Einstein–Hilbert action is the Fierz–Pauli kinetic form. De Donder gauge and its residual gauge freedom reduce it to the transverse-traceless action shown, with the $+$ and $\times$ polarizations. These Hessian calculations establish (C1)–(C4). None constructs the quantum data in (Q1)–(Q3), so no unconditional particle conclusion follows. □

Corollary 6.18a (conditional particle interpretation). If a branch in Theorem 6.18 is also supplied with its quantum-particle receipt, its positive-residue $\delta(\mu^2)$ term defines a massless particle pole with the displayed physical polarizations. Without that receipt, this paper claims only the corresponding classical or perturbative carrier mode. In particular, the Tannaka group alone does not imply a Maxwell kinetic term or deconfined phase, and the classical Einstein equation alone does not imply a graviton state.

Hard-term versus physical-spectrum boundary. The vanishing quadratic parameters $\mu^2_{Q,\mathrm{quad}}$, $\mu^2_{\mathrm{YM,quad}}$, and $\mu^2_{\mathrm{EH,TT}}$ in the displayed actions are action-level statements, not universal exact particle masses. Gauge or diffeomorphism redundancy can coexist with Higgs/Stueckelberg completions, plasma or medium self-energies, confinement, higher-derivative poles, bimetric sectors, or additional massive fields. Excluding those possibilities requires the declared field-content and phase hypotheses rather than the symmetry label alone.

2.2.12 Quotient-protected charge quantization

Theorem 6.19 (Conditional charge quantization for packet representations). If the maximal faithful matter image is

$$G_{\mathrm{packet}} = \frac{\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)}{Z_6},$$

as fixed in Proposition 6.6 by tensor descent on the derived projector pair, and if one separately assumes the electroweak branch with $Q = T_3 + Y$, then

Every color-singlet packet representation has integer Q-charge.

Equivalently, the conditional representation packet contains no color-singlet representation with charge such as ±1/3. This statement does not establish QFT realization or a physical particle spectrum.

Proof. The Z_6 quotient identifies the center elements (e^{2πi/3}, -1, e^{iπ/3}) ∈ SU(3) × SU(2) × U(1) with the identity. For a color-singlet state ($\tau = 0$), the SU(3) factor acts trivially. The resulting identification requires the SU(2) × U(1) quantum numbers to satisfy

$$(-1)^{2j} \cdot e^{i\pi n/3} = 1,$$

where j is the SU(2) spin and n = 6Y is the integer hypercharge label. This gives n ≡ -6j (mod 6), i.e., n ≡ 0 (mod 6) for integer j and n ≡ 3 (mod 6) for half-integer j. Equivalently: Y is integer when j is integer, and Y is half-integer when j is half-integer.

After electroweak breaking, Q = T_3 + Y. For integer j, $T_3 \in \mathbb{Z}$ and Y ∈ ℤ, so Q ∈ ℤ. For half-integer j, $T_3 \in \mathbb{Z} + 1/2$ and Y ∈ ℤ + 1/2, so Q = (half-integer) + (half-integer) ∈ ℤ. In both cases, Q ∈ ℤ. QED.

Experimental bound: No fractionally charged color-singlet particles have been observed. Three independent high-precision bounds confirm this:

1. **Neutrality of matter** (PDG 2024): The proton-electron charge sum satisfies

$$|q_p + q_e|/e < 1 \times 10^{-21},$$

   confirming charge quantization to 21 decimal places.

2. **Fractional charge searches in bulk matter**: Silicone oil drop experiments limit fractionally charged particle abundance to

$$(\text{fractionally charged particles})/\text{nucleon} \lesssim 10^{-22}.$$

3. **Collider searches** (CMS, PRL 134, 2025): Exclusions for stable particles with $q \in [e/3, 0.9e]$ up to masses $\sim$ 640 GeV (95% CL).

Existing searches are consistent with this conditional structural consequence.

The representation-theoretic implication is exact under its stated response and electroweak premises. The searches do not discharge those premises or establish the open QFT realization.

2.2.13 Coupling extraction from edge-sector probabilities

The edge-center completion (Theorem 2.3) yields sector probabilities p_α on collar boundaries. These probabilities encode the renormalized gauge coupling through a heat-kernel/Laplacian weighting law.

Abelian case (Z_n). For a Z_n gauge theory, the edge sectors are labeled by charge q ∈ {0, 1, ..., n-1}. The correct "Casimir" eigenvalue is the Laplacian eigenvalue of the boundary random walk:

$$\lambda_q = 4 \sin^2\left(\frac{\pi q}{n}\right).$$

Note: only in the limit n → ∞ and q ≪ n does λ_q ≈ $(2\pi q/n)^2 \propto q^2$. For finite n, the exact form is essential.

The sector probabilities follow a heat-kernel law:

$$p_q \propto e^{-t(\mu)\lambda_q},$$

where t(μ) is the dimensionless heat-kernel parameter encoding the renormalization scale. It is distinct from the BW modular parameter and from operational time. The extraction formula is:

$$t(\mu) = -\frac{\log(p_q/p_0)}{\lambda_q}, \qquad g^2_{\text{ent}}(\mu) = \frac{t(\mu)}{2\pi}.$$

Consistency requires that t extracted from different charges q agrees; this is verified numerically (see Section 6.14).

Electric-center measurement. The edge sectors are measured using the *electric-center* prescription. For a region A and boundary vertex v ∈ ∂A, define the restricted star operator:

$$Q_v^{(A)} = \prod_{\ell \in \text{star}(v) \cap A} X_\ell^{\pm 1},$$

where X_ℓ is the shift operator on link ℓ. The sector projectors are:

$$P_{v,q} = \frac{1}{n} \sum_{m=0}^{n-1} \omega^{-mq} \left(Q_v^{(A)}\right)^m, \qquad \omega = e^{2\pi i/n},$$

and the probabilities are p_{v,q} = ⟨P_{v,q}⟩. This electric-center operator, built from X's rather than Z's, correctly captures the boundary gauge charge/flux that labels entanglement edge sectors.

Non-abelian generalization. For SU(N) gauge theories, the edge sectors are labeled by irreducible representations with probabilities:

$$p_j \propto d_j \, e^{-t(\mu) C_2(j)},$$

where d_j is the dimension and C_2(j) the quadratic Casimir. Extraction:

$$t(\mu) = -\frac{\log(p_j/p_0)}{C_2(j)}, \qquad g^2_{\text{ent}}(\mu) = \frac{t(\mu)}{2\pi}.$$

Theoretical derivation. The fixed-cutoff same-overlap edge law is carried by the finite microphysics model once the declared overlap-sector projectors, reversible thermal branch, and quasi-local local-Gibbs generator of Theorem 2.6 are in hand. On this synthesis surface, the point is to summarize that fixed-cutoff model package and to state the refinement/Peter–Weyl lift boundary explicitly rather than to relocate the leaf-level finite calculation. The finite calculation is not a physical federation, source-instrument, or support-S^2 receipt. The Laplacian spectrum selection and the one-sided sector rank do not follow from MaxEnt, gauge invariance, and locality alone; both are named hypotheses of the theorem.

Theorem 6.20 (Fixed-cutoff edge-sector law and refinement lift boundary). Under the OPH axioms, the fixed-cutoff overlap-gauge realization of Sections 2.3 and 3.2, the quasi-local local-Gibbs form supplied by Theorem 2.6, and hypotheses (EH-1) and (EH-2) below, the regulated same-overlap edge-sector probability distribution satisfies:

$$p_R = \frac{d_R\, e^{-t\lambda_R}}{\sum_{R'} d_{R'}\, e^{-t\lambda_{R'}}}$$

where λ_R is the Laplacian eigenvalue on the R-isotypic component and t is determined by the collar Gibbs parameter.

Hypothesis (EH-1, edge Hamiltonian). The MaxEnt edge generator is the factorwise group Laplacian, $H_{\text{edge}} = t\,\Delta_G$ (for a product group $G = \prod_i G_i$, a positive combination $\sum_i c_i \Delta_{G_i}$); for finite groups, the Cayley-graph Laplacian with $\lambda_R = |S| - (1/d_R)\sum_{s\in S}\chi_R(s)$. This is a declared microphysics input, not a MaxEnt consequence: gauge invariance and the Gibbs form force only $H_{\text{edge}} = \sum_R h_R P_R$, and any function of the sector label (for instance $h_R = C_2(R)^2$, or a finite projector energy) is gauge invariant and realizable by a spatially local lattice Hamiltonian at fixed cutoff. Differential order on group space is not spatial finite-range locality, so no locality argument excludes those alternatives. Within second-order bi-invariant differential operators the Laplacian is unique up to the stated scale freedom, and the appendix supplies a model-specific mechanism that derives (EH-1) for one declared Markov chain: a proposal kernel with weighted symmetry and a Casimir acceptance rule has the stated law as its stationary distribution. That derivation covers that chain, not every OPH/MaxEnt edge branch; universality of (EH-1) across branches is open and tested numerically in Section 6.14.

Hypothesis (EH-2, one-sided edge algebra). The probability space is the one-sided electric-center edge algebra $\mathcal{A}_{\text{edge}} = \bigoplus_R B(V_R)$, the Gauss-law constrained algebra visible from one side of the cut, on which the sector projector has rank d_R. The convention is fixed here once: on the full two-sided space $L^2(G) \cong \bigoplus_R V_R \otimes V_R^*$ the R-isotypic projector has rank d_R^2, and a Gibbs trace taken there gives sector weights $\propto d_R^2 e^{-t\lambda_R}$ (the closed-surface partition sum $\sum_R d_R^2 e^{-tC_2(R)}$ is that two-sided object). The d_R law of the theorem is a statement about the one-sided algebra and is not obtained by tracing the two-sided Gibbs state, since the two-sided reduction keeps weight $d_R^2 e^{-t\lambda_R}$. (EH-2) declares the one-sided state assignment directly, consistent with the Markov normal form in which the edge factor on one side contributes $\log d_R$ to the entropy.

Proof.

Step 1 (Edge Hilbert space). From the fixed-cutoff overlap-gauge realization, the edge degrees of freedom at a boundary circle $\Sigma = \partial C$ live in a Hilbert space transforming under the gauge group G. Microscopically this is a finite-dimensional quantum-link edge space on the declared overlap interface; the effective representation-theoretic description used only for the refinement lift is modeled by $L^2(G)$ with its Peter–Weyl decomposition [50], and the physical one-sided carrier is the (EH-2) algebra.

Step 2 (Gauge invariance). The Gauss law constrains physical states. For an entanglement cut at Σ, the physical edge observables from one side decompose sectorwise as $\mathcal{A}_{\text{edge}} = \bigoplus_R B(V_R)$ by (EH-2).

Step 3 (Edge Hamiltonian). By (EH-1) the MaxEnt generator restricted to edge modes is $H_{\text{edge}} = \sum_R \lambda_R P_R$ with P_R the (EH-2) sector projector of rank d_R.

Step 4 (MaxEnt selection). MaxEnt with the (EH-1) generator selects the Gibbs state on the (EH-2) algebra:

$$\rho_{\text{edge}} = \frac{1}{Z} e^{-tH_{\text{edge}}} = \frac{1}{Z}\sum_R e^{-t\lambda_R} P_R.$$

Step 5 (Sector probabilities). The probability of sector R is $p_R = \operatorname{Tr}(\rho_{\text{edge}} P_R)$, and on the (EH-2) algebra $\operatorname{Tr} P_R = d_R$, so

$$p_R = \frac{d_R\, e^{-t\lambda_R}}{Z}.$$

QED.

Scope. The fixed-cutoff same-overlap derivation is the theorem-bearing part of the edge-law package, conditional on (EH-1) and (EH-2). The quasi-local local-Gibbs generator used here is derived from Theorem 2.6: if MaxEnt constraints are expectations of finitely many quasi-local operators, the entropy maximizer is automatically a Gibbs state with a quasi-local generator. The additional representation-theoretic step summarized here is the compact-group / Peter–Weyl refinement lift. The Casimir-ratio predictions, coupling extraction, α_U closure, and the 2D-YM/worldsheet bridge downstream of this law are valid on models satisfying (EH-1)/(EH-2), not as consequences of the listed axioms alone; the numerical tests of Section 6.14 probe (EH-1) on explicit models. Model-family uniqueness and large-edge continuations have their own claim boundaries.

Normalization anchor: 2D Yang-Mills. The parameter t can be exactly matched to a conventional coupling in 2D Yang-Mills, where the physical Hamiltonian is literally the group Laplacian:

$$H = \frac{g^2}{2}\Delta_G, \qquad \Delta_G \chi_R = -C_2(R)\chi_R \quad \Rightarrow \quad E_R = \frac{g^2}{2} C_2(R).$$

Euclidean evolution for "time" A (the area of a cylinder in 2D YM) gives weight(R) $\propto$ exp(-A E_R) = exp(-g^2 A C$_2$(R)/2). Comparing with the heat-kernel expansion K_t(U) = Σ_R d_R χ_R(U) e^{-t C$_2$(R)} yields the exact identification:

$$t_{\text{phys}} = \frac{g^2 A}{2} \quad \text{(in 2D YM, no ambiguity).}$$

This shows that the Laplacian + MaxEnt $\to$ heat-kernel structure is exact in a solvable two-dimensional Yang-Mills case. The coefficient in front of C$_2$ is fixed. This normalization anchor is a two-dimensional edge-theory check, separate from The Standard Model gauge paper's conditional support-visible compact-gauge repair-gap theorem. In any regime where the edge theory reduces to an effective 2D YM with known "Euclidean thickness" A_eff:

$$g^2(\mu) = \frac{2}{A_{\text{eff}}(\mu)} \cdot \frac{\Delta_R(\mu)}{C_2(R)},$$

and the RHS must be R-independent. This R-independence is an internal precision consistency test; the formula itself is the normalization map that connects t to the conventional gauge coupling.

2.2.14 Numerical validation of the heat-kernel law

The heat-kernel/Laplacian weighting of edge sectors is validated in explicit 2D Z_n gauge models on closed geometries.

Model. A 2×2 periodic lattice gauge theory (8 links) with Z_n link Hilbert spaces and Hamiltonian:

$$H = -K\sum_p \mathrm{Re}(B_p) - h\sum_\ell \mathrm{Re}(X_\ell) - \Gamma\sum_v \mathrm{Re}(A_v),$$

where X_ℓ is the Z_n shift on link ℓ, B_p is the oriented plaquette operator (product of Z's around plaquette p), and A_v is the oriented star/Gauss operator (outgoing X, incoming X†). With K = 1 and Γ = 5, the ground state satisfies $\langle$A_v$\rangle$ = 1 at all vertices to numerical precision.

Region and edge operator. Region A consists of links whose tail has x = 0 ("half-lattice" cut). At each boundary vertex v, the electric-center edge charge is the restricted star Q_v^{(A)} = ∏_{ℓ ∈ star(v) ∩ A} X_ℓ^{±1}.

Results for $\mathbf{Z_2}$. With $\lambda_1 = 4\sin^2(\pi/2) = 4$:

h	p_0	p_1	t	g_ent
0.5	0.8266	0.1734	0.391	0.249
1.0	0.9612	0.0388	0.803	0.357
2.0	0.9917	0.0083	1.194	0.436

Z_2 has a single nontrivial Laplacian eigenvalue, so a one-parameter fit of t can always be made exactly for any admissible nontrivial weight. The Z_2 table is therefore an implementation and convention check, not a test of the heat-kernel form.

Results for $\mathbf{Z_3}$ (symmetry-forced consistency check). With $\lambda_1 = \lambda_2 = 4\sin^2(\pi/3) = 3$:

h	p_0	p_1	p_2	t(q=1)	t(q=2)	g_ent	m_plaq
0.2	0.4395	0.2803	0.2803	0.1500	0.1500	0.154	2.22
0.5	0.7509	0.1245	0.1245	0.5989	0.5989	0.309	1.75
1.0	0.9606	0.0197	0.0197	1.2956	1.2956	0.454	4.07
1.5	0.9851	0.0074	0.0074	1.6288	1.6288	0.509	7.06
2.0	0.9921	0.0039	0.0039	1.8440	1.8440	0.542	10.10

Both equalities in this table are forced by symmetry before any heat-kernel ansatz is imposed. The equality $p_1 = p_2$ is exact charge-conjugation symmetry in Z_3, and the two nontrivial Fourier modes share the degenerate eigenvalue $\lambda_1 = \lambda_2 = 3$, so extracting t from q = 1 and q = 2 must give the same value for *any* conjugation-invariant distribution. The observed agreement (t_{q=1} ≈ 1.2956389318579 versus t_{q=2} ≈ 1.2956389318521 at h = 1.0, i.e. ~10^{-14}) therefore validates the implementation and the symmetry conventions only; it cannot distinguish the heat-kernel form from an arbitrary conjugation-symmetric distribution on Z_3, and it is not counted as an independent confirmation of the heat-kernel law. The substantive overconstrained tests are the Z_5 and S_3 calculations below, which involve two distinct nonzero eigenvalues.

Region-choice stability. At h = 1, the extracted g_ent is nearly independent of region size:

- 2 links (one vertex's outgoing links): g_ent ≈ 0.453
- 4 links (half-lattice): g_ent ≈ 0.454
- 6 links (three vertices): g_ent ≈ 0.453

This locality confirms that the coupling is dominated by physics near the cut, not global bookkeeping, exactly what is expected if this behaves like a local QFT observable.

Results for $\mathbf{Z_5}$ (golden ratio test). The Z_5 case provides a stringent test because the Laplacian eigenvalues have a distinctive ratio involving the golden ratio $\phi = (1+\sqrt{5})/2$:

$$\lambda_q = 4\sin^2\left(\frac{\pi q}{5}\right), \qquad \frac{\lambda_2}{\lambda_1} = \frac{\sin^2(72°)}{\sin^2(36°)} = \phi^2 \approx 2.618.$$

This ratio distinguishes the Laplacian law from naive alternatives: a linear model (λ_q ∝ q) would predict ratio 2, while a quadratic model (λ_q ∝ q^2) would predict ratio 4.

Direct diagonalization of the displayed Hamiltonian on the full 2×2-torus link basis (5^8 states, K = 1, Γ = 5), measuring the restricted-star edge charge at a boundary vertex of the half-lattice region, gives:

h	Measured ratio $\ln(p_2/p_0)/\ln(p_1/p_0)$	Deviation from ϕ^2
0.50	2.876	9.9%
0.20	2.750	5.1%
0.10	2.671	2.0%
0.05	2.643	1.0%

In the weak-field limit h → 0 the measured ratio converges to the golden ratio squared, and the stated limiting direction agrees with the tabulated sequence: the deviation decreases monotonically as h decreases. (The table reports the direct link-basis run, reproducible with the held-out validation script in the released edge-sector code; a dual/flux-basis parameterization carries a tabulated parameter that runs opposite to the h of the displayed Hamiltonian and is not used. At large h the ratio instead drifts toward the perturbative order-counting value 2, which is why the weak-field direction is the declared convergence regime for this model.) The test is genuinely overconstrained: Z_5 has two distinct nonzero eigenvalues, t is fitted on the q = 1 sector alone, the q = 2 weight is a held-out prediction, and the deviation column reports the held-out residual. Convergence of that residual confirms that the vacuum entanglement spectrum encodes the precise geometric structure of the gauge group Laplacian.

Significance. This validates the mathematical law (sector probabilities weighted by Laplacian eigenvalues) in explicit 2D gauge-invariant models with non-flat sector distributions. Because the nontrivial Z_3 eigenvalues are degenerate, it is the Z_5 test (and the S_3 test below) that is structurally identical to SU(2)/SU(3): two or more distinct nonzero eigenvalues overconstrain the slope, and held-out agreement confirms the mechanism works before jumping to nonabelian groups.

Results for S_3 (first nonabelian test). The abelian tests above use charge-sector projectors that reduce to Fourier modes. For nonabelian groups, the edge-sector projector must be generalized to character projectors:

$$P_{v,R} = \frac{d_R}{|G|} \sum_{h \in G} \chi_R(h^{-1}) Q_v^{(A)}(h),$$

where d_R is the dimension of irrep R, χ_R is its character, and Q_v^{(A)}(h) is the restricted gauge action at boundary vertex v acting only on links in region A.

For S_3 (the smallest nonabelian group, order 6), there are three irreps: trivial (d=1), sign (d=1), and standard (d=2). The Cayley-graph Laplacian eigenvalues for the transposition generating set are:

$$\lambda_{\text{triv}} = 0, \qquad \lambda_{\text{sign}} = 6, \qquad \lambda_{\text{std}} = 3.$$

Exact reduction on one plaquette. For the single-plaquette model (4 links), imposing Gauss's law at all vertices means the physical wavefunction depends only on the plaquette holonomy's conjugacy class. Since S_3 has exactly 3 conjugacy classes, the gauge-invariant Hilbert space is 3-dimensional, spanned by the character states {|χ_R⟩}. In this basis, the edge-sector probabilities are exactly p_R = |c_R|^2 where $|\psi_0\rangle$ = Σ_R c_R |χ_R⟩. This is an exact identity for the one-plaquette gauge-invariant sector.

The heat-kernel ansatz predicts p_R $\propto$ d_R exp(-t λ_R). Extracting t independently from the sign and standard irreps provides an overconstrained test: the ratio λ_sign/λ_std = 6/3 = 2 is a parameter-free prediction. Results from a single-plaquette S_3 lattice gauge model (K=1, Γ=5):

h	p_triv	p_sign	p_std	t (sign)	t (std)	Δt/t	log-ratio
0.5	0.909	0.0013	0.089	1.09	1.01	8.4%	2.17
1.0	0.980	7.5×10^{-5}	0.020	1.58	1.54	2.8%	2.06
2.0	0.996	4.3×10^{-6}	0.004	2.06	2.04	1.0%	2.02
5.0	0.9993	1.0×10^{-7}	0.00066	2.68	2.67	0.3%	2.006
12	0.9999	3.0×10^{-9}	0.00011	3.27	3.27	0.1%	2.002
100	1.0000	6.1×10^{-13}	2.0×10^{-6}	4.69	4.69	0.009%	2.0002

The "Δt/t" column shows the fractional difference (t_sign - t_std) / $\bar{t}$. The "log-ratio" column shows log(p_sign/p_0) / log(p_std/(2 p_0)), which should equal λ_sign/λ_std = 2 if the heat-kernel form holds exactly. Equivalently, t is fitted from the standard sector alone and the sign-sector weight is a held-out prediction; log-ratio minus 2 and Δt/t are the held-out residuals. Unlike Z_3, the two nonzero eigenvalues are distinct (6 versus 3), so no symmetry forces this agreement.

As h increases, both diagnostics converge: Δt/t drops below 10^{-4} and the log-ratio approaches 2.000. This is exactly the expected behavior: finite-size corrections are largest at strong coupling; the heat-kernel form becomes exact as the perturbative regime is approached.

This gives a nonabelian validation of the edge-sector extraction mechanism. The structure (character projectors, Laplacian eigenvalues from the group's Cayley graph, overconstrained t extraction) is identical to the SU(2) and SU(3) continuation setup.

Parameter-free predictions for SU(2) and SU(3). The heat-kernel law yields exact, parameter-free ratio predictions that require no scheme matching. Define the "Casimir log-gap":

$$\Delta_R \equiv \ln\left(\frac{p_0}{d_0}\right) - \ln\left(\frac{p_R}{d_R}\right) = t\, C_2(R).$$

Ratios of Δ_R cancel all unknowns (t, partition function):

$$\frac{\Delta_{R_1}}{\Delta_{R_2}} = \frac{C_2(R_1)}{C_2(R_2)} \quad \text{(exact, parameter-free).}$$

SU(2) predictions. Irreps labeled by spin j have d_j = 2j+1 and C_2(j) = j(j+1). The framework predicts:

- $\Delta_1/\Delta_{1}/_2$ = 2/(3/4) = **8/3 ≈ 2.667**
- $\Delta_3/_2/\Delta_1/_2$ = (15/4)/(3/4) = **5**
- $\Delta_3/_2/\Delta_1$ = (15/4)/2 = **15/8 = 1.875**

SU(3) predictions. Irreps labeled by Dynkin indices (p,q) have C_2(p,q) = (p^2 + q^2 + pq + 3p + 3q)/3. Using the fundamental **3** = (1,0) with C_2 = 4/3 as the reference:

- Δ_8/Δ_3 = 3/(4/3) = **9/4 = 2.25**
- Δ_6/Δ_3 = (10/3)/(4/3) = **5/2 = 2.5**
- Δ_{10}/Δ_3 = 6/(4/3) = **9/2 = 4.5**
- Δ_{15}/Δ_3 = (16/3)/(4/3) = **4**
- Δ_{27}/Δ_3 = 8/(4/3) = **6**

These are the SU(2)/SU(3) analogs of the Z_5 golden-ratio test: exact rational numbers fixed entirely by group theory, with no adjustable parameters.

Preliminary SU(3) results. A one-plaquette SU(3) "quantum link" model (finite truncated irrep basis, n_max = 12, κ = 2) extracts t from 14 different irreps simultaneously. The results show internal consistency at the 1-3% level:

bare g^2	extracted t (mean±std)	g_ent	gap
0.3	0.314±0.0005	0.224	1.92
0.5	0.539±0.0025	0.293	1.83
0.8	0.896±0.012	0.378	1.72
1.0	1.144±0.025	0.427	1.64

The standard deviation across irreps provides a built-in error estimate. This is a QCD proton-physics surrogate rather than a full proton calculation; it lacks dynamical quarks and operates on a single plaquette. It nevertheless demonstrates that the nonabelian extraction machinery produces self-consistent outputs without tuning.

Extracting the normalization factor A_eff. The 2D YM anchor (Section 6.13) gives t = g^2 A / 2, so the "effective Euclidean thickness" is

$$A_{\text{eff}} = \frac{2t}{g^2}.$$

Computing this from the SU(3) table:

bare g^2	extracted t	A_eff
0.3	0.314	2.093
0.5	0.539	2.156
0.8	0.896	2.240
1.0	1.144	2.288

Mean: A_eff ≈ 2.19 with point-to-point scatter ~4%.

Extrapolation to weak coupling. The systematic drift in A_eff shows fitting A_eff(g^2) = A_0 + a · g^2. A weighted linear fit gives:

$$A_0 = 2.004 \pm 0.012$$

with χ^2/dof ≈ 0.09, indicating excellent consistency. This strongly shows that, in this toy UV completion, the "missing normalization" converges to A_eff → 2 as $g^2 \to 0$.

The normalization factor behaves like a quasi-constant and extrapolates to a simple value (≈ 2) in the weak-coupling limit. This provides a concrete path to absolute coupling predictions: once A_eff is determined from microphysics, the conversion g^2 = 2t/A_eff fixes the gauge coupling without additional free parameters.

Internal validation summary. The finite-group calculations separate into symmetry-forced implementation checks and genuinely overconstrained tests:

- **Z_2, Z_3**: implementation and convention checks. Z_2 has one nontrivial eigenvalue (one parameter, one datum); in Z_3 the equalities $p_1 = p_2$ and t_{q=1} = t_{q=2} are forced by charge conjugation plus the degenerate eigenvalue $\lambda_1 = \lambda_2$, so the ~10^{-14} agreement is expected from symmetry alone

- **Z_5**: overconstrained golden-ratio test ($\lambda_2/\lambda_1 = \phi^2$, two distinct nonzero eigenvalues), held-out residual 1.0% at h = 0.05, decreasing monotonically as h → 0
- **S_3**: overconstrained nonabelian Casimir log-ratio test (λ_sign/λ_std = 2), held-out residual 0.01%
- **SU(3)**: 14-irrep simultaneous extraction, internal scatter 1-3%

The Z_3 equality is forced by symmetry and is retained only as a regression check of the implementation; it is not an independent confirmation of the heat-kernel law. The substantive validation rests on the Z_5, S_3, and SU(3) tests, each of which fits t using fewer spectral sectors than it predicts, involves two or more distinct nonzero eigenvalues, and reports held-out residuals that converge in the declared regime of the corresponding model (h → 0 for the Z_5 torus model; h large for the single-plaquette S_3 model). This provides internal validation of the mechanism "MaxEnt + Laplacian => heat-kernel sector weights" before applying it to physical gauge groups. A self-contained held-out refit script for the finite-group models is provided with the public code release; it fits t on a declared subset of sectors and prints held-out residuals for the remaining sectors.

2.2.15 Frozen IBM Quantum Cloud engineering archive

The former IBM Quantum Cloud program is frozen as an engineering and reproducibility archive. It checked whether OPH-motivated reduced-sector structures survived preparation and readout on real superconducting-qubit devices, and whether a programmed record-gated controller beat restricted controller nulls. Standard quantum mechanics predicts every prepared state and dynamic circuit used in that program. The runs therefore have no power to distinguish OPH from quantum mechanics, regardless of shot count or blinding quality. No new quantum-cloud campaign is promoted as OPH evidence until a source-closed observable with a different numerical QM prediction and an identifiable effect size is derived.

Detailed circuits, representative derived summaries, the complete blinded-run receipts, and archived rerun instructions are public in the frozen project write-up and the accompanying code and data bundle.

- **Stage 1 (Markov/recoverability benchmark).** On `ibm_marrakesh` and `ibm_fez`, the structured state reconstructs below both controls in conditional mutual information and above both controls in Petz fidelity. On `ibm_marrakesh`, the observed ordering is $0.2309 < 0.3890 < 0.9474$ for CMI and $0.9297 > 0.8649 > 0.6066$ for Petz fidelity; on `ibm_fez`, it is $0.1498 < 0.4992 < 0.9166$ and $0.9479 > 0.8117 > 0.6449$. This reproduces the programmed ordering on two real backends; the Petz calculation is offline and the result is not theory-discriminating.
- **Z_3 hardware sanity check.** The two-sector Z_3 extraction (whose two-sector agreement is expected a priori from the degenerate eigenvalues and charge-conjugation symmetry, so this is an implementation check, not an overconstrained test) passes cleanly on hardware: across prepared $t = 0.30, 0.60, 0.90$, the mean extracted t is within about 0.02 of the target, the two independent extractions agree to within 0.0142, 0.0034, and 0.0043, and leakage stays below 0.1%. This shows that the reduced-sector preparation and readout path is internally coherent on-device before sharper ratio claims are interpreted.
- **Z_5 exact-ratio preparation.** The programmed target is $\Delta_2/\Delta_1 = \varphi^2 \approx 2.618$. The measured values span approximately 2.407 to 2.878, and the focused high-shot $t = 0.90$ interval lies below the exact target. Because the target amplitudes are directly prepared, QM predicts the same ideal ratio.
- **S_3 nonabelian layout diagnostic.** The first hardware run revealed a real layout-dependent bias. Reversing the qubit layout moved the mitigated ratio from 1.8724 to 2.0299, and a repeat

returned 2.0657. Because the better mapping was selected after observing the bias, this is a post-hoc device diagnostic rather than a blinded confirmation.

These frozen runs document preparation, readout, layout, and feedback engineering on real devices. They add no evidence for OPH over standard quantum mechanics because the two descriptions make no different prediction for the implemented interventions. Their proper role is reproducibility and controller validation, not theory confirmation.

2.3 Gauge-chain status and falsification boundary

The finite carrier calculation derives the incidence polynomial for J, the target-blind response $R = -J$, its relative sector signs, and the equivariant current algebra $\mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{u}(1)$. Given the declared fermionic Spin category and rank-15 matter projector pair, the matter certificate derives anomaly balance, exact hypercharge up to conjugation, three colors, the common $\mathbb{Z}_6$ tensor kernel, and the maximal faithful matter image. These implications do not use Minimal Admissible Realization.

Four central global forms act consistently on the same local tensors: the cover and its $\mathbb{Z}_2$, $\mathbb{Z}_3$, and $\mathbb{Z}_6$ quotients. The kernel theorem identifies the maximal faithful local matter image; it does not select the physical spectrum of line operators or bundles. The physical global form is open. Physical fermion typing, laboratory-current attachment, equality between the finite current and the categorically reconstructed group, scalar attachment and multiplicity, physical family attachment, and QFT realization are also open.

The gauge result fails if the incidence identity, target-blind response, equivariant lift, anomaly scan, tensor-character enumeration, or refinement persistence receipt fails its registered negative controls. The categorical route fails outside its combined strictification, trivial-holonomy, localization, and refinement/fiber receipts.

A Technical Details for the Specialist Derivations

This appendix carries the structural gauge details that support the Standard Model derivation. It sharpens the product-group, baryogenesis-boundary, effective-worldsheet, and compact-gauge Yang–Mills surfaces without widening the recovered core beyond its stated branch conditions.

A.1 Product-Group Structural Corollaries

Theorem A.1 (Factorization equivalence)**.**

Factorizing edge weights $\Longleftrightarrow$ *Additive boundary Laplacian* $\Longleftrightarrow$ *Product gauge group.*

If

$$H_\partial = H_\partial^{(1)} + H_\partial^{(2)} + H_\partial^{(3)} \qquad \text{with} \qquad [H_\partial^{(i)}, H_\partial^{(j)}] = 0,$$

then

$$p(R_1, R_2, R_3) \propto \prod_{i=1}^{3} d_{R_i} e^{-t_i C_2(R_i)}.$$

Applied to the assumed transportable refinement-directed edge-sector colimit on the realized branch, together with the rigid symmetric C^*-tensor and faithful bosonic fiber-functor conditions of the compact-gauge theorem, Tannaka–Krein reconstruction first yields some compact group G. Under the MAR admissibility package, and once the admissible class includes one connected abelian charge factor, the selected connected Lie realization is $SU(3) \times SU(2) \times U(1)$ up to finite quotient.

Corollary A.2 (No simple-GUT X/Y gauge channel)**.** *With product gauge group, the adjoint representation of the full connected gauge group is*

$$(8, 1, 0) \oplus (1, 3, 0) \oplus (1, 1, 0),$$

equivalently the derived nonabelian adjoint is

$$(8, 1, 0) \oplus (1, 3, 0).$$

There are therefore no gauge generators in mixed representations $(3, 2, \pm 5/6)$, i.e. no simple-GUT X, Y bosons on the realized branch. This excludes the ordinary X/Y exchange channel, not every gauge-mediated ultraviolet mechanism. General proton stability and a proton lifetime require separate baryon-violating operator and ultraviolet data.

A.2 Finite-Quotient Baryogenesis Source Theorem and Gauge-Deck No-Go

At regulator r, let

$$\mathfrak{B}_r = (Q_r, C_r, E_r, \omega_{R,r}, L_{r,T}, p_{r,i}, \tau_r, T_r, \rho_{R,r})$$

be a finite baryogenesis source packet. Here Q_r is the physical quotient, C_r is its CP involution, $\omega_{R,r} : E_r \to \mathbb{Z}$ is a CP-odd integral winding cocycle, $L_{r,T}$ is a quotient-intrinsic continuous-time Markov generator, $p_{r,i}$ is the source initial law, τ_r and $T_r(\tau)$ supply the physical clock and temperature map, and $\rho_{R,r}$ assigns record charges r_ψ to left-handed Weyl multiplets. Let $\theta_0 = 2\pi/m_R$, where m_R is the primitive winding period.

Theorem A.3 (Finite-quotient anomaly-and-current theorem)**.** *If the record phase acts as $\psi \mapsto e^{ir_\psi \Theta_R}\psi$, then its electroweak topological coefficient and expected phase velocity are*

$$k_R = \sum_{\psi\,\mathrm{LH}} r_\psi \, 2T_2(R_\psi), \qquad \dot{\Theta}_R(T) = \theta_0 \sum_{q,q' \in Q_r} p_r(q, T) L_{r,T}(q, q') \omega_{R,r}(q, q').$$

The probability law satisfies $dp_r/d\tau = p_r L_{r,T(\tau)}$. If $p_{r,i}(C_r q) = p_{r,i}(q)$ and $L_{r,T}(C_r q, C_r q') = L_{r,T}(q, q')$, then

$$\dot{\Theta}_R(T) = Y_B = 0.$$

For the determinant/deck direction of the realized Standard Model branch, the natural central gauge attachment is hypercharge and

$$k_R^{YWW} = N_g(3Y_Q + Y_L) = 3\left(3 \cdot \frac{1}{6} - \frac{1}{2}\right) = 0.$$

The OPH $\mathbb{Z}_6$ gauge/deck phase therefore supplies no direct $\Theta_R W\widetilde{W}$ baryogenesis coupling. A conventional global $B+L$ attachment would instead give $k_R = 2N_g = 6$, conditional on a quotient-visible gauge-singlet record phase carrying that attachment.

Proof. The chiral measure gives the mixed $SU(2)_L^2 U(1)_R$ anomaly index, with an $SU(2)$ doublet contributing $r_\psi 2T_2(\mathbf{2}) = r_\psi$, including spectator multiplicities. The expected number of transitions $q \to q'$ in proper time $d\tau$ is $p_r(q,T)L_{r,T}(q,q')d\tau$; weighting by the integral winding increment gives the displayed current. Pairing every transition with its CP image cancels the current under a CP-symmetric law. The gauge/deck value is the Standard Model mixed hypercharge anomaly, which vanishes generation by generation. For $B+L$, the quark and lepton weak doublets contribute two units per generation. □

Source boundary. The determinant/deck cocycle supplies a candidate phase coordinate and its CP inversion. Normal-form settlement and the oriented 24-slot register supply no transition probabilities and select no baryon sign. A nonzero branch requires a distinct anomalous record attachment, a source-only $L_{r,T}$, a physical clock and initial or boundary law, a quotient-intrinsic CP-odd affinity, and cosmological domain coherence. On the conditional $B+L$ branch, the transport functional requires

$$\left\langle \frac{\dot{\Theta}_R}{T} \right\rangle_{\text{fo}} = (4.463 \pm 0.028) \times 10^{-9}.$$

This number is an evaluation target for the source generator. It is not an input from which the generator may be selected.

Proton-spin continuation benchmark. On the realized D8–D9 branch the color factor is SU(3), so the structural OPH input for proton-spin bookkeeping is

$$C_F = \frac{4}{3}, \qquad C_A = 3.$$

If one adds the QCD-dependent continuation ansatz that quark/gluon spin sharing equilibrates according to these Casimirs, then

$$\Delta\Sigma \approx \frac{C_F}{C_F + C_A} = \frac{4}{13} \approx 0.308.$$

Against a representative lattice benchmark $\Delta\Sigma \approx 0.286$, this lands within about 8%, but it is only a deferred continuation benchmark. A first-principles OPH derivation of proton spin would require the nonperturbative light-quark/hadron completion plus an explicit map from OPH data to renormalized proton matrix elements.

A.3 Controlled Worldsheet Effective Description

Exact OPH-to-2D-Yang–Mills bridge. On the compact-group heat-kernel branch, the edge-partition theorem proves the exact identity

$$Z_{\text{edge}}(t) = K_t(1)$$

by gluing the open-edge weights $p_R(t) \propto d_R e^{-tC_2(R)}$ into the closed partition sum

$$Z_{\text{edge}}(t) = \sum_R d_R^2 e^{-tC_2(R)}.$$

This is the precise theorem-level sense in which the OPH edge-sector partition reorganizes into the two-dimensional Yang–Mills heat-kernel surface, and the Chapman–Kolmogorov law supplies the matching collar-sewing rule. This bridge is a two-dimensional heat-kernel partition identity on the stated compact-group branch. The four-dimensional compact-gauge repair-gap theorem is the separate support-visible result in Theorem A.29.

Large-N_{edge} boundary. The large-N_{edge} worldsheet interpretation is carried only when one fixes a distinct large-N_{edge} sequence, with $N_{\mathrm{edge}} \neq N_c = 3$, a fixed-$\tau$ window for

$$\tau = tN_{\mathrm{edge}},$$

and the uniform genus-remainder control stated for that expansion. On that branch the edge free energy satisfies the Standard Model gauge paper's theorem-level criterion for a controlled genus expansion, so the standard Gross–Taylor rewriting becomes a controlled worldsheet effective description of edge dynamics.

External items. The exact bridge theorem uses the compact-group heat-kernel branch and the quadratic-Casimir normalization declared in the Standard Model gauge paper. The continuation theorem uses the declared large-N_{edge} regime and the imported Gross–Taylor large-N worldsheet dictionary for two-dimensional Yang–Mills. It does not identify a critical worldsheet CFT. Worldsheet supersymmetry, critical dimension, modular invariance, anomaly cancellation, GSO projection, and full massless-spectrum matching use separate continuation inputs beyond the compact recovered-core chain.

A.4 Support-Visible Yang–Mills Gap from Repair Dynamics

Assumption A.4 (Support-visible compact-gauge Yang–Mills branch). *The compact-gauge branch used in this subsection is the ordinary or central zero-obstruction compact-gauge sector with compact simple structure group G, a four-dimensional Euclidean scaling chart, and a cofinal tail carrying the compact-gauge refinement receipt of the compact-gauge refinement receipt stated in the main derivation. It also has a reflection-positive ordinary vacuum, topological angle $\theta = 0$, gauge-invariant local finite-constraint MaxEnt/Gibbs refinement, no additional relevant dimension-four pure-gauge operator on the branch besides the positive quadratic curvature invariant, active exact-Markov repair collars, bounded-color collar covers, and repair completeness. It is carried on the separated cofinal regulator system stated in the main derivation. No continuum cylinder state, GNS space, transfer generator, or vacuum projection is included in this assumption; those objects are constructed, or separately certified, below.*

Definition A.5 (Finite support-visible compact-gauge cylinder system). *At regulator r, make the following construction twice: for the finite four-dimensional Euclidean slab $\Lambda_r^{(4)} = (V_r^{(4)}, E_r^{(4)})$ cut out by the scaling chart, and for its time-zero spatial slice $\Lambda_r^{(0)} = (V_r^{(0)}, E_r^{(0)})$. In the formulas below, $\Lambda_r = (V_r, E_r)$ denotes either choice and the superscript is restored whenever the two must be distinguished. Let D_r be the finite collection of repaired collar records and zero-obstruction sector labels retained by the compact-gauge patch carrier. In the gauge-register presentation, let*

$$\widetilde{X}_r \subseteq G^{E_r} \times D_r$$

be the closed set satisfying the finite Gauss, overlap, and repaired-record constraints, let $\mathcal{G}_r = G^{V_r}$ act by endpoint gauge transformations, and let $\sim_{\mathrm{ov}}$ be the closed equivalence relation of support-visible overlap/presentation indistinguishability. The finite-stage support-visible configuration space is

$$X_r := (\widetilde{X}_r/\mathcal{G}_r)/{\sim_{\mathrm{ov}}}.$$

Thus X_r is compact Hausdorff. "Finite stage" means that Λ_r and D_r are finite; X_r need not be a finite set when G is a compact Lie group. A finite quantum-link truncation gives the same commuting Euclidean readout algebra after restriction to its joint cylinder spectrum; no noncommutative support-quotient claim is made here.

For each active collar C, let

$$\rho_{C,r} : X_r \longrightarrow Y_{C,r}$$

be the continuous complete repaired readback. On the classical compact-holonomy branch used by the repair lemma, the complementary conditional fiber contains every record or holonomy coordinate actually resampled by collar repair. It is either finite with its uniform conditional law or a standard atomless probability space with its conditioned MaxEnt/Haar law. The complete readback retains exactly the fixed repaired datum, not every pre-repair holonomy; otherwise nonconstant Wilson cylinders would incorrectly lie in the repair kernel.

For a bounded cell region $O \subset \Lambda_r$, let $\mathfrak{C}_r^{G,\mathrm{sv}}(O)$ be the unital *-algebra generated by

(a) *contracted Peter–Weyl spin-network functions, Wilson-loop characters, and any explicitly retained gauge-invariant boundary-carrier contractions supported in O;*

(b) *the overlap-sector projectors whose collars lie in O; and*

(c) *continuous functions of the complete repaired readbacks $\rho_{C,r}$ for those collars.*

Its uniform closure

$$\mathcal{A}_r^{G,\mathrm{sv}}(O) := \overline{\mathfrak{C}_r^{G,\mathrm{sv}}(O)}^{\|\cdot\|_\infty} \subseteq C(X_r)$$

is the finite-stage local cylinder algebra. Equivalently, start from the freely presented readout-generator *-algebra and quotient by the kernel of its evaluation homomorphism on X_r, then complete. That kernel is a closed two-sided overlap-trivial ideal fixed before a state is chosen. The relation $\sim_{\mathrm{ov}}$ identifies exactly the configurations on which all declared spin-network, sector, and repaired-readback generators agree. These generators are self-adjoint, contain the constants, and therefore separate points of X_r. Stone–Weierstrass gives

$$\mathcal{A}_r^{G,\mathrm{sv}}(\Lambda_r) = C(X_r).$$

Consequently, for any selected stage measure π_r, their GNS cylinder closure is the full $L^2(X_r, \pi_r)$ after the finite π_r-null support reduction.

Write the resulting spaces and algebras as

$$(X_r^{(4)}, \mathcal{A}_r^{G,\mathrm{sv},(4)}(O)) \quad \textit{and} \quad (X_r^{(0)}, \mathcal{A}_{r,0}^{G,\mathrm{sv}}(B)).$$

Restriction of a Euclidean history to the time-zero slice gives a continuous map

$$\tau_{0,r} : X_r^{(4)} \longrightarrow X_r^{(0)}$$

and the corresponding pullback embeds time-zero cylinders into four-dimensional cylinders.

For $r \preceq s$, the deterministic coarse shadow multiplies the ordered fine-edge holonomies lying over each coarse edge, restricts collar readbacks, and sends a fine sector label to its coarse shadow. Gauge covariance makes this a map

$$p_{sr} : X_s \longrightarrow X_r, \qquad p_{tr} = p_{sr} \circ p_{ts} \quad (r \preceq s \preceq t).$$

The four-dimensional and time-zero coarse shadows satisfy the finite commuting square

$$\tau_{0,r} p_{sr}^{(4)} = p_{sr}^{(0)} \tau_{0,s}.$$

On local cylinders it gives the unital *-map

$$\iota_{rs}^O : \mathcal{A}_r^{G,\mathrm{sv}}(O) \longrightarrow \mathcal{A}_s^{G,\mathrm{sv}}(O_s), \qquad \iota_{rs}^O(a) = a \circ p_{sr}.$$

On the compact-gauge refinement branch, the sector-projector part must agree with the separately supplied block-multiplicity injection $\jmath^B_{rs}$ of the compact-gauge refinement receipt; it is not inferred from the deterministic coarse-shadow map. The holonomy and readback parts are its declared compatible extension. The maps preserve isotony and obey

$$\iota^O_{rt} = \iota^{O_s}_{st}\iota^O_{rs}.$$

If every coarse configuration on the selected branch has a fine extension, p_{sr} is surjective and ι^O_{rs} is injective. Without that finite extendability receipt, the C^*-inductive limit quotients the common refinement-null kernel; the extraction proof below does not assume prelimit injectivity of the full cylinder map. That broader quotient extraction is not, by itself, the receipt-certified sector ladder of the refinement/fiber descent theorem in the main derivation. Write $\mathcal{A}^{G,\mathrm{sv}}_r := \mathcal{A}^{G,\mathrm{sv}}_r(\Lambda_r)$. Choose a countable cofinal regulator sequence and a countable bounded-region exhaustion of the four-dimensional chart. This is the cylinder family used below. For a non-countable regulator presentation, the same construction uses the full directed family and a subnet instead of this cofinal sequence.

Proposition A.6 (Projective compact-gauge cylinder extraction)**.** *For every regulator s, let $\mu^{(4)}_s$ be the declared finite quotient-Gibbs probability measure on $X^{(4)}_s$, let*

$$\pi_s := (\tau_{0,s})_\# \mu^{(4)}_s$$

be its stationary time-zero marginal on $X^{(0)}_s$, and put $\mu^{(0)}_s := \pi_s$. For $d \in \{4, 0\}$, write

$$\omega^{(d)}_s(a) := \int_{X^{(d)}_s} a \, d\mu^{(d)}_s, \qquad \omega^{(d)}_{s\downarrow r} := \omega^{(d)}_s \circ \iota^{(d)}_{rs}.$$

The statements below hold for both $d = 4$ and $d = 0$; the superscript is suppressed in the displayed proof. Then the following statements hold.

(i) *The marginals of any one fine state are exactly projective: for $q \preceq r \preceq s$,*

$$\omega_{s\downarrow r} \circ \iota_{qr} = \omega_{s\downarrow q}.$$

(ii) *There is a cofinal subsequence, or a cofinal subnet in the general directed case, and states $\overline{\omega}_r \in S(\mathcal{A}^{G,\mathrm{sv}}_r)$ such that*

$$\omega_{s_\alpha\downarrow r} \xrightarrow[\alpha]{} \overline{\omega}_r \quad \textit{weak-*}$$

on every fixed local cylinder algebra. The limit family is projective:

$$\overline{\omega}_r = \overline{\omega}_t \circ \iota_{rt} \qquad (r \preceq t).$$

(iii) *The compatible local states define a state ω_∞ on the algebraic inductive limit and extend uniquely to its local C^*-completion*

$$\mathcal{A}^{G,\mathrm{cyl}}_\infty(O) := \overline{\varinjlim_r \mathcal{A}^{G,\mathrm{sv}}_r(O_r)}.$$

The resulting local algebras are isotone and their union is dense in the global cylinder algebra.

(iv) *Let $\mathcal{N}_\omega(O)$ be the null ideal of the limiting cylinder measure,*

$$\mathcal{N}_\omega(O) := \{a \in \mathcal{A}_\infty^{G,\mathrm{cyl}}(O) : \omega_\infty(a^*a) = 0\}.$$

Because the Euclidean cylinder algebra is commutative, this is a closed two-sided ideal. On

$$\mathcal{A}_\infty^{G,\mathrm{supp}}(O) := \mathcal{A}_\infty^{G,\mathrm{cyl}}(O)/\mathcal{N}_\omega(O)$$

the induced state is faithful. The local GNS spaces glue isometrically, their local cylinder images have dense union, and their Hilbert direct limit is the GNS space $(K^{(d)}, \Pi^{(d)}, \mathbf{1}^{(d)})$ of $\omega_\infty^{(d)}$. Denote the four-dimensional Euclidean cylinder GNS space by $K^E := K^{(4)}$ and the time-zero repair GNS space by $K := K^{(0)}$. Each is a faithful cyclic GNS pair on its support quotient; the time-zero vector is identified as the physical vacuum only by the transfer/OS receipt below.

If, in addition, the finite four-dimensional quotient presentation satisfies the fiber-sum identity stated after the Local MaxEnt and Refinement Stability axiom, then $(p_{sr}^{(4)})_\# \mu_s^{(4)} = \mu_r^{(4)}$. The time-zero commuting square gives $(p_{sr}^{(0)})_\# \pi_s = \pi_r$, so both original state families are projective. Without that finite receipt, the bars on $\overline{\omega}_r$ are essential: compactness extracts a projective cluster family but does not turn the originally chosen coarse Gibbs measures into an exactly projective family. In the continuous compact-holonomy presentation, the corresponding receipt is the pushforward identity $(p_{sr}^{(4)})_\# \mu_s^{(4)} = \mu_r^{(4)}$, expressed through disintegration rather than a finite sum.

Proof. Functoriality of the coarse shadows gives

$$(\omega_s \circ \iota_{rs}) \circ \iota_{qr} = \omega_s \circ \iota_{qs},$$

which proves (i) without a continuum assumption.

For each fixed cylinder algebra, its state space is weak-* compact by Banach–Alaoglu. The algebras in the chosen cylinder exhaustion are separable: compact metrizable G has a countable Peter–Weyl test algebra, and only finitely many readbacks and sector projectors occur at one stage. Their state spaces are therefore weak-* metrizable. Starting with the first cylinder, extract a subsequence on which its marginals converge; from that subsequence extract one for the second cylinder, and continue. Choose the k-th diagonal index beyond the k-th regulator stage. The resulting diagonal subsequence is cofinal and converges on every cylinder. In the general directed presentation, compactness of the product of the local state spaces gives the same conclusion with a cofinal subnet: unresolved early coordinates may be filled by arbitrary states because every fixed coordinate is genuinely resolved on a cofinal tail. For $q \preceq r$, weak-* continuity of precomposition by ι_{qr}, together with (i), gives

$$\overline{\omega}_r \circ \iota_{qr} = \lim_\alpha \omega_{s_\alpha \downarrow r} \circ \iota_{qr} = \lim_\alpha \omega_{s_\alpha \downarrow q} = \overline{\omega}_q.$$

This proves (ii).

If $a \in \mathcal{A}_r^{G,\mathrm{sv}}$ represents an element $[a]_r$ of the algebraic inductive limit, set

$$\omega_\infty([a]_r) := \overline{\omega}_r(a).$$

Projectivity makes this independent of the representative. Positivity and normalization hold at the finite stage containing any given algebraic element. For every $t \succeq r$,

$$|\overline{\omega}_r(a)| = |\overline{\omega}_t(\iota_{rt}(a))| \le \|\iota_{rt}(a)\|.$$

Taking the infimum along the directed tail gives the C^*-inductive-limit norm bound, including when the maps are noninjective. Hence the state extends uniquely to the completion. The same construction region by region proves isotony and density, giving (iii).

By the Riesz representation theorem, a commutative local limiting state is integration against a Radon probability measure μ_O. Its null ideal consists precisely of the continuous functions vanishing on the measure support, so the quotient is canonically $C(\operatorname{supp}\mu_O)$ and the induced state is faithful. If $O \subseteq O'$ and $\iota_{OO'}$ is the local inclusion, state compatibility gives

$$a \in \mathcal{N}_\omega(O) \quad \Longleftrightarrow \quad \iota_{OO'}(a) \in \mathcal{N}_\omega(O').$$

Thus the local inclusions descend to injective maps of the support quotients, which therefore form an isotone faithful local net. For $r \preceq t$, projectivity makes

$$V_{rt} : K_{\overline{\omega}_r} \longrightarrow K_{\overline{\omega}_t}, \qquad V_{rt}[a]_r = [\iota_{rt}(a)]_t$$

well defined and isometric, because

$$\|V_{rt}[a]_r\|^2 = \overline{\omega}_t(\iota_{rt}(a^*a)) = \overline{\omega}_r(a^*a).$$

The maps compose, preserve $\mathbf{1}$, and intertwine left multiplication. Their Hilbert direct limit is therefore the GNS completion of the algebraic cylinder union and has dense local images. This proves (iv). Finally, suppressing the $d = 4$ superscript, the finite fiber-sum identity gives

$$Z_s = \sum_{x \in X_r} \sum_{x' : p_{sr}(x') = x} m_s(x') e^{-S_s(x')} = \alpha_{sr} Z_r,$$

and hence

$$(p_{sr})_\# \mu_s(x) = \frac{\alpha_{sr} m_r(x) e^{-S_r(x)}}{\alpha_{sr} Z_r} = \mu_r(x).$$

This is equivalent to $\omega_s^{(4)} \circ \iota_{rs}^{(4)} = \omega_r^{(4)}$; the commuting time-zero square gives the corresponding identity for $d = 0$. □

Remark A.7 (Two different quotients)**.** *The overlap-trivial presentation ideal in Definition A.5 is not the GNS null ideal $\mathcal{N}_\omega(O)$. The former removes data that no support-visible readout can distinguish before a state is selected; the latter removes directions that have zero norm in the extracted state. Faithful finite-stage states can have a nonfaithful weak-* limit, so quotienting only the overlap-trivial ideal would not prove faithfulness of the limiting pair.*

Assumption A.8 (Renormalized four-dimensional Yang–Mills identification receipt)**.** *On the selected four-dimensional cluster family, compatible renormalized holonomy cylinders obey uniform local Cauchy and regularity bounds. Their infinitesimal rectangle tests converge, in the declared distributional sense, to a connection/curvature pair*

$$U_{\mu\nu}(\varepsilon, x) = \mathbf{1} + \varepsilon^2 F_{\mu\nu}(x) + O(\varepsilon^3), \qquad F = dA + A \wedge A.$$

The renormalized local actions converge to the unique reflection-even positive dimension-four pure-gauge density on this branch, while the declared irrelevant remainders vanish. This is the regularity/universality receipt that identifies the projective generalized-holonomy state with the four-dimensional Yang–Mills cylinder state. The notation $e^{-S_E[A]} DA/G$ below abbreviates that cylinder family; it is not a literal infinite-dimensional Lebesgue measure.

Theorem A.9 (Conditional OPH four-dimensional Euclidean Yang–Mills form)**.** *Under Assumptions A.4 and A.8, on an extracted cylinder family from Proposition A.6 that lies on the declared local four-dimensional scaling branch, the continuum gauge-sector Euclidean action is*

$$S_E[A] = \frac{1}{4g^2} \int_{\mathbb{R}^4} \langle F_{\mu\nu}, F_{\mu\nu} \rangle \, d^4x, \qquad F = dA + A \wedge A,$$

with compact simple structure group G. The extracted gauge-quotient cylinder family is denoted

$$d\mu_{\mathrm{YM}}(A) = Z^{-1} e^{-S_E[A]} \, DA/G$$

in the OPH support-visible GNS representation. When Assumption A.22 also holds, its support-visible continuum transfer semigroup is the corresponding Euclidean Yang–Mills semigroup.

Proof. The proof spine is included here to make the branch target explicit. The present Yang–Mills branch takes compact simple G as branch data. When G is imported from OPH reconstruction, that import is conditional on the compact-gauge refinement receipt; on that tail the zero-obstruction transportable bosonic sector category and its faithful forgetful fiber functor reconstruct G. At fixed cutoff, the finite gauge-register / quantum-link presentation gives support-visible link holonomies and plaquette holonomies. In the refinement limit, the zero-obstruction gluing law makes infinitesimal rectangle holonomies multiplicative and path-local. The regularity and Cauchy bounds in Assumption A.8 upgrade this generalized-holonomy limit to a local connection A on the four-dimensional scaling chart, with infinitesimal plaquette defect

$$U_{\mu\nu}(\varepsilon, x) = \mathbf{1} + \varepsilon^2 F_{\mu\nu}(x) + O(\varepsilon^3), \qquad F = dA + A \wedge A.$$

The Euclideanized MaxEnt/local-Gibbs branch supplies a local finite-range action density built from support-visible gauge-invariant collar data. Gauge quotienting permits only class functions of the curvature and its covariant derivatives. Four-dimensional scaling, Euclidean rotation invariance, locality, and reflection positivity leave one relevant dimension-four positive quadratic invariant in the pure gauge sector:

$$\langle F_{\mu\nu}, F_{\mu\nu} \rangle.$$

The possible topological density $\langle F \wedge F \rangle$ is reflection odd and belongs to a separate topological-angle sector; it is absent on the ordinary reflection-positive zero-obstruction vacuum branch used here. Higher curvature powers and covariant-derivative terms are irrelevant under the declared continuum scaling; their disappearance in the extracted state is the remainder bound in Assumption A.8. Normalizing the unique positive quadratic invariant defines the coupling g. Definition A.5 identifies the finite-stage sources of the gauge cylinders, and Proposition A.6 proves their projective weak-$*$ / GNS extraction. Identification of its transfer generator is the separate dynamic statement in Theorem A.23; it is not a consequence of weak-* compactness alone. □

Prize-facing proof separation. The proof has two separate claims. Theorem A.9 conditionally identifies the four-dimensional Euclidean Yang–Mills form on the certified support-visible compact-gauge branch. The repair-dynamics theorem proves a spectral gap for that Hamiltonian.

Standing compact-gauge setup. Fix a compact simple gauge group G carried by an OPH compact-gauge zero-obstruction vacuum branch on a cofinal tail satisfying the compact-gauge refinement receipt, realized at fixed cutoff by the declared compact-gauge patch-carrier architecture. For each regulator r, let

$$(\mathcal{H}_r, \Omega_r, H_r), \qquad T_r(t) = e^{-tH_r},$$

be the physical Euclidean Hilbert space, vacuum, Hamiltonian, and transfer semigroup. Let $X_r := X_r^{(0)}$ be the support-visible time-zero compact-gauge quotient configuration space, and let π_r be the stationary time-zero measure from Proposition A.6, and set

$$K_r := L^2(X_r, \pi_r).$$

Write $\omega_r(a) := \int_{X_r} a\, d\pi_r$ for its state functional. All statements in K_r are automatically made after the finite π_r-null support reduction. Let $\mathcal{C}_r$ be the finite family of active repair collars. For each $C \in \mathcal{C}_r$, use the complete repaired visible datum $\rho_C := \rho_{C,r} : X_r \to Y_{C,r}$, and let

$$E_C : K_r \to K_r$$

be conditional expectation onto the ρ_C-measurable functions. This subsection stays on the ordinary or central zero-obstruction vacuum branch where the compact-gauge reconstruction ladder is carried by the explicit refinement receipt used in the Standard Model gauge paper.

Proposition A.10 (Local exact repair equals conditional expectation)**.** *For each active collar C, the exact-Markov repair map on the support-visible quotient is the π_r-preserving conditional expectation E_C.*

Proof. On the exact-Markov branch, repair preserves exactly the repaired visible datum ρ_C, changes only complementary invisible fiber data, and acts on the quotient-first physical algebra rather than on representatives. Let Φ_C be the Heisenberg repair map and let $\mathcal{N}_C$ be the repaired local fixed algebra. Exact repair semantics require $\mathcal{N}_C$ to be exactly the ρ_C-measurable subalgebra. Then

$$\Phi_C(a) = a \quad (a \in \mathcal{N}_C), \qquad \Phi_C(\mathcal{A}_r^{G,\mathrm{sv}}) \subseteq \mathcal{N}_C, \qquad \omega_r \circ \Phi_C = \omega_r,$$

and Φ_C is $\mathcal{N}_C$-bimodular. Hence, for $a \in \mathcal{N}_C$ and $x \in \mathcal{A}_r^{G,\mathrm{sv}}$,

$$\omega_r(a^* \Phi_C(x)) = \omega_r(a^* x).$$

Since $\Phi_C(x) \in \mathcal{N}_C$, this characterizes its class in K_r as the orthogonal projection of x onto the ρ_C-measurable subspace, uniquely modulo π_r-null functions. That projection is the π_r-preserving conditional expectation E_C. ☐

Lemma A.11 (Fiber-homogeneous orbit condition)**.** *Fix an active collar C and a repaired value $y \in Y_{C,r}$. Let $F_C(y) = \rho_{C,r}^{-1}(y)$, with conditional law $\nu_{C,y}$. Assume this is either a finite uniform probability space or a standard atomless probability space, and that the primitive collar relaxation is invariant under every $\nu_{C,y}$-preserving automorphism. Then the complete conditional fiber, including every holonomy coordinate actually resampled by repair, is homogeneous for the local repair receipt.*

Proof. This is the complete-fiber homogeneity receipt. Quotienting removes declared implementation labels, but the proof must also check that no remaining source constraint or relaxation rate distinguishes points in the conditional fiber. Under that check, the conditioned state and primitive relaxation carry the stated full measure-preserving symmetry. ☐

Lemma A.12 (Scalar relaxation on a homogeneous conditional fiber)**.** *Let (F, ν) be either a finite uniform probability space or a standard atomless probability space. Let E_F be expectation onto constants, and let D_F be a bounded positive self-adjoint Markov relaxation generator with*

$$\ker D_F = \mathrm{Ran}(E_F)$$

that commutes with every measure-preserving automorphism of (F, ν). Then $D_F = c_F(I - E_F)$ for a scalar $c_F > 0$.

Proof. For a finite uniform fiber this is the symmetric-group argument. For a standard atomless fiber, restrict first to step functions on a partition into n equal-measure pieces. Automorphisms within pieces and permutations of pieces force every bounded operator in the commutant to be scalar on the mean-zero step functions. Compatibility under equal-measure refinements gives one scalar, and such step functions are dense in $L^2_0(F,\nu)$. Thus the commutant on L^2_0 is scalar. Positivity and the stated kernel make that scalar strictly positive. □

Assumption A.13 (Finite ground-state-transform and cross-fiber receipt)**.** *For each regulator, the reflection-positive finite transfer matrix supplies a unitary $U_r : \mathcal{H}_r \to K_r$, with $U_r\Omega_r = \mathbf{1}_r$, whose ground-state-transformed generator decomposes as*

$$U_r H_r U_r^{-1} = \sum_{C\in\mathcal{C}_r} D_C.$$

Each D_C is the bounded positive self-adjoint detailed-balance Markov generator acting on the complete complementary conditional fiber of collar C, including every record or holonomy coordinate changed by repair, with fixed algebra exactly the $\rho_{C,r}$-measurable algebra. If the fiberwise scalar supplied by Lemma A.12 over repaired value y is $c_C(y)$, the finite receipt also verifies the cross-fiber equality $c_C(y) = c_C$ on the support of π_r. A Gibbs state alone does not imply this transfer-matrix/decomposition receipt.

Theorem A.14 (Exact Euclidean-consensus law under homogeneous fibers)**.** *Under the fiber-homogeneous orbit condition for every active collar and Assumption A.13, there are positive constants $c_C > 0$ such that the ground-state transformed physical Euclidean generator is exactly*

$$L_r^{\mathrm{EC}} = \sum_{C\in\mathcal{C}_r} c_C(I - E_C),$$

and therefore

$$U_r e^{-tH_r} U_r^{-1} = e^{-tL_r^{\mathrm{EC}}} \qquad (t \geq 0),$$

for a unitary $U_r : \mathcal{H}_r \to K_r$ with $U_r\Omega_r = \mathbf{1}_r$.

Proof. The finite receipt supplies the ground-state transform and a decomposition into primitive local pieces D_C. Each such piece is supported on collar C, preserves the repaired visible datum $\rho_{C,r}$, and relaxes the complete complementary conditional fiber. Thus

$$\ker D_C = \mathrm{Ran}(E_C).$$

Lemma A.11 gives full hidden-fiber permutation symmetry. Applying Lemma A.12 fiberwise gives $D_{C,y} = c_C(y)(I - E_{C,y})$. The cross-fiber part of the finite receipt makes $c_C(y) = c_C$, so

$$D_C = c_C(I - E_C)$$

for a collar-type constant $c_C > 0$. Summing over the active collars gives the displayed generator. Positivity and self-adjointness give the transfer-semigroup identity. □

Definition A.15 (Source-defined admissible atomic collar tower)**.** *For every regulator and allowed boundary condition b, suppose the finite source presentation gives an atomic register set V_r with finite local alphabets, a finite-range Gibbs law $\pi_{r,b}$, and, for each $v \in V_r$, a rooted active collar $C(v)$ and the heat-bath projection*

$$P_{v,r,b}f := \mathbb{E}_{\pi_{r,b}}[f \mid x_{V_r\setminus\{v\}}].$$

The source type *of* $C(v)$ *is the isomorphism class of the finite tuple consisting of its rooted interaction-radius neighbourhood, boundary/sector flags, local register alphabets, nonzero Gibbs-potential templates, complete repaired readback, conditional kernel, and normalized rate. The tower is admissible when the following data are printed by the source rather than inferred from the target gap:*

(G1) *one finite set* $\mathfrak{T}_{\mathrm{act}}$ *contains every active source type at every location, size, allowed boundary condition, and stage of the selected cofinal refinement tail;*

(G2) *the active family is exactly* $\mathcal{C}_r = \{C(v) : v \in V_r\}$, *with* $E_{C(v)} = P_{v,r,b}$*; type-equivalent collars have the same exact heat-bath kernel and rate, every atomic register has one active rooted collar, and* $c_{v,r,b} = c_{\tau(v)} \geq c_* > 0$*;*

(G3) *for the one-site conditional kernels, let* $a^{r,b}_{vu}$ *be the supremum total-variation change at root* v *when two exterior configurations differ only at* u*. Outward-rounded rational interval bounds printed in the finite type table prove*

$$\sup_{r,b,v} \sum_{u \neq v} a^{r,b}_{vu} \leq \eta_* < 1. \tag{INF}$$

The Dobrushin Poincaré comparison then derives, rather than assumes,

$$\mathrm{Var}_{\pi_{r,b}}(f) \leq \frac{1}{1-\eta_*} \sum_{v \in V_r} \|(I - P_{v,r,b})f\|^2_{2,\pi_{r,b}} \tag{AT}$$

uniformly over the family [59];

(G4) *coarse shadow sends each rooted fine collar to a rooted collar of the same declared type or to an explicitly listed type transition, preserves the conditional kernels on pulled-back coarse functions, and introduces no type outside* $\mathfrak{T}_{\mathrm{act}}$*.*

Collar CMI decay and qualitative finite-range Gibbs mixing are not item (G3)*.*

Proposition A.16 (Finite classification and uniform local-rate floor)**.** *For an admissible atomic collar tower, two active collars have the same source type exactly when their tuples in Definition A.15 are isomorphic. Thus the admissible active collar classes are precisely the nonempty fibers of* $\tau : \bigsqcup_r V_r \to \mathfrak{T}_{\mathrm{act}}$*; in particular there are at most* $|\mathfrak{T}_{\mathrm{act}}|$ *classes. Their rates obey*

$$c_{v,r,b} \geq c_* := \min_{t \in \mathfrak{T}_{\mathrm{act}}} c_t > 0$$

uniformly in location, boundary condition, system size, and cofinal refinement.

Proof. The tuple is a complete source signature, so equality of types is exactly rooted signature-isomorphism. Item (G1) makes its image finite. Item (G2) makes the rate a positive function on that finite image, and item (G4) prevents refinement from creating an unlisted or rate-degenerate signature. The displayed finite minimum is therefore positive and has all four stated uniformities. □

Theorem A.17 (Uniform collar-projection and transfer gap)**.** *On an admissible atomic collar tower satisfying the finite ground-state-transform receipt, define*

$$L^{\mathrm{col}}_{r,b} := \sum_{v \in V_r} c_{v,r,b}(I - P_{v,r,b}), \qquad P_{0,r,b} f := \pi_{r,b}(f)\mathbf{1}.$$

Then, with the single explicit modulus

$$\delta_* := c_*(1-\eta_*) > 0, \tag{GAP}$$

one has, simultaneously for all locations, allowed boundary conditions, system sizes, and stages of the cofinal refinement tower,

$$L^{\mathrm{col}}_{r,b} \geq \delta_*(I - P_{0,r,b}), \qquad \|e^{-tL^{\mathrm{col}}_{r,b}} - P_{0,r,b}\|_{2\to 2} \leq e^{-t\delta_*}.$$

Consequently the collar transfer operator has a positive projection/transfer gap at least δ_*. Under item (G2)*, $L^{\mathrm{col}}_{r,b} = L^{\mathrm{EC}}_r$.*

Proof. Conditional expectation is an orthogonal projection, hence

$$\langle f, (I - P_{v,r,b})f\rangle = \|(I - P_{v,r,b})f\|_2^2.$$

For $f \perp \mathbf{1}$, item (G3) gives (AT), and item (G2) gives

$$\langle f, L^{\mathrm{col}}_{r,b} f\rangle \geq c_* \sum_v \|(I - P_{v,r,b})f\|_2^2 \geq c_*(1-\eta_*)\|f\|_2^2.$$

This is the operator inequality. The spectral theorem gives the semigroup estimate. All constants are source-type constants and item (G4) carries the same table and receipt up the cofinal tower, so none depends on $r, b, |V_r|$, or the collar location. □

Proposition A.18 (Sharp hypothesis countermodels)**.** *Neither locality nor mixing may be deleted from the preceding theorem.*

(i) Without uniform mixing, take the faithful two-site law

$$\pi_\varepsilon(00) = \pi_\varepsilon(11) = \frac{1-\varepsilon}{2}, \qquad \pi_\varepsilon(01) = \pi_\varepsilon(10) = \frac{\varepsilon}{2}.$$

The exact spectrum of the two single-site heat-bath generator is $\{0, 2\varepsilon, 2(1-\varepsilon), 2\}$, so its gap is $2\varepsilon \to 0$. The limiting zero-gap example is therefore not doing any work hidden by a loss of faithfulness.

(ii) Without source locality, put the product fair-bit law on $\{0,1\}^m$, list its $N = 2^m$ states in cyclic Gray-code order, and let E_0, E_1 average over the two alternating perfect matchings of that cycle. The law has exact product mixing and zero separated CMI, but selecting which bit to change uses the entire configuration. For $L_m = (I - E_0) + (I - E_1)$, exact Fourier diagonalization of the cycle gives

$$\mathrm{gap}(L_m) = 1 - \cos(2\pi/2^m) \longrightarrow 0.$$

Both examples are finite at every displayed m; the second is a finite family with an explicit vanishing gap.

Remark A.19 (Claim boundary)**.** *Finite-range Gibbs form and collar-CMI decay alone do not prove* (AT)*, do not classify the kernel/rate signatures, and do not bound the Friedrichs angles between collar projections. Accordingly they do not imply a uniform repair gap. A numerical receipt is admissible only when outward-rounded intervals certify $c_*^{\mathrm{lo}} > 0$ and $\eta_*^{\mathrm{hi}} < 1$, in which case $\delta_*^{\mathrm{lo}} = c_*^{\mathrm{lo}}(1 - \eta_*^{\mathrm{hi}})$; a point estimate is not a certificate.*

Proposition A.20 (Executable finite calibration, not a physical receipt)**.** *The exact-rational four-dimensional Ising calibration table has 244 active collar types on the declared four-dimensional cofinal family $N_r = 3 \cdot 2^r$, for periodic, free, fixed-plus, and fixed-minus boundary conditions. Its complete deterministic table has*

$$c_* = 1, \qquad \eta_* \leq \frac{1}{2}, \qquad A_* \leq 2, \qquad \delta_* \geq \frac{1}{2}.$$

Thus its declared finite generator obeys

$$L_{r,b} \geq \frac{1}{2}(I - P_{0,r,b}), \qquad \|e^{-tL_{r,b}} - P_{0,r,b}\|_{2\to 2} \leq e^{-t/2}.$$

The bundled verifier recomputes every rational conditional-row variation and the receipt hash, and rejects binary floating-point inputs, nonpositive rate floors, and noncontractive influence rows.

Remark A.21 (Non-promotion boundary)**.** *Proposition A.20 validates the finite data model and the stated implication for a declared calibration family. It does not identify that Ising family with the compact-simple-gauge OPH tower. The intentionally uninstantiated physical manifest fails closed until quotient-first gauge collar types, source conditional kernels, rate floors, refinement closure, gauge and topological zero-mode handling, an independent ground-state-transform receipt, and the separate continuum/transfer/OS-noncollapse receipts are supplied.*

Assumption A.22 (Finite transfer and vacuum compatibility receipt)**.** *On the cofinal family selected in Proposition A.6, suppose the following finite-stage identities are certified.*

(T1) *The quotient ensembles obey the finite fiber-sum identity, or its continuous push-forward/disintegration analogue, so $(p^{(4)}_{sr})_{\#}\mu^{(4)}_s = \mu^{(4)}_r$ and, by time-zero restriction, $(p^{(0)}_{sr})_{\#}\pi_s = \pi_r$. Hence pullback gives coherent isometries*

$$J^K_{rs} : K_r \longrightarrow K_s, \qquad J^K_{rs} f = f \circ p^{(0)}_{sr}.$$

(T2) *There are coherent physical refinement isometries $J^{\mathcal{H}}_{rs} : \mathcal{H}_r \to \mathcal{H}_s$ satisfying*

$$J^K_{rs}\mathbf{1}_r = \mathbf{1}_s, \qquad J^{\mathcal{H}}_{rs}\Omega_r = \Omega_s, \qquad U_s J^{\mathcal{H}}_{rs} = J^K_{rs} U_r.$$

(T3) *With one refinement-independent physical Euclidean-time normalization, the finite semigroups are coherent:*

$$J^K_{rs} e^{-tL^{\mathrm{EC}}_r} = e^{-tL^{\mathrm{EC}}_s} J^K_{rs}, \qquad J^{\mathcal{H}}_{rs} e^{-tH_r} = e^{-tH_s} J^{\mathcal{H}}_{rs} \quad (t \geq 0).$$

(T4) *Let $\iota_{0,r}$ pull a time-zero observable into the four-dimensional cylinder algebra and let $\tau^{(r)}_t$ be finite Euclidean time translation. The finite time-zero Markov/transfer identity holds on the cylinder core:*

$$\omega^{(4)}_r\Big((\iota_{0,r} f)^* \tau^{(r)}_t (\iota_{0,r} g)\Big) = \langle f, e^{-tL^{\mathrm{EC}}_r} g\rangle_{K_r} = \langle U_r^{-1} f, e^{-tH_r} U_r^{-1} g\rangle_{\mathcal{H}_r}.$$

The reflection-positive finite OS time-zero quotient is therefore the same finite $(\mathcal{H}_r, \Omega_r, H_r)$, through U_r, rather than an unrelated Hilbert space.

These are finite commuting-diagram receipts. They are not consequences of the sector-category refinement functors or of weak- state compactness.*

Theorem A.23 (Continuum exact transfer identification)**.** *Under Assumption A.22, the finite repair and physical Hilbert spaces have Hilbert direct limits K and $\mathcal{H}$. Their semigroups induce strongly continuous self-adjoint contraction semigroups with nonnegative generators L^{rep} and H, and the finite unitaries induce a unitary $U : \mathcal{H} \to K$ such that*

$$Ue^{-tH}U^{-1} = e^{-tL^{\mathrm{rep}}} \qquad (t \geq 0),$$

and hence

$$UHU^{-1} = L^{\mathrm{rep}}.$$

The constant vectors and physical vacua define limit vectors $\mathbf{1} \in K$ and $\Omega \in \mathcal{H}$, with

$$U\Omega = \mathbf{1}, \qquad P_0^K = |\mathbf{1}\rangle\langle\mathbf{1}|, \qquad P_0^{\mathcal{H}} = |\Omega\rangle\langle\Omega|, \qquad UP_0^{\mathcal{H}}U^{-1} = P_0^K.$$

If $L_r^{\mathrm{EC}} \geq \delta_*(I - P_{0,r})$ for one $\delta_* > 0$, then

$$L^{\mathrm{rep}} \geq \delta_*(I - P_0^K), \qquad H \geq \delta_*(I - P_0^{\mathcal{H}}).$$

Proof. The fiber-sum identity makes J_{rs}^K isometric, since

$$\|f \circ p_{sr}^{(0)}\|_{K_s}^2 = \int_{X_r} |f|^2 \, d(p_{sr}^{(0)})_\# \pi_s = \|f\|_{K_r}^2.$$

It also makes $\overline{\omega}_r^{(0)} = \omega_r^{(0)}$, so these J_{rs}^K are exactly the time-zero GNS isometries $V_{rs}^{(0)}$ in Proposition A.6. Thus their Hilbert direct limit is canonically the support-visible GNS space K constructed there, not a second continuum repair space. Take the Hilbert direct limits under J_{rs}^K and $J_{rs}^{\mathcal{H}}$. On the dense finite-stage images define

$$S^K(t)J_r^K f := J_r^K e^{-tL_r^{\mathrm{EC}}} f, \qquad S^{\mathcal{H}}(t)J_r^{\mathcal{H}}\psi := J_r^{\mathcal{H}} e^{-tH_r}\psi.$$

The finite coherence identities make these definitions independent of the representative. They are contraction semigroups. Strong continuity holds on every finite-stage image and therefore, by density and uniform contractivity, on the whole direct limits. For each fixed t, the finite-stage inner-product identity makes the bounded extension symmetric on a dense union and hence symmetric everywhere; a bounded everywhere-defined symmetric operator is self-adjoint. Positivity also passes from the finite semigroups. Thus the spectral theorem gives nonnegative self-adjoint generators L^{rep} and H.

Define $UJ_r^{\mathcal{H}}\psi := J_r^K U_r \psi$. The finite unitary square makes this well defined and isometric; the same construction with U_r^{-1} shows that its range is dense and closed, so it is unitary. The finite transfer identity of Theorem A.14 gives $US^{\mathcal{H}}(t)U^{-1} = S^K(t)$ on a dense union and hence everywhere. Uniqueness of self-adjoint semigroup generators gives $UHU^{-1} = L^{\mathrm{rep}}$.

Unitality and the physical refinement identity make the constant and vacuum vectors coherent. For $f \in K_r$,

$$P_{0,s} J_{rs}^K f = \langle \mathbf{1}_s, J_{rs}^K f \rangle \mathbf{1}_s = J_{rs}^K P_{0,r} f,$$

so the rank-one vacuum projections induce the displayed limit projections; the same holds on the physical side and the unitary identifies them.

The finite operator inequality implies

$$\|e^{-tL_r^{\mathrm{EC}}}(I - P_{0,r})\| \leq e^{-\delta_* t}.$$

Passing this norm bound through the direct-limit definition gives $\|e^{-tL^{\mathrm{rep}}}(I - P_0^K)\| \leq e^{-\delta_* t}$. The spectral theorem is equivalent to $L^{\mathrm{rep}} \geq \delta_*(I - P_0^K)$, and conjugation by U^{-1} gives the physical bound. ∎

Remark A.24 (Asymptotic alternative)**.** *If the exact semigroup identities in Assumption A.22 are not available, the valid replacement is generalized Mosco convergence of both finite transfer forms, strong cylinder-core convergence of U_r and U_r^{-1}, and strong convergence of the vacuum projections. The Mosco theorem then supplies the two limit semigroups and the same intertwining conclusion. Weak-* convergence of cylinder states by itself supplies none of those dynamic statements.*

Assumption A.25 (OS regularity and noncollapse receipt)**.** *The selected finite cylinder family carries reflection-compatible refinement maps and exact finite reflection positivity. Its renormalized local Schwinger functions have uniform bounds sufficient to pass Euclidean covariance, locality, OS regularity, and clustering to the extracted family. There are gauge-invariant local cylinders A_r, transported from one fixed cylinder test, such that*

$$\inf_r \operatorname{Var}_{\pi_r}(A_r) > 0, \qquad \sup_r \langle A_r \mathbf{1}_r, L_r^{\mathrm{EC}} A_r \mathbf{1}_r \rangle < \infty.$$

The centered vectors $A_r\mathbf{1}_r - \omega_r(A_r)\mathbf{1}_r$ converge under the same cylinder/Hilbert identifications used in the extraction and transfer receipts. This is the separate OS/nontriviality receipt. It is not part of the weak- extraction theorem.*

Theorem A.26 (Osterwalder–Schrader reconstruction on the compact-gauge branch)**.** *Under Assumptions A.4, A.8, A.22, and A.25, with the extraction of Proposition A.6, the continuum support-visible compact-gauge cylinder family is Euclidean invariant, reflection positive, regular on gauge-invariant local cylinder observables, and cyclic for the vacuum sector. Thus Osterwalder–Schrader reconstruction gives a four-dimensional quantum Yang–Mills theory $(\mathcal{H}, \Omega, H, \mathcal{A}^G_{\mathrm{loc}})$ on the support-visible gauge-invariant local algebra, with $H \geq 0$ and e^{-tH} equal to the Euclidean transfer semigroup of Theorem A.9.*

Proof. For every finite list of positive-time cylinders, finite reflection positivity says that the associated reflected Gram matrix is positive semidefinite. Its entries converge along the extraction subnet, and the cone of positive semidefinite matrices is closed, so reflection positivity passes to the limit. Euclidean covariance, locality, regularity, and clustering pass by the compatible transformation maps and uniform bounds in Assumption A.25; weak-* compactness alone would not supply those bounds. The vacuum vector is cyclic for the GNS closure of the gauge-invariant local cylinder algebra. The finite time-zero identity (T4) passes to the limit on the dense cylinder core: the left side converges by four-dimensional cylinder extraction and the right side by the coherent semigroup limit. Consequently the OS time-zero inner product and translation semigroup are unitarily identical to the direct-limit $(\mathcal{H}, \Omega, e^{-tH})$ of Theorem A.23; the notation does not identify two unrelated Hamiltonians. The Osterwalder–Schrader reconstruction theorem then produces the Hilbert space, vacuum, local algebra, positive Hamiltonian, and Euclidean transfer semigroup [60, 61]. ☐

Remark A.27 (Scope of the compact-gauge OS theorem)**.** *The preceding reconstruction is restricted to the support-visible bosonic compact-gauge cylinder algebra on its declared branch. It falls short of the full chiral G_6 exact finite quantization, does not by itself instantiate the Standard-Model continuum observable sector, and supplies no resonance-sheet continuation. Those are separate typed inputs and implications in Appendix A.*

Proposition A.28 (Nontriviality of the support-visible compact-gauge theory)**.** *Under Assumptions A.22 and A.25, The support-visible compact-gauge local algebra on the zero-obstruction vacuum branch strictly contains the vacuum scalars and admits a non-vacuum finite-energy local excitation.*

Proof. The compact-gauge witness and physical-UV landing theorem supplies finite support-visible gauge-invariant local observables, for example nonconstant Wilson/plaquette cylinders. The uniform positive variance in Assumption A.25 says that the chosen centered cylinder does not enter the GNS null ideal in the limit, so it gives a nonzero vector orthogonal to the vacuum. The uniform form bound and lower-semicontinuity of the certified transfer-form limit give this vector finite continuum energy. Thus the support quotient has not erased the excitation that proves nontriviality. □

Theorem A.29 (Conditional positive compact-gauge Yang–Mills mass gap from OPH repair dynamics)**.** *Let G be a compact simple gauge group carried by a support-visible compact-gauge OPH vacuum branch satisfying the standing setup, Assumptions A.22 and A.25, the renormalized Yang–Mills identification receipt of Assumption A.8, the finite ground-state-transform and cross-fiber receipt of Assumption A.13, and all hypotheses of Theorem A.14, Lemma A.16, and Proposition A.17. The theory reconstructed in Theorem A.26 is nontrivial by Proposition A.28, and its continuum support-visible Hamiltonian H satisfies*

$$H \geq \delta_*(I - P_0^{\mathcal{H}}),$$

where $P_0^{\mathcal{H}}$ projects onto the physical vacuum. Therefore

$$\mathrm{Spec}(H) \cap (0, \delta_*) = \varnothing, \qquad \Delta_{\mathrm{YM}} \geq \delta_* > 0.$$

The Yang–Mills gap is the repair gap:

$$\Delta_{\mathrm{YM}} = \Delta_{\mathrm{rep}}.$$

Proof. Proposition A.17 gives

$$L_r^{\mathrm{EC}} \geq \delta_*(I - P_{0,r})$$

at every finite stage. Theorem A.23, using the explicit finite transfer/vacuum receipt, transports the bound to the support-visible continuum:

$$L^{\mathrm{rep}} \geq \delta_*(I - P_0^K).$$

Using $UHU^{-1} = L^{\mathrm{rep}}$, conjugation by U^{-1} gives the Hamiltonian lower bound. The spectral-gap statement follows immediately. Since U is unitary and maps the vacuum to the constant sector, it preserves the nonzero spectrum, so the Yang–Mills gap and repair gap are equal. □

Exact gap accounting. The theorem gives an identity stronger than a phenomenological estimate:

$$\mathrm{Spec}(H) \setminus \{0\} = \mathrm{Spec}(L^{\mathrm{rep}}) \setminus \{0\}.$$

Thus

$$\Delta_{\mathrm{YM}} = \inf\bigl(\mathrm{Spec}(H) \setminus \{0\}\bigr) = \inf\bigl(\mathrm{Spec}(L^{\mathrm{rep}}) \setminus \{0\}\bigr) = \Delta_{\mathrm{rep}}.$$

The finite-stage projection argument proves positivity of that same quantity.

Relation to the Clay/Jaffe–Witten statement. The Clay problem asks for a nontrivial quantum Yang–Mills theory on $\mathbb{R}^4$, for each compact simple G, satisfying axiomatic properties at least as strong as the stated Wightman or Osterwalder–Schrader references and possessing a positive mass gap [58, 57]. Theorem A.9 supplies the four-dimensional Euclidean Yang–Mills form on the OPH support-visible compact-gauge branch, and Theorem A.26 supplies the support-visible OS reconstruction on the gauge-invariant local algebra. Proposition A.28 supplies nontriviality. Theorem A.29 supplies the exact spectral gap accounting on that same branch. Proposition A.6 proves the algebra/state/GNS extraction only. Full Clay admissibility additionally requires the finite transfer/vacuum and OS/noncollapse receipts in Assumptions A.22 and A.25, together with the renormalized Yang–Mills identification, finite ground-state-transform/cross-fiber, and uniform-gap receipts; weak-* compactness is not a substitute for those inputs. The standalone Clay note is only a focused presentation of this same theorem surface; it does not enlarge the branch beyond the data stated here.

References

[1] B. Müller, D. Matscheko, and J. Hill, *Observation-Determined Normal Forms: Stability, Obstructions, and Refinement in Constraint and Rewrite Systems*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/observable_normal_forms.pdf.

[2] B. Müller, A. Osika, M. Poneder, K. Xue, P. Nguyen, M. A. Visser, and D. Matscheko, *Recovering Observer Spacetime and Einstein Dynamics from Overlap Consistency*, 2026. Available at GitHub PDF.

[3] B. Müller, A. Osika, M. Poneder, K. Xue, P. Nguyen, M. A. Visser, and D. Matscheko, *Deriving Standard Model Gauge Structure from Observer Overlap Consistency*, 2026. Available at GitHub PDF.

[4] B. Müller, K. Xue, K. A. Anirudha, D. Matscheko, and J. Hill, *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/reality_as_consensus_protocol.pdf.

[5] B. Müller, A. Osika, K. Xue, B. Cassie, M. A. Visser, and D. Matscheko, *Federated Echosahedral Screen Microphysics: Patch Hardware, Records, and Observer Synchronization in OPH*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/screen_microphysics_and_observer_synchronization.pdf.

[6] B. Müller, A. Osika, M. Poneder, K. Xue, M. A. Visser, and D. Matscheko, *Deriving the Particle Zoo from Observer Consistency*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/deriving_the_particle_zoo_from_observer_consistency.pdf.

[7] B. Müller, A. Osika, and D. Matscheko, *Theoretical Bounds on χ_ν in Observer-Patch Holography*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/chi_nu_susceptibility_bounds.pdf.

[8] B. Müller and D. Matscheko, *Observer-Patch Holography and the Dark Matter Phenomenon*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/cosmology/oph_dark_matter_paper.pdf.

[9] B. Müller, *Explaining the Yang–Mills Mass Gap with Observer-Patch Repair Dynamics*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/yang_mills_gap_clay_problem.pdf.

[10] B. Müller, *Observer-Patch Holography as a String-Vacuum Selector: Observers, Clocks, Edge Strings, and the Bouchard-Donagi Witness*, 2026. Available at https://github.com/Floatin

gPragma/observer-patch-holography/blob/main/extra/observer_patch_holography_as_string_vacuum_selector.pdf.

[11] B. Müller, *The de Sitter Time-Advance Sign from a Finite Screen with Fixed Capacity*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/de_sitter_time_advance_sign_from_fixed_screen_capacity.pdf.

[12] B. Müller, *OPH-FPE: finite simulator and receipt engine for Observer-Patch Holography physics experiments*, 2026. Available at https://github.com/muellerberndt/oph-physics-sim.

[13] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, and T. Schwetz, *NuFit-6.0: Updated global analysis of three-flavor neutrino oscillations*, JHEP **12** (2024) 216, arXiv:2410.05380, with the NuFIT 6.1 (2025) profile-table release at https://www.nu-fit.org/?q=node/309.

[14] M. H. A. Newman, "On theories with a combinatorial definition of 'equivalence'," *Ann. of Math. (2)* **43** (1942), 223–243.

[15] E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," *Commun. Math. Phys.* **28** (1972), 251–257.

[16] J. J. Bisognano and E. H. Wichmann, "On the duality condition for a Hermitian scalar field," *J. Math. Phys.* **16** (1975), 985–1007.

[17] J. J. Bisognano and E. H. Wichmann, "On the duality condition for quantum fields," *J. Math. Phys.* **17** (1976), 303–321.

[18] R. Brunetti, D. Guido, and R. Longo, "Modular Structure and Duality in Conformal Quantum Field Theory," *Commun. Math. Phys.* **156** (1993), 201–219, arXiv:funct-an/9302008.

[19] H.-W. Wiesbrock, "Half-Sided modular inclusions of von-Neumann-Algebras," *Commun. Math. Phys.* **157** (1993), 83–92.

[20] H. Araki and L. Zsidó, "Extension of the structure theorem of Borchers and its application to half-sided modular inclusions," *Rev. Math. Phys.* **17** (2005), 491–543, arXiv:math/0412061.

[21] D. Guido, R. Longo, and H.-W. Wiesbrock, "Extensions of Conformal Nets and Superselection Structures," *Commun. Math. Phys.* **192** (1998), 217–244, arXiv:hep-th/9703129.

[22] H.-W. Wiesbrock, "Modular Intersections of von-Neumann-Algebras in Quantum Field Theory," *Commun. Math. Phys.* **193** (1998), 269–285.

[23] W. G. Unruh, "Notes on black-hole evaporation," *Phys. Rev. D* **14** (1976), 870–892.

[24] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Phys. Rev. Lett.* **75** (1995), 1260–1263, doi:10.1103/PhysRevLett.75.1260, arXiv:gr-qc/9504004.

[25] T. Jacobson, "Entanglement equilibrium and the Einstein equation," *Phys. Rev. Lett.* **116** (2016), 201101, doi:10.1103/PhysRevLett.116.201101, arXiv:1505.04753.

[26] G. W. Gibbons and S. W. Hawking, "Cosmological event horizons, thermodynamics, and particle creation," *Phys. Rev. D* **15** (1977), 2738–2751, doi:10.1103/PhysRevD.15.2738.

[27] Y. Chen, D. Stanford, H. Tang, and Z. Yang, "Negative shocks versus static patch holography," arXiv:2607.14042 [hep-th] (2026). https://arxiv.org/abs/2607.14042

[28] Planck Collaboration, "Planck 2018 results. VI. Cosmological parameters," *Astron. Astrophys.* **641** (2020), A6, doi:10.1051/0004-6361/201833910, arXiv:1807.06209.

[29] R. Bousso, Z. Fisher, J. Koeller, S. Leichenauer, and A. C. Wall, "Proof of the quantum null energy condition," *Phys. Rev. D* **93** (2016), 024017, arXiv:1509.02542.

[30] R. Bousso, Z. Fisher, S. Leichenauer, and A. C. Wall, "Quantum focusing conjecture," *Phys. Rev. D* **93** (2016), 064044, arXiv:1506.02669.

[31] S. Balakrishnan, T. Faulkner, Z. U. Khandker, and H. Wang, "A general proof of the quantum null energy condition," *J. High Energy Phys.* **2019** (2019), 20, arXiv:1706.09432.

[32] E. H. Lieb and M. B. Ruskai, "Proof of the strong subadditivity of quantum-mechanical entropy," *J. Math. Phys.* **14** (1973), 1938–1941.

[33] E. H. Wichmann, "Density matrices arising from incomplete measurements," *J. Math. Phys.* **4** (1963), 884–896.

[34] F. Hiai, M. Ohya, and M. Tsukada, "Sufficiency, KMS condition and relative entropy in von Neumann algebras," *Pacific J. Math.* **96** (1981), 99–109.

[35] M. Takesaki, "Conditional expectations in von Neumann algebras," *J. Funct. Anal.* **9** (1972), 306–321.

[36] D. Petz, "Sufficient subalgebras and the relative entropy of states of a von Neumann algebra," *Commun. Math. Phys.* **105** (1986), 123–131.

[37] D. Petz, "Sufficiency of channels over von Neumann algebras," *Quart. J. Math.* **39** (1988), 97–108.

[38] O. Fawzi and R. Renner, "Quantum conditional mutual information and approximate Markov chains," *Commun. Math. Phys.* **340** (2015), 575–611, arXiv:1410.0664.

[39] P. Hayden, R. Jozsa, D. Petz, and A. Winter, "Structure of states which satisfy strong subadditivity of quantum entropy with equality," *Commun. Math. Phys.* **246** (2004), 359–374.

[40] S. Doplicher and J. E. Roberts, "A new duality theory for compact groups," *Invent. Math.* **98** (1989), 157–218.

[41] S. Doplicher and J. E. Roberts, "Why there is a field algebra with a compact gauge group describing the superselection structure in particle physics," *Commun. Math. Phys.* **131** (1990), 51–107.

[42] T. Tannaka, "Über den Dualitätssatz der nichtkommutativen topologischen Gruppen," *Tohoku Math. J.* **45** (1938), 1–12.

[43] M. G. Krein, "A principle of duality for a bicompact group and a square block algebra," *Dokl. Akad. Nauk SSSR* **69** (1949), 725–728.

[44] H. Georgi and S. L. Glashow, "Unity of all elementary-particle forces," *Phys. Rev. Lett.* **32** (1974), 438–441.

[45] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak interactions with lepton-hadron symmetry," *Phys. Rev. D* **2** (1970), 1285–1292.

[46] E. Witten, "An SU(2) anomaly," *Phys. Lett. B* **117** (1982), 324–328.

[47] S. Dimopoulos, S. Raby, and F. Wilczek, "Supersymmetry and the scale of unification," *Phys. Rev. D* **24** (1981), 1681–1683.

[48] U. Amaldi, W. de Boer, and H. Fürstenau, "Comparison of grand unified theories with electroweak and strong coupling constants measured at LEP," *Phys. Lett. B* **260** (1991), 447–455.

[49] D. J. Gross and W. Taylor, "Two-dimensional QCD is a string theory," *Nucl. Phys. B* **400** (1993), 181–208, arXiv:hep-th/9301068.

[50] F. Peter and H. Weyl, "Die Vollständigkeit der primitiven Darstellungen einer geschlossenen kontinuierlichen Gruppe," *Math. Ann.* **97** (1927), 737–755.

[51] A. Bullivant, M. Calçada, Z. Kádár, P. Martin, and J. Faria Martins, "Topological phases from higher gauge symmetry in 3+1D," *Phys. Rev. B* **95** (2017), 155118, arXiv:1606.06639.

[52] S. Chandrasekharan and U.-J. Wiese, "Quantum link models: A discrete approach to gauge theories," *Nucl. Phys. B* **492** (1997), 455–471, arXiv:hep-lat/9609042.

[53] W. Donnelly and A. C. Wall, "Entanglement entropy of electromagnetic edge modes," *Phys. Rev. Lett.* **114** (2015), 111603, arXiv:1412.1895.

[54] F. Pastawski, B. Yoshida, D. Harlow, and J. Preskill, "Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence," *JHEP* **06** (2015), 149, arXiv:1503.06237.

[55] M. A. Levin and X.-G. Wen, "String-net condensation: A physical mechanism for topological phases," *Phys. Rev. B* **71** (2005), 045110, arXiv:cond-mat/0404617.

[56] D. Laghi, G. Carullo, J. Veitch, and W. Del Pozzo, "Quantum black hole spectroscopy: probing the quantum nature of the black hole area using LIGO-Virgo ringdown detections," *Class. Quantum Grav.* **38** (2021), 095005, arXiv:2011.03816.

[57] Clay Mathematics Institute, "Yang–Mills & the Mass Gap." Available at https://www.claymath.org/millennium/yang-mills-the-maths-gap/.

[58] A. Jaffe and E. Witten, "Quantum Yang–Mills Theory," official Clay Mathematics Institute problem description. Available at https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf.

[59] L. Wu, "Poincaré and transportation inequalities for Gibbs measures under the Dobrushin uniqueness condition," *Annals of Probability* **34** (2006), 1960–1989, arXiv:math/0611635.

[60] K. Osterwalder and R. Schrader, "Axioms for Euclidean Green's functions," *Communications in Mathematical Physics* **31** (1973), 83–112.

[61] K. Osterwalder and R. Schrader, "Axioms for Euclidean Green's functions II," *Communications in Mathematical Physics* **42** (1975), 281–305.

[62] Particle Data Group, "Review of Particle Physics," *Phys. Rev. D* **110** (2024), 030001. Available at https://pdg.lbl.gov/2024/download/db2024.pdf.

[63] Particle Data Group, "Electroweak Model and Constraints on New Physics," Available at https://pdg.lbl.gov/2024/reviews/rpp2024-rev-standard-model.pdf.