Recovering Relativity and the Standard Model from Observer Overlap Consistency

B. Müller Alexander Osika Kai Xue Peter Nguyen

June 8, 2026

Paper release: `r1465` **Released:** June 8, 2026

Abstract

Observer-Patch Holography starts from finite screen geometry following the spherical symmetries of observer-accessible cuts. Neighboring observer patches must agree on their overlap. The screen is a regulator and symmetry chart; in the broader OPH stack the relevant geometry includes A_5-type icosahedral patch data and E_8-type exceptional organization where those branches are invoked. From that rule, the recovered core yields a derived Lorentz/null-modular/Einstein branch and a bosonic compact-gauge reconstruction branch. On the realized minimal one-Higgs matter branch, the same stack yields the realized Standard Model quotient $(\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1))/\mathbb{Z}_6$ with the exact hypercharge lattice, the structural electroweak gauge content $\mathrm{SU}(2)_L\times\mathrm{U}(1)_Y\to\mathrm{U}(1)_Q$, the realized color triplet $N_c=3$, and the generation count $N_g=3$. On the support-visible compact-gauge branch, the same compact-gauge reconstruction and four-dimensional scaling data yield the Euclidean Yang–Mills form; exact repair dynamics identifies Euclidean transfer with a positive repair generator and gives a compact-gauge Yang–Mills mass gap for compact simple groups carried by that branch. Fermions and chirality enter through the explicit realized matter-sector input; the compact paper leaves the super-Tannakian reconstruction theorem outside its scope. The basis used here is quantum-algebraic: patch algebras, states, trace/Born probabilities where measurement records are invoked, and generalized entropy are part of the starting formalism. OPH's reconstruction claim is that this algebraic-information basis can support a consistent theory-of-everything program by recovering the observed effective universe from overlap consistency.

The quantitative completion is formulated as a zero-input closure program with two constants rather than fitted parameters. The local pixel fixed point is $P_\star=\varphi+\sqrt{\pi}/A_T(P_\star)$. The global capacity fixed point is the cosmic record-closure fixed point $N_{\mathrm{CRC}}=F(N_{\mathrm{CRC}})$, where

$$F(N)=\mathrm{Cap}_{\mathrm{read}}\big(\mathrm{Obs}(\mathrm{nf}(\mathfrak{U}_N))\big)$$

is the active horizon capacity read back by observers inside the universe supplied with capacity N. The finite-count representation is

$$N_\star=\mathrm{MAR}\arg\max_N(\log|\Omega_N^{\mathrm{sc}}|-N).$$

A contraction certificate for F, or equivalently a derivative-sign certificate for $\ell(N)=\log|\Omega_N^{\mathrm{sc}}|-N$, gives the unique stable fixed point. Informally, P links outside screen-cell area to inside electromagnetic interaction/observation strength, while N_{CRC} links outside total horizon capacity to inside observer-accessible public record capacity. The universe must be able to reconstruct its own boundary: observers inside infer geometry, horizons, entropy, Λ, history, and records from information available inside the universe. Thus N_{CRC} is the unique global capacity at which the universe reads back its own boundary without deficit or slack. The cosmological readout is $\Lambda_{\mathrm{CRC}}=3\pi/(GN_{\mathrm{CRC}})$. On the observed branch this fixed point is the de Sitter entropy capacity $N_{\mathrm{CRC}}\simeq 3.31\times10^{122}$. This compact paper isolates the recovered core and keeps it distinct from the quantitative particle and continuation branches built on top of the same closure data.

1 Introduction

Physics is usually presented in two boxes. One box contains spacetime and gravity. The other contains the particle world. This paper asks whether both can be recovered from a smaller starting point: a finite screen whose neighboring observer patches must agree on their overlaps.

That single requirement drives a long structural chain. It explains why the Lorentz group appears, why the low-energy spacetime description follows a Jacobson-type Einstein branch, why the gauge structure is

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

why the observed three colors and three generations are realized, why strong/color, weak-isospin, hypercharge, and electric-charge assignments take their observed form, what fixes the electroweak and flavor scales, and what the framework can and cannot say about the cosmological constant and possible dark-sector continuations.

This paper presents the compact reconstruction program based on Observer-Patch Holography (OPH). Physical data are organized patchwise: each observer has access only to a patch algebra, neighboring observers must agree on overlaps, and the reconstruction is driven by those consistency conditions. Under the five axioms and the branch hypotheses and theorems stated below, that structure feeds a Lorentzian modular-phase classification, a Jacobson-type Einstein branch, compact gauge reconstruction in a bosonic internal-gauge sector, color and electroweak charge quantization, and the low-energy architecture fixed on the realized MAR gauge branch. Black-hole spectroscopy, Page-curve/island claims, dark-sector models, and particle-family continuations belong to companion or continuation surfaces.

The paper separates exact structural theorems, scaling-limit branches, quantitative outputs, and phenomenological continuations. Several longer arguments are written in compressed proof form so the claim boundary stays readable.

The framework is presented as a reconstruction program in which general relativity and the Standard Model appear as effective sectors derived from the stated quantum-algebraic axiom set, explicit refinement constructions, and the support-visible BW scaling theorem. This theory-of-everything framing is a claim about recovering the observed effective universe from that basis and its declared branch hypotheses, not a demand that every mathematical ingredient be derived from a blank starting point. In the gauge lane, the logical split is explicit: overlap and holonomy data classify transportable sectors and obstruction routing; Doplicher–Roberts/Tannaka reconstruction yields a compact group from the zero-obstruction bosonic sector category; MAR plus the realized one-Higgs chiral matter package selects the Standard Model quotient on the realized branch.

Prioritized achievements

The results emphasized in this paper are ordered below by a combination of impact and defensibility.

1. A Jacobson-type Einstein branch is recovered from observer-overlap modular geometry, the support-visible BW scaling theorem on the prime geometric subnet, and the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations on the realized cap-label-preserving MaxEnt family; in the null modular bridge, quasi-local propagation, exact-or-controlled strip additivity, endpoint-Lipschitz renormalized half-line control, the resulting weak tail generator, and the half-sided modular pair are supplied internally by the null-interval structure and the scaling-limit geometric action, after which Borchers–Wiesbrock yields the explicit positive null-translation generator on its Stone domain together with the

affine half-line modular relation and the half-line generator/charge identification with the local null-stress charge on that same family; the Einstein side then adds the internal small-ball bridge from the geometric cap generator together with bounded-interval transport through the separate projective branch, and the later tensor upgrade states its all-directions/all-reference-states conditions explicitly on that scaling branch.

2. The exact Standard Model gauge structure

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

is recovered on the realized MAR-admissible branch of the bosonic compact-gauge theorem; together with the one-Higgs package it fixes the structural electroweak force content $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_Q$ and the $W^{\pm}/Z/A$ gauge-boson basis. This is a realized-branch gauge result with explicit matter-package and admissibility inputs; numerical W/Z masses and mixing are D10 quantitative readouts.

3. On the support-visible compact-gauge branch, OPH derives the four-dimensional Euclidean Yang–Mills form from compact-gauge holonomy data and the local MaxEnt/Gibbs continuum limit. Local repair collars generate Euclidean time through conditional expectations. A uniform positive repair gap transfers exactly to the compact-gauge Hamiltonian, so the Yang–Mills gap is the repair gap for every compact simple group carried by that OPH branch. The Clay-facing reading is the companion theorem surface in which support-visible continuum extraction and Osterwalder–Schrader reconstruction supply the claimed four-dimensional axiomatic construction on that declared branch.

4. On the realized one-generation chiral matter plus one-Higgs package, anomaly constraints and Yukawa invariance fix the hypercharge lattice, the realized color triplet fixes $N_c = 3$, and the same one-Higgs quark branch then fixes $N_g = 3$ from CKM phase counting, weak-sector asymptotic freedom, and MAR.

5. The quantitative implementation used here exposes two fixed-point closure values: the local pixel fixed point $P_\star$ and the cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$.

6. Downstream matter-sector continuations, including the charged-lepton exact centered-readback / common-shift frontier, are visible but explicitly outside this SM/GR derivation paper's recovered-core theorem package.

7. The same Einstein branch closes globally at the cosmic record-capacity fixed point, tying the cosmological-constant branch to global screen capacity rather than local vacuum energy.

8. Exclusion of gauge-mediated proton decay follows from product-group structure; any stronger proton-lifetime claim is UV/EFT/QCD-dependent.

9. Dark-sector, heuristic baryogenesis continuations, string, and spectroscopy branches are kept visible only as phenomenological continuations or, where additional ansätze are involved, conjectural phenomenology.

10. A bare overlap net is a finite constraint code: codewords are the globally overlap-consistent states $C = \Phi^{-1}(0)$. Overlap repair admits a unique schedule-independent normal form from each fixed initial physical quotient state on the stated finite patch-net hypotheses, where Lyapunov descent gives termination and the local-diamond plus repair-completeness clauses give

confluence. Same-boundary uniqueness additionally requires a preserved boundary/sector map whose consistent quotient fiber has a unique extension. The same package induces refinement-limit normal-form and holonomy classes on separated cofinal refinement systems with compatible projections, and makes reconciliation commute with coarse-graining up to explicitly declared normal-form and obstruction defects. QECC distance, min-cut resilience, exponential convergence, BFT wall-clock liveness, and hardware speedup are separate certified branches, not consequences of a generic overlap graph.

11. The fixed-cutoff topological UV package closes on the ordinary branch, the central-defect branch, and the genuinely noncentral branch through a compact crossed-module higher-gauge collar theorem; the support-visible continuum BW/geometric lift is supplied by Theorem 6.8, while the realized compact-gauge branch is supplied by the compact-gauge witness and physical-UV landing theorem below.

On that same fixed-cutoff consensus surface, the status split is sharp: normal-form computation is a finite-state decidable problem with the Lyapunov step bound from the consensus paper, the only automatic approximate-stability inputs are the collar-local splice and record estimates carried there. Long-run noisy approximate consensus is theorem-grade only on the separate fair-block contraction branch, where the implementation supplies $(\lambda, \varepsilon, A, \beta, L)$ for distance to the exact quotient normal-form set. Computational expressiveness for growing patch-net families is a separate consensus-paper complexity boundary rather than a recovered-core dependency. The same firewall applies to coding language: the default code is a finite constraint code, while topological-code distance/min-cut and Knill–Laflamme resilience require explicit topological-code and error-model certificates. The imported consensus package is the fixed-cutoff theorem stack: asynchronous confluence for the declared recovery-derived repair law, with termination separated from the local-diamond and repair-completeness clauses, the cycle-obstruction / higher-gauge defect package, gauge-quotient descent, observable-level confluence on the declared physical observable algebras, the refinement-limit normal-form / holonomy theorem, the controlled coarse-graining / reconciliation square, the no-free-min-cut boundary for bare graphs, and the record-algebra theorem on the observer-accessible surface. The imported consensus branch conditions are repair completeness and the stated Petz support/CPTP clause where that branch is used. The scaling-limit bridge from that fixed-cutoff patch-net package to the Lorentz, Einstein, and compact-gauge branches is separate; the consensus paper supplies the intermediate inverse-limit normal-form and holonomy theorem once the cofinal refinement projections commute with finite-stage normal forms and holonomy maps, and it supplies the approximate RG/reconciliation comparison once the chosen coarse-graining channel has controlled normal-form and obstruction defects.

2 Overview of Results

This section summarizes the claim tiers used throughout. The table below is the formal dependency ledger for this SM/GR derivation paper. The recovered core contains fixed-cutoff overlap repair and collar structure, the Lorentz/null-modular/Einstein branch, bosonic compact gauge reconstruction, and the realized-branch Standard Model quotient with exact hypercharge, structural electroweak force content, $N_c = 3$, and $N_g = 3$ under the explicit matter-package and admissibility inputs. The quantitative closures contain the screen-capacity closure of the same Einstein branch and the pixel-driven electroweak/gauge-coupling branch. The continuations contain flavor details beyond the stated theorem surfaces, dark-sector proposals, heuristic baryogenesis branches, spectroscopy, hadrons, and string/worldsheet effective-description branches. The canonical five-axiom basis, the quantitative quantities listed next, the support-visible BW scaling theorem, and the theorem-produced gauge chain used below are fixed in Section 3. Each row names a node in the reconstruction program, records its immediate internal parents, separates imported standard mathematics from declared external inputs, and states the resulting claim tier.

Theorem key for the table. For the recovered core, the scaling/BW step is Theorem 6.8, the support-visible BW scaling theorem. The compact-gauge nontriviality step is handled by the compact-gauge witness and physical-UV landing theorem. The fixed-cap generalized-entropy stationarity result is Theorem 6.28. Transportability, the fixed-cutoff bosonic-category construction, and refinement/fiber descent are theorem-produced by Theorems 5.8, 5.11, and 7.2.

Table 1: Dependency checklist for the OPH SM/GR reconstruction program.

Node	Output	Immediate internal ingredients	Standard mathematics used	Branch-local inputs / external data	Claim tier
D1	Unique schedule-independent normal form from each fixed initial physical quotient state, boundary-conditioned uniqueness under a preserved boundary/sector map with unique consistent extension, observable-level confluence on the declared fixed-cutoff physical algebras for the recovery-derived repair relation, inverse-limit normal-form / holonomy classes on separated cofinal refinement systems, and controlled reconciliation/coarse-graining compatibility	overlap-consistency problem on a finite patch net together with the declared fixed-cutoff collar recovery step, its touched-overlap local-fit acceptance contract, support-local disjoint commutation, restriction-compatible union-collar gluing on the physical quotient, the fixed-point / quotient-local physical observable algebras carried by that same collar package, cofinal refinement projections compatible with finite-stage normal forms and holonomy maps, and a coarse-graining channel with declared normal-form and obstruction defects	local-to-global Lyapunov descent on the inconsistency potential; local-diamond completion from overlap-associative center-sector / higher-gauge union-collar gluing; Newman's lemma; inverse limits of separated finite quotient systems; pseudometric control of commuting coarse-graining diagrams	repair completeness; stated Petz support/CPTP control where that branch is used; for same-boundary uniqueness, preservation of the boundary/sector map and at most one consistent quotient extension in the fiber; normal-form naturality, holonomy naturality, visible separation, and explicit coarse-graining defect bounds for the refinement system	Phase I structural theorem
D2	Collar Markov/entropy split and G-dictionary setup (Theorem 5.2, Propositions 5.18, 5.19, 5.25, Definition 5.26)	Axioms 3.1–3.4	HJPW Markov structure theorem; Petz and Fawzi–Renner recovery theory	exact block-factorized identities require exact Markovity or an explicitly stated idealized limit; the $A/(4G)$ step is the coarse-grained dictionary of Definition 5.26	Phase I structural lemma/proposition layer
D3	Lorentz branch $\mathrm{Conf}^+(S^2) \cong \mathrm{SO}^+(3,1)$ (Theorem 6.8, Corollary 6.16)	Axioms 3.1–3.4, Theorem 6.2, and the fixed-cutoff collar/consensus package	Bisognano–Wichmann modular geometry; conformal-group classification; weak-$*$/GNS extraction	support-visible regularized modular transport on the prime geometric cap net; support-readable modular covariance; ordered cut-pair rigidity; the scaling-limit observer algebra may leave the finite type-I regulator class, so the theorem is automorphism-level there and $K_C = 2\pi B_C$ is only the special type-I limit form	Phase I scaling-limit theorem

Table 1: Dependency checklist for the OPH SM/GR reconstruction program.

Node	Output	Immediate internal ingredients	Standard mathematics used	Branch-local inputs / external data	Claim tier
D4	Null modular bridge to T_{kk}, including exact-or-controlled strip additivity, endpoint-Lipschitz renormalized half-line families, the weak tail generator, the derived half-sided modular inclusion, the explicit positive half-line null-translation generator on its Stone domain with affine half-line modular relations, and the exact half-line generator/charge identification with the local null-stress charge (Propositions 6.17, 6.20; Corollaries 6.18, 6.22; Theorem 6.21, Lemma 6.23; Theorem 6.24)	D2+D3 and Axioms 3.1–3.4	Borchers–Wiesbrock positive-generator theorem for standard half-sided inclusions; Stone's theorem; distributional differentiation on half-lines	the theorem-local null-cut center transfer, the inherited left/right strip-split package used for the spatial-collar-type tensor decomposition, and the exact-or-controlled Markov hypotheses of Proposition 6.17 and Corollary 6.18; quasi-local propagation and endpoint-Lipschitz control are supplied internally by Axiom 3.3 on the local finite-constraint branch, the geometric scaling action on the null half-line blow-up net derives the half-sided modular pair, and bounded-interval formulas use the separate interval-preserving projective branch	Phase I bridge theorem
D5	Jacobson-type Einstein branch: rest-frame relation plus tensor upgrade (Theorem 6.28, Lemma 6.30, Theorem 6.32, Corollary 6.33)	D3+D4 and Axioms 3.1–3.4	fixed-volume area-variation identity for small geodesic balls; local quadratic-polarization argument for the tensor upgrade	Theorem 6.28, which fixes the admissible fixed-cap MaxEnt variation class and derives generalized-entropy stationarity on the realized cap-label-preserving MaxEnt family; the internal small-ball bridge of Lemma 6.30, which uses the geometric cap generator together with the D4 half-line generator/charge identification and the separate interval-preserving projective branch; locally Lorentzian $d = 4$ scaling regime; small-ball constancy assumptions; D4 carried remainder negligible at order ℓ^4; all local directions/reference states for the tensor upgrade	Phase I scaling-limit theorem/corollary
D6	Local/global theorem stack for Λ: null-invisible metric ambiguity, screen-capacity closure of the same Einstein branch, the derived de Sitter static-patch parameter package, and the cosmic record-closure fixed point on the observed branch, plus a legacy capacity-level neutrino side estimate (Theorem 6.45)	D5 local Einstein recovery leaves a $+\Lambda g_{ab}$ ambiguity	de Sitter entropy relation, static-patch radius/time formulas, finite self-closing normal-form density/readback map, and dimensional analysis	cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$ and its observed de Sitter entropy readout	Phase II zero-input global closure

Table 1: Dependency checklist for the OPH SM/GR reconstruction program.

Node	Output	Immediate internal ingredients	Standard mathematics used	Branch-local inputs / external data	Claim tier
D7	Construction of the refinement-stable bosonic sector category plus compact gauge reconstruction (Theorem 7.2, Theorem 7.3, Theorem 7.4), the four-dimensional Euclidean Yang–Mills form (Theorem A.17), and the support-visible compact-gauge repair-gap theorem (Theorem A.28)	Axioms 3.1–3.4	directed-colimit descent for monoidal C^*-categories; Doplicher–Roberts / Tannaka reconstruction; holonomy-to-curvature expansion; finite-dimensional conditional expectations and spectral comparison for commuting collar projections	Theorems 5.8, 5.11, and 7.2 on the ordinary or central-defect bosonic zero-obstruction branch; the compact-gauge witness theorem supplies nonempty realized MAR-admissible witness data, while the Standard Model selection step belongs to D8–D9; the Euclidean-form and repair-gap results additionally use the four-dimensional scaling chart, reflection-positive ordinary vacuum, active exact-Markov repair collars, bounded-color collar covers, repair completeness, and support-visible continuum extraction; the genuinely noncentral fixed-cutoff branch is handled separately by Theorems 5.3–5.6 and is not itself the ordinary compact-group reconstruction theorem	Phase I structural theorem
D8	Product gauge structure up to finite quotient (Lemmas 7.7–7.12, Theorem 7.13)	D7 + Axiom 3.5	compact Lie representation classification; Schur's lemma	the same ordinary or central-defect bosonic branch as D7, together with a connected positive-dimensional Lie admissible class, one connected abelian factor, and faithful action on the minimal coupled carrier	Phase I realized-branch theorem
D9	Realized Standard Model quotient, hypercharges, structural electroweak force content, $N_c = 3$, $N_g = 3$, and product-group corollaries (Theorem 7.24, Theorem 7.14, Corollaries 7.15, 7.16, 7.22, Proposition 7.19)	D8	anomaly-cancellation algebra; Witten global-anomaly argument; CKM CP counting; stabilizer computation for the neutral Higgs direction	realized one-generation chiral matter plus one Higgs package, with the CP-capability and weak-sector UV clauses contained in Axiom 3.5; no extra post-MAR selector is used	Phase I realized-branch theorem/corollary chain
D10	Forward gauge-coupling closure and declared electroweak readout of the integrated D10 quantitative-closure package	D9 + pixel ratio P	printed RG evolution, matching, and scheme-conversion conventions	pixel constraint and the source-only forward transmutation solve $\mathcal{F}(\alpha_U; P) = 0$; the fixed-cutoff edge heat-kernel / Casimir theorem on the microphysics surface together with the compact-group / Peter–Weyl lift used on the D10 lane; printed beta-function, threshold, and scheme-conversion conventions	Phase II quantitative-closure sector
D12	Charged-lepton exact centered-readback / common-shift frontier, strong-CP branch, texture, dark-sector, heuristic baryogenesis continuations, black-hole spectroscopy, proton-spin, proton-lifetime estimates beyond the gauge-channel exclusion, controlled large-N_{edge} string/worldsheet effective descriptions, conjectural critical-superstring extensions, and other continuations	various subsets of D6, D9, and D10	branch-specific EFT and phenomenological manipulations	additional ansätze such as the uniform $\mathbb{Z}_6$ center-label ensemble, discrete texture choices, dark-sector response assumptions, discrete-horizon assumptions, a distinct large-N_{edge} regime with fixed $\tau = tN_{\text{edge}}$ window and uniform genus-remainder control, or additional worldsheet/CFT assumptions	Phase III phenomenological continuation branch

The recovered core of this SM/GR derivation paper is D1–D5 together with D7–D9. Within that package, D7 is the bosonic sector-category construction plus compact-gauge reconstruction stage, while D8–D9 are realized-branch results that keep the realized one-generation chiral matter plus one-Higgs package and MAR admissibility inputs explicit. D6 sits outside that recovered core because the cosmological-capacity closure is the separate global completion of the Einstein branch. D5 leaves the metric term undetermined, D6 closes that same branch at the cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$, and the same D6 stack emits the de Sitter static-patch parameter package. The density $\log|\Omega_N^{\mathrm{sc}}| - N$ is the finite-count representation of the same target after division by the full screen Hilbert-space dimension e^N. Informally, that point is where the universe reads back its own boundary without deficit or slack. D10 is the integrated quantitative-closure branch, and D12 collects phenomenological continuations. The modular-anomaly dark-sector branch belongs to D12: it imports the D6 static-patch scale only as an IR benchmark and does not derive MOND/RAR-style response laws, a controlled galaxy-scale source/response law, or galaxy/cluster phenomenology. The two continuous closure values are $P_\star$ and N_{CRC}; descendants use their branch-appropriate readouts.

Cosmology continuations above D6 obey the same firewall. The flat-FLRW reading is allowed only as a conditional holonomy statement: spatial curvature is a visible scalar holonomy on a declared FLRW branch, and zero curvature is the zero-visible-holonomy branch when no independent curvature charge is included in the preserved boundary datum. Inflation replacement, CMB scalar or parity kernels, H_0/S_8 prediction, dark/anomaly growth kernels, and baryogenesis require additional continuation theorems or likelihood contracts. They are not promoted to recovered-core SM/GR outputs by the D6 capacity relation.

3 Five Axioms, the Phase-II Pixel Fixed Point, the Capacity Target, and Theorem Checklist

This section states the theorem checklist and dependency map for this SM/GR derivation paper. The basis used below consists of the five core axioms, the quantitative quantities listed next, the support-visible BW scaling theorem, and the theorem-produced bosonic gauge-reconstruction chain. This is an algebraic-information basis: the screen-net, state, trace/probability, and generalized-entropy structures are starting ingredients of the OPH framework. The purpose is to test whether this basis supports a consistent and comprehensive theory-of-everything reconstruction of the observed effective universe. On the fixed-cutoff local finite-constraint MaxEnt branch, the local-Gibbs form, quasi-local dynamics, Lieb–Robinson propagation control, and endpoint-Lipschitz interval control are internal consequences. Any Dobrushin/local-Gibbs/mixing language used below names only a sufficient fixed-cutoff collar-recoverability mechanism on that branch. It is not an infinite-volume uniqueness claim for the refinement-limit theory. What Axiom 3.3 fixes under refinement is the realized state-side MaxEnt branch itself. Theorem 6.2 fixes the operational/geometric split at fixed cutoff, and Theorem 6.8 fixes the support-visible scaling-limit cap automorphism, its 2π normalization, and the Lorentz-facing geometric cap action on that subnet. Corollary 6.22 then derives the null half-sided modular pair after the null blow-up step, and Lemma 6.23 promotes that pair to the explicit positive null-translation generator on its Stone domain with the affine half-line modular relation. Theorem 6.28 internalizes the fixed-cap generalized-entropy stationarity step for admissible fixed-cap MaxEnt variations on the realized cap-label-preserving MaxEnt family. The compact-gauge realized branch is supplied by Theorem 7.23. The bounded-interval transport/projective branch is a branch-local construction only where those interval formulas are invoked. Gluing-side path-independent transport is Theorem 5.8. On the ordinary/central branch the zero obstruction is $[z]_\Sigma = 0$; on the genuinely noncentral branch the zero obstruction is $q_\Sigma = 0$, and $q_\Sigma \neq 0$ is routed to the fixed-cutoff higher-gauge sector rather than to ordinary compact-group reconstruction.

The formulation uses one Phase-II pixel variable together with one external continuous input:

$$P \equiv a_{\mathrm{cell}}/\ell_P^2, \tag{1}$$

$$N_{\mathrm{scr}} \equiv \log \dim \mathcal{H}_{\mathrm{tot}}. \tag{2}$$

In the quantitative implementation used here, P is the local particle-physics scale variable fixed by the outer/inner closure program summarized in the synthesis paper *Observers Are All You Need* [1]; this compact paper uses it only as the quantity carried by the forward particle-physics map. That closure writes the outside detuning as

$$P = \varphi + \alpha_{\mathrm{in}}(P)\sqrt{\pi},$$

where the inner side is the electromagnetic observation scale emitted by the same cell on the declared quantitative branch. The fixed-point readout is

$$\alpha^{-1}(0) = 137.035999177(21), \qquad \alpha \simeq 0.00729735256433, \qquad P \simeq 1.6309682094.$$

The source computation runs in the order

$$P \mapsto M_U(P) \mapsto \alpha_U(P) \mapsto \alpha_i(m_Z; P) \mapsto a_0(P) \mapsto A_T(P),$$

where $A_T(P) = \alpha_{\mathrm{em}}^{-1}(0; P)$ is the Ward-projected Thomson endpoint. The realized cell solves $P = \varphi + \sqrt{\pi}/A_T(P)$. The fine-structure value is forced on this branch because the same local pixel must satisfy the outer entropy-detuning equation and the inner electromagnetic endpoint equation with no remaining free local scale. A separate hardware note reports an optical-cavity check of the same fixed-point geometry; this is treated as corroborating engineering evidence. The technical endpoint table in Section 9 records the source-side audit trunk and residual package. N_{scr} is inferred from the observed cosmological horizon and enters the cosmological-capacity branch. The capacity normalization uses the Gibbons–Hawking de Sitter entropy [13]: if

$$N_{\mathrm{patch}} = \left(\frac{r_{\mathrm{dS}}}{\ell_P}\right)^2$$

denotes the bare horizon area ratio, then

$$N_{\mathrm{scr}} = S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4\ell_P^2} = \pi N_{\mathrm{patch}} = \frac{3\pi}{\Lambda \ell_P^2}.$$

Using the Planck-2018 cosmological benchmark [14] gives $r_{\mathrm{dS}} \simeq 1.66 \times 10^{26}$ m, $N_{\mathrm{patch}} \simeq 1.05 \times 10^{122}$, $N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}$, and $\Lambda \ell_P^2 \simeq 2.85 \times 10^{-122}$. Any retained capacity-level neutrino side estimate on that input is legacy bookkeeping rather than the weighted-cycle theorem lane.

Axiom 3.1 (Screen Net)**.** *Physical data is organized on a horizon screen S^2 carrying a net of local algebras*

$$P \longmapsto \mathcal{A}(P)$$

for connected patches $P \subset S^2$, with isotony

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

Axiom 3.2 (Overlap Consistency)**.** *For overlapping patches $P_1 \cap P_2 \neq \varnothing$, the local states induced on the shared algebra agree:*

$$\omega_{P_1}|_{\mathcal{A}(P_1\cap P_2)} = \omega_{P_2}|_{\mathcal{A}(P_1\cap P_2)}.$$

Axiom 3.3 (Local MaxEnt and Refinement Stability)**.** *At the regulator scale $\ell_{\rm UV}$, the realized branch is selected by maximizing entropy subject to a finite family of gauge-invariant local constraints*

$$\mathcal{C}_{\ell_{\rm UV}} = \{O_a(x)\}_{a=1}^{N_{\rm con}},$$

where each $O_a(x)$ is supported in a ball of radius $O(\ell_{\rm UV})$ and the label set a is independent of the number of regulator cells. Under refinement, the same finite constraint family is preserved, so the realized states at different cutoffs belong to one common finite-dimensional MaxEnt family. The realized low-energy branch is the refinement-stable branch of that family; symmetry-allowed relevant operators are therefore held at zero on that branch only by symmetry or because they lie in the explicitly retained constraint family. This is a statement about branch persistence, not a universal entropy-ordering theorem for arbitrary phases away from that branch.

Derived local-Gibbs branch. On the finite regulator realization of Axiom 3.1, the standard finite-dimensional Lagrange-multiplier argument applied to Axiom 3.3 gives

$$\omega_{\ell_{\rm UV}}(\lambda) = Z_{\ell_{\rm UV}}(\lambda)^{-1}\exp\big(-K_{\ell_{\rm UV}}(\lambda)\big),$$

with

$$K_{\ell_{\rm UV}}(\lambda) = \sum_x \sum_{a=1}^{N_{\rm con}} \lambda_a O_a(x) + \sum_b \mu_b Q_b,$$

where the Q_b are the finitely many optional global conserved-charge constraints. Thus the logarithm of the selected state is itself a quasi-local UV generator built from the same bounded-support densities that define the branch.

Derived quasi-local propagation and interval control. Fix the regulator graph metric coming from cell adjacency and let $\tau_t^{\ell_{\rm UV}}$ denote the automorphism group generated by $K_{\ell_{\rm UV}}(\lambda)$, or more generally by any branch generator lying in the norm-closed algebra generated by the same bounded-support local densities. Standard Lieb–Robinson estimates for finite-range interactions on the finite regulator net then give constants $C, \xi, v_{\ell_{\rm UV}} < \infty$ such that for local observables A_X, B_Y, [5]

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

So finite-velocity quasi-local propagation is branch-internal: it is the local finite-constraint MaxEnt branch written in propagation form, not an extra selector.

The same locality statement yields the bounded-interval endpoint-Lipschitz control used by the compact-gauge branch in the null modular chain. Whenever later sections use a fixed-cutoff interval or collar generator on this branch, it is the corresponding restriction of the same local density list, with the central endpoint term separated off exactly as in the later edge-center bookkeeping. Changing an interval I to I' then alters only an $O(|I'\Delta I|)$ collar of local terms. Hence on any fixed local-energy-bounded domain one has

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

which is precisely the endpoint-Lipschitz / finite-variation control used by the compact-gauge branch in the null modular chain. This paper does *not* separately derive a second independent microscopic generator beyond this quasi-local branch generator; if a UV completion carries one, the only later requirement is that it lie in the same bounded-support algebraic closure so that the same propagation constants control the selected branch.

This local-Gibbs/Lieb–Robinson package, together with any Dobrushin-type estimate used by the compact-gauge branch, should be read only as fixed-cutoff collar-local recoverability, support control, and carried-error bookkeeping on the selected branch. One may impose such estimates on a chosen finite collar model without deciding the refinement-limit gauge phase. The package is therefore not a proof that the refinement limit lands in a trivial thermodynamic phase and not a proof that the realized branch is nontrivial. Refinement persistence and bosonic fiber descent for zero-obstruction sectors are supplied instead by Theorem 7.2. Its theorem-level role is precisely the collar-local mixing and support/error control carried into the later fixed-cutoff and scaling-limit arguments.

Internal refinement notion. Choose any family of refinement channels $\Phi_{\ell\to L}$ compatible with Axiom 3.3. Because the same finite local constraint family is preserved, the regulator-scale realized states are parameterized by one common finite-dimensional multiplier space λ, and 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 that multiplier space. The refinement-stable realized branch is therefore the persistent trajectory or invariant subset selected inside this finite-dimensional space, not an imported regularity package. This state-side refinement notion is enough to compare realized states across cutoffs. Any Dobrushin/local-mixing hypothesis used by the compact-gauge branch is logically separate: it supplies a collar estimate on a fixed finite-dimensional model, not information about the refinement-limit gauge phase. Whenever later sections speak of a "refinement-stable directed colimit" of zero-obstruction edge sectors, the monoidal-refinement/fiber clause is Theorem 7.2; Axiom 3.3 supplies the realized state branch along which the theorem-produced sector witness persists.

Axiom 3.4 (Recoverable Generalized Entropy)**.** *A generalized entropy functional exists on caps,*

$$S_{\text{gen}}(C) = S_{\text{bulk}}(C) + \langle L_C\rangle,$$

where L_C is a positive edge-center entropy functional. In the semiclassical scaling branch, its leading coarse-grained contribution is identified with $A(\partial C)/(4G)$. The functional obeys quantum focusing on null generators, and collar tripartitions have small CMI with controlled recovery maps [16, 17, 18, 19].

Axiom 3.5 (Minimal Admissible Realization)**.** *Among realized 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 loop-coherent, anomaly-free, refinement-stable with light chiral matter, single-Higgs Yukawa-completable with one connected abelian charge factor acting nontrivially on the coupled carrier, intrinsically quark-sector CP-capable, and weak-sector UV-completable on that same one-Higgs branch, the realized package is the lexicographically minimal one under*

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

$\mathfrak{S}$ is the sector package on which MAR acts; it is not the bare tensor category alone. MAR is therefore an explicit structural-economy axiom on admissible realized low-energy branches, not a

theorem derived from the preceding axioms. χ_{cpl} is the coupled carrier dimension: the smallest unitary carrier containing a common irreducible block on which the admissible pseudoreal and complex nonabelian charge types both act nontrivially. This is stronger than the abstract minimal faithful representation dimension.

Definition 3.6 (MAR realization space and order)**.** *Let $\mathfrak{A}_{\mathrm{MAR}}$ be the set of isomorphism classes of finite low-energy sector packages*

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

on the ordinary or central zero-obstruction bosonic branch, together with the explicit one-Higgs chiral matter package used below. A package is MAR-admissible exactly when it satisfies the six predicates in Axiom 3.5. Two packages are physically equivalent when a compact-group isomorphism and a fiber-compatible symmetric monoidal equivalence preserve the observer-visible representations, Yukawa invariants, anomaly polynomial, normalized hypercharge lattice, and one-Higgs branch, modulo generation relabeling, charge-conjugation convention, gauge-center quotienting, implementation hiding, and inert ancillary stabilization. MAR orders packages lexicographically by $C(\mathfrak{S}) \in \mathbb{N}^4$ and then quotients ties by this physical equivalence.

Proposition 3.7 (Well-founded MAR minima)**.** *Every nonempty MAR-admissible class has at least one MAR-minimal package. The minimal packages are exactly those whose complexity vector is the lexicographically least element of $C(\mathfrak{A}) \subseteq \mathbb{N}^4$ for the chosen nonempty admissible class $\mathfrak{A}$.*

Proof. Lexicographic order on $\mathbb{N}^4$ is well-founded: minimize the first coordinate, then the second on that fiber, then the third, then the fourth. Each step minimizes a nonempty subset of $\mathbb{N}$. □

Remark 3.8 (Meaning of MAR uniqueness)**.** *MAR uniqueness means uniqueness of the observer-visible low-energy package modulo Definition 3.6. It is not uniqueness of a microscopic regulator representative. The later D8–D9 lemmas prove that, within the connected positive-dimensional Lie admissible class with one connected abelian factor and faithful action on the minimal coupled carrier, all MAR-minimal representatives realize the same connected SM gauge image, normalized hypercharge lattice, structural electroweak force content, $N_c = 3$, and $N_g = 3$.*

Remark 3.9 (Recovered-core theorem sources)**.** *The scaling/BW step is Theorem 6.8, the support-visible BW scaling theorem. The fixed-cutoff realized presentation, the local-Gibbs form, quasi-local Lieb–Robinson propagation, endpoint-Lipschitz interval control, the induced finite-dimensional refinement branch, the fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations on the realized cap-label-preserving MaxEnt family, and the operational/geometric split of Theorem 6.2 are established internally from Axioms 3.1–3.3. Transportability, the fixed-cutoff bosonic category, the refinement/fiber ladder, and the realized compact-gauge witness are supplied by Theorems 5.8, 5.11, 7.2, and 7.23. The genuinely noncentral fixed-cutoff branch is controlled by Theorems 5.3–5.6: $q_\Sigma = 0$ strictifies to the ordinary compact-group transportable case, while $q_\Sigma \neq 0$ remains a higher-gauge sector and is not itself the ordinary compact-group reconstruction theorem. Fermionic signs and chirality belong to the later fermionic/super-Tannakian matter lift. The classification/selection split is packaged in Theorem 7.5, and the observer-visible selected endpoint is Theorem 7.24.*

Remark 3.10 (Log convention)**.** *Unless an explicit base is written, logarithms and entropies in this paper use natural logs (nats). Base-2 logarithms are written as $\log_2$ and quoted in bits.*

Dependency DAG

The table below is the dependency map for this SM/GR derivation paper. The "Immediate OPH ingredients" column records only parents internal to the OPH program. Imported standard mathematics, branch-local constructions, and external inputs are stated separately rather than hidden inside the word "derived."

Node	Theorem labels and output	Immediate OPH ingredients	Standard mathematics used	Branch-local inputs / external data	Status
D1	Theorems 4.4 and 4.7: unique quotient normal form and schedule independence from a fixed initial quotient state, plus same-boundary uniqueness under unique consistent extension	overlap-consistency problem on a finite patch net together with the declared fixed-cutoff collar recovery step, its touched-overlap local-fit acceptance contract, support-local disjoint commutation, and restriction-compatible union-collar gluing on the physical quotient	local-to-global Lyapunov descent on the inconsistency potential; local-diamond completion from overlap-associative center-sector / higher-gauge union-collar gluing; Newman's lemma	repair completeness; stated Petz support/CPTP control where that branch is used; for same-boundary uniqueness, boundary/sector preservation and at most one consistent quotient extension in the boundary fiber	structural theorem
D2	Theorem 5.2, Propositions 5.18, 5.19, 5.25, Definition 5.26: collar Markov/entropy layer	Axioms 3.1–3.4	HJPW Markov structure theorem; Petz and Fawzi–Renner recovery theory	exact Markov identities require exact Markovity or an explicitly stated idealized recoverability limit; the $A/(4G)$ identification is the coarse-grained dictionary of Definition 5.26	structural lemma/ proposition layer
D3	Theorem 6.8, Corollary 6.16: support-visible BW scaling theorem and Lorentz branch	Axioms 3.1–3.4, Theorem 6.2, and the fixed-cutoff collar/consensus package	Bisognano–Wichmann modular geometry; conformal classification on S^2; weak-$*$/GNS extraction	support-visible regularized modular transport on the prime geometric cap net; support-readable modular covariance; ordered cut-pair rigidity; the scaling-limit observer algebra may leave the finite type-I regulator class, so the theorem is automorphism-level there and $K_C = 2\pi B_C$ is only the special type-I limit form	scaling-limit theorem

Node	Theorem labels and output	Immediate OPH ingredients	Standard mathematics used	Branch-local inputs / external data	Status
D4	Propositions 6.17, 6.20; Corollaries 6.18, 6.22; Theorem 6.21, Lemma 6.23; Theorem 6.24: null modular bridge to T_{kk}, including exact-or-controlled strip additivity, endpoint-Lipschitz renormalized half-line families, the weak tail generator, the derived half-sided modular inclusion, the explicit positive half-line null-translation generator on its Stone domain with affine half-line modular relations, and the exact half-line generator/charge identification with the local null-stress charge	D2+D3 and Axioms 3.1–3.4	Borchers–Wiesbrock positive-generator theorem for standard half-sided inclusions; Stone's theorem; distributional differentiation on half-lines	the theorem-local null-cut center transfer, the inherited left/right strip-split package used for the spatial-collar-type tensor decomposition, and the exact-or-controlled Markov hypotheses of Proposition 6.17 and Corollary 6.18; quasi-local propagation and endpoint-Lipschitz control are supplied internally by Axiom 3.3 on the local finite-constraint branch, the geometric scaling action on the null half-line blow-up net derives the half-sided modular pair, and bounded-interval formulas use the separate interval-preserving projective branch	bridge theorem
D5	Theorem 6.28, Lemma 6.30, Theorem 6.32, Corollary 6.33: Jacobson-type Einstein branch via a rest-frame relation plus tensor upgrade	D3+D4 and Axioms 3.1–3.4	fixed-volume area-variation identity for small geodesic balls; local quadratic-polarization argument for the tensor upgrade	Theorem 6.28, which fixes the admissible fixed-cap MaxEnt variation class and derives generalized-entropy stationarity on the realized cap-label-preserving MaxEnt family; the internal small-ball bridge of Lemma 6.30, which uses the geometric cap generator together with the D4 half-line generator/charge identification and the separate interval-preserving projective branch; locally Lorentzian $d = 4$ scaling regime; small-ball constancy assumptions; D4 carried remainder negligible at order ℓ^4; all local directions/reference states for the tensor upgrade	scaling-limit theorem/corollary

Node	Theorem labels and output	Immediate OPH ingredients	Standard mathematics used	Branch-local inputs / external data	Status
D6	Theorem 6.45: local/global theorem stack for Λ, namely the null-invisible metric ambiguity, the screen-capacity closure of the same Einstein branch, the derived de Sitter static-patch parameter package, and the cosmic record-closure readback fixed point, plus a legacy neutrino side estimate	D5 local Einstein recovery, which leaves the $+\Lambda g_{ab}$ ambiguity	de Sitter entropy relation, static-patch radius/time formulas, finite self-closing normal-form counting/readback, and dimensional analysis	cosmic record-capacity fixed point $N_{\rm CRC} = F(N_{\rm CRC})$ and its observed de Sitter entropy readout	zero-input global closure
D7	Theorem 7.2, Theorem 7.3, Theorem 7.4: construction of the refinement-stable bosonic sector category and compact gauge reconstruction; Theorem A.17: four-dimensional Euclidean Yang–Mills form; Theorem A.28: support-visible compact-gauge Yang–Mills repair gap	Axioms 3.1–3.4	directed-colimit descent for monoidal C^*-categories; Doplicher–Roberts / Tannaka reconstruction; holonomy-to-curvature expansion; finite-dimensional conditional expectations and spectral comparison for commuting collar projections	Theorem 5.8, Theorem 5.11, and Theorem 7.2 on the ordinary or central-defect bosonic zero-obstruction branch; Theorem 7.23 supplies nonempty realized MAR-admissible witness data, while the Standard Model selection step belongs to D8–D9; the Euclidean-form and repair-gap theorems additionally use the four-dimensional scaling chart, reflection-positive ordinary vacuum, active exact-Markov repair collars, bounded-color collar covers, repair completeness, and support-visible continuum extraction; the genuinely noncentral fixed-cutoff branch is handled separately by Theorems 5.3–5.6 and is not itself the ordinary compact-group reconstruction theorem	structural theorem
D8	Lemmas 7.7–7.12, Theorem 7.13: product gauge structure up to finite quotient	D7 + Axiom 3.5	compact Lie representation classification; Schur's lemma	the same ordinary or central-defect bosonic branch as D7, together with a connected positive-dimensional Lie admissible class, one connected abelian factor, and faithful action on the minimal coupled carrier	realized-branch theorem
D9	Theorem 7.24, Theorem 7.14, Corollaries 7.15, 7.16, 7.22, Proposition 7.19: realized Standard Model quotient, hypercharges, structural electroweak force content, $N_c = 3$, $N_g = 3$, and product-group corollaries	D8	anomaly-cancellation algebra; Witten global-anomaly argument; CKM CP counting; stabilizer computation for the neutral Higgs direction	realized one-generation chiral matter plus one Higgs package, with the CP-capability and weak-sector UV clauses contained in Axiom 3.5; no extra post-MAR selector is used	realized-branch theorem/corollary chain

Node	Theorem labels and output	Immediate OPH ingredients	Standard mathematics used	Branch-local inputs / external data	Status
D10	forward gauge-coupling closure and declared electroweak readout of Section 7	D9 + pixel ratio P	printed RG evolution, matching, and scheme-conversion conventions	pixel constraint and the source-only forward transmutation solve $\mathcal{F}(\alpha_U; P) = 0$; the fixed-cutoff edge heat-kernel / Casimir theorem on the microphysics surface together with the compact-group / Peter–Weyl lift used on the D10 lane; printed beta-function, threshold, and scheme-conversion conventions	quantitative-closure sector
D12	Sections 8–10 continuations: charged leptons, strong CP, proton spin/lifetime, and controlled string/worldsheet effective-description branches	various subsets of D6, D9, and D10	branch-specific EFT and phenomenological manipulations	additional ansätze such as the uniform $\mathbb{Z}_6$ center-label ensemble, texture choices, dark-sector response assumptions, discrete-horizon assumptions, a distinct large-N_{edge} regime with fixed $\tau = tN_{\text{edge}}$ window and uniform genus-remainder control, or additional worldsheet/CFT assumptions	phenomenological continuation

The dependency DAG records the immediate internal parents, imported mathematics, branch-local constructions, and external data for each node.

Theorem 3.11 (Summary theorem for the recovered relativity-plus-Standard-Model core)**.** *This theorem packages the recovered-core nodes D1–D5 and D7–D9 of the dependency DAG above. Node D6 sits outside this recovered-core theorem because the cosmic record-capacity fixed point is the separate global closure of the Einstein branch recovered in item (iii); D10 and D12 are outside this theorem's scope. Assume Axioms 3.1–3.4, Axiom 3.5, and the hypotheses of Theorems 4.4, 4.7, 6.8, 6.24, 6.32, 7.2, 7.3, 7.4, 7.13, and 7.14, together with Corollaries 7.16, 7.15, and Proposition 7.19. Then, on the support-visible scaling branch of Theorem 6.8:*

(i) overlap repair admits a unique schedule-independent normal form from each fixed initial physical quotient state, and from fixed boundary/sector data when the consistent quotient extension in that fiber is unique;

(ii) cap modular flow on the extracted prime geometric cap pair is geometric and yields the connected Lorentz group

$$\mathrm{Conf}^+(S^2) \cong \mathrm{SO}^+(3,1);$$

(iii) the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations on the realized cap-label-preserving MaxEnt family together with the null modular bridge and the bounded-interval projective branch yield the Jacobson-type rest-frame relation

$$\delta(G_{00} + \Lambda g_{00}) = 8\pi G\, \delta\langle T_{00}\rangle$$

at the cap center in the diamond rest frame; if that rest-frame relation holds for all local directions and reference states in the scaling regime, then the null-to-tensor upgrade gives

$$G_{ab} + \Lambda g_{ab} = 8\pi G \langle T_{ab} \rangle;$$

(iv) on the theorem-produced ordinary or central-defect bosonic zero-obstruction branch, the edge-sector category reconstructs a compact gauge group, and the realized connected gauge structure has the form

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\Gamma}$$

for some finite central subgroup Γ*.*

After the hypercharge lattice, color-count, generation-count, and trivial-action quotient steps, the finite quotient is fixed to $\Gamma = \mathbb{Z}_6$*, so the realized gauge structure on the realized MAR-admissible branch is*

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_c = 3, \qquad N_g = 3,$$

with the exact Standard Model hypercharge lattice on the realized matter package. The same package also fixes the structural electroweak force content: the weak $\mathrm{SU}(2)_L$ *factor, the hypercharge* $\mathrm{U}(1)_Y$ *factor, the neutral-Higgs stabilizer* $\mathrm{U}(1)_Q$*, and the* $W^{\pm}/Z/A$ *gauge-boson basis. Here the count statements are internal to the same realized branch: the D8 minimal coupled-carrier theorem stack fixes the color triplet and hence* $N_c = 3$*; CKM phase counting and weak-sector asymptotic freedom then bound* N_g *on that same one-Higgs branch; MAR fixes the smallest realized value* $N_g = 3$*; and Witten's* $\mathrm{SU}(2)$ *anomaly is retained only as a consistency check on the resulting triplet-doublet package.*

Proof. Items (i)–(iii) are collected from Theorem 4.4, Corollary 6.16, Theorem 6.32, and Corollary 6.33. On the gauge side, Theorem 5.8 supplies the strict zero-obstruction transport criterion, Theorem 5.11 constructs the fixed-cutoff bosonic collar-sector categories, and Theorem 7.2 constructs the monoidal refinement transport and compatible finite multiplicity fibers. Theorem 7.3 descends those data to the refinement-limit category and fiber functor, and compact gauge reconstruction is then Theorem 7.4. Theorem 7.5 separates this classification stage from the MAR selection stage. The realized nontrivial branch is supplied by Theorem 7.23. The product gauge structure up to finite quotient is Theorem 7.13; and the final identification of the exact Standard Model quotient and structural electroweak content is Theorem 7.24, Theorem 7.14 together with Corollaries 7.15, 7.16, 7.22, and Proposition 7.19. □

Theorem 3.11 therefore packages only the recovered relativity-plus-Standard-Model core: D1–D5 together with D7–D9. D6 is not a detached second gravity lane: it is the global closure of the same Einstein branch at the cosmic record-capacity fixed point. The quantitative-closure branch and phenomenological continuations are outside the theorem's scope.

Usually Postulated Elsewhere, Recovered or Continued Here with Explicit Status

Structure	Common treatment	OPH treatment
Lorentz kinematics	background symmetry or starting axiom	scaling-limit branch from geometric modular flow on screen caps
Einstein dynamics	fundamental field equation	first-variation relation with tensor upgrade from generalized entropy, null modular data, and derived fixed-cap generalized-entropy stationarity
Gauge group	model input	construction of the refinement-limit bosonic edge-sector category, then compact group reconstruction from it; exact SM quotient selected on the realized MAR-admissible branch with the realized one-generation/one-Higgs package
Hypercharge lattice	matter-assignment input	solved from anomaly cancellation and Yukawa invariance on the realized one-generation chiral matter plus one-Higgs package
Color and generation count	empirical input	color triplet fixed by the D8 minimal coupled carrier; generation count fixed by CKM phase counting, weak-sector asymptotic freedom, and MAR; Witten parity retained as a consistency check
Flavor hierarchy unit	model-dependent small parameter	phenomenological ansatz $\varepsilon = 1/6$ motivated by a uniform $\mathbb{Z}_6$ center-label ensemble
Charged-lepton continuation	Koide-style phenomenological fit	exact centered readback plus the closed common-shift no-go above $\widehat{C}_e^{\text{cand}}$; any $\delta = 2/9$ phase relation is continuation-only
Cosmological constant	local vacuum-energy puzzle	global screen-capacity term fixed by $N_{\text{CRC}} = F(N_{\text{CRC}})$, with $\Lambda_{\text{CRC}} = 3\pi/(GN_{\text{CRC}})$

4 Observer-Overlap Consistency as a Fixed-Point Problem

The discrete overlap problem is the mathematical skeleton of the framework.

Definition 4.1 (Patch Net)**.** *Let $G = (V, E)$ be a finite connected graph. Each vertex $i \in V$ carries a finite local state space S_i. For each edge $e = \{i, j\}$ let I_e be an interface alphabet and let*

$$\pi_{i,e} : S_i \to I_e, \qquad \pi_{j,e} : S_j \to I_e$$

be interface projections. The global state space is

$$\Sigma = \prod_{i \in V} S_i,$$

and the consistency set is

$$C = \{s \in \Sigma : \pi_{i,e}(s_i) = \pi_{j,e}(s_j) \text{ for all } e = \{i, j\} \in E\}.$$

This finiteness is a regulator-level assumption for the discrete patch-net normal-form theorem; it is not the continuum limit statement.

Define the inconsistency potential

$$\Phi(s) = \sum_{e=\{i,j\}\in E} w_e\, d_e\big(\pi_{i,e}(s_i), \pi_{j,e}(s_j)\big),$$

with $w_e > 0$ and $d_e(a, b) = 0$ iff $a = b$.

Let

$$q_{\text{phys}} : \Sigma \to Q$$

be the physical quotient map identifying hidden representatives with the same observer-facing overlap data, and write

$$C_Q := q_{\text{phys}}(C).$$

Accepted repairs are required to descend to quotient repairs $\overline{T}_i : Q \to Q$. Write $x \to y$ on Q when $y = \overline{T}_i(x) \neq x$ for some accepted repair, and let $\Phi_Q : Q \to \mathbb{R}_{\geq 0}$ be the induced finite mismatch functional. When no hidden representative quotient is present, take $Q = \Sigma$.

Remark 4.2 (Default code claim)**.** *At this level the overlap net is a finite constraint code and nothing stronger: the codewords are the states in C, equivalently the zero set $C = \Phi^{-1}(0)$. A graph min-cut, by itself, does not determine the distance of this code, since the same graph can carry trivial constant readouts with distance 1 or repetition constraints with distance $|V|$. QECC distance, topological-code resilience, Knill–Laflamme correction, exponential convergence, and BFT wall-clock liveness are therefore read only on the corresponding certified branches, not from the bare overlap graph.*

Definition 4.3 (Recovery-derived local repair law)**.** *A law λ is a family of local repair maps*

$$T_i^{\lambda} : \Sigma \to \Sigma$$

changing only patch i or a bounded neighborhood of i. On the fixed-cutoff collar branch these are not free rewrite primitives: each accepted local update is read from exact Markov splice or a declared Petz/Fawzi–Renner recovery move on a collar chart around i, then lifted back to the finite patch presentation, and committed only when it strictly lowers the touched-overlap contribution to Φ on the overlaps it can modify. Disjoint accepted repairs therefore commute directly by support locality, and when two accepted repairs touch a common overlap they are read inside one finite union collar whose physical glued state is parenthesization-independent on the quotient and whose local decoders are restriction-compatible on nested collars.

Theorem 4.4 (Quotient normal-form uniqueness)**.** *Assume:*

(i) *finite descent: $x \to y$ implies $\Phi_Q(y) < \Phi_Q(x)$;*

(ii) *quotient-local diamond: every one-step peak $y \leftarrow x \to z$ in Q has a common descendant w with $y \to^* w$ and $z \to^* w$;*

(iii) *repair completeness: terminal quotient states are exactly C_Q.*

Then every initial quotient state $x \in Q$ has a unique terminal normal form

$$N_Q(x) \in C_Q,$$

and $N_Q(x)$ is independent of asynchronous repair order.

Proof. Finite descent on finite Q gives termination. The quotient-local diamond gives local confluence. Newman's lemma [4] says that a terminating locally confluent rewrite system is confluent. If two terminal states y, z are reachable from the same x, confluence gives a common descendant w. Since y and z are terminal, $y = w = z$. Thus the terminal state reachable from x is unique. Repair completeness identifies that terminal state with an element of C_Q, so $N_Q : Q \to C_Q$ is well defined and schedule-independent. □

Remark 4.5 (Confluence is an extra obligation)**.** *The descent argument above proves termination only. OPH does not infer uniqueness from termination. Uniqueness and order-independence enter through the local diamond on the physical quotient, Newman's lemma, and repair completeness. If two accepted schedules from the same initial state terminate in different observer-facing quotient normal forms, with no declared holonomy/higher-gauge obstruction and no mere hidden-representative difference, then the proposed repair law fails the fixed-cutoff consensus criterion.*

Corollary 4.6 (Schedule Independence)**.** *If observables factor through Q and are evaluated on $N_Q(q_{\text{phys}}(s))$, the resulting physical law is independent of update order from the fixed initial quotient state $q_{\text{phys}}(s)$.*

Theorem 4.7 (Boundary-conditioned uniqueness)**.** *Let*

$$B : Q \to \mathcal{B}$$

record fixed external boundary data, conserved charge, root packet, holonomy sector, or task input. Assume B is preserved by accepted repairs:

$$x \to y \implies B(x) = B(y).$$

If for each $b \in \mathcal{B}$ the consistent quotient fiber

$$C_b := \{x \in C_Q : B(x) = b\}$$

has at most one element, then all initial quotient states with boundary value b settle to the same observer-facing normal form. Equivalently, $B(x) = B(x')$ implies

$$N_Q(x) = N_Q(x').$$

Proof. By Theorem 4.4, $N_Q(x)$ and $N_Q(x')$ exist and lie in C_Q. Boundary preservation along repair sequences gives $B(N_Q(x)) = B(x)$ and $B(N_Q(x')) = B(x')$. If $B(x) = B(x') = b$, then both normal forms lie in C_b. The unique consistent extension assumption says C_b has at most one element, hence the two normal forms are equal. □

Theorem 4.8 (Cycle Obstruction / Holonomy Criterion)**.** *Let A be an abelian group. For each oriented edge $e : u \to v$ assign $b_e \in A$ and consider the affine overlap equations*

$$x_v - x_u = b_e.$$

A global solution exists if and only if the signed sum of b_e vanishes on every cycle.

Proof. Summing the edge equations around a cycle gives necessity. For sufficiency, fix a root and define each x_v by summing labels along a path from the root to v. Vanishing cycle sums make this independent of the chosen path. □

Definition 4.9 (Physical repair law and representative lift)**.** *Suppose a local gauge group $\Gamma = \prod_i \Gamma_i$ acts on Σ while leaving all interface data invariant. Write*

$$q : \Sigma \to \Sigma/\Gamma, \qquad q(s) = [s].$$

A physical repair law is the family of local quotient maps induced by the recovery-derived collar updates just described:

$$\overline{T}_i^\lambda : \Sigma/\Gamma \to \Sigma/\Gamma$$

on the overlap-invariant quotient. A representative repair family is any family

$$T_i^\lambda : \Sigma \to \Sigma$$

with

$$q \circ T_i^\lambda = \overline{T}_i^\lambda \circ q.$$

Proposition 4.10 (Representative lifts descend to the quotient)**.** *For any representative repair family of Definition 4.9,*

$$q(T_i^\lambda(\gamma \cdot s)) = q(T_i^\lambda(s)) \qquad \forall\, i,\ \forall\, \gamma \in \Gamma,\ \forall\, s \in \Sigma.$$

In particular the repair relation descends to Σ/Γ without any extra gauge-covariance axiom.

Proof. Because $q(\gamma \cdot s) = q(s)$,

$$q(T_i^\lambda(\gamma \cdot s)) = \overline{T}_i^\lambda(q(\gamma \cdot s)) = \overline{T}_i^\lambda(q(s)) = q(T_i^\lambda(s)).$$

So gauge-equivalent inputs induce the same repaired physical state. □

Proposition 4.11 (Gauge Quotient)**.** *Under Definition 4.9, Proposition 4.10, and Theorem 4.4 with $Q = \Sigma/\Gamma$, the normal-form map is a quotient map:*

$$\overline{N}_\lambda : \Sigma/\Gamma \to q(C), \qquad [s] \mapsto N_Q([s]).$$

Hence gauge-invariant observables are unique on the quotient, and the fixed-cutoff physical observable algebra has representative-independent terminal expectations on that carrier.

Proof. By Proposition 4.10, every repair step has quotient image determined only by the current orbit. Iterating, the quotient image of any repair sequence depends only on the initial orbit. Theorem 4.4 gives a unique terminal quotient normal form $N_Q([s])$ from that orbit. This makes $\overline{N}_\lambda$ well-defined. □

On the fixed-cutoff quantum lift, the same quotient-local carrier does more than fix the terminal orbit: for every declared fixed-cutoff physical algebra, the terminal expectation functional is the same on any two representative lifts whose regional collar data lie in one quotient-local glued state, even when the microscopic representatives differ by gauge or sector relabelings on that carrier.

Theorem 4.12 (Refinement-limit consensus classes)**.** *Let R be a directed cofinal refinement set. For each $r \in R$, let $Q_r = \Sigma_r/\Gamma_r$ be the finite physical quotient state space, let $n_r : Q_r \to Q_r$ be the finite quotient normal-form map, and let $h_r : Q_r \to \mathcal{H}_r$ be the finite holonomy or higher-gauge obstruction map. Suppose that for $r \preceq s$ there are restriction maps*

$$\rho_{sr} : Q_s \to Q_r, \qquad \chi_{sr} : \mathcal{H}_s \to \mathcal{H}_r$$

forming directed inverse systems and satisfying

$$\rho_{sr} n_s = n_r \rho_{sr}, \qquad \chi_{sr} h_s = h_r \rho_{sr}.$$

Assume visible separation: compatible families in the inverse limits are equal whenever they agree on a cofinal subset of finite stages. Then

$$n_\infty((x_r)_r) = (n_r(x_r))_r, \qquad h_\infty((x_r)_r) = (h_r(x_r))_r$$

define a unique schedule-independent refinement-limit normal-form class and a refinement-limit holonomy class. The finite normal forms and holonomies converge to those classes in the inverse-limit topology. A nonzero limiting holonomy has a finite-stage witness, and agreement of the normal-form and holonomy projections on a cofinal tail gives the same refinement-limit consensus class.

Proof. Compatibility of the x_r gives $\rho_{sr}(x_s) = x_r$. Normal-form naturality then gives

$$\rho_{sr}(n_s(x_s)) = n_r(\rho_{sr}(x_s)) = n_r(x_r),$$

so the normal forms form a compatible inverse-limit family. Holonomy naturality gives the same calculation for $h_s(x_s)$. Finite-stage schedule independence comes from Theorem 4.4; visible separation upgrades agreement of all cofinal finite projections to uniqueness of the inverse-limit class. A nonzero inverse-limit holonomy must have a nonzero projection at some finite stage by visible separation. □

Proposition 4.13 (Coarse-graining / reconciliation compatibility)**.** *Let $r \preceq s$ be two refinement stages in the D1 consensus system, with coarse-graining maps*

$$\rho_{sr} : Q_s \to Q_r, \qquad \chi_{sr} : \mathcal{H}_s \to \mathcal{H}_r.$$

Equip Q_r and $\mathcal{H}_r$ with the pseudometrics used for macroscopic readout at stage r. If the chosen coarse-graining channel has normal-form and obstruction defects ε^n_{sr} and ε^h_{sr}, meaning

$$d^Q_r(\rho_{sr} n_s(x), n_r \rho_{sr}(x)) \le \varepsilon^n_{sr}, \qquad d^{\mathcal{H}}_r(\chi_{sr} h_s(x), h_r \rho_{sr}(x)) \le \varepsilon^h_{sr}$$

for every $x \in Q_s$, then reconciling at stage s and coarse-graining to r gives the same macroscopic normal-form and obstruction readout as coarse-graining first and reconciling at r, up to

$$\max\{\varepsilon^n_{sr}, \varepsilon^h_{sr}\}.$$

In the exact separated cofinal system of Theorem 4.12, these defects are zero. If the defects vanish on cofinal tails, the two procedures define the same macroscopic inverse-limit consensus class.

Proof. The displayed inequalities are exactly the two components of the claimed product-readout bound. Exact naturality in Theorem 4.12 is the zero-defect case. Cofinal vanishing gives convergence of every fixed coarse cylinder value in the inverse-limit topology. □

This removes the extra gauge-covariance axiom and keeps repair on its actual quotient-local carrier: recovery dynamics on fixed-cutoff collars. On the declared fixed-cutoff branch, the theorem-local branch condition is repair completeness, together with the Petz support/CPTP clause where that branch is used. The touched-overlap local-fit contract supplies Lyapunov Φ-descent on accepted moves, and the support-local commutation and union-collar compatibility package belongs to the declared repair law itself. Stability of that package under refinement or branch change is handled by the compatibility clauses of Theorem 4.12 when a separated cofinal refinement system is supplied. Proposition 4.13 is the corresponding RG-facing statement: it does not say that arbitrary coarse-graining maps commute with repair, only that the selected coarse-graining branch commutes with reconciliation to the extent that the displayed defects are controlled.

Definition 4.14 (Ancilla stabilization)**.** *Fix a finite-cutoff OPH realization*

$$\mathfrak{U} = (\{\mathcal{H}_P, \mathcal{A}(P), \omega_P\}_{P\in\mathcal{P}}, \Gamma, T^\lambda).$$

Choose finite-dimensional ancillary factors K_P *with product state* $\eta = \bigotimes_{P\in\mathcal{P}} \eta_P$*. The associated ancilla stabilization is the realization*

$$\mathfrak{U}^\eta = (\{\mathcal{H}_P^\eta, \mathcal{A}^\eta(P), \omega_P^\eta\}_{P\in\mathcal{P}}, \Gamma, (T^\lambda)^\eta),$$

with

$$\mathcal{H}_P^\eta := \mathcal{H}_P \otimes K_P, \qquad \mathcal{A}^\eta(P) := \mathcal{A}(P) \otimes \mathbf{1}_{K_P}, \qquad \omega_P^\eta := \omega_P \otimes \eta_P,$$

and lifted repair dynamics acting trivially on the ancillas.

Proposition 4.15 (Ancilla-stable UV underdetermination)**.** *Let* $\mathfrak{U}^\eta$ *be an ancilla stabilization of a finite-cutoff OPH realization* $\mathfrak{U}$*. Then:*

1. *observable expectations on the physical subalgebras are unchanged:*

$$\omega_P^\eta(a \otimes \mathbf{1}_{K_P}) = \omega_P(a) \qquad \forall\, a \in \mathcal{A}(P);$$

2. *the interacting local MaxEnt branch on the physical subalgebra is unchanged, since*

$$\omega_{\ell_{\mathrm{UV}}}^\eta(\lambda) = \omega_{\ell_{\mathrm{UV}}}(\lambda) \otimes \eta;$$

3. *for every collar split* $A : B : D$*,*

$$I(A : D \mid B)_{\rho\otimes\eta} = I(A : D \mid B)_\rho,$$

   so the Fawzi–Renner remainder $r_{\mathrm{FR}}(\varepsilon)$ *and the collar Markov modulus* $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$ *are unchanged;*

4. *if the ancillas are inert under repair, then they remain hidden representative data and the quotient normal form on physical observables is identical:*

$$N_{Q^\eta}\Big(q^\eta_{\mathrm{phys}}(s \otimes k)\Big) = N_Q(q_{\mathrm{phys}}(s))\,;$$

5. *if the ancillas carry only trivial neutral sectors, the MAR-selected realized sector package is unchanged.*

Hence $\mathfrak{U}$ *and* $\mathfrak{U}^\eta$ *are OPH-indistinguishable although they are different microscopic regulator realizations.*

Proof. Item 1 is immediate from the definition of $\mathcal{A}^\eta(P)$ and ω_P^η. Item 2 is the product-state form of the same fixed-cutoff MaxEnt branch. Item 3 follows from additivity of entropy for product ancillas, which cancels in the conditional mutual-information combination and therefore leaves both r_{FR} and δ^{M} unchanged. Item 4 holds because the lifted repair maps do not act on the ancillary factors. Item 5 is immediate when the ancillas are neutral and carry no extra realized sector data. □

Definition 4.16 (OPH-stable UV equivalence)**.** *Two finite-cutoff realizations* $\mathfrak{U}$ *and* $\mathfrak{U}'$ *are* OPH-stably equivalent*, written*

$$\mathfrak{U} \sim_{\mathrm{OPH}} \mathfrak{U}',$$

if after finite ancilla stabilizations they are related by local gauge-covariant **-isomorphisms intertwining the observable patch net, overlap maps, local states, and repair dynamics on the physical subalgebras.*

Corollary 4.17 (Unique physical UV branch only modulo OPH-stable equivalence)**.** *Under the OPH axiom language, the UV invariant determined by the theory is the class* $[\mathfrak{U}]_{\mathrm{OPH}}$*, not a unique microscopic regulator presentation. Literal microscopic UV uniqueness is therefore not an OPH invariant.*

Proof. Propositions 4.10 and 4.11 fix the schedule-independent physical branch on the quotient, while Proposition 4.15 shows that inert ancillary refinements leave every OPH observable invariant. So the physical branch is fixed only modulo $\sim_{\mathrm{OPH}}$, not at the level of one microscopic representative. ∎

These three results encode much of the eventual physical interpretation. Objectivity becomes confluence, gauge symmetry becomes quotient invariance induced by the physical overlap algebra, and stable defects are represented by nontrivial overlap holonomy classes. The physically relevant uniqueness statement is therefore quotient uniqueness together with ancilla-stable equivalence: OPH fixes a unique gauge-invariant physical branch modulo boundary redundancy, implementation hiding, and inert ancillary stabilization, not a unique microscopic regulator presentation.

5 Collars, Edge Centers, and Generalized Entropy

The fixed-point picture explains why overlap consistency produces gauge quotients: once repair is read on the physical overlap algebra, descent to the quotient is automatic. The gravity and consensus arguments later use a more specific collar fact: after edge-center completion, interior observables become insensitive to compatible exterior substitutions, and modular additivity becomes exact in the Markov normal form. The point of this section is therefore twofold. First, at fixed regulator scale, we derive the collar block decomposition from overlap consistency itself on the ordinary or central-defect branch and then state its genuinely noncentral higher-gauge replacement. Second, we state precisely when the approximate recoverability clause of Axiom 3.4 is allowed to converge to the exact HJPW normal form used by the later spatial and null collar theorems.

The sharp claim boundary is as follows. Small conditional mutual information always gives a constructive recovered comparison state with $O(\varepsilon^{1/2})$ observable error. It does *not* by itself give a universal one-shot trace-norm bound to an exact Markov state. The exact Markov normal form is recovered either when $I(A:D\mid B)=0$ holds literally, or for a controlled family on one fixed finite-dimensional collar model, or after pullback to such a model, where $\varepsilon\to 0$. On that fixed collar one gets a collar-dependent modulus $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon)\to 0$ measuring distance to the exact Markov set, and this is the error that must be carried into later exact splice or modular-additivity identities.

Proposition 5.1 (Derived regulator gluing datum)**.** *Fix a regulator-scale finite patch cover of a cap neighborhood and let R be any finite union of regulator cells. Assume only the finite patch-net presentation implicit in Definition 4.1 together with overlap consistency on common interfaces. Then:*

1. *each regulator cell i with finite local state space S_i can be Hilbertized as $\tilde{\mathcal{H}}_i\cong\mathbb{C}^{|S_i|}$, so the extended algebra before overlap quotienting is the finite type-I algebra*

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

2. *for each connected boundary component $\Sigma\subset\partial R$, after choosing a reference cut presentation $\tilde{\mathcal{H}}_\Sigma$, every overlap-consistent change of local chart along Σ is implemented by a unitary on*

$\tilde{\mathcal{H}}_\Sigma$*, and the compact closure of the subgroup generated by all such recharting unitaries is a compact boundary redundancy group*

$$K_\Sigma \subset U(\tilde{\mathcal{H}}_\Sigma);$$

before the choice of reference chart, the same data form a compact unitary groupoid of overlap-preserving transitions;

3. *if triple-overlap defects are central, the projective composition law of item 2 lifts to a direct action of a compact central extension* $\widehat{K}_\Sigma$*; a genuinely noncentral defect is the only obstruction to reducing the transition system to an ordinary compact group action;*

4. *in the invariant-state realization of the quotient, the chart-independent endomorphisms of the lifted presentation form the boundary-invariant algebra*

$$\mathcal{A}_{\mathrm{inv}}(R) = \widetilde{\mathcal{A}}(R)^{\widehat{K}_{\partial R}}, \qquad \widehat{K}_{\partial R} := \prod_{\Sigma \subset \partial R} \widehat{K}_\Sigma,$$

with the convention $\widehat{K}_\Sigma = K_\Sigma$ *whenever no central extension is needed; the physical state space is the corresponding invariant subspace, and on collars the sector-preserving algebra induced on that subspace is the block-diagonal algebra used in Theorem 5.2.*

Items 1 and the lifted fixed-point presentation of item 4 are the fixed-cutoff realized data used by the compact-gauge branch; they are not extra axioms beyond the finite regulator presentation plus overlap redundancy.

Proof sketch. Item 1 is the Hilbertization of finite sets. For item 2, each overlap-consistent recharting acts on a finite-dimensional matrix algebra, and every *-automorphism of such an algebra is inner. After transporting all local overlap presentations to a chosen reference chart, each recharting is therefore implemented by a unitary on the cut Hilbert space. The subgroup generated by those unitaries has compact closure inside a finite-dimensional unitary group, which yields K_Σ; without fixing the reference chart one has the corresponding compact unitary groupoid. Item 3 is the usual projective-versus-central-extension lift for a central 2-cocycle. Item 4 records the lifted fixed-point presentation: chart-independent endomorphisms of the unreduced boundary data are exactly the fixed points of the derived boundary action, while the collar theorem uses the invariant-state realization and the sector-preserving algebra induced on the matched-cut subspace. If the defect is genuinely noncentral, the quotient is well defined, but the correct bookkeeping object is the crossed-module or higher-gauge data rather than an ordinary compact group. □

Theorem 5.2 (Derived edge-center collar decomposition and exact Markov normal form)**.** *Let A-B-D be a collar tripartition realized at fixed regulator scale on the ordinary or central-defect branch of Proposition 5.1. Write* $B = B_L \cup B_R$ *with common interface* $\Sigma = \partial C$*, and let* $\widehat{K}_\Sigma$ *be the derived compact boundary action attached to that cut, with* $\widehat{K}_\Sigma = K_\Sigma$ *on the ordinary branch. In the invariant-state realization,*

$$\mathcal{H}_B = (\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R})^{\widehat{K}_\Sigma}.$$

Decompose the left boundary data as

$$\tilde{\mathcal{H}}_{B_L} \cong \bigoplus_\alpha W_\alpha \otimes \mathcal{H}_{b_L^\alpha},$$

with W_α irreducible $\widehat{K}_\Sigma$-modules. Because the right half-collar carries the inverse overlap transport across the same cut, it decomposes contragrediently:

$$\tilde{\mathcal{H}}_{B_R} \cong \bigoplus_\beta W_\beta^* \otimes \mathcal{H}_{b_R^\beta}.$$

Then

$$\mathcal{H}_B \cong \bigoplus_\alpha \mathcal{H}_{b_L^\alpha} \otimes \mathcal{H}_{b_R^\alpha},$$

and the sector-preserving collar algebra

$$\mathcal{A}_{\mathrm{EC}}(B) := \bigoplus_\alpha \mathcal{B}(\mathcal{H}_{b_L^\alpha}) \otimes \mathcal{B}(\mathcal{H}_{b_R^\alpha})$$

has center

$$Z(\mathcal{A}_{\mathrm{EC}}(B)) = \bigoplus_\alpha \mathbb{C}\,\mathbf{1}_\alpha.$$

If, in addition, the reference state on $A \cup B \cup D$ is exact Markov,

$$I(A : D|B)_\rho = 0,$$

then the state decomposes as

$$\rho_{ABD} = \bigoplus_\alpha p_\alpha\, \rho^{(\alpha)}_{Ab_L^\alpha} \otimes \rho^{(\alpha)}_{b_R^\alpha D}.$$

The same formula holds in any limit state that is exact Markov on this fixed collar model. Thus the collar-center block decomposition is forced by overlap consistency plus the derived regulator package on the ordinary or central-defect branch; exact Markov factorization is an additional state condition, not a consequence of the block decomposition alone.

Proof. By Proposition 5.1, one may realize the physical collar states as the $\widehat{K}_\Sigma$-invariant subspace of $\tilde{\mathcal{H}}_{B_L} \otimes \tilde{\mathcal{H}}_{B_R}$. Complete reducibility of finite-dimensional unitary representations gives the displayed decomposition of $\tilde{\mathcal{H}}_{B_L}$, and the right half-collar carries the inverse transport law across the same cut, hence the contragredient decomposition of $\tilde{\mathcal{H}}_{B_R}$. Therefore

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

Taking $\widehat{K}_\Sigma$-invariants gives

$$\mathcal{H}_B \cong \bigoplus_{\alpha,\beta} (W_\alpha \otimes W_\beta^*)^{\widehat{K}_\Sigma} \otimes (\mathcal{H}_{b_L^\alpha} \otimes \mathcal{H}_{b_R^\beta}).$$

By Schur's lemma,

$$(W_\alpha \otimes W_\beta^*)^{\widehat{K}_\Sigma} \cong \begin{cases} \mathbb{C}, & \alpha = \beta, \\ 0, & \alpha \neq \beta, \end{cases}$$

which yields the displayed block decomposition. The sector-preserving collar algebra therefore has the displayed block-diagonal form, so its center is generated by the block projectors. Within each block, observables from $A \cup B$ act only on the left factor and observables from $B \cup D$ act only on the right factor. If the state is exact Markov, the standard HJPW structure theorem gives the blockwise factorized normal form [20]. The same formula is used in the idealized recoverability limit when exact identities are taken literally. □

Theorem 5.3 (Derived higher-gauge cut datum)**.** *Fix a connected cut Σ and a finite regulator chart $\{P_i\}_{i\in I_\Sigma}$ meeting along Σ, with finite nerve N_Σ. On the genuinely noncentral branch, the weak gluing data are implemented by a compact crossed module*

$$\mathbb{K}_\Sigma = (H_\Sigma \xrightarrow{\partial_\Sigma} G_\Sigma, \triangleright)$$

together with maps

$$g_{ij} : P_{ij} \to G_\Sigma, \qquad h_{ijk} : P_{ijk} \to H_\Sigma,$$

obeying

$$g_{ij}g_{jk} = \partial_\Sigma(h_{ijk})\, g_{ik}, \qquad h_{jkl}h_{ijl} = (g_{ij} \triangleright h_{ikl})\, h_{ijk}.$$

The corresponding higher-gauge change system on the cut is the compact semidirect product

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

with multiplication

$$(\eta, u) \cdot (\eta', u') = (\eta\, (u \triangleright \eta'),\ uu'),$$

and standard coboundary action

$$g_{ij} \mapsto u_i\, \partial_\Sigma(\eta_{ij})\, g_{ij}\, u_j^{-1},$$

$$h_{ijk} \mapsto (u_i \triangleright h_{ijk})\, \eta_{ij}\, (g_{ij} \triangleright \eta_{jk})\, \eta_{ik}^{-1}.$$

The fixed-cutoff physical sector on the genuinely noncentral branch is therefore the orbit class

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

and no external continuous input P or N_{scr} enters this theorem package.

Proof sketch. On pair overlaps, every regulator-scale recharting acts on a finite-dimensional matrix algebra and is therefore inner, so it is implemented by a unitary on the cut Hilbert space. On triple overlaps, genuinely noncentral failure of strict composition is recorded by unitary associators rather than scalars. These 1- and 2-morphisms form a compact unitary 2-group of rechartings. Skeletal strictification of a compact 2-group is equivalent to crossed-module data, which yields the displayed compact $\mathbb{K}_\Sigma$. Because the nerve is finite, $C^1(N_\Sigma, H_\Sigma)$ and $C^0(N_\Sigma, G_\Sigma)$ are finite products of compact groups, hence $\mathcal{T}_\Sigma$ is compact. The displayed formulas are the standard crossed-module coboundary action, and the orbit class is exactly the nonabelian Čech 2-class of the defect data; see [32] for an explicit higher-lattice-gauge realization of the same crossed-module structures. □

Theorem 5.4 (Higher-gauge edge-center completion)**.** *Let $B = B_L \cup B_R$ be a fixed-cutoff collar around a connected cut Σ on the genuinely noncentral branch of Theorem 5.3. Then, as unitary $\mathcal{T}_\Sigma$-modules,*

$$\widetilde{\mathcal{H}}_{B_L} \cong \bigoplus_{\lambda\in\Lambda_\Sigma} W_\lambda \otimes \mathcal{H}_{b_L^\lambda}, \qquad \widetilde{\mathcal{H}}_{B_R} \cong \bigoplus_{\mu\in\Lambda_\Sigma} W_\mu^* \otimes \mathcal{H}_{b_R^\mu},$$

for a finite set Λ_Σ of irreducible unitary representations of $\mathcal{T}_\Sigma$ that actually occur. Hence the physical higher-gauge collar space is

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

and the collar algebra and its center are

$$\mathcal{A}_{2g}(B) \cong \bigoplus_{\lambda\in\Lambda_\Sigma} \mathcal{B}(\mathcal{H}_{b_L^\lambda}) \otimes \mathcal{B}(\mathcal{H}_{b_R^\lambda}), \qquad Z(\mathcal{A}_{2g}(B)) = \bigoplus_{\lambda\in\Lambda_\Sigma} \mathbb{C}\,\mathbf{1}_\lambda.$$

Proof sketch. Because $\mathcal{T}_\Sigma$ is compact and the half-collar spaces are finite-dimensional, both $\widetilde{\mathcal{H}}_{B_L}$ and $\widetilde{\mathcal{H}}_{B_R}$ split orthogonally into irreducible unitary $\mathcal{T}_\Sigma$-modules with finite multiplicity. The right side sees inverse transport across the same cut, hence the contragredient module. Therefore

$$\widetilde{\mathcal{H}}_{B_L} \otimes \widetilde{\mathcal{H}}_{B_R} \cong \bigoplus_{\lambda,\mu} (W_\lambda \otimes W_\mu^*) \otimes (\mathcal{H}_{b_L^\lambda} \otimes \mathcal{H}_{b_R^\mu}).$$

Taking $\mathcal{T}_\Sigma$-invariants leaves

$$(W_\lambda \otimes W_\mu^*)^{\mathcal{T}_\Sigma} \cong \mathrm{Hom}_{\mathcal{T}_\Sigma}(W_\mu, W_\lambda) \cong \begin{cases} \mathbb{C}, & \lambda = \mu, \\ 0, & \lambda \neq \mu, \end{cases}$$

by Schur's lemma. This yields the matched block decomposition and the center formula. □

Theorem 5.5 (Higher-gauge Markov collar and carried errors)**.** *Let ρ_{ABD} be a faithful state on a collar whose middle region B carries the higher-gauge block decomposition of Theorem 5.4. If*

$$I(A : D \mid B)_\rho = 0,$$

then

$$\rho_{ABD} = \bigoplus_{\lambda \in \Lambda_\Sigma} p_\lambda \, \rho_{Ab_L^\lambda} \otimes \rho_{b_R^\lambda D}.$$

If instead $I(A : D \mid B)_\rho \leq \varepsilon$ on one fixed faithful collar model, then the same carried Fawzi–Renner and fixed-collar Markov errors as in the ordinary branch continue to hold:

$$r_{\mathrm{FR}}(\varepsilon) = 2\sqrt{1 - e^{-\varepsilon}} \leq 2\sqrt{\varepsilon}, \qquad 4\lambda_*^{-1} \, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

Proof sketch. Theorem 5.4 shows that the higher-gauge collar algebra is a finite direct sum of type-I tensor blocks. HJPW depends only on that finite-dimensional algebraic structure, not on whether the block label arose from an ordinary compact group, a central extension, or a crossed-module gauge system. Likewise, the Fawzi–Renner remainder and the fixed-collar Markov modulus are finite-dimensional entropy and recovery statements and are blind to the origin of the block label. □

Theorem 5.6 (Higher-gauge defect transportability)**.** *With the genuinely noncentral cut data of Theorem 5.3, define*

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

Then:

1. *q_Σ is invariant under all local rechartings and higher-gauge changes $(\eta, u) \in \mathcal{T}_\Sigma$.*

2. *Two genuinely noncentral representatives describe the same fixed-cutoff physical cut sector if and only if they determine the same class q_Σ.*

3. *The defect is removable if and only if $q_\Sigma = 0$.*

Thus the genuinely noncentral branch carries the same transportability logic as the ordinary or central branch, but in crossed-module Čech degree 2 rather than ordinary group cohomology.

Proof sketch. Items 1–3 are the crossed-module version of ordinary Čech transportability: coboundaries change representatives while preserving classes, the trivial class strictifies the cocycle, and a nontrivial class cannot be gauged away. In any explicit crossed-module lattice realization, the same sectors are equivalently the connected components of the 2-flat configuration groupoid, or homotopy classes of maps into the classifying space $B\mathbb{K}_\Sigma$. □

Definition 5.7 (Overlap-transport groupoid and strict transportability)**.** *Fix a connected cut* Σ *and a finite regulator charting nerve* N_Σ*. The* overlap-transport groupoid $\Pi_1^{\mathrm{ov}}(N_\Sigma)$ *has the regulator cut charts as objects and words in overlap-preserving rechartings as morphisms, modulo insertion and deletion of immediate inverse rechartings. On the ordinary or central branch, an oriented edge* $i \to j$ *is represented by a unitary recharting* U_{ij} *on the cut presentation, with* $U_{ji} = U_{ij}^{-1}$ *after choosing the inverse chart. On the genuinely noncentral branch, the same notation denotes the* G_Σ*-part of the crossed-module data of Theorem 5.3, with the* H_Σ*-valued associators retained as* 2*-morphisms.*

A collar charge at chart i *is a simple edge-center summand, equivalently an irreducible boundary module* $W_\alpha^{(i)}$ *together with its multiplicity/intertwiner block in the collar decomposition. For a path*

$$p = (i_0 \to i_1 \to \cdots \to i_m)$$

define the transported charge by applying the composite recharting

$$U_p := U_{i_{m-1} i_m} \cdots U_{i_0 i_1}$$

to the boundary module and its intertwiner block. Transport is strictly path-independent *when for any two overlap paths* p, p' *with the same endpoints the induced transport functors on the collar charge blocks are naturally equal after the allowed local recharting gauge changes. Equivalently, every closed overlap path acts trivially on the sector class and on its transported intertwiner block, rather than only projectively or up to a noncentral* 2*-morphism.*

Theorem 5.8 (TransportabilityFromOverlapGluing)**.** *Assume the fixed-cutoff regulator gluing package of Proposition 5.1 on the ordinary or central branch, and assume the crossed-module cut datum of Theorem 5.3 on the genuinely noncentral branch. Then path-independent transportability is a theorem-level zero-obstruction criterion, not a separate input.*

(i) ***Ordinary branch.*** *If the overlap rechartings are strict on triple overlaps, then transport of a collar charge along an overlap path is computed by the path composite* U_p*. It is strictly path-independent if and only if every closed overlap loop has trivial holonomy on the collar-sector block. In the tree-cover case this condition is automatic; on a cover with loops it is exactly the ordinary loop-coherence obstruction.*

(ii) ***Central branch.*** *Suppose the only failure of strict triple-overlap composition is central:*

$$U_{ij} U_{jk} = z_{ijk} U_{ik}, \qquad z_{ijk} \in Z.$$

Let $[z]_\Sigma \in \check{H}^2(N_\Sigma, Z)$ *denote the resulting central loop-coherence class, including the central multipliers accumulated by elementary triangle moves between overlap paths. Then strict path-independent transport of collar charges exists if and only if* $[z]_\Sigma = 0$*. When* $[z]_\Sigma = 0$*, one may choose a central* 1*-cochain* a_{ij} *and replace* U_{ij} *by* $a_{ij}^{-1} U_{ij}$*, after which the rechartings compose strictly and the transport functor depends only on endpoints. When* $[z]_\Sigma \neq 0$*, transport remains only projective: the path dependence is exactly the central multiplier represented by* $[z]_\Sigma$*.*

(iii) ***Genuinely noncentral branch.*** *Let*

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

be the crossed-module class of Theorem 5.6. Strict path-independent transport of ordinary collar charges exists if and only if $q_\Sigma = 0$*. When* $q_\Sigma = 0$*, a crossed-module coboundary*

(η, u) strictifies the weak gluing data to a genuine G_Σ-valued 1-cocycle, and the strict transport construction of the ordinary branch applies. When $q_\Sigma \neq 0$, there is no gauge in which all path comparisons are ordinary equalities: the fixed-cutoff sector is instead a higher-gauge sector labelled by q_Σ, and transport is higher transport in the crossed-module 2-groupoid rather than ordinary path-independent DHR transport.

Thus overlap gluing constructs the transport operation, and the complete obstruction to strict path independence is $[z]_\Sigma$ on the ordinary/central branch and q_Σ on the genuinely noncentral branch.

Proof. For a path $p = (i_0 \to \cdots \to i_m)$, Definition 5.7 gives the transport functor by composing the finite-dimensional recharting implementers. This construction uses only the derived overlap unitary transition system of Proposition 5.1; no transportability assumption is invoked.

On the ordinary branch, strict triple-overlap coherence says that elementary replacements of the form $(i \to j \to k) \rightsquigarrow (i \to k)$ do not change the composite recharting. Any two paths with the same endpoints in the finite nerve differ, after inserting or deleting immediate backtracks, by a finite sequence of such elementary moves together with closed loop moves. Backtracks contribute the identity because $U_{ji} = U_{ij}^{-1}$. Therefore the only possible path dependence is the ordinary loop holonomy. The transport is strictly path-independent exactly when that loop holonomy acts trivially on the collar-sector block.

On the central branch, the same elementary triangle replacement changes the composite by the central multiplier z_{ijk}. For a finite sequence of elementary moves between two paths, the total discrepancy is the product of the corresponding z's over the filling 2-chain. The quadruple-overlap coherence identity is precisely the Čech cocycle identity, so this product depends only on the class $[z]_\Sigma$. If $[z]_\Sigma = 0$, choose a central 1-cochain a_{ij} with $z_{ijk} = a_{ij}a_{jk}a_{ik}^{-1}$. Replacing U_{ij} by $a_{ij}^{-1}U_{ij}$ kills the central multiplier on every triangle, so all elementary path moves preserve the transport functor and the ordinary argument gives strict path independence. Conversely, if strict path-independent transport exists, every elementary triangular comparison has trivial central multiplier after an allowed recharting gauge change; hence z is a coboundary and $[z]_\Sigma = 0$.

On the genuinely noncentral branch, elementary triangle comparisons are non-scalar H_Σ-valued 2-morphisms h_{ijk}, and the consistency of different fillings of a path homotopy is the crossed-module cocycle identity of Theorem 5.3. A local higher-gauge change (η, u) changes representatives by the crossed-module coboundary action displayed there, while preserving the class q_Σ by Theorem 5.6. If $q_\Sigma = 0$, the data are equivalent to a strict representative: the h_{ijk} can be gauged away and $g_{ij}g_{jk} = g_{ik}$ holds. The ordinary path-composite construction then gives strict path-independent transport. Conversely, strict ordinary path-independent transport gives a representative in which every triangular 2-comparison is the identity and every closed surface comparison is trivial; this is exactly a trivial crossed-module Čech 2-class, so $q_\Sigma = 0$. If $q_\Sigma \neq 0$, Theorem 5.6 says the defect is not removable, so no strict ordinary transport functor can exist. □

Corollary 5.9 (Transportability from overlap gluing)**.** *The ordinary/central compact-gauge branch may invoke only those finite-cutoff collar sectors that satisfy the zero-obstruction conclusion of Theorem 5.8. The condition is the strict-transport corollary of overlap gluing. If the genuinely noncentral class q_Σ is nonzero, the branch is a higher-gauge fixed-cutoff sector rather than an ordinary compact-group DR reconstruction sector; if $q_\Sigma = 0$, it strictifies to the ordinary transportable case.*

Definition 5.10 (Fixed-cutoff zero-obstruction collar sectors)**.** *Fix a regulator cutoff r and a connected overlap collar $B = B_L \cup B_R$ around a cut Σ on the ordinary or central-defect branch. Let $\widehat{K}_{\Sigma,r}$ be the compact boundary action supplied by the fixed-cutoff overlap gluing theorem, with*

$\widehat{K}_{\Sigma,r} = K_{\Sigma,r}$ *on the ordinary branch. Edge-center completion gives*

$$\widetilde{\mathcal{H}}_{B_L} \cong \bigoplus_{\alpha\in A_{\Sigma,r}} W_\alpha \otimes \mathcal{H}_{b_L^\alpha}, \qquad \widetilde{\mathcal{H}}_{B_R} \cong \bigoplus_{\alpha\in A_{\Sigma,r}} W_\alpha^* \otimes \mathcal{H}_{b_R^\alpha},$$

and hence

$$\mathcal{H}_B = (\widetilde{\mathcal{H}}_{B_L} \otimes \widetilde{\mathcal{H}}_{B_R})^{\widehat{K}_{\Sigma,r}} \cong \bigoplus_{\alpha\in A_{\Sigma,r}} \mathcal{H}_{b_L^\alpha} \otimes \mathcal{H}_{b_R^\alpha}.$$

The minimal central projector onto the α-summand is denoted P_α. A simple fixed-cutoff collar charge is the zero-obstruction overlap-transport orbit of such a pair (P_α, W_α). Here zero obstruction means the strict path-independence criterion of Theorem 5.8: $[z]_\Sigma = 0$ on the ordinary/central branch, or $q_\Sigma = 0$ if a genuinely noncentral representative has strictified to the ordinary case. For objects X, Y obtained from the W_α's by finite direct sums, tensor products, and subobjects, define

$$\mathrm{Hom}_r(X,Y) := \mathrm{Hom}_{\widehat{K}_{\Sigma,r}}(X,Y),$$

after transporting all localizations to one reference collar by the zero-obstruction transport functor. Different reference collars give canonically unitarily equivalent Hom spaces by Theorem 5.8.

Theorem 5.11 (FixedCutoffBosonicSectorCategory)**.** *At every fixed regulator cutoff r, on the ordinary or central-defect zero-obstruction branch and in the bosonic internal-gauge sector of the $3+1$-dimensional EFT regime, the construction of Definition 5.10 produces a semisimple rigid symmetric C^*-tensor category*

$$\mathsf{Sect}_r^{\mathrm{bos}}.$$

Its simple objects are the transportable edge-center collar charges (P_α, W_α), its tensor product is collar concatenation, its duals are orientation reversal / charge conjugation, its *-operation and C^*-norm are inherited from the finite-dimensional lifted collar algebras, and its symmetry is the bosonic spacelike-exchange symmetry of the $3+1$-dimensional EFT branch.

Proof. By edge-center completion, every simple localized charge visible on one reference collar is represented by a minimal central summand P_α together with an irreducible boundary carrier W_α of $\widehat{K}_{\Sigma,r}$. Since $\widehat{K}_{\Sigma,r}$ is compact and acts unitarily on finite-dimensional collar data, its finite-dimensional representations are completely reducible. The additive Karoubi envelope generated by the W_α's is therefore semisimple, with finite-dimensional Hom spaces.

Transport between collars is not assumed. It is the path-composite construction of Theorem 5.8. On the branch used here the obstruction vanishes, so transport of P_α, W_α, and their intertwiners is independent of the overlap path. Hence the Hom spaces defined after moving all localizations to a reference collar do not depend on the chosen path or reference collar, up to canonical transported unitary identification.

For two objects X and Y, choose spacelike separated or nested collar representatives and concatenate the collars. On the lifted finite-dimensional presentation, the boundary carrier of the concatenated charge is the diagonal tensor product $X \otimes Y$ of $\widehat{K}_{\Sigma,r}$-modules. Complete reducibility decomposes $X \otimes Y$ into simple edge-center summands, so the product stays inside the generated category. The vacuum collar sector is the trivial one-dimensional module and is the tensor unit. Associativity is the parenthesization-independence of union-collar gluing; on the lifted Hilbert presentation it is implemented by the standard finite-dimensional tensor associator, and the pentagon identity says that both paths are the same rebracketing of a fourfold union collar.

Duals are obtained by reversing the collar orientation, which swaps the half-collar boundary actions and sends W_α to W_α^*. Evaluation and coevaluation are the standard $\widehat{K}_{\Sigma,r}$-invariant pairings.

Their zig-zag identities are the usual finite-dimensional duality identities, equivalently cap/cup cancellation for collar gluing.

The *-structure is inherited from the lifted finite-dimensional collar algebra. Each $\mathrm{Hom}_r(X, Y)$ is a closed subspace of operators between finite-dimensional Hilbert spaces, adjoints are again $\widehat{K}_{\Sigma,r}$-intertwiners, composition is operator composition, and the norm is the operator norm. Thus $\|f^*f\| = \|f\|^2$, direct sums and subobjects are implemented by orthogonal projections, and the category is a C^*-category.

In the 3+1-dimensional bosonic internal-gauge EFT branch, two localized bosonic collar charges can be exchanged by a spacelike isotopy of their codimension-two collar supports. This exchange induces the canonical unitary flip $c_{X,Y} : X \otimes_r Y \to Y \otimes_r X$. Naturality and the hexagon identities are the corresponding isotopy identities, and the double exchange is isotopic to the identity on the bosonic branch. Hence $c_{Y,X}c_{X,Y} = \mathrm{id}$, so the braiding is symmetric. 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. □

Corollary 5.12 (Fixed-cutoff category from sector construction)**.** *The fixed-cutoff collar-sector packages are bosonic symmetric C^*-tensor categories by Theorem 5.11. The refinement functors and compatible finite bosonic fibers are supplied by Theorem 7.2. The realized low-energy branch is then supplied by Theorem 7.23, so there is no separate fixed-cutoff category input.*

Proposition 5.13 (Higher-gauge interacting compatibility)**.** *In a fixed-cutoff crossed-module lattice realization of Theorem 5.3, one may add local collar terms*

$$K^{2g}_{\mathrm{collar}} = \sum_{v\in\Sigma_0} \beta_v(1 - A_v) + \sum_{e\in\Sigma_1} \beta_e(1 - B_e) + \sum_{c\in\mathcal{C}_\Sigma} \beta_c(1 - C_c),$$

where A_v, B_e, and C_c are the local gauge and fake-flat projectors of the chosen higher-gauge realization. These terms are bounded-support and commute in the standard higher-gauge projector realization [32], so the full fixed-cutoff generator is quasi-local and inside the Axiom 3.3 class.

Proof sketch. The additional terms are local projector constraints supported on finitely many collar cells, so they preserve bounded support and quasi-locality. In the standard higher-gauge projector realization they commute, so the same finite-constraint MaxEnt/Lieb–Robinson bookkeeping used on the ordinary branch continues to apply. □

Remark 5.14 (Status of the regulator package and the precise Markov boundary)**.** *Proposition 5.1, Theorem 5.2, and Theorems 5.3–5.6 together with Proposition 5.13 show that edge-center completion is a theorem of overlap consistency plus the finite regulator presentation on the ordinary, central-defect, and genuinely noncentral higher-gauge branches. No preferred quantum-link realization is assumed on the ordinary or central branch, and the genuinely noncentral branch is handled by the compact crossed-module replacement rather than by one more ordinary-group lemma. The fixed-cutoff topological UV package is closed on all three branches. The continuum modular/geometric lift on the support-visible geometric subnet is Theorem 6.8; the compact-gauge realized branch is handled by Theorem 7.23. The distinct state-control issue is that exact splice and modular-additivity identities require either literal exact Markovity or a controlled family that converges to the exact Markov set on one fixed finite-dimensional collar.*

Definition 5.15 (Exact Markov set and collar-distance modulus)**.** *Fix the finite-dimensional collar Hilbert space of Theorem 5.2. Let*

$$\mathfrak{M}_{A:B:D} := \{\sigma_{ABD} : \ I(A : D \mid B)_\sigma = 0\}$$

be the set of exact Markov states on that fixed collar. Equivalently, by HJPW and Theorem 5.2, $\sigma \in \mathfrak{M}_{A:B:D}$ iff

$$\sigma_{ABD} = \bigoplus_{\alpha} q_\alpha \, \sigma^{(\alpha)}_{Ab^\alpha_L} \otimes \sigma^{(\alpha)}_{b^\alpha_R D}$$

for the fixed edge-center decomposition of B.

For a general state ρ_{ABD}, define its distance to the exact Markov set by

$$d_{\mathrm{M}}(\rho) := \inf_{\sigma \in \mathfrak{M}_{A:B:D}} \|\rho - \sigma\|_1,$$

and the fixed-collar exact-Markov modulus by

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) := \sup \left\{ d_{\mathrm{M}}(\rho) : I(A:D \mid B)_\rho \le \varepsilon \right\}.$$

Proposition 5.16 (Controlled exact Markov-collar limit)**.** *On a fixed finite-dimensional collar, the exact Markov set $\mathfrak{M}_{A:B:D}$ is compact and*

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \longrightarrow 0 \qquad \text{as } \varepsilon \downarrow 0.$$

Consequently, if $\rho^{(n)}_{ABD}$ is any sequence of states on that same collar with

$$I(A:D \mid B)_{\rho^{(n)}} \le \varepsilon_n, \qquad \varepsilon_n \to 0,$$

then there exist exact Markov states $\sigma^{(n)} \in \mathfrak{M}_{A:B:D}$ such that

$$\|\rho^{(n)} - \sigma^{(n)}\|_1 \le \delta^{\mathrm{M}}_{A:B:D}(\varepsilon_n) \longrightarrow 0.$$

Thus approximate recoverability converges to the exact HJPW normal form used by the compact-gauge branch.

Proof. The full state space on a fixed finite-dimensional Hilbert space is compact in trace norm. Conditional mutual information is continuous there because von Neumann entropy is continuous in finite dimension. Hence $\mathfrak{M}_{A:B:D}$ is closed, so it is compact as a closed subset of a compact space.

Suppose $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \not\to 0$. Then there exist $\eta > 0$, a sequence $\varepsilon_n \downarrow 0$, and states $\rho^{(n)}$ with $I(A:D \mid B)_{\rho^{(n)}} \le \varepsilon_n$ but $d_{\mathrm{M}}(\rho^{(n)}) \ge \eta$ for all n. By compactness, pass to a trace-norm convergent subsequence $\rho^{(n_k)} \to \rho^*$. Continuity of conditional mutual information gives

$$I(A:D \mid B)_{\rho^*} = \lim_{k\to\infty} I(A:D \mid B)_{\rho^{(n_k)}} = 0,$$

so $\rho^* \in \mathfrak{M}_{A:B:D}$. But then

$$d_{\mathrm{M}}(\rho^{(n_k)}) \le \|\rho^{(n_k)} - \rho^*\|_1 \longrightarrow 0,$$

contradicting $d_{\mathrm{M}}(\rho^{(n_k)}) \ge \eta$. Thus $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \to 0$. The final statement follows by choosing $\sigma^{(n)} \in \mathfrak{M}_{A:B:D}$ within $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon_n) + 1/n$ of $\rho^{(n)}$. □

Theorem 5.17 (Finite-collar Markov replacement stability)**.** *Fix one finite-dimensional collar model $A : B : D$ and a faithful floor $\lambda_* > 0$ for the collar state. Let*

$$\mathcal{S}_{\lambda_*} := \{\rho_{ABD} : \rho_{ABD} \ge \lambda_* \mathbf{1}\}$$

inside the affine state space on that fixed collar, and let

$$\mathfrak{M}^{\lambda_*}_{A:B:D} := \mathfrak{M}_{A:B:D} \cap \mathcal{S}_{\lambda_*}.$$

Assume this faithful exact-Markov class is nonempty. There are constants $C_{A:B:D,\lambda_*} > 0$ *and* $\theta_{A:B:D,\lambda_*} > 0$*, depending on the chosen collar model and floor but not on a refinement family, such that every* $\rho \in \mathcal{S}_{\lambda_*}$ *satisfies*

$$d_{\mathrm{M}}^{\lambda_*}(\rho) := \inf_{\sigma \in \mathfrak{M}_{A:B:D}^{\lambda_*}} \|\rho - \sigma\|_1 \leq C_{A:B:D,\lambda_*}\, I(A:D \mid B)_\rho^{\theta_{A:B:D,\lambda_*}}.$$

Equivalently, on that fixed faithful collar class,

$$\delta_{A:B:D}^{\mathrm{M},\lambda_*}(\varepsilon) \leq C_{A:B:D,\lambda_*}\, \varepsilon^{\theta_{A:B:D,\lambda_*}}.$$

The constants are collar-local; this is not a dimension-free stability theorem for arbitrary tripartite systems.

Proof. On $\mathcal{S}_{\lambda_*}$ the von Neumann entropy terms are real analytic functions of the matrix entries, hence

$$F(\rho) := I(A:D \mid B)_\rho$$

is a nonnegative real analytic function. By equality in strong subadditivity, or equivalently the HJPW structure theorem, the zero set of F is precisely $\mathfrak{M}_{A:B:D}^{\lambda_*}$. The compact-set Lojasiewicz inequality for a real analytic nonnegative function gives constants $c > 0$ and $q > 0$ on this fixed compact collar class such that

$$F(\rho) \geq c\, d_{\mathrm{M}}^{\lambda_*}(\rho)^q.$$

Taking $C = c^{-1/q}$ and $\theta = 1/q$ gives the displayed estimate. □

Proposition 5.18 (Approximate collar recovery)**.** *Let A-B-D be a collar tripartition satisfying*

$$I(A:D|B)_\rho \leq \varepsilon,$$

and let $\mathcal{R}_{B\to BD}$ *be a recovery map such that*

$$F(\rho_{ABD}, (\mathrm{id}_A \otimes \mathcal{R}_{B\to BD})(\rho_{AB})) \geq e^{-\varepsilon/2}.$$

Define the recovered comparison state

$$\rho_{ABD}^{\mathrm{rec}} := (\mathrm{id}_A \otimes \mathcal{R}_{B\to BD})(\rho_{AB}), \qquad r_{\mathrm{FR}}(\varepsilon) := 2\sqrt{1 - e^{-\varepsilon}} \leq 2\sqrt{\varepsilon}.$$

Then

$$\|\rho_{ABD} - \rho_{ABD}^{\mathrm{rec}}\|_1 \leq r_{\mathrm{FR}}(\varepsilon).$$

Consequently, for every CPTP map $\Lambda_{BD\to B'D'}$ *and every bounded observable* X *supported on* $A \cup B'$*,*

$$|\mathrm{Tr}[X(\mathrm{id}_A \otimes \Lambda)(\rho_{ABD} - \rho_{ABD}^{\mathrm{rec}})]| \leq \|X\|_\infty \|(\mathrm{id}_A \otimes \Lambda)(\rho_{ABD} - \rho_{ABD}^{\mathrm{rec}})\|_1 \leq \|X\|_\infty\, r_{\mathrm{FR}}(\varepsilon).$$

Proof. The fidelity lower bound and the Fuchs–van de Graaf inequality give

$$\|\rho_{ABD} - \rho_{ABD}^{\mathrm{rec}}\|_1 \leq 2\sqrt{1 - F(\rho_{ABD}, \rho_{ABD}^{\mathrm{rec}})^2} \leq 2\sqrt{1 - e^{-\varepsilon}} = r_{\mathrm{FR}}(\varepsilon).$$

The estimate $r_{\mathrm{FR}}(\varepsilon) \leq 2\sqrt{\varepsilon}$ follows from $1 - e^{-x} \leq x$. Contractivity of trace norm under CPTP maps gives the transported bound, and the observable estimate is the duality between trace norm and operator norm [19]. The bound is constructive and dimension-free. What it controls is closeness to a recovered comparison state; no claim is made that $\rho_{ABD}^{\mathrm{rec}}$ is itself exact Markov. □

Proposition 5.19 (Exact splice theorem and controlled collar approximation)**.** *Under Theorem 5.2, let $\sigma_{ABD} \in \mathfrak{M}_{A:B:D}$ have exact Markov form*

$$\sigma_{ABD} = \bigoplus_{\alpha} p_\alpha\, \sigma^{(\alpha)}_{Ab_L^\alpha} \otimes \sigma^{(\alpha)}_{b_R^\alpha D}.$$

Let $\tau^{(\alpha)}_{b_R^\alpha D'}$ be any normalized family of states compatible with the same right-boundary sectors, and define

$$\sigma'_{ABD'} = \bigoplus_{\alpha} p_\alpha\, \sigma^{(\alpha)}_{Ab_L^\alpha} \otimes \tau^{(\alpha)}_{b_R^\alpha D'}.$$

Define the common left algebra

$$\mathcal{A}_{Ab_L} := \bigoplus_{\alpha} \mathcal{B}(\mathcal{H}_{Ab_L^\alpha}) \otimes \mathbf{1}_{b_R^\alpha},$$

canonically represented on both direct-sum Hilbert spaces. Then:

1. *for every $X \in \mathcal{A}_{Ab_L}$,*
$$\mathrm{Tr}(X\sigma'_{ABD'}) = \mathrm{Tr}(X\sigma_{ABD});$$
2. *if ρ_{ABD} is any state on the same fixed collar with $I(A:D \mid B)_\rho \le \varepsilon$, and if $\sigma_\varepsilon \in \mathfrak{M}_{A:B:D}$ is chosen so that*
$$\|\rho_{ABD} - \sigma_\varepsilon\|_1 \le \delta^{\mathrm{M}}_{A:B:D}(\varepsilon),$$
then the corresponding exact splice σ'_ε satisfies
$$\left|\mathrm{Tr}(X\rho_{ABD}) - \mathrm{Tr}(X\sigma'_\varepsilon)\right| \le \|X\|_\infty\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$$
for all $X \in \mathcal{A}_{Ab_L}$.

Hence the exact splice identity used by the compact-gauge branch is justified either at exact Markovity or along any controlled collar family for which $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon_\delta) \to 0$.

Proof. For item 1, blockwise factorization gives

$$\mathrm{Tr}(X\sigma'_{ABD'}) = \sum_{\alpha} p_\alpha \mathrm{Tr}\Big(X_\alpha \sigma^{(\alpha)}_{Ab_L^\alpha}\Big)\, \mathrm{Tr}\Big(\tau^{(\alpha)}_{b_R^\alpha D'}\Big).$$

Here X_α denotes the α-block of X in $\mathcal{A}_{Ab_L}$. Each $\tau^{(\alpha)}$ is normalized, so the second factor is 1. The same computation for σ_{ABD} gives the same value because the original right factor is likewise normalized.

For item 2, apply item 1 to σ_ε:

$$\mathrm{Tr}(X\sigma'_\varepsilon) = \mathrm{Tr}(X\sigma_\varepsilon).$$

Therefore

$$\left|\mathrm{Tr}(X\rho_{ABD}) - \mathrm{Tr}(X\sigma'_\varepsilon)\right| = |\mathrm{Tr}[X(\rho_{ABD} - \sigma_\varepsilon)]| \le \|X\|_\infty\, \|\rho_{ABD} - \sigma_\varepsilon\|_1 \le \|X\|_\infty\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

□

Proposition 5.20 (Transport of modular-additivity errors from the exact Markov reference)**.** *Let $K_X(\omega) := -\log \omega_X$ for a faithful reduced state on region X, and define the modular defect*

$$\Delta K(\omega) := K_{ABD}(\omega) - K_{AB}(\omega) - K_{BD}(\omega) + K_B(\omega).$$

Fix a collar and let ρ be a faithful state with $I(A : D \mid B)_\rho \le \varepsilon$. Choose $\sigma_\varepsilon \in \mathfrak{M}_{A:B:D}$ with

$$\|\rho - \sigma_\varepsilon\|_1 \le \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

Assume that the four marginals ρ_X and $(\sigma_\varepsilon)_X$ for $X \in \{ABD, AB, BD, B\}$ all satisfy

$$\rho_X \ge \lambda_* \mathbf{1}, \qquad (\sigma_\varepsilon)_X \ge \lambda_* \mathbf{1}$$

on their supports for some $\lambda_* > 0$. Then

$$\|\Delta K(\rho) - \Delta K(\sigma_\varepsilon)\|_\infty \le 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

Since $\Delta K(\sigma_\varepsilon)$ is central or blockwise constant, every matrix element of the modular defect of ρ is within $4\lambda_*^{-1}\delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$ of the exact Markov value.

Proof. By monotonicity of trace distance under partial trace,

$$\|\rho_X - (\sigma_\varepsilon)_X\|_1 \le \|\rho - \sigma_\varepsilon\|_1 \le \delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$$

for each $X \in \{ABD, AB, BD, B\}$. On the interval $[\lambda_*, 1]$, the function $\log x$ has derivative bounded by λ_*^{-1}. By the standard integral representation of the operator logarithm, this implies the operator-norm Lipschitz bound

$$\|K_X(\rho) - K_X(\sigma_\varepsilon)\|_\infty = \|\log(\sigma_\varepsilon)_X - \log \rho_X\|_\infty \le \lambda_*^{-1}\, \|\rho_X - (\sigma_\varepsilon)_X\|_\infty \le \lambda_*^{-1}\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

Summing the four region contributions gives

$$\|\Delta K(\rho) - \Delta K(\sigma_\varepsilon)\|_\infty \le 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

For $\sigma_\varepsilon \in \mathfrak{M}_{A:B:D}$, the exact HJPW block factorization implies that $\Delta K(\sigma_\varepsilon)$ is central or blockwise constant, so the same bound controls all matrix elements relative to the exact Markov modular-additivity value. ☐

Theorem 5.21 (Finite-stage modular-defect propagation)**.** *Consider a fixed finite branch calculation that uses N collar or strip modular-additivity identities, followed by finitely many sums, products, commutators with bounded operators, bounded functional-calculus operations on a fixed spectral interval, and bounded-time modular transports. Replace the j-th exact identity by its finite-stage controlled form with errors*

$$r_{\mathrm{FR}}(\varepsilon_j), \qquad \delta^{\mathrm{M}}_j(\varepsilon_j), \qquad \eta^{\mathrm{reg}}_j,$$

where η^{reg}_j denotes the regularized support-visible transport remainder when no common full-algebra floor is used. Then every bounded downstream modular observable O produced by that calculation obeys

$$|\langle O\rangle_{\mathrm{finite}} - \langle O\rangle_{\mathrm{exact}}| \le \mathcal{P}_O\Big(\{r_{\mathrm{FR}}(\varepsilon_j)\}_{j=1}^N, \{\delta^{\mathrm{M}}_j(\varepsilon_j)\}_{j=1}^N, \{\eta^{\mathrm{reg}}_j\}_{j=1}^N\Big),$$

for a polynomial-continuity modulus $\mathcal{P}_O$ whose coefficients depend only on the fixed collar models, the bounded-time interval, the operator norms, and the declared faithful floors or regularization schedules. In particular $\mathcal{P}_O \to 0$ as all listed errors tend to zero.

Proof. Induct over the expression tree defining O. Sums and products are controlled by triangle inequalities and submultiplicativity on the bounded operator class. Commutators satisfy

$$\|[A, B] - [A', B']\| \leq \|A - A'\| \, \|B\| + \|A'\| \, \|B - B'\| + \|B - B'\| \, \|A - A'\|,$$

with all norms bounded on the fixed calculation. Bounded functional calculus on a fixed spectral interval is uniformly continuous, and it is Lipschitz for the logarithmic comparisons when the floor λ_* is present. For a bounded-time modular transport generated by K and K', Duhamel's formula gives

$$\|e^{itK} X e^{-itK} - e^{itK'} X e^{-itK'}\| \leq 2|t| \, \|X\| \, \|K - K'\| + O(\|K - K'\|^2)$$

uniformly for t in the chosen compact interval. When $K - K'$ is not an unregularized full-algebra bounded operator, Theorem 6.12 supplies the support-visible matrix-element replacement and its remainder η_j^{reg}. Composing these finitely many continuity estimates gives the displayed polynomial modulus. □

Theorem 5.22 (Collar-locality of dimension-dependent constants)**.** *All constants used in exact-Markov replacement, logarithmic modular comparison, regularized modular transport, and the finite-stage propagation of modular defects are functions only of the declared fixed local collar model, its support-visible dimension, the chosen faithful floor or regularization schedule, the bounded observable class, and the bounded modular-time interval. No estimate used in the BW/Lorentz/null-modular/Einstein branch requires a dimension-free trace-norm stability theorem from small conditional mutual information to exact Markov normal form for arbitrary tripartite quantum systems.*

Proof. The Fawzi–Renner term r_{FR} is the only one-shot constructive recoverability estimate used here, and it compares the state to a recovered comparison state rather than to the exact Markov set. The exact-Markov replacement modulus δ^{M} is defined after pullback to one fixed finite collar model; Theorem 5.17 gives a rate only on a fixed faithful collar class. Logarithmic comparison uses the local floor λ_*, while the BW branch avoids a refinement-uniform full-algebra floor by using the regularized support-visible transport theorem. The propagation theorem above then composes only those collar-local estimates. □

Remark 5.23 (What is and is not carried forward)**.** *There are therefore two distinct quantitative controls, and this paper keeps them separate:*

1. *the constructive one-shot recoverability bound $r_{\mathrm{FR}}(\varepsilon) = O(\varepsilon^{1/2})$, which compares the physical state to a recovered comparison state and is enough for bounded-observable statements at one regulator stage;*

2. *the fixed-collar exact-Markov modulus $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) \to 0$, which compares the physical state to the exact Markov set and is the correct quantity for justifying exact splice or modular-additivity identities in a controlled limit, with an additional factor $4\lambda_*^{-1}$ when logarithms are involved.*

Later spatial and null collar theorems therefore use exact identities only in two regimes: literal exact Markovity, or a controlled refinement or scaling family in which the relevant $\delta^{\mathrm{M}}(\varepsilon_\delta)$ tends to zero after transport to one fixed finite-dimensional collar model. The manuscript does not *claim a universal dimension-free one-shot trace-norm estimate from small conditional mutual information directly to an exact Markov state.*

Remark 5.24 (Floor transfer to the exact-Markov reference)**.** *On one fixed finite-dimensional collar model, Proposition 5.20 does not require a second independent faithfulness input on the exact-Markov comparison family. Suppose the transported physical marginals satisfy*

$$\rho_X \geq \bar{\lambda}_{m,\delta}\mathbf{1}$$

eventually for each $X \in \{ABD, AB, BD, B\}$*, and choose exact-Markov replacements* $\sigma_{n,m,\delta}$ *with*

$$\|\rho_X - (\sigma_{n,m,\delta})_X\|_1 \leq \delta^{\mathrm{M}}_{m,\delta}(\varepsilon_{n,m,\delta}) \longrightarrow 0.$$

Then $\|\rho_X - (\sigma_{n,m,\delta})_X\|_\infty \leq \|\rho_X - (\sigma_{n,m,\delta})_X\|_1$*, so for all sufficiently large* n*,*

$$(\sigma_{n,m,\delta})_X \geq \tfrac{1}{2}\bar{\lambda}_{m,\delta}\mathbf{1}.$$

Thus a lower spectral bound, when available on a fixed collar model, is inherited by the exact-Markov reference once the exact-Markov modulus tends to zero. The support-visible theorem below does not require such a bound uniformly across the full refining matrix algebra.

Proposition 5.25 (Generalized-entropy split)**.** *If the reduced cap state takes the block form*

$$\rho_C = \bigoplus_\alpha p_\alpha \left(\rho^{(\alpha)}_{\mathrm{bulk},C} \otimes \frac{\mathbf{1}^{(\alpha)}_{\mathrm{edge}}}{d_\alpha} \right),$$

then

$$S(\rho_C) = S_{\mathrm{bulk}}(C) + \mathrm{Tr}(\rho_C L_C),$$

where

$$S_{\mathrm{bulk}}(C) = H(p_\alpha) + \sum_\alpha p_\alpha S(\rho^{(\alpha)}_{\mathrm{bulk},C}), \qquad L_C = \sum_\alpha (\log d_\alpha) P_\alpha.$$

Proof. The entropy of a direct sum is

$$S\left(\bigoplus_\alpha p_\alpha \sigma_\alpha \right) = H(p_\alpha) + \sum_\alpha p_\alpha S(\sigma_\alpha).$$

Apply this identity with

$$\sigma_\alpha = \rho^{(\alpha)}_{\mathrm{bulk},C} \otimes \frac{\mathbf{1}^{(\alpha)}_{\mathrm{edge}}}{d_\alpha}.$$

Since

$$S(\sigma_\alpha) = S(\rho^{(\alpha)}_{\mathrm{bulk},C}) + \log d_\alpha,$$

the first statement follows. □

Definition 5.26 (Effective Newton constant in the refinement dictionary)**.** *In the refinement-scaling regime, if*

$$\mathrm{Tr}(\rho_C L_C) \approx N_\Sigma \bar{\ell}(t), \qquad A(\partial C) \approx N_\Sigma a_{\mathrm{cell}},$$

then matching $\mathrm{Tr}(\rho_C L_C)$ *to the area term* $A(\partial C)/(4G)$ *defines the effective Newton constant*

$$G := \frac{a_{\mathrm{cell}}}{4\bar{\ell}(t)}.$$

This is a continuum dictionary identification, not part of the exact algebraic proposition.

Proposition 5.27 (Shared edge-entropy identity on the realized product-group branch)**.** *On the realized product-group branch*

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

let

$$R = R_3 \boxtimes R_2 \boxtimes q$$

be a lifted product presentation of one cut sector on the quotient branch, and let the R-sector contribution of the collar edge-center entropy operator be

$$(L_C)\,|_R = \log d_R.$$

Then

$$(L_C)\,|_R = L_C^{(3)} + L_C^{(2)}, \qquad L_C^{(3)} := \log d_{R_3}, \qquad L_C^{(2)} := \log d_{R_2},$$

because every irreducible $\mathrm{U}(1)$ *representation is one-dimensional and contributes* $\log 1 = 0$*. Consequently, on the product heat-kernel branch*

$$\bar{\ell}_{\text{shared}} = \langle L_C \rangle = \bar{\ell}_{\mathrm{SU}(3)}(t_{3,\text{run}}) + \bar{\ell}_{\mathrm{SU}(2)}(t_{2,\text{run}}).$$

If the same branch satisfies the D10 pixel law

$$\bar{\ell}_{\mathrm{SU}(2)}(t_{2,\text{run}}) \; + \; \bar{\ell}_{\mathrm{SU}(3)}(t_{3,\text{run}}) = P/4,$$

then

$$\bar{\ell}_{\text{shared}} = P/4.$$

Proof. For the lifted product presentation $R = R_3 \boxtimes R_2 \boxtimes q$ of a sector on the quotient branch,

$$d_R = d_{R_3} d_{R_2} d_q.$$

Every irreducible $\mathrm{U}(1)$ representation is one-dimensional, so $d_q = 1$. Therefore

$$\log d_R = \log d_{R_3} + \log d_{R_2} + \log d_q = \log d_{R_3} + \log d_{R_2}.$$

This proves the sector identity

$$(L_C)\,|_R = L_C^{(3)} + L_C^{(2)}.$$

On the product heat-kernel branch, expectation values add, so

$$\bar{\ell}_{\text{shared}} = \langle L_C \rangle = \bar{\ell}_{\mathrm{SU}(3)}(t_{3,\text{run}}) + \bar{\ell}_{\mathrm{SU}(2)}(t_{2,\text{run}}).$$

If the same branch satisfies the D10 pixel law, the right-hand side equals $P/4$. □

Proposition 5.28 (Local familiar-unit readout package on the gravity row)**.** *On the declared local extension surface, let* a_{cell} *be the microscopic cell-area datum and define*

$$G_{\text{nat}} := \frac{a_{\text{cell}}}{4\bar{\ell}_{\text{shared}}}.$$

If Proposition 5.27 gives $\bar{\ell}_{\text{shared}} = P/4$*, then*

$$G_{\text{nat}} = \frac{a_{\text{cell}}}{P} = \ell_P^2.$$

In SI units $\ell_P^2 = \hbar G_{\mathrm{SI}}/c^3$. *Therefore*

$$G_{\mathrm{SI}} = \frac{c^3 a_{\mathrm{cell}}}{\hbar P}.$$

More generally, if $\widehat{L}(P)$, $\widehat{T}(P)$, $\widehat{E}(P)$, *and* $\widehat{\Theta}(P)$ *denote dimensionless branch outputs on that same local branch, then*

$$L_{\mathrm{loc}} = \sqrt{a_{\mathrm{cell}}}\,\widehat{L}(P), \qquad t_{\mathrm{loc}} = \frac{\sqrt{a_{\mathrm{cell}}}}{c}\,\widehat{T}(P),$$

$$E_{\mathrm{loc}} = \frac{\hbar c}{\sqrt{a_{\mathrm{cell}}}}\,\widehat{E}(P), \qquad \Theta_{\mathrm{loc}} = \frac{\hbar c}{k_B\sqrt{a_{\mathrm{cell}}}}\,\widehat{\Theta}(P).$$

Thus, at fixed P, *local lengths scale with* $a_{\mathrm{cell}}^{1/2}$ *alone, local times use that same* $a_{\mathrm{cell}}^{1/2}$ *scale divided by the structural Lorentz output* c, *and local mass/energy and temperature rows use the inverse* $a_{\mathrm{cell}}^{1/2}$ *scale dressed only by the familiar-unit display conventions* $\hbar$ *and* k_B.

Proof. By Proposition 5.27,

$$G_{\mathrm{nat}} = \frac{a_{\mathrm{cell}}}{4(P/4)} = \frac{a_{\mathrm{cell}}}{P}.$$

Since $P = a_{\mathrm{cell}}/\ell_P^2$, this equals ℓ_P^2. The SI relation

$$\ell_P^2 = \frac{\hbar G_{\mathrm{SI}}}{c^3}$$

then gives

$$G_{\mathrm{SI}} = \frac{c^3 a_{\mathrm{cell}}}{\hbar P}.$$

The following formulas are dimensional bookkeeping on the same declared local branch. Because P is dimensionless and a_{cell} is the only local microscopic area datum, every local length readout is $\sqrt{a_{\mathrm{cell}}}$ times a dimensionless branch quantity. The structural Lorentz output c converts that same local ruler to seconds, while the familiar-unit display constants $\hbar$ and k_B convert the inverse local ruler to energy/mass rows and to Kelvin rows. □

Proposition 5.19 is the algebraic version of interior invariance under compatible exterior substitutions. Proposition 5.20 is the precise bridge from small collar conditional mutual information to the exact modular identities later used by the spatial and null branches: at finite stage one carries a remainder, and exact identities appear only when that remainder is zero or tends to zero in the controlled collar limit. Proposition 5.25 isolates the exact entropy split, while Definition 5.26 supplies the continuum identification that turns the edge center into a gravitational coupling. Proposition 5.27 identifies the shared edge entropy on the realized product-group branch with the same $\mathrm{SU}(2) + \mathrm{SU}(3)$ quantity used by the D10 pixel law, and Proposition 5.28 makes the familiar-unit readout explicit relative to the declared microscopic datum a_{cell}.

6 Modular Geometry, Relativity, and Einstein Dynamics

6.1 Geometric modular flow

Before stating the continuum modular theorem, fix the algebraic target. At fixed cutoff, the full operational patch algebra contains geometric overlap data, declared record summaries, and compare/write/verify pointer layers. The BW question should be asked only on the overlap-generated geometric subnet, not on the entire operational algebra.

Definition 6.1 (Operational total algebra, geometric subnet, and auxiliary/record sector)**.** *At fixed cutoff n and patch P, let $\mathcal{A}_n^{\mathrm{tot}}(P)$ denote the operational patch algebra consisting of the physical patch algebra together with the declared overlap-readable summaries used by the compare/write/verify surface. Let $B_{IJ} \subset P$ range over overlap collars incident on P, let $\Pi_\alpha^{(IJ)}$ denote the overlap sector projectors, and let $Q_a^{(IJ)}$ denote the boundary observables carrying the geometric cut data. Define the geometric subnet by*

$$\mathcal{A}_n^{\mathrm{geo}}(P) := W^*\Big\langle \bigcup_{B_{IJ}\subset P} \iota_{IJ\to P}\big(\{\Pi_\alpha^{(IJ)}\}_\alpha \cup \{Q_a^{(IJ)}\}_a\big)\Big\rangle_{\mathrm{repair,isotony}} \Big/ \mathcal{N}_n(P),$$

where $\mathcal{N}_n(P)$ is the maximal repair-invariant overlap-trivial kernel, namely the maximal part whose restriction to every overlap algebra inside P is scalar. The complementary operational sector generated by record summaries, pointer registers, and other interface-inert observables is denoted $\mathcal{A}_n^{\mathrm{aux/rec}}(P)$.

Theorem 6.2 (Geometric subnet and operational primeness)**.** *For each fixed-cutoff patch P, the operational algebra admits a canonical quotient*

$$\pi_{n,P}^{\mathrm{geo}} : \mathcal{A}_n^{\mathrm{tot}}(P) \twoheadrightarrow \mathcal{A}_n^{\mathrm{geo}}(P)$$

with the following properties:

- (i) *$\mathcal{A}_n^{\mathrm{geo}}(P)$ is generated by the overlap sector projectors and geometric boundary observables, and is closed under repair and isotony;*
- (ii) *the record-summary block, pointer algebra, and interface-inert auxiliary observables lie in the complementary operational sector $\mathcal{A}_n^{\mathrm{aux/rec}}(P)$ and do not enlarge $\mathcal{A}_n^{\mathrm{geo}}(P)$;*
- (iii) *if $N \subset \mathcal{A}_n^{\mathrm{geo}}(P)$ is a repair-invariant subalgebra whose restriction to every overlap algebra inside P is scalar, then $N = \mathbb{C}\,\mathbf{1}$.*

Proof. Items (i) and (ii) are the content of Definition 6.1. For item (iii), if such an N were nontrivial, then its pullback under $\pi_{n,P}^{\mathrm{geo}}$ would define a nontrivial repair-invariant overlap-trivial component of $\mathcal{A}_n^{\mathrm{tot}}(P)$, contradicting the maximality of $\mathcal{N}_n(P)$. Thus no nontrivial overlap-trivial factor survives inside $\mathcal{A}_n^{\mathrm{geo}}(P)$. □

The collar analysis proves a fixed-cutoff statement on the finite type-I regulator net: the reduced cap state has a literal density matrix, its modular Hamiltonian exists, and its nonadditive part is confined to a shrinking collar up to carried errors. The Lorentz claim is therefore not a literal statement about the finite regulator matrices. The target object is a support-visible realized scaling-limit geometric cap pair

$$(\mathcal{A}_\infty^{\mathrm{geo}}(C), \omega_\infty^{\mathrm{geo},C})$$

for each round cap $C \subset S^2$. Axiom 3.3 controls the state-side realized branch across refinement through one common finite-dimensional MaxEnt family, and Theorem 6.2 removes the overlap-trivial auxiliary/record layers from the BW target. The theorem below uses support-visible regularized modular transport, weak-$*$/GNS extraction, support-readable modular covariance, ordered cut-pair rigidity, and KMS/BW normalization to close the cap automorphism statement needed downstream.

Fix a cap $C \subset S^2$ and a shrinking collar family $(A_\delta, B_\delta, D_\delta)$ around ∂C, with $\delta \downarrow 0$ and $A_\delta \cup B_\delta \cup D_\delta$ covering a neighborhood of the cut. Write

$$\varepsilon_\delta := I(A_\delta : D_\delta \mid B_\delta)_\omega, \qquad r_{\mathrm{FR}}(\varepsilon_\delta) := 2\sqrt{1 - e^{-\varepsilon_\delta}} \le 2\sqrt{\varepsilon_\delta},$$

and, on each fixed faithful collar model,

$$\eta_\delta^{\mathrm{M}} := 4\lambda_*^{-1}\, \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta),$$

where $\lambda_* > 0$ is the lower spectral bound used to compare modular Hamiltonians. Every exact collar identity below is therefore to be read in one of two ways:

1. literal exactness when the reference collar state is exact Markov; or
2. a controlled collar family for which

$$\delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta) \to 0 \qquad (\delta \downarrow 0),$$

   with the finite-stage errors $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η_δ^{M} carried explicitly.

For each cap C, let $\lambda_C(s) \subset \mathrm{Conf}^+(S^2)$ denote the standard cap-preserving conformal one-parameter subgroup, normalized so that the null blow-up near a smooth cut acts by $v \mapsto e^{-s}v$, and let $\alpha_{\lambda_C(s)}$ denote the induced automorphism of the scaling-limit cap net.

Remark 6.3 (BW is a conditional scaling-limit branch)**.** *The Lorentz/BW statement is not "finite cells imply Lorentz invariance." The branch theorem is the conditional implication*

$$\begin{gathered} \textit{support-visible cap-pair extraction} + \textit{ regularized modular transport} \\ + \textit{ support-readable modular covariance} \\ + \textit{ round-cap rigidity} + \textit{ KMS/BW normalization} \Longrightarrow \sigma_t^{\omega_\infty^{\mathrm{geo},C}} = \alpha_{\lambda_C(2\pi t)}. \end{gathered}$$

All four inputs are part of the branch certificate. Finite type-I regulators supply collar control and regularized support-visible matrix elements; they do not make the full finite operational algebra, pointer/record sectors, or off-support directions Lorentz covariant.

Proposition 6.4 (Support-visible cap-pair extraction on the local GNS support quotient)**.** *Fix a round cap $C \subset S^2$. Under Axioms 3.1–3.4, Theorem 6.2, and the derived fixed-cutoff regulator/collar/consensus package, consider the transported geometric cap-local test family on $\mathcal{A}_n^{\mathrm{geo}}(C_n)$, its projectively compatible transported marginals, the asymptotic transport-equivalence certificate, and a cutoff schedule satisfying Theorem 6.12. Along a cofinal refinement subnet:*

(i) *the transported cap marginals admit local weak-$*$ limits on every support-visible geometric cap subalgebra;*

(ii) *the compatible local limits glue to a state on $\mathcal{A}_\infty^{\mathrm{geo,sv}}(C)$;*

(iii) *the GNS representation gives a cap pair*

$$(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C})$$

faithful on the local support quotient, and the regularized support-visible modular matrix elements converge to its modular automorphism group.

Proof. For each fixed local collar model, the finite-stage state spaces are weak-$*$ compact. The transported marginal family is projectively compatible up to the refinement-equivalence certificate supplied by the consensus package, so the usual diagonal subnet argument gives compatible local weak-$*$ limits on the cap-local test family. Theorem 6.12 controls the regularized modular matrix elements on every bounded support-visible collar observable under the stated cutoff schedule. Passing to the GNS representation of the glued state quotients the null directions and leaves a faithful representation on the local support. This gives the displayed support-visible cap pair and its modular automorphism group. □

Lemma 6.5 (Support-readable modular covariance)**.** *Let $(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C})$ be the cap pair of Proposition 6.4. Assume the extracted support-visible prime geometric subnet satisfies outer regularity / minimal support: for every cap-local observable O, the intersection of all connected regions P with $O \in \mathcal{A}_\infty^{\mathrm{geo,sv}}(P)$ is again a connected region, denoted* $\mathrm{supp}(O)$*. For every connected cap-local region $R \subset C$, define*

$$f_t^C(R) := \bigcup_{O \in \mathcal{A}_\infty^{\mathrm{geo,sv}}(R)} \mathrm{supp}\Big(\sigma_t^{\omega_\infty^{\mathrm{geo},C}}(O)\Big).$$

Then the modular automorphism group induces a one-parameter support map on cap-local regions:

$$\sigma_t^{\omega_\infty^{\mathrm{geo},C}}(\mathcal{A}_\infty^{\mathrm{geo,sv}}(R)) = \mathcal{A}_\infty^{\mathrm{geo,sv}}\Big(f_t^C(R)\Big).$$

Proof. The modular group is an automorphism group of the cap algebra. By outer regularity / minimal support, the support of each cap-local observable is readable from the net labeling after GNS quotienting of null directions. Thus the displayed definition of $f_t^C(R)$ is intrinsic to the support-visible net. The inclusion "$\subseteq$" follows by construction. Applying the same argument to the inverse automorphism $\sigma_{-t}^{\omega_\infty^{\mathrm{geo},C}}$ gives the reverse inclusion, hence equality. ☐

Theorem 6.6 (Projective Markov replacement compatibility)**.** *Let ρ_n be a cofinal refinement family, and let $R_{n\to m}$ denote restriction to a fixed local collar model m. Suppose that for each fixed m*

$$I(A_m : D_m \mid B_m)_{R_{n\to m}\rho_n} \leq \varepsilon_{n,m}, \qquad \delta_m^{\mathrm{M}}(\varepsilon_{n,m}) \to 0 \quad (n \to \infty),$$

and suppose the finite-stage restriction maps commute with the physical quotient and normal-form maps up to the declared refinement-equivalence certificate. Then, after passing to a cofinal subnet, one can choose exact Markov replacements $\sigma_{n,m} \in \mathfrak{M}_m$ such that for every fixed $k \leq m$,

$$\|R_{m\to k}\sigma_{n,m} - \sigma_{n,k}\|_1 \to 0.$$

Hence the exact-Markov comparison states define a projective support-visible Markov comparison class in the scaling limit.

Proof. For each fixed m, Proposition 5.16 gives exact Markov replacements within $\delta_m^{\mathrm{M}}(\varepsilon_{n,m}) + o(1)$ of $R_{n\to m}\rho_n$. The exact Markov sets $\mathfrak{M}_m$ are compact on fixed collars, so a diagonal subnet has limits on every fixed m. Restriction maps are continuous and, by hypothesis, commute with the physical quotient and normal form up to the same refinement-equivalence certificate used by the cap-local test family. Therefore the restriction of the m-limit to a smaller fixed collar k agrees with the k-limit. Pulling the chosen approximants along the diagonal subnet gives the displayed asymptotic compatibility. ☐

Proposition 6.7 (Round-cap rigidity from surviving cut data)**.** *Let $C \subset S^2$ be a round cap, and let the cap pair be the support-visible extracted pair of Proposition 6.4. Assume the modular support map of Lemma 6.5 acts by orientation-preserving conformal maps on the support-visible cap geometry. The geometric cut observables descending through Theorem 6.2 determine the boundary circle and two ordered null endpoint classes. Any orientation-preserving conformal one-parameter subgroup preserving this data is $\lambda_C(\kappa s)$ for a unique normalization constant $\kappa > 0$.*

Proof. Choose a Möbius map carrying C to the upper half-plane and the ordered endpoint classes to 0 and ∞. The orientation-preserving conformal maps of the upper half-plane preserving 0, ∞, and the order are exactly $z \mapsto e^s z$, $s \in \mathbb{R}$. Transporting this subgroup back to C gives the cap-preserving hyperbolic subgroup. Only its parameter normalization remains, hence the factor $\kappa > 0$. ☐

Theorem 6.8 (Support-visible BW scaling theorem on the prime geometric subnet)**.** *Assume Axioms 3.1–3.4, Theorem 6.2, and the derived fixed-cutoff regulator/collar/consensus package established above. Assume also the support-readable net regularity of Lemma 6.5 and the branch condition that the resulting modular support maps act by orientation-preserving conformal maps on the support-visible cap geometry. For each OPH-realized observer-supporting refinement branch and each round cap $C \subset S^2$, the support-visible prime geometric cap net admits a weak-$*$/GNS scaling-limit cap pair*

$$(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C}).$$

Then:

(i) *At each finite regulator stage, the cap algebra is type I, the reduced cap state has a density matrix $\rho_C^{(\delta)}$, and its modular Hamiltonian*

$$K_C^{(\delta)} := -\log \rho_C^{(\delta)}$$

has nonadditive part confined to the shrinking collar up to the carried errors measured by $r_{\mathrm{FR}}(\varepsilon_\delta)$, the fixed-collar replacement modulus δ^{M}, and the regularized support-visible modular transport bound of Theorems 6.12 and 6.13. In particular, on each fixed collar model the physical collar state differs from its constructive recovered comparison state by $r_{\mathrm{FR}}(\varepsilon_\delta)$ in trace norm, and the support-visible modular matrix elements converge under the displayed cutoff schedule independently of the exact-Markov replacement sequence.

(ii) *For every bounded support-visible cap-local observable O, the regularized modular matrix elements converge in the local weak-$*$/GNS support-quotient topology supplied by Proposition 6.4. The convergence is controlled by the explicit error budget*

$$r_{\mathrm{FR}}(\varepsilon_\delta), \qquad \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta), \qquad \frac{\Delta_\delta}{a_\delta}, \qquad d_{m,\delta}a_\delta, \qquad \Delta_\delta|\log a_\delta|,$$

which vanishes under the stated Markov-replacement and cutoff schedule.

(iii) *On the support-visible scaling-limit cap pair, the modular automorphism group is geometric:*

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

Equivalently, the modular parameter t and the geometric cap-dilation parameter s are related by

$$s = 2\pi t.$$

(iv) *No separate cap-isotropy/*$\mathrm{SO}(2)$*-equivariance selector, finite-cell Lorentz-invariance premise, or unregularized full-algebra common floor is used in this theorem. The conformal support-map regularity is the explicit scaling-limit branch condition above; it is not a hidden finite-regulator Euclidean-regularity premise. The target algebra is the support-visible prime geometric subnet. Proposition 6.7 fixes the geometric cap subgroup on that extracted prime pair, $\lambda_C(s)$ is normalized by the null blow-up $v \mapsto e^{-s}v$, and the modular KMS normalization fixes the coefficient 2π.*

(v) *If the scaling-limit cap algebra happens to be type I, item* (iii) *may be written as the operator identity*

$$K_C = 2\pi B_C.$$

In the generic continuum case, where the scaling-limit cap algebra has left the regulator class and is expected in QFT examples to be non-type-I, the correct theorem statement is item (iii) *itself, and the geometric modular action is generally outer rather than inner.*

Consequently, replacing the finite-stage modular action by the geometric cap-dilation action incurs only the carried collar and support-visible regularization errors, and these vanish in the refinement limit. The theorem does not assert the false stronger statement that every off-support direction of the full finite matrix algebra has a refinement-uniform unregularized spectral floor, and it does not assert Lorentz covariance for record/pointer or interface-inert auxiliary registers outside the extracted geometric subnet.

Proof. Step 1: fixed-cutoff collar control. On the finite type-I regulator net, Propositions 5.19, 5.20, and 5.25 localize the modular defect to the shrinking collar and quantify the discrepancy between the physical collar state, its constructive recovered comparison state, and its exact-Markov reference by the carried errors $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η_δ^{M}.

Step 2: support-visible regularization. Proposition 6.11 shows that a refinement-uniform common floor on the full finite matrix algebra is unavailable in general. Theorem 6.12 supplies the needed replacement: for $K_a(\rho) = -\log(\rho + a\mathbf{1})$ and bounded support-visible collar observables, the modular matrix-element difference is bounded by

$$\|O\| \left(\frac{4\Delta_n}{a} + d_{m,\delta} a + 4\Delta_n |\log a| \right).$$

Choosing $a_n \downarrow 0$ with $\Delta_n / a_n \to 0$, $d_{m,\delta} a_n \to 0$, and $\Delta_n |\log a_n| \to 0$ makes the right-hand side vanish on every fixed local collar model. Theorems 6.6 and 6.13 then make the exact-Markov comparison family projectively compatible and replacement-independent on support-visible observables. Thus the theorem concerns the support-visible scaling limit rather than off-support full-algebra directions.

Step 3: scaling-limit cap-pair extraction. Theorem 6.2 fixes the target algebra as the overlap-generated prime geometric subnet, with overlap-trivial auxiliary/record factors removed. Proposition 6.4 supplies the support-visible weak-$*$/GNS cap pair from the transported cap-local test family, the projectively compatible marginal family, the asymptotic transport-equivalence certificate, and the regularized modular transport cutoff schedule.

Step 4: support-readable geometric identification on the prime subnet. Lemma 6.5 makes the modular automorphism group induce a support map on cap-local regions of the extracted prime geometric subnet. Under the stated conformal support-map branch condition, Proposition 6.7 identifies the only orientation-preserving conformal one-parameter subgroup preserving the surviving round-cap cut data as the standard cap-preserving hyperbolic subgroup up to normalization. Therefore there is a constant $\kappa_C > 0$ such that

$$\sigma_t^{\omega_\infty^{\mathrm{geo},C}} = \alpha_{\lambda_C(\kappa_C t)}.$$

Step 5: normalization. The modular group is KMS at inverse temperature 1 by definition. Since $\lambda_C(s)$ is normalized by unit null dilation under the blow-up $v \mapsto e^{-s} v$, the Bisognano–Wichmann/Rindler normalization fixes

$$\kappa_C = 2\pi,$$

hence

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

Step 6: operator versus automorphism form. If the limit cap algebra is type I, one can represent the modular automorphism group by a modular Hamiltonian $K_C = -\log \rho_C$ and recover $K_C =$

$2\pi B_C$. If the limit algebra leaves the regulator class, the modular automorphism group is well defined while the inner operator representative need not exist inside $\mathcal{A}_\infty^{\mathrm{geo,sv}}(C)$. In that case the automorphism identity is the full theorem statement. This is exactly the observer-facing content needed for Lorentz kinematics, the null half-sided modular bridge, and the local Einstein branch. □

Definition 6.9 (BW-branch observer-relative time reading)**.** *On the branch satisfying the hypotheses of Theorem 6.8, OPH uses the modular automorphism parameter t of the extracted cap pair $(\mathcal{A}_\infty^{\mathrm{geo}}(C), \omega_\infty^{\mathrm{geo},C})$ as that cap observer's relative time parameter. This is a declared BW-branch reading of the geometric modular flow derived in Theorem 6.8; it is not an additional proof that arbitrary operational clocks, global time, or the full problem of time have been derived from the axioms.*

Remark 6.10 (BW-side UV scaffold after the geometric-subnet split)**.** *Theorem 6.2 fixes the target algebra: the continuum Lorentz claim is asked only on the support-visible extracted prime geometric subnet, not on the full operational algebra. Record/pointer/interface-inert observables are outside that geometric subnet, and off-support directions that disappear in the limiting GNS support are not observer-facing data.*

The transported geometric cap-local system consists of the cap-local test family on $\mathcal{A}_n^{\mathrm{geo}}(C_n)$, the projectively compatible transported marginal family, and the asymptotic transport-equivalence certificate. On each fixed local collar model, the regularized estimate of Theorem 6.12 replaces the unavailable unregularized full-algebra floor. Proposition 6.4 gives the local GNS support quotient cap pair, Lemma 6.5 reads modular flow as a support map on cap-local regions, and Proposition 6.7 identifies the residual cap-preserving conformal freedom with the standard hyperbolic subgroup. This is the theorem-side content of the BW lift used below.

Proposition 6.11 (Recoverability is not modular geometry: common-floor collapse countermodel)**.** *Exact or asymptotically exact Markov recovery at finite cutoff does not imply the eventual common floor used in the unregularized BW/geometric cap-pair extraction. In $M_2(\mathbb{C})$, the faithful states*

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

have $\lambda_{\min}(\rho_n) \to 0$. Tensoring this family with any fixed finite exact-Markov collar factor gives a full-rank exact-Markov collar family, but no refinement-uniform lower spectral floor survives. On an off-diagonal matrix unit the modular generator carries a logarithmic gap of order n, so unregularized modular transport can fail even though every finite stage is faithful and Markov. Thus the full-algebra common-floor route is the wrong target; Theorem 6.12 supplies the support-visible replacement used by Theorem 6.8.

Theorem 6.12 (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, and $\Delta_n = \|\rho_n - \widehat{\rho}_n\|_1$. For $a > 0$, set $K_a(\rho) = -\log(\rho + a\mathbf{1})$. For every bounded support-visible collar observable O,*

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

If $a_n \downarrow 0$, $\Delta_n/a_n \to 0$, $d_{m,\delta}a_n \to 0$, and $\Delta_n|\log a_n| \to 0$, then the regularized support-visible modular matrix elements converge on that fixed collar model.

Proof. The integral representation of the operator logarithm gives the displayed Lipschitz control above the spectral cutoff a. Splitting the comparison into the visible support above a, the a-tail, and the trace-distance error gives the three terms in the bound, which vanish under the stated cutoff schedule. ☐

Theorem 6.13 (Support-visible modular convergence from controlled Markov collars)**.** *Let* $(A_\delta, B_\delta, D_\delta)$ *be a shrinking collar family around a cap cut. Suppose that after pullback to every fixed local collar model,*

$$r_{\mathrm{FR}}(\varepsilon_\delta) \to 0, \qquad \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta) \to 0,$$

and suppose the regularized cutoff schedule in Theorem 6.12 is satisfied. Then for every bounded support-visible geometric collar observable O*,*

$$\lim_{\delta\downarrow 0} \operatorname{Tr} \rho_\delta O\big(K_{a_\delta}(\rho_\delta) - K_{a_\delta}(\sigma_\delta)\big) = 0$$

for any projectively compatible exact-Markov replacement family σ_δ*. Consequently, the scaling-limit modular automorphism on the extracted geometric cap pair is independent of the particular exact-Markov replacement sequence used to audit the finite-stage calculation.*

Proof. Theorem 6.12 gives the displayed matrix element bound with $\Delta_\delta = \|\rho_\delta - \sigma_\delta\|_1$. The controlled Markov replacement hypotheses make $\Delta_\delta \to 0$ on each fixed collar model, and the chosen a_δ schedule sends

$$\Delta_\delta / a_\delta, \qquad d_{m,\delta} a_\delta, \qquad \Delta_\delta |\log a_\delta|$$

to zero. The projective compatibility theorem ensures that different compatible replacement choices determine the same local weak-$*$ limits after GNS quotienting of null directions. Therefore the support-visible cap modular automorphism emitted in the scaling limit is a property of the physical cap pair, not of an arbitrary replacement choice. ☐

Remark 6.14 (Support-visible closure status)**.** *The BW/geometric branch is closed at the support-visible regularized level. The theorem does not need, and does not claim, a full-algebra unregularized common floor on directions that collapse out of the limiting observer support. The observer-facing content is the automorphism statement of Theorem 6.8, which is exactly what the Lorentz, null-modular, and local Einstein branches use.*

Remark 6.15 (Stronger BW convergence statements not claimed)**.** *Theorem 6.8 proves convergence of regularized support-visible modular matrix elements on fixed collar models and identifies the resulting scaling-limit cap modular automorphism. It does not assert uniform-on-compact-time convergence of regularized modular automorphism groups, for example*

$$\sup_{|t|\le T} \left|\omega_n\Big(X_n^* [\sigma_{t,n}^{a_n}(O_n) - \alpha_{\lambda_C(2\pi t)}(O)_n] Y_n\Big)\right| \to 0,$$

without an additional equicontinuity, strong-resolvent, or Trotter-type lemma for the regularized modular groups. Nor does it invoke a black-box AQFT BW theorem from conformal-net axioms. Such an independent route would require verifying the relevant scaling-limit cap-net axioms, such as isotony, additivity, locality, Haag or wedge duality, standardness of the vacuum, positive-energy covariance, and suitable modular inclusions. Those would be useful validation routes, but they are extra analytic certificates rather than hidden premises of Theorem 6.8.

Corollary 6.16 (Lorentz Kinematics)**.** *Under the hypotheses of Theorem 6.8,*

$$\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1).$$

The cap modular flows are therefore the standard one-parameter Lorentz boost/dilation subgroups in the celestial-sphere realization, and the induced local kinematic group is the connected Lorentz group.

Proof. Orientation-preserving conformal maps of S^2 are exactly the Möbius transformations, so

$$\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1).$$

By Theorem 6.8, on the extracted prime geometric subnet the realized scaling-limit cap modular automorphism groups are the standard cap-preserving conformal dilations with the 2π normalization fixed internally. Hence the local kinematics induced by modular flow is the connected Lorentz group. □

6.2 The null modular bridge

The fixed-cutoff strip analysis begins from the transferred cut-center data and the inherited-strip Markov structure. At fixed regulator scale one obtains exact or controlled four-term strip additivity together with endpoint-Lipschitz control of the renormalized half-line family; on the same scaling-limit geometric-cap branch used in Theorem 6.8, the null half-line blow-up net then inherits geometric dilation and therefore half-sided modular inclusion. Standard Borchers–Wiesbrock theory [9] applies only after that derived half-sided modular pair is in hand, yielding positivity and unitary implementation of null translations. The downstream boundary is narrower: the separate interval-preserving projective branch is needed to transport the half-line result to bounded null intervals, and the tensor upgrade carries the standard null-invisible metric ambiguity. The half-line generator itself is identified below with the effective local null-stress charge.

The present subsection therefore establishes the following chain:

1. null-strip center transfer and the inherited left/right split hypothesis;
2. exact strip additivity on the exact inherited Markov model, and a carried defect operator on controlled exact-Markov replacements on one fixed strip model;
3. endpoint-Lipschitz matrix elements for the renormalized half-line family and the resulting weak tail generator;
4. derived half-sided modular inclusion on the null half-line blow-up net from the declared scaling-limit geometric-cap branch, and the resulting Borchers positive translation generator.

The only explicit downstream boundary retained below is the bounded-interval transport/projective branch together with the later tensor upgrade; the half-line generator/charge identification is proved inside the bridge itself. Lemma 6.23 states the positive null-translation generator explicitly: the Borchers unitary group is constructed on its Stone domain, positivity is theorem-level rather than implicit, and the half-line modular Hamiltonians are tied to it by explicit affine covariance and quadratic-form identities. The density-upgrade template of Theorem 6.21 is recorded separately and does not carry the operator-identification burden.

Proposition 6.17 (Null-strip center transfer and inherited split)**.** *Fix a null generator* Ω *and a regulated tripartition*

$$I_- = (v_1, v_2), \qquad J = (v_2, v_3), \qquad I_+ = (v_3, v_4),$$

with cuts

$$\Gamma_- := \{v = v_2\}, \qquad \Gamma_+ := \{v = v_3\}.$$

Assume the derived fixed-cutoff regulator/collar package established above, and assume that the two null cuts inherit the ordinary or central-defect boundary-redundancy data of Proposition 5.1 in the following precise sense:

(i) *each cut* $\Gamma_\pm$ *carries a compact derived boundary action* $\widehat{K}_\pm$ *on a reference cut Hilbert space, with* $\widehat{K}_\pm = K_\pm$ *on the ordinary branch;*

(ii) *in a compatible type-I regulator presentation, the adjacent regions carry inverse transport across each cut, so for irreducible* $\widehat{K}_\pm$*-modules* $W_{\alpha_\pm}$ *one has decompositions*

$$\tilde{\mathcal{H}}_{I_-} \cong \bigoplus_{\alpha_-} W_{\alpha_-} \otimes \mathcal{H}_{i_-^{\alpha_-}},$$

$$\tilde{\mathcal{H}}_J \cong \bigoplus_{\alpha_-,\alpha_+} W^*_{\alpha_-} \otimes \mathcal{H}_{j^{\alpha_-,\alpha_+}} \otimes W_{\alpha_+},$$

$$\tilde{\mathcal{H}}_{I_+} \cong \bigoplus_{\alpha_+} W^*_{\alpha_+} \otimes \mathcal{H}_{i_+^{\alpha_+}};$$

(iii) *the gauge-invariant strip algebra is the commutant of the transported boundary action on* $\tilde{\mathcal{H}}_J$*, so that the pair* (α_-, α_+) *is a central cut label;*

(iv) *for each* (α_-, α_+)*, the multiplicity space* $\mathcal{H}_{j^{\alpha_-,\alpha_+}}$ *admits a chosen factorization*

$$\mathcal{H}_{j^{\alpha_-,\alpha_+}} \cong \mathcal{H}_{j_L^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{j_R^{\alpha_-,\alpha_+}}$$

such that, blockwise, $\mathcal{A}(I_- \cup J)$ *acts only on* $\mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j_L^{\alpha_-,\alpha_+}}$ *and* $\mathcal{A}(J \cup I_+)$ *acts only on* $\mathcal{H}_{j_R^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{i_+^{\alpha_+}}$*.*

Then the strip algebra has the inherited central decomposition

$$\mathcal{A}(J) \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{B}(\mathcal{H}_{j^{\alpha_-,\alpha_+}}), \qquad Z(\mathcal{A}(J)) = \bigoplus_{\alpha_-,\alpha_+} \mathbb{C}\, P_{\alpha_-,\alpha_+},$$

and, under item (iv),

$$\mathcal{A}(J) \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{B}(\mathcal{H}_{j_L^{\alpha_-,\alpha_+}}) \otimes \mathcal{B}(\mathcal{H}_{j_R^{\alpha_-,\alpha_+}}).$$

Moreover the glued tripartition carries the block decomposition

$$\mathcal{H}_{I_- \cup J \cup I_+} \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j_L^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{j_R^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{i_+^{\alpha_+}}.$$

Thus the two null cuts transfer the same center data as the spatial collar branch, while the left/right split of the strip multiplicity spaces is exactly the extra decomposition-inheritance hypothesis and is not forced by center transfer alone.

Proof. By complete reducibility of the finite-dimensional unitary actions $\widehat{K}_\pm$, the displayed decompositions of $\tilde{\mathcal{H}}_{I_-}$, $\tilde{\mathcal{H}}_J$, and $\tilde{\mathcal{H}}_{I_+}$ exist. Under item (iii), the strip algebra is the commutant of the transported $\widehat{K}_- \times \widehat{K}_+$ action on

$$\tilde{\mathcal{H}}_J = \bigoplus_{\alpha_-,\alpha_+} W_{\alpha_-}^* \otimes \mathcal{H}_{j^{\alpha_-,\alpha_+}} \otimes W_{\alpha_+}.$$

Write an operator $X \in \mathcal{B}(\tilde{\mathcal{H}}_J)$ in matrix form relative to that direct sum. The block

$$X_{(\alpha_-,\alpha_+),(\beta_-,\beta_+)}$$

intertwines

$$W_{\alpha_-}^* \otimes W_{\alpha_+} \longrightarrow W_{\beta_-}^* \otimes W_{\beta_+}$$

for the product action of $\widehat{K}_- \times \widehat{K}_+$. By Schur's lemma, such an intertwiner is zero unless $(\alpha_-,\alpha_+) = (\beta_-,\beta_+)$, and when the sector pair agrees it is scalar on the representation factors. Therefore the commutant is exactly

$$\mathcal{A}(J) \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{B}(\mathcal{H}_{j^{\alpha_-,\alpha_+}}),$$

and the center is generated by the corresponding block projectors P_{α_-,α_+}.

For the glued tripartition, tensor the three regulator Hilbert spaces and decompose:

$$\tilde{\mathcal{H}}_{I_-} \otimes \tilde{\mathcal{H}}_J \otimes \tilde{\mathcal{H}}_{I_+} \cong \bigoplus_{\alpha_-,\beta_-,\alpha_+,\beta_+} (W_{\alpha_-} \otimes W_{\beta_-}^*) \otimes \mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j^{\beta_-,\beta_+}} \otimes (W_{\beta_+} \otimes W_{\alpha_+}^*) \otimes \mathcal{H}_{i_+^{\alpha_+}}.$$

Taking $\widehat{K}_- \times \widehat{K}_+$-invariants and applying Schur's lemma to each cut gives

$$(W_{\alpha_-} \otimes W_{\beta_-}^*)^{\widehat{K}_-} \cong \begin{cases} \mathbb{C}, & \alpha_- = \beta_-, \\ 0, & \alpha_- \neq \beta_-, \end{cases} \qquad (W_{\beta_+} \otimes W_{\alpha_+}^*)^{\widehat{K}_+} \cong \begin{cases} \mathbb{C}, & \beta_+ = \alpha_+, \\ 0, & \beta_+ \neq \alpha_+, \end{cases}$$

so only matching cut sectors survive and one obtains

$$\mathcal{H}_{I_- \cup J \cup I_+} \cong \bigoplus_{\alpha_-,\alpha_+} \mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{i_+^{\alpha_+}}.$$

If item (iv) holds, substitute the chosen factorization

$$\mathcal{H}_{j^{\alpha_-,\alpha_+}} \cong \mathcal{H}_{j_L^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{j_R^{\alpha_-,\alpha_+}}$$

to obtain the displayed inherited split and the stated action of the left and right union algebras. Nothing in the Schur-lemma argument itself forces item (iv); it is exactly the extra structural condition required to reproduce the same blockwise tensor pattern as Theorem 5.2. □

Corollary 6.18 (Exact or controlled four-term null modular relation on an inherited strip model)**.** *Fix one finite-dimensional regulated strip model satisfying Proposition 6.17, and define*

$$\mathfrak{M}_{I_-:J:I_+} := \{\sigma : \ I(I_- : I_+ \mid J)_\sigma = 0\},$$

$$\delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon) := \sup\left\{ \inf_{\sigma \in \mathfrak{M}_{I_-:J:I_+}} \|\rho - \sigma\|_1 : I(I_- : I_+ \mid J)_\rho \leq \varepsilon \right\}.$$

For any state η on this strip model, write

$$\Delta K_J(\eta) := K_{I_- \cup J \cup I_+}(\eta) - K_{I_- \cup J}(\eta) - K_{J \cup I_+}(\eta) + K_J(\eta).$$

If the reference state ω is exact Markov on this inherited decomposition, then

$$\Delta K_J(\omega) \in Z(\mathcal{A}(J)),$$

and in the canonical blockwise identification of Proposition 6.17 one may take

$$\Delta K_J(\omega) = 0.$$

Equivalently, without fixing that normalization there exists a central operator

$$K_{\partial,J}(\omega) \in Z(\mathcal{A}(J))$$

such that

$$K_{I_- \cup J \cup I_+}(\omega) = K_{I_- \cup J}(\omega) + K_{J \cup I_+}(\omega) - K_J(\omega) + K_{\partial,J}(\omega).$$

If instead $I(I_- : I_+ \mid J)_\omega \le \varepsilon$, then on this fixed strip model one may choose an exact Markov state $\widetilde{\omega}_J \in \mathfrak{M}_{I_-:J:I_+}$ with

$$\|\omega - \widetilde{\omega}_J\|_1 \le \delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon).$$

Define the carried defect operator

$$\mathfrak{D}_J(\omega, \widetilde{\omega}_J) := \Delta K_J(\omega) - \Delta K_J(\widetilde{\omega}_J).$$

Then

$$K_{I_- \cup J \cup I_+}(\omega) = K_{I_- \cup J}(\omega) + K_{J \cup I_+}(\omega) - K_J(\omega) + K_{\partial,J}(\widetilde{\omega}_J) + \mathfrak{D}_J(\omega, \widetilde{\omega}_J),$$

where

$$K_{\partial,J}(\widetilde{\omega}_J) := \Delta K_J(\widetilde{\omega}_J) \in Z(\mathcal{A}(J)).$$

For every bounded observable O in $\mathcal{A}(I_- \cup J)$ or $\mathcal{A}(J \cup I_+)$,

$$\big|\omega(O) - \widetilde{\omega}_J(O)\big| \le \|O\|_\infty \, \delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon).$$

If the marginals entering ΔK_J for ω and $\widetilde{\omega}_J$ are uniformly faithful with lower spectral bound $\lambda_* > 0$, then

$$\|\mathfrak{D}_J(\omega, \widetilde{\omega}_J)\|_\infty \le 4\lambda_*^{-1} \, \delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon).$$

Independently, the Fawzi–Renner recovery map supplies a constructive comparison state with trace-norm error

$$r_{\mathrm{FR}}(\varepsilon) = 2\sqrt{1 - e^{-\varepsilon}} \le 2\sqrt{\varepsilon}.$$

Hence the exact four-term strip relation is available only at exact Markovity, or along a controlled strip family on one fixed inherited strip model for which

$$\delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon_J) \to 0.$$

In the controlled case, $\mathfrak{D}_J$ is carried explicitly at each finite stage.

Proof. Assume first that ω is exact Markov on the inherited split of Proposition 6.17. Then the HJPW structure theorem gives, on each sector pair $\alpha = (\alpha_-, \alpha_+)$,

$$\omega_{I_-JI_+}|_\alpha = p_\alpha\, \omega^{(\alpha)}_{I_-j_L} \otimes \omega^{(\alpha)}_{j_R I_+}.$$

Accordingly,

$$K_{I_-\cup J\cup I_+}(\omega)\big|_\alpha = (-\log p_\alpha)\mathbf{1} + K^{(\alpha)}_{I_-j_L} \otimes \mathbf{1}_{j_R I_+} + \mathbf{1}_{I_-j_L} \otimes K^{(\alpha)}_{j_R I_+}.$$

Tracing over I_+, I_-, or both gives the marginal block formulas

$$K_{I_-\cup J}(\omega)\big|_\alpha = (-\log p_\alpha)\mathbf{1} + K^{(\alpha)}_{I_-j_L} \otimes \mathbf{1}_{j_R} + \mathbf{1}_{I_-j_L} \otimes K^{(\alpha)}_{j_R},$$

$$K_{J\cup I_+}(\omega)\big|_\alpha = (-\log p_\alpha)\mathbf{1} + K^{(\alpha)}_{j_L} \otimes \mathbf{1}_{j_R I_+} + \mathbf{1}_{j_L} \otimes K^{(\alpha)}_{j_R I_+},$$

$$K_J(\omega)\big|_\alpha = (-\log p_\alpha)\mathbf{1} + K^{(\alpha)}_{j_L} \otimes \mathbf{1}_{j_R} + \mathbf{1}_{j_L} \otimes K^{(\alpha)}_{j_R}.$$

Subtracting the last three expressions from the first gives zero on every block. Thus, in the canonical blockwise identification supplied by Proposition 6.17,

$$\Delta K_J(\omega) = 0.$$

If one chooses to keep the blockwise endpoint-label bookkeeping explicit rather than absorbing those constants into the canonical identification, the remainder is a direct sum of block constants and is therefore central. This yields the stated central form with $K_{\partial,J}(\omega) \in Z(\mathcal{A}(J))$.

For the controlled statement, apply Proposition 5.16 on this fixed finite-dimensional strip model, relabeling A, B, D there as I_-, J, I_+. This yields an exact Markov replacement $\widetilde{\omega}_J \in \mathfrak{M}_{I_-:J:I_+}$ with the displayed trace-norm bound. The identity

$$\Delta K_J(\omega) = \Delta K_J(\widetilde{\omega}_J) + \mathfrak{D}_J(\omega, \widetilde{\omega}_J)$$

is tautological, so the displayed four-term formula follows from the exact-Markov case applied to $\widetilde{\omega}_J$. The observable estimate is immediate from

$$\big|\omega(O) - \widetilde{\omega}_J(O)\big| \le \|O\|_\infty\, \|\omega - \widetilde{\omega}_J\|_1.$$

Finally, Proposition 5.20, again relabeled $A, B, D \mapsto I_-, J, I_+$, yields the operator-norm bound on $\mathfrak{D}_J$ whenever the relevant marginals have the stated lower spectral bound. The Fawzi–Renner estimate is the same one-shot recoverability bound used elsewhere and is independent of the exact-Markov replacement modulus. ☐

Definition 6.19 (Renormalized null modular functional)**.** *For a null interval I on generator Ω, let $K_\partial(I, \Omega)$ denote the central endpoint-label term singled out by Corollary 6.18 on the inherited strip model, or by its controlled exact-Markov replacement when that model is used as reference. Define*

$$\widetilde{K}[I, \Omega] := K[I, \Omega] - K_\partial(I, \Omega).$$

For a null half-line $H_a := (a, \infty)$, abbreviate

$$\widetilde{K_a}(\Omega) := \widetilde{K}[H_a, \Omega].$$

Proposition 6.20 (Endpoint-Lipschitz null modular families and weak tail generator)**.** *Under the derived quasi-local propagation and bounded-interval endpoint-Lipschitz control internal to Axiom 3.3, the renormalized interval family obeys the following matrix-element bounds on every fixed local-energy-bounded domain. For bounded intervals*

$$I = (a, b), \qquad I' = (a', b'),$$

contained in one compact endpoint window,

$$|\langle \psi, (\widetilde{K}[I', \Omega] - \widetilde{K}[I, \Omega])\phi\rangle| \le C_{\psi,\phi,\Omega} \left(|a' - a| + |b' - b|\right).$$

In particular the matrix elements of $\widetilde{K}[I, \Omega]$ are jointly Lipschitz in the two endpoints and therefore have finite variation there.

For null half-lines $H_a = (a, \infty)$, one likewise has

$$|\langle \psi, (\widetilde{K}_{a'}(\Omega) - \widetilde{K}_a(\Omega))\phi\rangle| \le C_{\psi,\phi,\Omega} \, |a' - a|$$

for a, a' in every compact window. Hence, for fixed ψ, ϕ,

$$f_{\psi,\phi}(a) := \langle \psi, \widetilde{K}_a(\Omega)\phi\rangle$$

is locally Lipschitz and therefore absolutely continuous. Its distributional derivative defines a locally L^∞ sesquilinear-form density

$$\langle \psi, q(a, \Omega)\phi\rangle := -\frac{1}{2\pi} \, \partial_a f_{\psi,\phi}(a),$$

and for $a < b$,

$$\langle \psi, (\widetilde{K}_b(\Omega) - \widetilde{K}_a(\Omega))\phi\rangle = -2\pi \int_a^b \langle \psi, q(v, \Omega)\phi\rangle \, dv.$$

The object $q(a, \Omega)$ is the weak tail generator rather than a local operator-valued density; its identification with the positive self-adjoint Borchers generator occurs only after the derived half-sided modular inclusion of Corollary 6.22 and Lemma 6.23.

Proof. The bounded-interval endpoint-control statement derived from Axiom 3.3 applies precisely to the renormalized family obtained after removal of the central endpoint term. Therefore, for intervals contained in one fixed compact endpoint window,

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

where $I' \Delta I$ is the symmetric difference of the two intervals. If

$$I = (a, b), \qquad I' = (a', b'),$$

then

$$|I' \Delta I| \le |a' - a| + |b' - b|,$$

which gives the displayed joint endpoint-Lipschitz bound.

For half-lines $H_a = (a, \infty)$ and $H_{a'} = (a', \infty)$, the symmetric difference is the bounded interval between the two endpoints, so

$$|H_{a'} \Delta H_a| = |a' - a|.$$

Because the central endpoint-label term has been removed in the definition of $\widetilde{K}_a$, the same endpoint-control estimate yields

$$|\langle \psi, (\widetilde{K}_{a'}(\Omega) - \widetilde{K}_a(\Omega))\phi \rangle| \le C_{\psi,\phi,\Omega}\, |a' - a|.$$

Hence, on every compact a-interval, the function

$$f_{\psi,\phi}(a) := \langle \psi, \widetilde{K}_a(\Omega)\phi \rangle$$

is Lipschitz. A Lipschitz function on a compact interval is absolutely continuous, so it has an almost-everywhere derivative in L^∞ and obeys the fundamental theorem of calculus:

$$f_{\psi,\phi}(b) - f_{\psi,\phi}(a) = \int_a^b \partial_v f_{\psi,\phi}(v)\, dv.$$

Define

$$\langle \psi, q(v, \Omega)\phi \rangle := -\frac{1}{2\pi}\, \partial_v f_{\psi,\phi}(v)$$

distributionally. Substituting this definition into the previous identity gives

$$\langle \psi, (\widetilde{K}_b(\Omega) - \widetilde{K}_a(\Omega))\phi \rangle = -2\pi \int_a^b \langle \psi, q(v, \Omega)\phi \rangle\, dv.$$

This is exactly the claimed weak-tail formula. □

Theorem 6.21 (Downstream density-upgrade template)**.** *Fix a null generator Ω and the renormalized half-line family $\widetilde{K}_a(\Omega)$ of Proposition 6.20. For fixed vectors ψ, ϕ in a common dense domain, let*

$$f_{\psi,\phi}(a) := \langle \psi, \widetilde{K}_a(\Omega)\phi \rangle, \qquad q_{\psi,\phi}(a) := \langle \psi, q(a, \Omega)\phi \rangle = -\frac{1}{2\pi}\, \partial_a f_{\psi,\phi}(a).$$

Assume in addition that $q_{\psi,\phi}$ is weakly differentiable in a, that

$$\lim_{a\to+\infty} f_{\psi,\phi}(a) = 0, \qquad \lim_{a\to+\infty} q_{\psi,\phi}(a) = 0,$$

and define distributionally

$$\langle \psi, p(a, \Omega)\phi \rangle := -\partial_a q_{\psi,\phi}(a) = \frac{1}{2\pi}\, \partial_a^2 f_{\psi,\phi}(a).$$

Then

$$q_{\psi,\phi}(a) = \int_a^\infty \langle \psi, p(v, \Omega)\phi \rangle\, dv,$$

and

$$\langle \psi, \widetilde{K}_a(\Omega)\phi \rangle = 2\pi \int_a^\infty (v - a)\, \langle \psi, p(v, \Omega)\phi \rangle\, dv.$$

This statement is not part of the present fixed-cutoff bridge; it records the exact additional input/output package carried into the downstream density-upgrade task.

Null half-line blow-up net. Fix a smooth cut point and choose affine coordinate v on generator Ω so that the cut sits at $v = 0$. For $a \geq 0$, write

$$H_a := (a, \infty), \qquad \mathcal{M}_a(\Omega) := \overline{\bigvee_{a<c<d<\infty} \mathcal{A}((c,d),\Omega)}.$$

Thus $\mathcal{M}_a(\Omega)$ is the half-line blow-up algebra generated by bounded null interval algebras inside H_a. By isotony of the null interval net,

$$a \leq b \quad \Longrightarrow \quad \mathcal{M}_b(\Omega) \subseteq \mathcal{M}_a(\Omega).$$

Corollary 6.22 (Derived half-sided modular inclusion on null half-lines)**.** *On the scaling-limit geometric-cap branch of Theorem 6.8, the blow-up modular action near a smooth entangling cut acts on the null coordinate by*

$$v \mapsto e^{-2\pi t} v.$$

Therefore, for every $a \geq 0$,

$$\sigma_t^\omega(\mathcal{M}_a(\Omega)) = \mathcal{M}_{e^{-2\pi t}a}(\Omega).$$

Hence, for every $a > 0$,

$$\sigma_t^\omega(\mathcal{M}_a(\Omega)) \subseteq \mathcal{M}_a(\Omega) \qquad (t \leq 0),$$

so the inclusion

$$\mathcal{M}_a(\Omega) \subset \mathcal{M}_0(\Omega)$$

is half-sided modular. Equivalently, after the harmless convention change $t \mapsto -t$, one recovers the standard positive-time half-sided-inclusion form. The half-sided-inclusion step is therefore derived from the null-interval structure, isotony, and the scaling-limit geometric action rather than imported as an additional modular-QFT ingredient.

Proof. By Theorem 6.8, on every controlled collar family whose carried remainder vanishes in the refinement limit, the cap modular flow converges to the exact geometric cap dilation with 2π normalization. Blow up the cap near a smooth cut and restrict to the chosen null generator. In that tangent limit, the cap-preserving flow becomes the null dilation $v \mapsto e^{-2\pi t} v$. For every bounded interval $(c, d) \subset H_a$, the blow-up action therefore sends

$$\mathcal{A}((c,d),\Omega) \longmapsto \mathcal{A}((e^{-2\pi t}c, e^{-2\pi t}d),\Omega).$$

Taking the von Neumann closure of the algebra generated by all such intervals gives

$$\sigma_t^\omega(\mathcal{M}_a(\Omega)) = \mathcal{M}_{e^{-2\pi t}a}(\Omega).$$

If $t \leq 0$, then $e^{-2\pi t}a \geq a$, hence

$$H_{e^{-2\pi t}a} \subseteq H_a.$$

Isotony of the null interval net therefore implies

$$\mathcal{M}_{e^{-2\pi t}a}(\Omega) \subseteq \mathcal{M}_a(\Omega).$$

This is exactly the half-sided modular-inclusion property for the pair $\mathcal{M}_a(\Omega) \subset \mathcal{M}_0(\Omega)$. Reparametrizing the modular parameter by $t \mapsto -t$ gives the more common positive-time convention if desired. □

Lemma 6.23 (Positive null-translation generator)**.** *For the derived half-sided modular inclusion*

$$\mathcal{M}_a(\Omega) \subset \mathcal{M}_0(\Omega) \qquad (a > 0)$$

of Corollary 6.22, let $\Delta_0(\Omega)$ *be the modular operator of the standard pair* $(\mathcal{M}_0(\Omega), \omega)$ *and define*

$$K_0(\Omega) := -\log \Delta_0(\Omega).$$

Then Borchers–Wiesbrock yields a unique strongly continuous one-parameter unitary group

$$U_\Omega(a) = e^{iaP_\Omega}, \qquad a \in \mathbb{R},$$

such that

$$U_\Omega(a)\omega = \omega, \qquad U_\Omega(a)\, \mathcal{M}_b(\Omega)\, U_\Omega(a)^* = \mathcal{M}_{a+b}(\Omega) \quad (a, b \geq 0),$$

and whose generator P_Ω *is positive and self-adjoint on the Stone domain*

$$D(P_\Omega) := \left\{ \psi \in \mathcal{H} : \lim_{a \to 0} \frac{U_\Omega(a)\psi - \psi}{ia} \ \textit{exists} \right\}.$$

Moreover,

$$\Delta_0(\Omega)^{it} U_\Omega(a) \Delta_0(\Omega)^{-it} = U_\Omega(e^{-2\pi t}a), \qquad \Delta_0(\Omega)^{it} P_\Omega \Delta_0(\Omega)^{-it} = e^{-2\pi t} P_\Omega,$$

and on the common invariant analytic core of $K_0(\Omega)$ *and* P_Ω*,*

$$[K_0(\Omega), P_\Omega] = -\, i\, 2\pi P_\Omega.$$

For each $a \geq 0$*, let* $\Delta_a(\Omega)$ *be the modular operator of the translated standard pair* $(\mathcal{M}_a(\Omega), \omega)$ *and define*

$$K_a(\Omega) := -\log \Delta_a(\Omega).$$

Then

$$\Delta_a(\Omega) = U_\Omega(a)\Delta_0(\Omega)U_\Omega(a)^*, \qquad K_a(\Omega) = U_\Omega(a)K_0(\Omega)U_\Omega(a)^*,$$

hence

$$K_a(\Omega) = K_0(\Omega) - 2\pi a\, P_\Omega$$

as a quadratic-form identity on $D(K_0(\Omega)) \cap D(P_\Omega)$*. Equivalently,*

$$\langle \psi, (K_b(\Omega) - K_a(\Omega))\phi \rangle = -2\pi(b - a)\, \langle \psi, P_\Omega \phi \rangle$$

for all $\psi, \phi \in D(K_0(\Omega)) \cap D(P_\Omega)$*.*

In the canonical blockwise normalization of Corollary 6.18, where the half-line endpoint term has been absorbed into the central part, the weak endpoint derivative of Proposition 6.20 agrees with the Borchers generator:

$$\langle \psi, P_\Omega \phi \rangle = -\frac{1}{2\pi} \frac{d}{da}\Big|_{a=0^+} \langle \psi, \widetilde{K}_a(\Omega)\phi \rangle$$

for the same domain vectors. The half-sided-inclusion step proves only this affine half-line pair $(K_0(\Omega), P_\Omega)$*; bounded-interval modular-Hamiltonian formulas are downstream consequences of the separate interval-preserving projective branch recorded in Theorem 6.24.*

Proof. Corollary 6.22 gives the required half-sided modular inclusion. The standard Borchers–Wiesbrock theorem for a standard half-sided inclusion [9] then yields a unique strongly continuous unitary group $U_\Omega(a)$ with $U_\Omega(a)\omega = \omega$, the translation law for the nested half-line net, and a positive self-adjoint generator P_Ω; Stone's theorem gives the displayed domain formula.

The affine covariance relation

$$\Delta_0(\Omega)^{it} U_\Omega(a) \Delta_0(\Omega)^{-it} = U_\Omega(e^{-2\pi t} a)$$

is the Borchers relation. Differentiating at $a = 0$ gives

$$\Delta_0(\Omega)^{it} P_\Omega \Delta_0(\Omega)^{-it} = e^{-2\pi t} P_\Omega.$$

Since $\Delta_0(\Omega)^{it} = e^{-itK_0(\Omega)}$, differentiating at $t = 0$ on the common analytic core yields

$$-i[K_0(\Omega), P_\Omega] = -2\pi P_\Omega,$$

hence

$$[K_0(\Omega), P_\Omega] = -\, i\, 2\pi P_\Omega.$$

Because $U_\Omega(a)\omega = \omega$ and

$$\mathcal{M}_a(\Omega) = U_\Omega(a) \mathcal{M}_0(\Omega) U_\Omega(a)^*,$$

the modular operator of the translated standard pair is

$$\Delta_a(\Omega) = U_\Omega(a) \Delta_0(\Omega) U_\Omega(a)^*,$$

so

$$K_a(\Omega) = U_\Omega(a) K_0(\Omega) U_\Omega(a)^*.$$

Differentiating this identity in a on $D(K_0(\Omega)) \cap D(P_\Omega)$ gives

$$\frac{d}{da} K_a(\Omega) = i[P_\Omega, K_a(\Omega)] = -2\pi P_\Omega,$$

because the commutator relation just proved is invariant under conjugation by $U_\Omega(a)$. Integrating from 0 to a yields the displayed quadratic-form identity, and taking matrix elements gives the equivalent form. Proposition 6.20 had produced the weak endpoint derivative of the renormalized half-line family; in the canonical normalization of Corollary 6.18 the central term has been removed, so that weak derivative is exactly the Borchers generator. The final sentence is the explicit claim boundary: bounded intervals require the additional projective interval action not supplied by half-sided inclusion alone. □

Theorem 6.24 (The half-line generator is the local null-stress charge)**.** *Fix a null generator* Ω *and work in the canonical blockwise normalization of Corollary 6.18. Let* $\widetilde{K}_a(\Omega)$ *denote the renormalized modular Hamiltonian of the half-line* H_a*, and let* P_Ω *be the positive Borchers generator of Lemma 6.23. Then on*

$$D_\Omega := D(K_0(\Omega)) \cap D(P_\Omega)$$

one has

$$\langle \psi, P_\Omega \phi \rangle = -\frac{1}{2\pi} \frac{d}{da}\Big|_{a=0^+} \langle \psi, \widetilde{K}_a(\Omega) \phi \rangle \qquad (\psi, \phi \in D_\Omega),$$

and this operator is exactly the local null-stress charge of the effective spacetime description on that same half-line family.

Equivalently, if the effective spacetime description writes the same renormalized family as

$$\widetilde{K}_a^{\mathrm{eff}}(\Omega) = 2\pi\, Q_a^{\mathrm{mod}}(\Omega) + K_\partial(a,\Omega),$$

with $Q_a^{\mathrm{mod}}(\Omega)$ *the standard local null modular-energy term and* $K_\partial(a,\Omega)$ *central and endpoint-supported, then*

$$Q_{kk}(\Omega) := -\frac{1}{2\pi}\frac{d}{da}\Big|_{a=0^+}\widetilde{K}_a^{\mathrm{eff}}(\Omega)$$

satisfies

$$P_\Omega = Q_{kk}(\Omega)$$

as a quadratic-form identity on D_Ω.

Hence the renormalized half-line null generator is not a separate EFT-side input: it is the same operator that the effective spacetime description calls the local null-stress charge of the null half-line.

Proof. Lemma 6.23 proves the first displayed identity: in the canonical normalization of Corollary 6.18, the weak endpoint derivative of the renormalized half-line family is exactly the Borchers generator. Passing to the effective spacetime description does not change the operator family; it rewrites the same $\widetilde{K}_a(\Omega)$ in local continuum variables appropriate to the local Lorentzian regime. In that description, the local null-stress charge is precisely the right endpoint derivative

$$-\frac{1}{2\pi}\frac{d}{da}\Big|_{a=0^+}\widetilde{K}_a^{\mathrm{eff}}(\Omega)$$

of the same half-line modular family. Because $\widetilde{K}_a^{\mathrm{eff}}(\Omega)$ and $\widetilde{K}_a(\Omega)$ represent the same renormalized modular Hamiltonians, their right derivatives agree on D_Ω. Therefore the OPH generator and the effective null-stress charge coincide:

$$P_\Omega = Q_{kk}(\Omega).$$

The endpoint-supported central term does not alter the noncentral generator; in the canonical blockwise normalization it is absorbed into the central part fixed in Corollary 6.18. This proves the identification. □

Imported input and downstream boundary. The only continuum input used above is the standard local modular-Hamiltonian form on the effective description of that same half-line family. It is used only to name, in continuum language, the operator fixed by the OPH half-line derivative. There is no separate "null-stress identification" assumption. The downstream boundary is only:

(i) transport of the half-line statement to bounded null intervals, which requires the separate interval-preserving projective branch and its affine-covariant kernel $g_I(v)$; and

(ii) reconstruction of a full local symmetric tensor from the directional charges, which is determined only up to the null-invisible metric term ϕg_{ab}.

So the gap closed here is the generator/charge identification itself: inside the null bridge and on the declared half-line family, P_Ω is exactly the local null-stress charge.

Lemma 6.25 (Null data determine the stress tensor modulo a metric term)**.** *Let* X_{ab} *be a symmetric tensor. If*

$$X_{ab}k^a k^b = 0$$

for every null vector k *at a point, then*

$$X_{ab} = \phi g_{ab}$$

for some scalar ϕ.

Proof. Write

$$X_{ab} = Y_{ab} + \phi g_{ab}$$

with Y_{ab} traceless. Since $g_{ab}k^a k^b = 0$ on null vectors, the hypothesis implies $Y_{ab}k^a k^b = 0$ for all null k. In local inertial coordinates let $k = (1, \hat{n})$ with $|\hat{n}| = 1$. Then

$$Y_{ab}k^a k^b = Y_{00} + 2\hat{n}^i Y_{0i} + \hat{n}^i \hat{n}^j Y_{ij}.$$

Oddness under $\hat{n} \mapsto -\hat{n}$ forces $Y_{0i} = 0$. Averaging that relation over S^2 gives

$$Y_{00} + \frac{1}{3}\delta^{ij} Y_{ij} = 0.$$

Tracelessness of Y_{ab} gives

$$-Y_{00} + \delta^{ij} Y_{ij} = 0,$$

so both Y_{00} and the spatial trace $\delta^{ij}Y_{ij}$ vanish. The condition is therefore

$$\hat{n}^i \hat{n}^j Y_{ij}^{\mathrm{TF}} = 0 \qquad \text{for all } \hat{n} \in S^2,$$

where Y_{ij}^{TF} is the traceless spatial part. The $\ell = 2$ spherical harmonics are linearly independent, so this implies $Y_{ij}^{\mathrm{TF}} = 0$. Hence $Y_{ab} = 0$ and therefore $X_{ab} = \phi g_{ab}$. □

Remark 6.26 (Claim boundary in the null bridge)**.** *The c.12-local bridge ends with explicit theorem-level objects: the positive self-adjoint null-translation generator P_Ω on its Stone domain, its Borchers affine covariance relation, the half-line modular identities above, and the exact half-line generator/charge identification of Theorem 6.24. The downstream boundary is transport to bounded null intervals through the separate interval-preserving projective branch and the tensor upgrade modulo the null-invisible metric term. The density-upgrade template of Theorem 6.21 is recorded separately and does not carry the operator-identification burden.*

6.3 Generalized entropy and the Einstein equation

From this point through Corollary 6.33 we specialize to the physical $d = 4$ scaling branch. The claim boundary is explicit. The only standard geometric input used below, beyond D3 and D4, is the fixed-volume area-variation identity for a small geodesic ball. The small-ball modular kernel used below is *not* imported from a separate EFT law: it is the local expression of the geometric cap generator B_C from Theorem 6.8, evaluated using the half-line null-stress identification of Theorem 6.24 together with the separate interval-preserving projective branch that transports it to bounded null intervals.

For a reference state ω and a small cap C,

$$\delta S_C = \delta \langle K_C \rangle.$$

By Theorem 6.8,

$$\delta S_C = 2\pi\, \delta \langle B_C \rangle.$$

By Proposition 5.25,

$$S_{\mathrm{gen}}(C) = \mathrm{Tr}(\rho L_C) + S_{\mathrm{bulk}}(C).$$

Together with Definition 5.26, this identifies the area term as the coarse-grained form of the edge entropy density rather than an independent postulate. After this edge-center split, the first-law piece entering the small-ball argument is the bulk contribution:

$$\delta S_{\mathrm{bulk}}(C) = 2\pi\, \delta \langle B_C \rangle.$$

Definition 6.27 (Admissible fixed-cap MaxEnt variation)**.** *Fix a cap C and a realized reference state ω_C on the cap-reduced realized MaxEnt branch. An admissible fixed-cap MaxEnt variation is a first-order variation*

$$\delta\rho_C = \frac{d}{ds}\Big|_{s=0}\rho_C(s)$$

along a C^1 curve $s \mapsto \rho_C(s)$ such that:

1. *every $\rho_C(s)$ lies in the same cap-label-preserving block class of Proposition 5.25, so the same edge-center operator L_C and sector projectors P_α apply;*
2. *the cap C, its size/volume, and the optional conserved charges Q_b are held fixed; and*
3. *the Axiom 3.3 constraint values are fixed to first order:*

$$\frac{d}{ds}\Big|_{s=0}\mathrm{Tr}(\rho_C(s)O_a(x)) = 0, \qquad \frac{d}{ds}\Big|_{s=0}\mathrm{Tr}(\rho_C(s)Q_b) = 0$$

for every retained local constraint $O_a(x)$ supported in the fixed cap neighborhood and every retained conserved charge Q_b.

Shape or null-cut deformations are denoted separately by δ_{shape} or ∂_λ; they are not part of this fixed-cap variation class.

Theorem 6.28 (Derived fixed-cap generalized-entropy stationarity)**.** *Let C be a fixed cap and let ω_C be a realized reference state on the cap-reduced realized MaxEnt branch of Axiom 3.3. For every admissible fixed-cap MaxEnt variation δ of Definition 6.27,*

$$\delta S_{\text{gen}}(C) = 0.$$

Equivalently, the realized reference state is stationary for generalized entropy on the allowed fixed-cap variation class selected by the realized MaxEnt family.

Proof. Axiom 3.3 selects the realized branch by entropy maximization on the finite-dimensional constraint surface cut out by the retained local constraints and optional conserved charges. Definition 6.27 restricts to first-order tangent directions that stay on that same fixed-cap constraint surface. Hence the first variation of the cap entropy vanishes at the realized reference state:

$$\delta S(\rho_C)\big|_{\omega_C} = 0.$$

For the same cap-label-preserving block class, Proposition 5.25 gives

$$S(\rho_C) = S_{\text{bulk}}(C) + \mathrm{Tr}(\rho_C L_C) = S_{\text{gen}}(C).$$

Differentiating this identity along the admissible curve yields

$$\delta S_{\text{gen}}(C) = \delta S(\rho_C)\big|_{\omega_C} = 0.$$

□

Remark 6.29 (Claim boundary of the fixed-cap generalized-entropy stationarity theorem)**.** *Theorem 6.28 internalizes only the fixed-cap, cap-label-preserving, constraint-preserving first-order variation class selected by Axiom 3.3. It does not claim generalized-entropy stationarity for arbitrary shape deformations, arbitrary null-cut deformations, or arbitrary off-branch perturbations outside that realized MaxEnt family.*

Lemma 6.30 (Internal d=4 small-ball bridge)**.** *Assume the hypotheses of Theorem 6.24, together with the separate interval-preserving projective branch invoked there for bounded intervals. Let p be the cap center and let u^a be the future-directed unit tangent of the local diamond rest frame at p. In local inertial coordinates (t, x^i) adapted to u^a, let D_ℓ be the causal diamond whose $t = 0$ slice is the Euclidean ball*

$$B_\ell = \{t = 0,\ r < \ell\}, \qquad r^2 = \delta_{ij} x^i x^j.$$

Then

$$\delta S_{\text{bulk}}(C) = 2\pi \int_{B_\ell} \frac{\ell^2 - r^2}{2\ell}\, \delta\langle T_{00}\rangle\, d^3x + \delta\langle E^{(\eta)}_{C,\ell}\rangle.$$

If, moreover, $\delta\langle T_{00}\rangle$ is approximately constant across B_ℓ and

$$\delta\langle E^{(\eta)}_{C,\ell}\rangle = o(\ell^4)$$

along the same scaling family, then in $d = 4$,

$$\delta S_{\text{bulk}}(C) = \frac{8\pi^2 \ell^4}{15}\, \delta\langle T_{00}\rangle + O(\ell^5 \partial T) + o(\ell^4).$$

Proof. Work in the local Lorentzian scaling regime around p. On the geometric-subnet branch of Theorem 6.8, the cap modular flow is the geometric flow of the diamond-preserving conformal Killing field. In the tangent diamond D_ℓ, that field is

$$\xi_{D_\ell} = \frac{1}{2\ell}\Big((\ell^2 - r^2 - t^2)\, \partial_t - 2t\, x^i \partial_i\Big),$$

so on the $t = 0$ slice one has

$$\xi_{D_\ell} \cdot n = \frac{\ell^2 - r^2}{2\ell},$$

with $n^a = u^a$ the slice normal at the center.

The same kernel is the bounded-interval kernel supplied by that separate projective branch. Along each null generator of the diamond, Theorem 6.24 identifies the half-line generator constructed above with the local null-stress charge, and the interval-preserving projective branch transports that identification to the affine-covariant interval weight that vanishes at the two diamond endpoints. On the $t = 0$ slice that weight is exactly $(\ell^2 - r^2)/(2\ell)$. Therefore the null-stress bridge reconstructs the bulk modular charge of the geometric generator as

$$\delta\langle B_C\rangle = \int_{B_\ell} \frac{\ell^2 - r^2}{2\ell}\, \delta\langle T_{00}\rangle\, d^3x + \frac{1}{2\pi}\, \delta\langle E^{(\eta)}_{C,\ell}\rangle.$$

Multiplying by the first-law factor 2π gives the displayed formula for $\delta S_{\text{bulk}}(C)$.

If $\delta\langle T_{00}\rangle$ is approximately constant across B_ℓ, then

$$\int_{B_\ell} \frac{\ell^2 - r^2}{2\ell}\, \delta\langle T_{00}\rangle\, d^3x = \frac{4\pi\ell^4}{15}\, \delta\langle T_{00}\rangle + O(\ell^5 \partial T).$$

The additional hypothesis $\delta\langle E^{(\eta)}_{C,\ell}\rangle = o(\ell^4)$ then yields

$$\delta S_{\text{bulk}}(C) = 2\pi \cdot \frac{4\pi\ell^4}{15}\, \delta\langle T_{00}\rangle + O(\ell^5 \partial T) + o(\ell^4),$$

which is exactly the stated coefficient. No separate EFT small-ball first law has been used: the kernel comes from the geometric cap generator together with the D4 null-stress bridge. □

Theorem 6.31 (Finite-cutoff Einstein remainder bound)**.** *On the Lorentz/null-modular/Einstein branch, before taking the controlled collar and small-ball limits, the small-ball rest-frame relation has the form*

$$\delta(G_{00} + \Lambda g_{00}) = 8\pi G\, \delta\langle T_{00}\rangle + \mathcal{E}_{\ell,\delta},$$

where, on one fixed faithful collar model and for a bounded support-visible observable class,

$$|\mathcal{E}_{\ell,\delta}| \leq C_1 r_{\mathrm{FR}}(\varepsilon_\delta) + C_2 \delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta) + C_3 \eta^{\mathrm{reg}}_\delta + C_4 \ell\, \|\partial T\| + o_\delta(1) + o_\ell(1).$$

Here $\eta^{\mathrm{reg}}_\delta$ is the regularized support-visible modular transport remainder, and the constants depend only on the fixed collar model, the bounded observable class, and the local small-ball chart. If the controlled collar limit removes the first three terms and the small-ball limit removes the derivative and geometric remainders, the exact rest-frame relation of Theorem 6.32 follows.

Proof. Lemma 6.30 expresses the bulk entropy variation as the geometric cap-kernel stress term plus the carried modular/collar operator $E^{(\eta)}_{C,\ell}$ and the ordinary long-wavelength derivative remainder. The finite-stage modular-defect propagation theorem bounds expectations of $E^{(\eta)}_{C,\ell}$ by the Fawzi–Renner observable error, the fixed-collar Markov replacement modulus, and the support-visible regularized modular-transport remainder. Inserting that estimate into fixed-cap generalized-entropy stationarity and dividing by the small-ball coefficient gives the displayed $\mathcal{E}_{\ell,\delta}$. □

Theorem 6.32 (Jacobson-type rest-frame relation from derived fixed-cap generalized-entropy stationarity)**.** *Assume Axioms 3.1–3.4, Assumption 3.9, and the hypotheses of Theorem 6.24 and Lemma 6.30. Then, for every sufficiently small cap centered at p and every admissible fixed-cap MaxEnt variation δ about a realized reference state ω in the locally Lorentzian $d = 4$ scaling regime, if u^a is the future-directed unit tangent of the local diamond rest frame at p, one has*

$$\delta\Big[(G_{ab} + \Lambda g_{ab}) u^a u^b\Big] = 8\pi G\, u^a u^b\, \delta\langle T_{ab}\rangle$$

at p. In the adapted rest frame $u^a = (1, 0, 0, 0)$, this is

$$\delta(G_{00} + \Lambda g_{00}) = 8\pi G\, \delta\langle T_{00}\rangle.$$

Proof. Theorem 6.28 gives the fixed-cap generalized-entropy stationarity condition

$$0 = \delta S_{\mathrm{gen}}(C) = \delta S_{\mathrm{bulk}}(C) + \frac{\delta A}{4G}.$$

By Lemma 6.30,

$$\delta S_{\mathrm{bulk}}(C) = \frac{8\pi^2 \ell^4}{15}\, u^a u^b\, \delta\langle T_{ab}\rangle + O(\ell^5 \partial T) + o(\ell^4),$$

where in the adapted rest frame the contraction is $\delta\langle T_{00}\rangle$.

For a fixed-volume small geodesic ball in $d = 4$, the standard area-variation identity gives

$$\delta A\big|_{V,\Lambda} = -\frac{4\pi \ell^4}{15}\, \delta\Big[(G_{ab} + \Lambda g_{ab}) u^a u^b\Big].$$

Therefore

$$0 = \frac{8\pi^2 \ell^4}{15}\, u^a u^b\, \delta\langle T_{ab}\rangle - \frac{\pi \ell^4}{15G}\, \delta\Big[(G_{ab} + \Lambda g_{ab}) u^a u^b\Big] + O(\ell^5 \partial T) + o(\ell^4).$$

Dividing by ℓ^4 and taking the small-ball limit removes the $O(\ell^5 \partial T) + o(\ell^4)$ terms and yields

$$\delta\Big[(G_{ab} + \Lambda g_{ab}) u^a u^b\Big] = 8\pi G\, u^a u^b\, \delta\langle T_{ab}\rangle.$$

This is precisely the desired rest-frame scalar relation. □

Corollary 6.33 (Internal tensor upgrade in the scaling regime)**.** *If the first-variation relation of Theorem 6.32 holds for all local observer four-velocities and all reference states in a connected scaling branch, then*

$$G_{ab} + \Lambda g_{ab} = 8\pi G \langle T_{ab} \rangle$$

throughout that branch. The only local ambiguity is the metric term isolated by Lemma 6.25, and that ambiguity is exactly the same branch constant Λ*.*

Proof. Define

$$Y_{ab} := G_{ab} + \Lambda g_{ab} - 8\pi G \langle T_{ab} \rangle.$$

At any point p, Theorem 6.32 gives

$$u^a u^b \, \delta Y_{ab} = 0$$

for the unit future-directed four-velocity u^a of the local diamond rest frame. By the Lorentz branch of Theorem 6.8 and overlap consistency across observers through the same point, these local rest frames exhaust all timelike directions. Hence

$$u^a u^b \, \delta Y_{ab} = 0 \qquad \text{for every unit timelike } u^a.$$

Choose local inertial coordinates at p and write

$$u^a = \gamma(1, v^i), \qquad |\vec{v}| < 1, \qquad \gamma = (1 - |\vec{v}|^2)^{-1/2}.$$

Then

$$0 = u^a u^b \, \delta Y_{ab} = \gamma^2 \Big(\delta Y_{00} + 2 v^i \delta Y_{0i} + v^i v^j \delta Y_{ij} \Big)$$

for every $|\vec{v}| < 1$. The quadratic polynomial in v^i therefore vanishes on an open ball, so all of its coefficients vanish:

$$\delta Y_{00} = 0, \qquad \delta Y_{0i} = 0, \qquad \delta Y_{ij} = 0.$$

Thus

$$\delta Y_{ab} = 0$$

as a full tensor.

Along any path of admissible fixed-cap MaxEnt variations inside the connected scaling branch, Y_{ab} is therefore constant. The null bridge isolated the only purely local freedom as a metric term via Lemma 6.25; fixing that term on one maximally symmetric reference state determines the same branch constant Λ. Equivalently,

$$Y_{ab} = 0,$$

so

$$G_{ab} + \Lambda g_{ab} = 8\pi G \langle T_{ab} \rangle$$

throughout the branch. □

6.4 Cosmological-constant / screen-capacity closure

Lemma 6.34 (Vacuum-Energy Blindness)**.** *For any null vector* k *and any vacuum-energy contribution* $T^{\mathrm{vac}}_{ab} = -\rho_{\mathrm{vac}} g_{ab}$*,*

$$T^{\mathrm{vac}}_{kk} = 0.$$

Proposition 6.35 (Structural Separation of Λ)**.** *Local null-modular data fix* T_{ab} *only up to* ϕg_{ab}*. Therefore* Λ *is not determined by local overlap consistency alone and requires a global capacity closure; that zero-input closure is the cosmic record-closure readback fixed point of Definition 6.40.*

Corollary 6.36 (Cosmological Constant from Capacity)**.** *If the cosmic record-capacity fixed point is identified with the de Sitter static-patch entropy on the observed branch,*

$$N_{\mathrm{CRC}} = S_{\mathrm{dS}},$$

and the standard de Sitter entropy relation

$$S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4G} = \frac{3\pi}{G\Lambda},$$

then

$$\Lambda_{\mathrm{CRC}} = \frac{3\pi}{GN_{\mathrm{CRC}}}.$$

Corollary 6.37 (Global screen-capacity closure of the Einstein branch)**.** *Assume the hypotheses of Corollaries 6.33 and 6.36. Then on the same connected scaling branch,*

$$G_{ab} + \frac{3\pi}{GN_{\mathrm{CRC}}} g_{ab} = 8\pi G \langle T_{ab} \rangle.$$

Thus the local Einstein recovery closes globally at the cosmic record-capacity fixed point N_{CRC}.

Proof. Corollary 6.33 gives

$$G_{ab} + \Lambda g_{ab} = 8\pi G \langle T_{ab} \rangle$$

on the connected scaling branch. Corollary 6.36 identifies

$$\Lambda_{\mathrm{CRC}} = \frac{3\pi}{GN_{\mathrm{CRC}}}$$

at the cosmic record-capacity fixed point. Substituting that value of Λ yields the stated equation. □

Corollary 6.38 (De Sitter static-patch parameter package on the D6 branch)**.** *Assume the hypotheses of Corollary 6.36. Then the same cosmic record-capacity fixed point fixes the coherent static-patch package*

$$S_{\mathrm{dS}} = N_{\mathrm{CRC}}, \qquad A_{\mathrm{dS}} = 4GN_{\mathrm{CRC}}, \qquad r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_{\Lambda} = \frac{r_{\mathrm{dS}}}{c}.$$

Together with Corollary 6.37, this packages the D6 branch as the global closure of the same Einstein branch plus its derived de Sitter entropy/radius/timescale readout.

Proof. Corollary 6.36 fixes Λ from N_{CRC}. The de Sitter entropy-area relation gives

$$S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4G} = N_{\mathrm{CRC}},$$

and the standard static-patch formulas give the stated r_{dS} and t_{Λ}. □

Remark 6.39 (Capacity normalization)**.** *The entropy capacity and the bare horizon area ratio differ by the factor π. In Planck units,*

$$N_{\mathrm{patch}} = \left(\frac{r_{\mathrm{dS}}}{\ell_P}\right)^2 = \frac{3}{\Lambda \ell_P^2}, \qquad N_{\mathrm{scr}} = \pi N_{\mathrm{patch}} = \frac{3\pi}{\Lambda \ell_P^2}.$$

For $\Lambda \ell_P^2 \simeq 2.85 \times 10^{-122}$, this gives $N_{\mathrm{patch}} \simeq 1.05 \times 10^{122}$ and $N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}$. The D6 branch uses the second quantity because N_{scr} is defined as the de Sitter entropy capacity.

Definition 6.40 (Cosmic record-closure capacity). *For a candidate entropy capacity N, let $\mathfrak{U}_N$ be the OPH universe candidate with supplied active boundary capacity N, let* nf *be the quotient normal-form map on the declared terminating/confluent repair surface, let* Obs *select the stable self-reading observer sector, and let* $\mathrm{Cap}_{\mathrm{read}}$ *return the capacity reconstructed by that sector. The capacity readback map is*

$$F(N) = \mathrm{Cap}_{\mathrm{read}}\big(\mathrm{Obs}(\mathrm{nf}(\mathfrak{U}_N))\big).$$

The cosmic record-closure capacity is the fixed point

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}), \qquad \Lambda_{\mathrm{CRC}} = \frac{3\pi}{GN_{\mathrm{CRC}}}.$$

The screen-normalized count representation is as follows. Let Ω_N^{sc} denote the terminal OPH normal forms at capacity N that are closed under repair, support at least one stable observer/checkpoint subfederation, carry the recovered local package, and whose own horizon record surface reads back capacity N. Define

$$\Pi(N) := \frac{|\Omega_N^{\mathrm{sc}}|}{\dim \mathcal{H}_{\partial,N}} = |\Omega_N^{\mathrm{sc}}|e^{-N}, \qquad \log \dim \mathcal{H}_{\partial,N} = N.$$

The corresponding input-free selector target is

$$N_\star = \mathrm{MAR}\arg\max_N \left[\log|\Omega_N^{\mathrm{sc}}| - N\right].$$

Equivalently, with $\ell(N) = \log|\Omega_N^{\mathrm{sc}}| - N$, the OPH-derived stationarity map

$$T_\eta(N) = N + \eta\ell'(N)$$

has a unique stable fixed point under the derivative-sign certificate stated on the synthesis surface. Informally, N_{CRC} is the single capacity where the universe reads back its own boundary without deficit or slack.

Remark 6.41 (Capacity fixed-point reading). *Definition 6.40 gives the precise D6 screen-capacity fixed point: the readback map F is defined on the OPH finite normal-form grammar, and the full screen dimension supplies the non-tunable normalization e^{-N} in the count representation. In differentiable form, the pressure certificate is the unique fixed-point condition $T_\eta(N) = N$, equivalently*

$$\frac{d}{dN}\left[\log|\Omega_N^{\mathrm{sc}}| - N\right] = 0.$$

On the observed branch the read-off is

$$N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}.$$

The observed branch identifies this fixed point with the de Sitter entropy capacity used by Theorem 6.45.

Remark 6.42 (Observed-age benchmark status). *The branch quantity t_Λ is the de Sitter timescale fixed by the same D6 closure. The observed cosmic age t_0 is not an additional D6 theorem output. It is a downstream FLRW comparison quantity once one chooses a cosmological model above the local/global D5→D6 stack. On the flat ΛCDM benchmark,*

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_\Lambda}}\sinh^{-1}\left(\sqrt{\frac{\Omega_\Lambda}{\Omega_m}}\right).$$

Lemma 6.43 (FLRW curvature as visible scalar holonomy)**.** *On a homogeneous-isotropic spatial slice with constant sectional curvature K, the small-loop holonomy in the u-v plane satisfies*

$$\mathrm{Hol}_{\square_{uv}} = \exp\Big(K A_\square J_{uv} + O(A_\square^{3/2})\Big),$$

where $A_\square$ is the loop area and J_{uv} is the infinitesimal rotation generator in that two-plane. On a visibly separated OPH refinement system, the refinement-limit scalar spatial holonomy vanishes if and only if $K = 0$. Thus the flat FLRW branch is the zero-visible spatial holonomy branch.

Proof. The displayed formula is the standard infinitesimal holonomy expansion for a connection with constant sectional curvature. If $K = 0$, every area-normalized small-loop curvature readout vanishes. Conversely, if $K \neq 0$, the area-normalized holonomy converges to KJ_{uv} in each visible two-plane. Visible separation forbids quotienting this nonzero family away as a mere gauge representative difference. Homogeneity and isotropy leave no independent scalar spatial-curvature obstruction beyond K, so vanishing refinement-limit scalar holonomy forces $K = 0$. □

Remark 6.44 (Flatness is not a D6 output)**.** *Lemma 6.43 only names the OPH meaning of an FLRW flat branch. A selection statement needs an additional cosmological continuation hypothesis: the preserved cosmological boundary datum must contain no independent curvature charge, and the same-boundary or MAR selector must choose the extension with minimal visible geometric obstruction. The lemma by itself does not solve the inflationary flatness or horizon problems and does not supply a CMB kernel.*

Theorem 6.45 (Cosmological-constant / screen-capacity closure stack)**.** *Assume the hypotheses of Corollary 6.33, the cosmic record-capacity fixed point*

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}),$$

its observed-branch de Sitter entropy readout

$$N_{\mathrm{CRC}} = S_{\mathrm{dS}},$$

the standard de Sitter entropy relation

$$S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4G} = \frac{3\pi}{G\Lambda},$$

and the standard de Sitter static-patch formulas

$$r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\mathrm{dS}}}{c}.$$

Then:

(i) *the local null-modular data determine the Einstein branch only modulo a metric term Λg_{ab};*

(ii) *the same branch closes globally as*

$$G_{ab} + \frac{3\pi}{G N_{\mathrm{CRC}}}\, g_{ab} = 8\pi G\, \langle T_{ab}\rangle;$$

(iii) *the same D6 closure fixes the static-patch package*

$$S_{\mathrm{dS}} = N_{\mathrm{CRC}}, \qquad A_{\mathrm{dS}} = 4GN_{\mathrm{CRC}}, \qquad r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\mathrm{dS}}}{c};$$

(iv) *the observed cosmic age t_0 is not an additional theorem output of this stack and is a downstream FLRW benchmark.*

Thus the cosmological-constant package is one local/global theorem stack: the local null-modular branch leaves exactly the null-invisible metric freedom, and the cosmic record-capacity fixed point closes that same Einstein branch globally.

Proof. Item (i) is Proposition 6.35, whose local ambiguity statement rests on the null-data blindness isolated by Lemma 6.34 together with the D5 Einstein branch. Item (ii) is Corollary 6.37. Item (iii) is Corollary 6.38. Item (iv) is Remark 6.42. □

Theorem 6.45 packages the local/global D5→D6 handoff. The D6 hypotheses are exactly the local Einstein branch, the cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$, its observed-branch de Sitter entropy readout $N_{\mathrm{CRC}} = S_{\mathrm{dS}}$, the standard de Sitter entropy relation, and the standard static-patch radius/time formulas; none of these are hidden inside the local null-modular reconstruction. The local theorem stack does not derive the global capacity from bare null data; the OPH-derived readback closure fixes it. The vacuum-energy problem is reorganized by separating local null data from the global capacity fixed point.

7 Gauge Reconstruction and Standard Model Structure

7.1 Compact gauge reconstruction

At any fixed UV cutoff, edge-center completion equips collars with finitely many sector labels, boundary charge carriers, and finite-dimensional intertwiner spaces. These are fixed-cutoff collar data, not refinement-limit assumptions. The local MaxEnt / collar-mixing package established earlier controls only fixed-cutoff recoverability, modular-support localization, and carried error terms on the realized branch. It is logically separate from whether transportable edge sectors survive refinement, and it supplies no theorem that the refinement-limit sector category is trivial or nontrivial.

On the ordinary or central-defect branch, path-independent movement of collar charges is supplied by the overlap-gluing theorem. Theorem 5.8 constructs transport from paths in the overlap recharting groupoid and proves the exact obstruction criterion: the ordinary/central branch is strictly path-independent exactly when $[z]_\Sigma = 0$. The same theorem handles the genuinely noncentral branch by the crossed-module criterion $q_\Sigma = 0$; if $q_\Sigma \neq 0$, the fixed-cutoff sector is a higher-gauge sector labelled by q_Σ, not an ordinary compact-group DR sector.

For the ordinary or central-defect bosonic zero-obstruction branch, the fixed-cutoff categories $\mathsf{Sect}_r^{\mathrm{bos}}$ are theorem-produced and the refinement/fiber ladder is constructed below. The resulting directed colimit

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

is the category on which Doplicher–Roberts / Tannaka reconstruction is applied. The theorem below is neutral about whether Sect_∞ is trivial or nontrivial: if only the tensor unit persists, the reconstructed compact group is the trivial group.

7.1.1 Classification is not realization

Proposition 7.1 (Obstruction neutrality of the Standard Model selection step)**.** *The ordinary zero-obstruction condition, the central condition $[z]_\Sigma = 0$, and the strictified noncentral condition $q_\Sigma = 0$*

are transportability conditions. They permit an ordinary transportable bosonic sector category on the corresponding branch; they do not select

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

On a zero-obstruction branch, DR/Tannaka reconstruction returns

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

for the constructed sector category and fiber functor. If the category is trivial, G is trivial; in general G is whatever compact group that tensor-fiber data reconstructs. The Standard Model quotient enters only after Axiom 3.5 is applied to a nonempty realized one-Higgs chiral sector package.

Proof. Theorem 5.8 identifies vanishing obstruction with strict path-independent transport, and with no more structure than that. Theorems 5.11, 7.2, and 7.3 use that transport condition to build the ordinary bosonic sector category and its finite-dimensional fiber functor. Theorem 7.4 then reconstructs G from $(\mathsf{Sect}_\infty,\mathcal{F})$. None of those steps names a weak doublet, color triplet, Higgs doublet, hypercharge lattice, generation count, or finite $\mathbb{Z}_6$ kernel. Those data are supplied by the MAR-realized branch through Theorem 7.24. □

Theorem 7.2 (RefinementFunctorAndFiberDescent)**.** *Let R be the separated cofinal OPH refinement system used by Theorem 4.12. For each $r\in R$, let $\mathsf{Sect}_r^{\mathrm{bos}}$ be the fixed-cutoff zero-obstruction bosonic collar-sector category constructed in Theorem 5.11, with strict transport supplied by Theorem 5.8. Then for every $r\preceq s$ there is a faithful monoidal *-functor*

$$U_{rs}:\mathsf{Sect}_r^{\mathrm{bos}}\longrightarrow\mathsf{Sect}_s^{\mathrm{bos}}$$

carrying $\mathbf{1}_r$ to $\mathbf{1}_s$, preserving zero-obstruction transportable sector classes on cofinal tails, and satisfying coherent composition isomorphisms

$$U_{st}\circ U_{rs}\cong U_{rt}\qquad(r\preceq s\preceq t).$$

Moreover the fixed-cutoff edge carriers and intertwiner multiplicity blocks define faithful objectwise finite-dimensional bosonic fiber functors

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

with monoidal natural isomorphisms

$$F_s\circ U_{rs}\cong F_r$$

compatible with refinement composition.

Proof. Fix $r\preceq s$ and a connected collar B around a cut Σ. OPH refinement subdivides B into a refined collar B_s and carries overlap records by the deterministic projection maps of the separated cofinal refinement system. Dualizing the finite refinement projection gives an injective unital *-homomorphism on the sector-preserving collar algebras,

$$\jmath_{rs}^{B}:\mathcal{A}_{\mathrm{EC},r}(B)\hookrightarrow\mathcal{A}_{\mathrm{EC},s}(B_s).$$

Since this map preserves centers, a central edge block $P_{\alpha,r}$ is sent to the sum of fine central blocks with coarse shadow α:

$$\jmath_{rs}^{B}(P_{\alpha,r})=\sum_{\beta:\,\beta\downarrow r=\alpha}P_{\beta,s}.$$

The same refinement projection gives a coarse-shadow homomorphism on the compact overlap recharting systems,

$$\pi^{\Sigma}_{sr} : \widehat{K}_{\Sigma,s} \to \widehat{K}_{\Sigma,r},$$

where $\widehat{K} = K$ on the ordinary branch and $\widehat{K}$ is the compact central extension on the central branch. Refinement associativity gives

$$\jmath^{B}_{rt} = \jmath^{B_s}_{st} \circ \jmath^{B}_{rs}, \qquad \pi^{\Sigma}_{tr} = \pi^{\Sigma}_{sr} \circ \pi^{\Sigma}_{ts}.$$

A simple object of $\mathsf{Sect}^{\mathrm{bos}}_r$ is represented by a zero-obstruction collar charge package $(P_{\alpha,r}, W_{\alpha,r})$, where $W_{\alpha,r}$ is the finite-dimensional boundary carrier. Define

$$U_{rs}(P_{\alpha,r}, W_{\alpha,r}) := (\jmath^{B}_{rs}(P_{\alpha,r}), (\pi^{\Sigma}_{sr})^{*} W_{\alpha,r}).$$

Here $(\pi^{\Sigma}_{sr})^{*} W_{\alpha,r}$ is the same finite-dimensional Hilbert space, viewed as a representation of the refined recharting system through π^{Σ}_{sr}. If the refined support decomposes into fine minimal central projectors, the image object is the corresponding direct sum in the additive Karoubi envelope.

For a morphism $f \in \mathrm{Hom}_{\widehat{K}_{\Sigma,r}}(W_X, W_Y)$, let $U_{rs}(f)$ be the same linear map regarded as an intertwiner between pulled-back refined carriers. This is well defined because f commutes with the $\widehat{K}_{\Sigma,r}$-action, hence also with the $\widehat{K}_{\Sigma,s}$-action through π^{Σ}_{sr}. The identities

$$U_{rs}(g \circ f) = U_{rs}(g) \circ U_{rs}(f), \qquad U_{rs}(f^{*}) = U_{rs}(f)^{*}$$

are identities of finite-dimensional operators. Since a nonzero finite-dimensional operator remains nonzero after pullback, and equivalently since $\jmath^{B}_{rs}$ is injective on the overlap-visible collar algebra, U_{rs} is faithful.

The tensor product in $\mathsf{Sect}^{\mathrm{bos}}_r$ is collar concatenation. Pullback of boundary representations commutes with tensor products,

$$(\pi^{\Sigma}_{sr})^{*}(W_X \otimes W_Y) \cong (\pi^{\Sigma}_{sr})^{*} W_X \otimes (\pi^{\Sigma}_{sr})^{*} W_Y,$$

and $\jmath^{B}_{rs}$ sends a concatenated collar support to the concatenation of the refined supports. Therefore there are canonical natural isomorphisms

$$J^{rs}_{X,Y} : U_{rs}(X \otimes_r Y) \xrightarrow{\sim} U_{rs}(X) \otimes_s U_{rs}(Y).$$

They satisfy the pentagon and unit identities because union-collar gluing is parenthesization-independent. The trivial boundary carrier pulls back to the trivial boundary carrier and $\jmath^{B}_{rs}(1) = 1$, so $U_{rs}\mathbf{1}_r = \mathbf{1}_s$.

For $r \preceq s \preceq t$, the displayed identities for $\jmath$ and π identify $U_{st}U_{rs}$ with U_{rt} on objects and morphisms. These are the required coherent composition isomorphisms.

By Theorem 5.8, strict transport is exactly the zero-obstruction condition: trivial loop holonomy on the ordinary branch, $[z]_{\Sigma} = 0$ on the central branch, and $q_{\Sigma} = 0$ after strictification of a genuinely noncentral representative. A refined closed path acts on $U_{rs}X$ through its coarse shadow, so a trivial coarse loop action remains trivial on the pulled-back refined image. In central language, a central 1-cochain that kills z_r pulls back to a central 1-cochain that kills the cocycle on the refined image. In the strictified crossed-module case, the pullback of the trivial class is again trivial. Thus a zero-obstruction sector cannot lose strict transportability under refinement. It also cannot vanish because $\jmath^{B}_{rs}$ is injective and unital. A split into two inequivalent persistent sector tails with the same cofinal overlap-visible support would contradict visible separation in Theorem 4.12; if the support differs, the difference is a new finite-stage sector datum rather than a silent change of the sector.

For the fibers, choose at each fixed cutoff a finite simple skeleton $\{X_{i,r}\}_{i\in I_r}$ and let $V_{i,r}$ be the finite-dimensional boundary carrier of $X_{i,r}$. Define

$$F_r(X) := \bigoplus_{i\in I_r} \mathrm{Hom}_r(X_{i,r}, X) \otimes V_{i,r}, \qquad F_r(f)(h \otimes v) := (f \circ h) \otimes v.$$

This Hilbert space is finite-dimensional because the fixed-cutoff category is semisimple with finite-dimensional Hom spaces and finite-dimensional boundary carriers. The functor is faithful: if $f \neq 0$, semisimplicity gives a simple summand map $h : X_{i,r} \to X$ with $f \circ h \neq 0$, so $F_r(f) \neq 0$.

The monoidal structure on F_r is the finite Clebsch–Gordan decomposition supplied by fixed-cutoff collar concatenation:

$$X_{i,r} \otimes X_{j,r} \cong \bigoplus_{k\in I_r} N_{ij}^k X_{k,r}, \qquad V_{i,r} \otimes V_{j,r} \cong \bigoplus_{k\in I_r} \mathbb{C}^{N_{ij}^k} \otimes V_{k,r}.$$

Associativity, unit, and symmetry of these isomorphisms are exactly the collar-gluing associativity and bosonic spacelike-exchange identities proved at fixed cutoff.

Finally, the carrier of $U_{rs}X_{i,r}$ is the pullback of $V_{i,r}$, hence the same finite-dimensional Hilbert space with refined action through π^{Σ}_{sr}. The Hom-space pullbacks are the maps U_{rs} above. Therefore the summandwise carrier and Hom identifications define monoidal natural isomorphisms

$$\theta_{rs} : F_s \circ U_{rs} \xrightarrow{\sim} F_r.$$

For $r \preceq s \preceq t$, the identity $\theta_{rt} = \theta_{rs} \circ (\theta_{st} U_{rs})$ follows from the associativity identities for $\jmath$, π, and union-collar gluing. This proves the theorem. □

Theorem 7.3 (Construction of the refinement-stable bosonic sector category)**.** *Assume Axioms 3.1–3.4. On the ordinary or central-defect bosonic zero-obstruction branch, let the fixed-cutoff sector categories* $\mathsf{Sect}_r^{\mathrm{bos}}$ *be those constructed in Theorem 5.11, and let the refinement functors and stagewise finite-dimensional bosonic fibers be those constructed in Theorem 7.2. Then the directed colimit*

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

inherits a well-defined semisimple rigid symmetric C^**-tensor structure. Its tensor unit is the persistent vacuum sector, its tensor product is induced from collar concatenation, its duals are induced from orientation reversal / charge conjugation, and its braiding is the induced bosonic spacelike-exchange symmetry of the* 3+1*-dimensional branch. The finite-dimensional stagewise fibers descend to a faithful bosonic fiber functor*

$$\mathcal{F} : \mathsf{Sect}_\infty \to \mathsf{Hilb}_{\mathrm{fd}}.$$

The theorem is neutral about nontriviality: if only the tensor unit persists, then Sect_∞ *is the trivial bosonic tensor category.*

Proof. The fixed-cutoff categorical structure is supplied by Theorem 5.11. The faithful monoidal *-refinement functors, their composition coherence, their preservation of zero-obstruction transportable sector classes on cofinal tails, and the compatible finite-dimensional stagewise fibers are constructed in Theorem 7.2.

The directed colimit category has objects represented by eventually compatible refinement tails and morphisms represented by eventual intertwiner classes. Faithfulness of the U_{rs} makes the morphism equivalence relation separated: a nonzero finite-stage intertwiner cannot become zero on a cofinal tail. The tensor unit on each fixed-cutoff category is carried to the tensor unit by every U_{rs}, so the vacuum tail defines $\mathbf{1}_\infty$.

For objects represented by tails X_r and Y_r, define

$$[X] \otimes [Y] := [X_r \otimes_r Y_r]$$

at any sufficiently fine common stage. This is independent of the chosen stage because the monoidal structure maps of U_{rs} identify

$$U_{rs}(X_r \otimes_r Y_r) \cong U_{rs}(X_r) \otimes_s U_{rs}(Y_r).$$

The same argument descends the associators, unitors, duality evaluation and coevaluation maps, and the bosonic symmetry. Since each fixed-cutoff category is rigid, symmetric, semisimple, and C^*, and since the refinement functors preserve *, direct sums, and subobjects, the colimit inherits the same structure on persistent tails.

For the fiber functor, choose a representative X_r of a colimit object $[X]$ and define

$$\mathcal{F}([X]) := F_r(X_r)$$

modulo the canonical monoidal natural isomorphisms

$$F_s(U_{rs}X_r) \cong F_r(X_r)$$

from Theorem 7.2. The composition compatibility of those isomorphisms makes this independent of representative. On morphisms, $\mathcal{F}$ is induced by the eventual action of the corresponding finite-stage F_r. Objectwise finite dimensionality, monoidality, and faithfulness descend from the stagewise F_r's. Thus $\mathcal{F} : \mathsf{Sect}_\infty \to \mathsf{Hilb}_{\mathrm{fd}}$ is a faithful bosonic fiber functor. □

Theorem 7.4 (Compact gauge reconstruction in the bosonic branch)**.** *Assume Axioms 3.1–3.4 and work on the ordinary or central-defect bosonic zero-obstruction branch. Let*

$$\mathcal{F} : \mathsf{Sect}_\infty \to \mathsf{Hilb}_{\mathrm{fd}}$$

be the faithful bosonic fiber functor constructed in Theorem 7.3 from Theorems 5.8, 5.11, and 7.2. Then

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

is a compact group and

$$\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G)$$

as a symmetric C^*-tensor category with fiber functor. In particular G is uniquely determined up to isomorphism by the constructed pair $(\mathsf{Sect}_\infty, \mathcal{F})$.

Proof. The hypotheses required by Doplicher–Roberts / Tannaka reconstruction are supplied by theorems rather than by a refinement assumption. Theorem 5.8 supplies the strict zero-obstruction transport criterion. Theorem 5.11 constructs the fixed-cutoff bosonic symmetric C^*-tensor categories. Theorem 7.2 constructs the faithful monoidal *-refinement functors and compatible finite-dimensional fibers, and Theorem 7.3 descends them to $(\mathsf{Sect}_\infty, \mathcal{F})$.

Fix a small skeleton of Sect_∞. For each object X, a monoidal natural automorphism $\eta \in \mathrm{Aut}_\otimes(\mathcal{F})$ has a unitary component $\eta_X \in U(\mathcal{F}(X))$. Hence

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

Naturality imposes the closed relations

$$\mathcal{F}(f)\eta_X = \eta_Y \mathcal{F}(f) \qquad (f : X \to Y),$$

and monoidality imposes the closed relations

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

Thus $G = \mathrm{Aut}_\otimes(\mathcal{F})$ is a closed subgroup of a product of compact unitary groups, and is compact.

Doplicher–Roberts / Tannaka reconstruction applies to the rigid symmetric C^*-tensor category with faithful finite-dimensional bosonic fiber functor constructed above, giving

$$\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G).$$

Uniqueness follows because the reconstructed compact group is the tensor automorphism group of the fiber functor. If Sect_∞ is trivial, then G is the trivial compact group; the realized nontrivial branch is supplied by Theorem 7.23, not hidden as a hypothesis in this reconstruction theorem. □

If a fermionic sign object is present, the correct reconstruction statement is super-Tannakian rather than purely Tannakian. This paper does not prove the full fermionic/chiral extension; it works only in the bosonic internal-gauge branch once the fixed-cutoff sector packages and their monoidal refinement transport have been constructed.

Theorem 7.5 (Gauge-sector classification–selection factorization)**.** *On the OPH compact-gauge lane, the realized Standard Model claim factors as*

$$\textit{overlap/gluing data} \longrightarrow \textit{obstruction class} \longrightarrow \mathsf{Sect}_\infty^{\mathrm{bos}} \longrightarrow G = \mathrm{Aut}_\otimes(\mathcal{F}) \longrightarrow \mathfrak{S}_{\mathrm{MAR}}.$$

The first arrow computes the ordinary, central, or crossed-module obstruction. The second keeps only the ordinary transportable zero-obstruction sector branch. The third reconstructs the compact group from the persistent tensor category and fiber functor. These are classification and reconstruction steps. The final arrow is the realization/selection step: MAR acts on realized admissible sector packages, not on obstruction classes alone.

Proof. The first two arrows are Theorem 5.8 and the branch split into $[z]_\Sigma = 0$, $q_\Sigma = 0$, or genuinely higher-gauge $q_\Sigma \neq 0$ sectors. The third arrow is Theorem 7.4. Proposition 7.1 shows that these arrows are neutral about the Standard Model quotient. The final arrow is Axiom 3.5 applied to the nonempty witness class of Theorem 7.23, with uniqueness summarized in Theorem 7.24. □

7.2 MAR minimality and the Standard Model quotient

Throughout this subsection, χ_{cpl} means the dimension of the minimal coupled carrier supporting the required weak-type and color-type nonabelian charges on a common block. It is not the abstract minimal faithful representation dimension of the final gauge group.

Definition 7.6 (Weak-type and color-type nonabelian roles)**.** *A weak-type nonabelian charge is a pseudoreal nonabelian doublet role. A color-type nonabelian charge is an intrinsically complex nonabelian role, meaning that its nonabelian factor itself acts by an irreducible representation not equivalent to its conjugate. Abelian twisting of a pseudoreal nonabelian irreducible representation does not count as a color-type role.*

The derivation of the Standard Model quotient separates into five steps:

1. existence of a compact reconstructed group;
2. identification of the minimal nonabelian sector content;
3. identification of the minimal coupled carrier;
4. uniqueness of the connected abelian factor once admitted on that carrier;
5. determination of product gauge structure up to finite quotient, then of the global finite quotient from the realized matter spectrum.

Lemma 7.7 (Minimal nonabelian sector content)**.** *Any admissible low-energy sector must contain both a weak-type pseudoreal nonabelian charge role and an intrinsically complex color-type nonabelian charge role.*

Proof. Light chiral matter requires left-handed multiplets and right-handed singlets to carry inequivalent gauge data. A pseudoreal doublet structure is the minimal way to realize the weak sector without immediately producing vectorlike masses. A genuinely complex nonabelian representation is required to distinguish quarks from antiquarks. A single nonabelian simple factor cannot simultaneously supply both a minimal weak-type role and an intrinsically complex color-type role; abelian twisting does not count toward the color-type requirement. □

Lemma 7.8 (Connected 2D pseudoreal image classification)**.** *Let H be a connected compact group with a faithful irreducible* 2*-dimensional pseudoreal unitary representation. Then the connected derived subgroup of its image is conjugate to* $\mathrm{SU}(2)$*. Equivalently, the nonabelian factor acting in that role is* $\mathrm{SU}(2)$ *up to finite central quotient.*

Proof sketch. Irreducibility places the image inside $\mathrm{U}(2)$ with scalar commutant. Its connected derived subgroup is therefore a connected compact subgroup of $\mathrm{SU}(2)$. The connected compact subgroups of $\mathrm{SU}(2)$ are either tori or $\mathrm{SU}(2)$ itself. The representation is nonabelian and pseudoreal, so the torus case is excluded. Hence the derived subgroup is conjugate to $\mathrm{SU}(2)$, with only a finite central kernel left invisible in the representation. □

Lemma 7.9 (Connected 3D intrinsically complex image classification)**.** *Let H be a connected compact group with a faithful irreducible intrinsically complex* 3*-dimensional unitary representation. Then the connected derived subgroup of its image is conjugate to* $\mathrm{SU}(3)$*. Equivalently, the nonabelian factor acting in that role is* $\mathrm{SU}(3)$ *up to finite central quotient.*

Proof sketch. Because the representation is irreducible on $\mathbb{C}^3$, the semisimple part of the Lie algebra of the image acts irreducibly on $\mathbb{C}^3$. A product of two nontrivial simple factors would force a tensor-product decomposition of dimension at least 4, so only one simple factor can occur. Among compact simple Lie algebras, the only ones with nontrivial irreducible representations of dimension at most 3 are $\mathfrak{su}(2)$ and $\mathfrak{su}(3)$. The 3-dimensional $\mathfrak{su}(2)$ representation is real, not intrinsically complex, whereas the fundamental $\mathfrak{su}(3)$ representation is intrinsically complex. Therefore the connected derived subgroup is conjugate to $\mathrm{SU}(3)$, again up to finite central kernel. □

Lemma 7.10 (Minimal coupled carrier under MAR)**.** *Within the connected positive-dimensional Lie admissible class used by MAR, and under the weak-type / intrinsically complex color-type criteria above, the minimal coupled carrier supporting both roles on a common block has the form*

$$V = \mathbb{C}^3 \otimes \mathbb{C}^2, \qquad \chi_{\mathrm{cpl}} = 6.$$

Proof. By Lemma 7.8, the minimal weak-type nonabelian role is the pseudoreal doublet of SU(2), hence has dimension 2. By Lemma 7.9, the minimal intrinsically complex color-type nonabelian role is the triplet of SU(3), hence has dimension 3.

By definition of χ_{cpl}, the weak-type and color-type roles must act nontrivially on one common irreducible nonabelian block. Disconnected direct-sum summands do not qualify. Commuting irreducible actions on such a common block therefore require dimension at least the product of the minimal weak and color dimensions:

$$\chi_{\mathrm{cpl}} \geq 2 \cdot 3 = 6.$$

Equality is achieved by the tensor-product carrier $\mathbb{C}^3 \otimes \mathbb{C}^2$. The block-diagonal representation $\mathbb{C}^3 \oplus \mathbb{C}^2$ of $S(U(3) \times U(2))$ is faithful of dimension 5, but it is not coupled and hence is not the MAR minimizer. The explicit color-type criterion excludes connected non-semisimple counterexamples of $U(2)$ type. □

Remark 7.11. *The dimension comparisons used in Lemma 7.10 apply to the positive-dimensional connected Lie image carrying the admissible nonabelian charges, equivalently to the identity component of the reconstructed group on that sector. Finite, disconnected, or connected non-semisimple counterexamples exist. They are outside the admissible class fixed by the explicit weak-type / color-type criteria.*

MAR supplies existence of a connected abelian factor; the next lemma proves only uniqueness and identification on the minimal coupled carrier.

Lemma 7.12 (Uniqueness of the connected abelian factor on the minimal coupled carrier)**.** *Inside* U(6)*, the commutant of* $\mathrm{SU}(3) \times \mathrm{SU}(2)$ *acting on* $\mathbb{C}^3 \otimes \mathbb{C}^2$ *is exactly* U(1)*.*

Proof. The SU(3) action is irreducible on the first tensor factor and trivial on the second. The SU(2) action is irreducible on the second tensor factor and trivial on the first. By Schur's lemma, any operator commuting with both actions is scalar on each irreducible factor. Hence the full commutant is a single U(1). □

Theorem 7.13 (Product Gauge Structure up to Finite Quotient)**.** *Assume Axiom 3.5, the hypotheses of Theorem 7.4 on the ordinary or central-defect realized low-energy branch, the weak-type / color-type MAR inputs above, and that the realized connected gauge image acts faithfully on the minimal coupled carrier. Then the realized connected gauge structure has the form*

$$G_{\mathrm{phys}} = \frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\Gamma}$$

for some finite central subgroup Γ*.*

Proof. Theorem 7.4 yields some compact group G. Lemma 7.7 identifies the minimal admissible sector content, Lemmas 7.8 and 7.9 identify the connected nonabelian factors realizing the minimal weak-type and color-type roles, and Lemma 7.10 fixes the minimal coupled carrier. The tensor-product structure of that carrier implies commuting weak and color actions, so the connected semisimple part of the realized image is locally $\mathrm{SU}(3) \times \mathrm{SU}(2)$. Axiom 3.5 supplies the existence of one connected abelian charge factor acting nontrivially on the coupled carrier, and Lemma 7.12 shows that any such connected abelian factor is necessarily a single U(1). A compact connected Lie group with Lie algebra $\mathfrak{su}(3) \oplus \mathfrak{su}(2) \oplus \mathfrak{u}(1)$ is a quotient of $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ by a finite central subgroup. Faithfulness rules out an invisible extra torus, so the only ambiguity is the finite central quotient Γ. □

Theorem 7.14 (Hypercharge lattice on the realized matter package)**.** *Assume gauge group* $\mathrm{SU}(N_c) \times \mathrm{SU}(2) \times \mathrm{U}(1)_Y$*, one generation of chiral matter* (Q, u^c, d^c, L, e^c)*, one Higgs doublet* H*, and Yukawa terms*

$$QHu^c, \qquad QH^\dagger d^c, \qquad LH^\dagger e^c.$$

Then anomaly cancellation and Yukawa invariance determine the hypercharge ratios up to an overall $\mathrm{U}(1)_Y$ *normalization:*

$$Y_L = -N_c Y_Q, \quad Y_H = N_c Y_Q, \quad Y_u = -(N_c+1)Y_Q, \quad Y_d = (N_c-1)Y_Q, \quad Y_e = 2N_c Y_Q.$$

Fixing the normalization by $Q = T_3 + Y$ *and* $Q(\nu_L) = 0$ *gives*

$$Y_Q = \frac{1}{2N_c}.$$

For $N_c = 3$ *one recovers the exact Standard Model lattice*

$$Y_Q = \frac{1}{6}, \quad Y_L = -\frac{1}{2}, \quad Y_u = -\frac{2}{3}, \quad Y_d = \frac{1}{3}, \quad Y_e = 1, \quad Y_H = \frac{1}{2}.$$

Proof. Yukawa invariance gives

$$Y_u = -(Y_Q + Y_H), \qquad Y_d = -Y_Q + Y_H, \qquad Y_e = -Y_L + Y_H.$$

The mixed anomalies yield

$$N_c Y_Q + Y_L = 0,$$

and

$$2N_c Y_Q + N_c Y_u + N_c Y_d + 2Y_L + Y_e = 0.$$

Solving the first anomaly gives

$$Y_L = -N_c Y_Q.$$

Substituting the Yukawa relations and $Y_L = -N_c Y_Q$ into the mixed gravitational anomaly then yields

$$Y_H = N_c Y_Q, \qquad Y_u = -(N_c+1)Y_Q, \qquad Y_d = (N_c-1)Y_Q, \qquad Y_e = 2N_c Y_Q.$$

With these relations in place, the $\mathrm{SU}(N_c)^2\mathrm{U}(1)$ and $\mathrm{U}(1)^3$ anomalies cancel identically. The normalization of Y_Q is fixed by the electric charge operator $Q = T_3 + Y$ together with $Q(\nu_L) = 0$, which implies

$$Y_L = -\frac{1}{2} = -N_c Y_Q, \qquad \text{hence} \qquad Y_Q = \frac{1}{2N_c}.$$

□

Corollary 7.15 (Three Colors on the realized MAR branch)**.** *Under the hypotheses of Theorem 7.13, the realized color-type factor acts in the fundamental triplet on the coupled carrier. Hence the realized quark doublet carries exactly*

$$N_c = 3.$$

Proof. Lemma 7.9 identifies the minimal intrinsically complex color-type role in the connected Lie admissible class as the 3-dimensional fundamental of SU(3). Lemma 7.10 then shows that MAR realizes this role on the coupled block $\mathbb{C}^3 \otimes \mathbb{C}^2$. The color multiplicity of the realized weak doublet Q is therefore fixed by the same D8 carrier to be 3. No later admissibility selector is needed to promote oddness to the specific value 3. □

Corollary 7.16 (Three Generations on the realized MAR branch)**.** *Under the hypotheses of Theorem 7.13, with $N_c = 3$ from Corollary 7.15, the CP-capability and weak-sector UV clauses contained in Axiom 3.5 force*

$$N_g = 3.$$

Proof. On the realized one-Higgs chiral branch, intrinsic CKM CP capability requires

$$\frac{(N_g - 1)(N_g - 2)}{2} > 0,$$

hence $N_g \geq 3$. For one Higgs doublet, the one-loop weak-sector coefficient is

$$b_2 = \frac{11}{3}C_A - \frac{2}{3}\sum_f T(R_f) - \frac{1}{3}\sum_s T(R_s) = \frac{22}{3} - \frac{N_g(N_c+1)}{3} - \frac{1}{6} = \frac{43}{6} - \frac{N_g(N_c+1)}{3}.$$

Asymptotic freedom therefore implies

$$N_g(N_c + 1) < \frac{43}{2}.$$

With $N_c = 3$, this gives

$$4N_g < \frac{43}{2}, \qquad \text{hence} \qquad N_g \leq 5.$$

These are not extra post-MAR selectors: they are the explicit numerical consequences of the CP-capability and weak-sector UV clauses present in Axiom 3.5 on the same realized branch. Once the first three MAR components are fixed by the D8 color/weak structure, the fourth component N_g is minimized on the allowed set $\{3, 4, 5\}$, so $N_g = 3$. □

Corollary 7.17 (MAR-minimal SM package is unique up to physical equivalence)**.** *On the ordinary or central zero-obstruction bosonic branch, assume the MAR-admissible class is nonempty and contains the explicit realized one-Higgs chiral matter package used in Theorem 7.13 through Corollary 7.16. Then every MAR-minimal representative has the same observer-visible Standard Model package,*

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_c = 3, \qquad N_g = 3,$$

with the hypercharge lattice of Theorem 7.14. The only residual freedom is physical equivalence in the sense of Definition 3.6.

Proof. Proposition 3.7 supplies MAR minima. Lemmas 7.8–7.12 fix the minimal weak, color, coupled-carrier, and connected abelian roles at the least first three complexity entries. Corollary 7.15 fixes $N_c = 3$, Corollary 7.16 fixes $N_g = 3$, Theorem 7.14 fixes the charge lattice, and Proposition 7.19 fixes the finite kernel. Definition 3.6 removes only relabelings, gauge-center conventions, and inert implementation data. □

Remark 7.18 (Witten parity on the realized color branch)**.** *With $N_c = 3$, each generation contributes $N_c + 1 = 4$ left-handed $\mathrm{SU}(2)$ doublets, so Witten's global anomaly constraint is automatically satisfied generation by generation. In this theorem stack the anomaly is therefore a consistency check on the realized triplet-doublet package, not the step that creates the color count [27].*

Proposition 7.19 (Global quotient from the realized matter spectrum)**.** *If the realized matter spectrum carries the hypercharges of Theorem 7.14 with $N_c = 3$, then the subgroup of $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ acting trivially on all realized states is exactly $\mathbb{Z}_6$.*

Proof. Write $q := 6Y \in \mathbb{Z}$ and $\eta := e^{i\pi/3}$, so the U(1) factor acts by η^q on charge-q multiplets. The element

$$g_0 := (\omega_3, -1, \eta)$$

acts trivially on the realized matter and Higgs multiplets:

$$Q:\ \omega_3(-1)\eta = 1, \qquad u^c:\ \omega_3^{-1}\eta^{-4} = 1, \qquad d^c:\ \omega_3^{-1}\eta^2 = 1,$$
$$L:\ (-1)\eta^{-3} = 1, \qquad e^c:\ \eta^6 = 1, \qquad H:\ (-1)\eta^3 = 1.$$

Its powers therefore generate a subgroup of order six acting trivially on all realized states. Conversely, any central element acting trivially on e^c must have U(1) part η^n, and triviality on Q then fixes the SU(3) and SU(2) center factors uniquely. Hence every trivial central element is a power of g_0, so the kernel is exactly $\mathbb{Z}_6$. □

Remark 7.20 (Arithmetic toy model for finite quotient labels)**.** *DULA's congruence-grading construction gives a simple number-theoretic analogue of this bookkeeping [3]. For integers coprime to* 6*, the number of prime factors congruent to* 5 *modulo* 6*, counted with multiplicity and reduced modulo* 2*, is an additive finite label that determines whether the product is* 1 *or* 5 *modulo* 6*. OPH does not use this as evidence or as a physics input. It is a toy model for the same formal pattern used here: local factor or sector data accumulate in a finite abelian label, and that label determines the global quotient or obstruction class.*

Corollary 7.21 (Standard Model Gauge Group)**.** *Under Theorem 7.13, Theorem 7.14, Corollaries 7.16 and 7.15, and Proposition 7.19,*

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

Proof. Theorem 7.13 yields the product structure up to finite quotient, and Proposition 7.19 fixes that quotient to $\mathbb{Z}_6$ once the realized hypercharge lattice and $N_c = 3$ are in place. □

Corollary 7.22 (Structural electroweak force and charges)**.** *On the realized one-Higgs branch of Theorem 7.14 and Corollary 7.17, 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 neutral direction with

$$Q_H = T_3 + Y = 0,$$

so the unbroken generator is

$$Q = T_3 + Y,$$

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. Thus the recovered 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 the numerical W/Z masses belong to the D10 running/matching surface rather than to this recovered-core structural corollary.*

Proof. The D8 result fixes the weak-type role to the SU(2) doublet and the connected abelian role to $\mathrm{U}(1)_Y$. Theorem 7.14 fixes $Y_H = 1/2$ and the matter hypercharge lattice. Choosing the neutral Higgs component has $T_3 = -1/2$, hence $T_3 + Y = 0$ on the selected vacuum direction. Therefore precisely the $Q = T_3 + Y$ direction stabilizes that branch. The two noncommuting charged SU(2) generators are broken directions and combine into $W^{\pm}$; the orthogonal neutral broken direction is the Z carrier, while the unbroken stabilizer is the photon A. The global $\mathbb{Z}_6$ quotient fixes the compatible charge lattice used in Proposition 7.19. □

Theorem 7.23 (Realized compact-gauge witness and physical UV landing)**.** *On the ordinary or central zero-obstruction bosonic branch, the OPH compact reference architecture admits an OPH-realizable cofinal heat-kernel edge-sector witness*

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

with

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

Every witness sector has positive realized support on the compact heat-kernel branch, and the witness is MAR-admissible. Hence the MAR-admissible class used above is nonempty. Moreover every OPH-admissible microscopic UV completion of this realized branch has observer-visible low-energy package, modulo physical equivalence in Definition 3.6,

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

with the hypercharge lattice of Theorem 7.14.

Proof. Use the fixed-cutoff compact-gauge patch-carrier architecture on the zero-obstruction branch. The microphysics carrier may be represented by finite gauge-register, quantum-link, or federated echosahedral regulator charts; only the declared overlap sector algebra and refinement-compatible compact labels are used here. At a sufficiently fine compact cutoff, retain the Peter–Weyl labels listed in $\mathcal{W}_{\mathrm{SM}}$. They are finite-dimensional unitary representations of $\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)$, and with integer normalization $q = 6Y$ Proposition 7.19 shows that the generator $g_0 = (\omega_3, -1, e^{i\pi/3})$ acts trivially on exactly the displayed matter and Higgs multiplets. The witness therefore descends to the quotient $(\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1))/\mathbb{Z}_6$.

The fixed-cutoff edge heat-kernel law on the microphysics surface gives sector weights

$$p_R(t) \propto d_R e^{-tC_2(R)}$$

on the compact branch. For finite t, $d_R > 0$ and $e^{-tC_2(R)} > 0$, so every witness projector has positive support. The refinement functors of Theorem 7.2 carry those zero-obstruction labels along the cofinal tail, and Theorem 7.3 places the corresponding objects in Sect_∞.

The witness is loop-coherent by the zero-obstruction branch. Theorem 7.14 supplies anomaly cancellation and one-Higgs Yukawa completeness on the displayed chiral package; Lemmas 7.8–7.10 give the weak-type and color-type roles on the minimal coupled carrier; Corollary 7.15 gives $N_c = 3$; and Corollary 7.16 gives $N_g = 3$ from the CP-capability and weak-sector UV clauses in Axiom 3.5. Thus $\mathcal{W}_{\mathrm{SM}}$ is an occupied MAR-admissible package.

Proposition 3.7 supplies minima for the resulting nonempty admissible class. Corollary 7.17 and Proposition 7.19 identify every MAR-minimal observer-visible package with the displayed Standard

Model package. Finally, Remark 3.8 and Definition 3.6 remove only microscopic regulator choices, implementation hiding, gauge-center conventions, generation relabeling, charge-conjugation convention, and inert ancillary stabilization. Thus any OPH-admissible UV completion of the realized branch lands on the same observer-visible low-energy package, not on a unique microscopic representative. □

Theorem 7.24 (Realized Standard Model branch)**.** *Assume:*

(1) *the ordinary or central zero-obstruction bosonic branch of Theorem 5.8;*

(2) *the fixed-cutoff bosonic sector category, refinement/fiber descent, and compact-group reconstruction of Theorems 5.11, 7.2, 7.3, and 7.4;*

(3) *a nonempty realized one-Higgs chiral MAR-admissible class, witnessed by $\mathcal{W}_{\rm SM}$ in Theorem 7.23;*

(4) *the weak-type and intrinsically complex color-type roles on a common coupled carrier;*

(5) *one connected abelian charge factor acting nontrivially on that carrier; and*

(6) *the anomaly-free, one-Higgs Yukawa-complete, CP-capable, and weak-sector UV clauses of Axiom 3.5.*

Then MAR minima exist, and every MAR-minimal observer-visible package in that class is physically equivalent to

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

where the matter package is

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

with

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

Proof. Proposition 3.7 gives existence of MAR minima because the admissible class is nonempty and $C(\mathfrak{S})\in\mathbb{N}^4$ is well ordered lexicographically. Lemmas 7.8 and 7.9 identify the minimal weak-type and color-type connected nonabelian roles as SU(2) and SU(3). Lemma 7.10 puts those roles on the minimal coupled carrier $\mathbb{C}^3\otimes\mathbb{C}^2$, not on the uncoupled $\mathbb{C}^3\oplus\mathbb{C}^2$, and Lemma 7.12 identifies the unique connected abelian commutant as U(1). Theorem 7.13 therefore gives product structure up to finite central quotient. Theorem 7.14 fixes the hypercharge lattice, Corollary 7.15 fixes $N_c=3$, Corollary 7.16 fixes $N_g=3$ as a MAR-branch consequence rather than as anomaly cancellation alone, and Proposition 7.19 fixes the global kernel. Corollary 7.17 then identifies all minima modulo the physical equivalence relation of Definition 3.6. □

Corollary 7.25 (Product-Group Consequences)**.** *The adjoint representation of the full connected gauge group contains only*

$$(8,1,0)\oplus(1,3,0)\oplus(1,1,0).$$

Equivalently, the adjoint representation of the derived gauge group contains only

$$(8,1,0)\oplus(1,3,0).$$

There are no mixed $(3,2,\pm5/6)$ gauge bosons. Hence gauge-mediated proton decay is absent. The unbroken gauge symmetries forbid elementary gauge-boson mass terms: the photon is massless, while gluons are massless gauge fields of the unbroken color symmetry.

Remark 7.26. *The same product-group structure removes the usual grand-unification monopole sector. In the present framework, any retained coupling-unification discussion does not proceed by embedding the gauge group into a simple group that is later broken. It proceeds, if at all, only through the separate D10 calibration branch built from shared edge heat-kernel data together with the printed running/matching/threshold/scheme conventions and extra calibration assumptions. Absence of gauge-mediated proton decay and unification-like running therefore coexist without tension on that branch. Any confinement statement is a separate infrared dynamical issue, not a corollary of the product-group adjoint argument alone.*

7.3 Fundamental scales

The quantitative D10 branch is a forward quantitative-closure sector, not part of the recovered-core theorem package. Its particle-physics scale variable is

$$P \equiv a_{\mathrm{cell}}/\ell_P^2,$$

together with the realized discrete gauge data of the structural branch, in particular

$$N_c = 3, \qquad \beta_{\mathrm{EW}} = N_c + 1 = 4.$$

The logical direction is therefore

$$P \longmapsto \big(M_U(P), E_{\mathrm{cell}}(P)\big) \longmapsto \alpha_U(P) \longmapsto \big(t_U(P), t_{\mathrm{tr}}(P)\big) \longmapsto \big(t_2(P), t_3(P), v(P)\big) \longmapsto \alpha_i(\mu_*; P).$$

The edge-law input on this lane is split explicitly by claim tier. The fixed-cutoff same-overlap thermal/Casimir theorem is carried on the separate microphysics surface. This compact paper imports that closure through the declared merge boundary: one fixed-cutoff edge-law package below, then the compact-group / Peter–Weyl lift used in the D10 pixel constraint above it. The compact-group lift is a declared D10 branch handoff, while large-N_{edge} and critical-string claims belong to the continuation lane. The finite-cutoff Casimir branch itself is not a missing input. No hardware evidence, private run log, or laboratory module claim is imported into this D10 theorem step. Measured electroweak quantities are not used to infer the internal transmutation parameters. They enter only at the end as compare-only validation targets for the printed D10 running/matching/threshold/scheme package. The forward transmutation certificate on the live D10 lane makes that logical order explicit: the source-only basis reconstructs the same $\alpha_U(P)$, hence the same unified diffusion parameter

$$t_U(P) = 4\pi^2 \alpha_U(P),$$

and transmutation exponent

$$t_{\mathrm{tr}}(P) = \frac{2\pi}{(N_c+1)\alpha_U(P)}$$

as the pixel-closure solve itself. Here the paper-side transmutation factor is the same $\beta_{\mathrm{EW}} = N_c + 1$ used below; legacy overloaded β-ratios are compare-only diagnostic readouts and are not part of the theorem contract.

The scales fixed directly from P are

$$M_U(P) = \frac{E_P}{e^{2\pi}}\, P^{1/6}, \qquad E_{\mathrm{cell}}(P) = \frac{E_P}{\sqrt{P}}.$$

The factor $e^{-2\pi}$ is the same modular normalization used on the Lorentz/BW side, and $P^{1/6}$ is the cell-area scaling relation carried by the present D10 implementation.

Golden-ratio equilibrium benchmark. The total/bulk/edge hierarchy has one exact self-similar balance point. Writing

$$x(C) := \frac{S_{\text{gen}}(C)}{S_{\text{bulk}}(C)} = 1 + \frac{\langle L_C \rangle}{S_{\text{bulk}}(C)},$$

the exact self-similar balance condition

$$\frac{S_{\text{gen}}(C)}{S_{\text{bulk}}(C)} = \frac{S_{\text{bulk}}(C)}{\langle L_C \rangle}$$

gives

$$x = \frac{1}{x-1}, \qquad x^2 - x - 1 = 0.$$

Hence the unique positive equilibrium point is

$$x = \varphi := \frac{1+\sqrt{5}}{2}.$$

Equivalently, the equilibrium-breaking order parameter

$$A_\varphi(x) := x - 1 - \frac{1}{x}$$

vanishes exactly at $x = \varphi$. In that sense φ is the exact self-similar balance point rather than a numerological comparison constant. The synthesis paper gives the outer/inner closure relation that fixes the realized value of P by matching that detuning to the inner electromagnetic observation scale emitted by the same cell. The role of the equilibrium theorem is to explain why the realized value sits close to φ. Exact equilibrium is too symmetric to support durable records, structure, and dynamics, so the realized branch sits at a small equilibrium-breaking detuning away from it. The technical question for the quantitative branch is the size of that detuning together with the reduced-residual/root-control analysis of the printed solve.

The one-dimensional internal variable solved on the D10 branch is the unified coupling α_U. For a trial value of α_U, define

$$v(\alpha_U, P) := E_{\text{cell}}(P) \exp\left(-\frac{2\pi}{\beta_{\text{EW}}\alpha_U}\right).$$

Run the one-loop D10 couplings from

$$\alpha_1(M_U) = \alpha_2(M_U) = \alpha_3(M_U) = \alpha_U$$

down to a scale μ using

$$\alpha_i^{-1}(\mu; \alpha_U, P) = \alpha_U^{-1} + \frac{b_i}{2\pi} \log \frac{M_U(P)}{\mu},$$

with the printed one-loop coefficients b_i. Let $\mu_* = \mu_*(\alpha_U, P)$ be the fixed point determined by the tree-level Z-mass relation

$$\mu_* = \frac{v(\alpha_U, P)}{2} \sqrt{g_2(\mu_*)^2 + g_Y(\mu_*)^2}, \qquad \alpha_Y = \frac{3}{5}\,\alpha_1.$$

At that fixed point define

$$t_2(\alpha_U, P) = 4\pi^2 \alpha_2(\mu_*; \alpha_U, P), \qquad t_3(\alpha_U, P) = 4\pi^2 \alpha_3(\mu_*; \alpha_U, P).$$

Let $\bar{\ell}_{\mathrm{SU}(2)}(t)$ and $\bar{\ell}_{\mathrm{SU}(3)}(t)$ denote the nonabelian edge-entropy functions appearing in the D10 pixel constraint. The pixel-closure functional is

$$\mathcal{F}(\alpha_U; P) := \bar{\ell}_{\mathrm{SU}(2)}(t_2(\alpha_U, P)) + \bar{\ell}_{\mathrm{SU}(3)}(t_3(\alpha_U, P)) - \frac{P}{4}.$$

The printed D10 package takes $\alpha_U(P)$ to be the branch value selected by solving

$$\mathcal{F}(\alpha_U; P) = 0.$$

Only after that forward solve are the internal transmutation parameters fixed:

$$t_U(P) := 4\pi^2 \alpha_U(P), \qquad t_{\mathrm{tr}}(P) := \frac{2\pi}{(N_c + 1)\alpha_U(P)},$$

and with them the downstream scale data

$$t_2(P) := t_2(\alpha_U(P), P), \qquad t_3(P) := t_3(\alpha_U(P), P), \qquad v(P) := v(\alpha_U(P), P).$$

Thus the branch runs from OPH input P to the internal transmutation data $(t_U(P), t_{\mathrm{tr}}(P))$, then to the downstream local scale data $(t_2(P), t_3(P), v(P))$, and only then to the low-energy couplings. It does not use measured $\alpha_i(m_Z)$ to infer those internal t-parameters.

Every quoted D10 electroweak number is downstream of this forward map. In particular, the source-locked running-family anchor $a_0(P) = \alpha_{\mathrm{em}}^{-1}(m_Z^2; P)$, the declared electromagnetic transport family $\alpha_{\mathrm{em}}^{-1}(q^2; P)$ and $\sin^2\theta_W(q^2; P)$, and the frozen public compare-only W/Z validation rows sit on the printed quantitative surface. The Thomson endpoint

$$\alpha_{\mathrm{Th}}^{-1}(P) = \lim_{q^2 \to 0} \alpha_{\mathrm{em}}^{-1}(q^2; P)$$

on the Ward-projected $\mathrm{U}(1)_Q$ lane is declared through the source-spectral reduction theorem and remains gated by the populated source spectral measure payload, same-scheme remainder, and interval certificate. When compared with observation, these rows serve as validation of that printed implementation on the quantitative-closure branch recorded in the theorem checklist. The repository carries a numerical witness for the outer/inner fixed-point closure, and the public endpoint value is kept out of the closure solve.

8 Relation to Holography and Existing UV Frameworks

OPH belongs to the holographic family of ideas, but its primitive data are different from those used in asymptotic-boundary constructions. The distinction matters because several of the central claims of this paper, including the treatment of Λ, the handling of factorization, and the emergence of gauge structure, depend on it.

8.1 Static-patch screen versus asymptotic boundary

In asymptotic-boundary holography one starts from a dual theory at infinity and reconstructs the bulk inward. Here one starts from a finite-capacity screen equipped with a net of local patch algebras and reconstructs the effective bulk from overlap consistency. The primitive object is therefore a family of observer patches with nontrivial overlaps, not a single global boundary theory.

This shift has two technical consequences. First, subsystem factorization is treated from the beginning as a gluing problem with centers and sector labels. Second, the physical role of the horizon is local and operational: every observer has direct access only to a patch algebra, and global law is whatever survives reconciliation on overlaps.

8.2 Positive Λ and finite capacity

The D5→D6 cosmological-capacity stack is explicit: the null-modular reconstruction of Section 6 determines the stress tensor only up to a metric term, so local null data do not fix Λ. The global value is the cosmic record-capacity fixed point

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$$

together with its observed-branch de Sitter entropy readout

$$N_{\mathrm{CRC}} = S_{\mathrm{dS}},$$

the standard de Sitter entropy relation

$$S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4G} = \frac{3\pi}{G\Lambda},$$

On that same branch this yields

$$\Lambda_{\mathrm{CRC}} = \frac{3\pi}{GN_{\mathrm{CRC}}}$$

and fixes the static-patch package

$$S_{\mathrm{dS}} = N_{\mathrm{CRC}}, \qquad r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\mathrm{dS}}}{c},$$

while the observed cosmic age is a downstream FLRW benchmark rather than an additional D6 theorem output. The compact paper is therefore formulated in a de Sitter-first language: the positive cosmological constant is tied to finite capacity rather than to a deformation of a negative-Λ starting point. The input-free global closure is the cosmic record-closure readback fixed point

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}), \qquad F(N) = \mathrm{Cap}_{\mathrm{read}}\big(\mathrm{Obs}(\mathrm{nf}(\mathfrak{U}_N))\big).$$

The screen-normalized self-closure density is its count representation:

$$\Pi(N) = |\Omega_N^{\mathrm{sc}}|e^{-N}.$$

In differentiable form, the OPH-derived stationarity map $T_\eta(N) = N + \eta\, d\log\Pi(N)/dN$ has a unique stable fixed point under the derivative-sign certificate. On the observed branch this fixed point is the de Sitter entropy capacity.

8.3 Factorization, gauge structure, and the string sector

The factorization problem of gauge theory and gravity is often treated as a technical nuisance. Here it is part of the architecture. Edge-center completion turns the collar center into a first-class object, and the entropy split

$$S_{\mathrm{gen}} = \langle L_C\rangle + S_{\mathrm{bulk}}$$

is then a structural consequence rather than an added prescription.

In the bosonic EFT branch, compact gauge reconstruction follows after Theorem 5.8 supplies the strict zero-obstruction transport criterion, Theorem 5.11 constructs the fixed-cutoff bosonic collar-sector categories, and Theorem 7.2 constructs the refinement functors and finite bosonic fiber descent. Crossed-module data handle the genuinely noncentral fixed-cutoff branch separately when $q_\Sigma \neq 0$; the ordinary compact-group theorem does not. This obstruction calculus is a classification and routing theorem, not a Standard Model selection theorem. DR/Tannaka reconstruction yields some compact group $G = \mathrm{Aut}_{\otimes}(\mathcal{F})$ from the zero-obstruction transportable sector category. On the realized MAR-admissible branch with the explicit realized one-generation chiral matter plus one-Higgs package, the Standard Model quotient is then selected by minimal admissibility. This differs sharply from approaches in which the gauge group is specified in advance.

The worldsheet relation is similarly reversed. The edge partition function

$$Z_{\mathrm{edge}}(t) = \sum_R d_R^2 e^{-tC_2(R)}$$

is the closed partition function obtained after gluing the open-edge weights

$$p_R(t) \propto d_R e^{-tC_2(R)}.$$

Theorem 8.1 (OPH-to-2D-Yang–Mills edge partition theorem)**.** *Assume the compact-group heat-kernel branch supplied by the fixed-cutoff edge-sector law together with the compact-group / Peter–Weyl lift used in the bosonic gauge lane, in the same quadratic-Casimir normalization as the compact heat kernel. Equivalently, assume the open-edge weights satisfy*

$$p_R(t) \propto d_R e^{-tC_2(R)}$$

on a compact gauge group G, and define the closed edge partition function by gluing the two edge boundaries:

$$Z_{\mathrm{edge}}(t) = \sum_R d_R^2 e^{-tC_2(R)}.$$

Then:

(i) *the closed edge partition function is exactly the compact-group heat kernel at the identity,*

$$Z_{\mathrm{edge}}(t) = K_t(1);$$

(ii) *the same sum is therefore the standard compact-group two-dimensional Yang–Mills heat-kernel partition function at identity, equivalently the closed-surface heat-kernel partition sum on that branch;*

(iii) *the Chapman–Kolmogorov law for K_t is the corresponding collar-sewing rule for the OPH edge partition.*

Thus the OPH edge-sector partition reorganizes exactly into the two-dimensional Yang–Mills heat-kernel form before any large-N_{edge} continuation is invoked.

Proof. The fixed-cutoff edge-sector theorem gives the heat-kernel weights $d_R e^{-tC_2(R)}$ on the compact-group branch. Gluing the two edge boundaries contributes a second factor of d_R, so the closed partition function is the displayed sum. Peter–Weyl identifies

$$\sum_R d_R \chi_R(g) e^{-tC_2(R)}$$

with the compact-group heat kernel $K_t(g)$, and evaluating at the identity $g = 1$ gives

$$K_t(1) = \sum_R d_R \chi_R(1) e^{-tC_2(R)} = \sum_R d_R^2 e^{-tC_2(R)} = Z_{\mathrm{edge}}(t),$$

since $\chi_R(1) = d_R$. The Chapman–Kolmogorov law for K_t is exactly the semigroup gluing law, so it gives the collar-sewing rule on the same branch. □

Scope boundary. The imported inputs are exactly the fixed-cutoff edge heat-kernel / Casimir law, the compact gauge-group / Peter–Weyl lift, the quadratic-Casimir heat-kernel normalization on that lift, and the Peter–Weyl heat-kernel identity. The external content is any large-N_{edge} regime, the Gross–Taylor dictionary on that regime, and every further critical-string ingredient. The theorem above proves the OPH-to-2D-Yang–Mills partition reorganization itself; it does not by itself produce a worldsheet genus expansion. The four-dimensional compact-gauge repair-gap theorem is the separate support-visible repair-dynamics result in Theorem A.28.

Theorem 8.2 (Criterion for a controlled large-N_{edge} worldsheet effective description)**.** *Assume Theorem 8.1 and a large-N_{edge} realization of the compact-group heat-kernel branch, with $N_{\text{edge}} \neq N_c = 3$, for which the 't Hooft-style variable*

$$\tau := tN_{\text{edge}}$$

is kept in a compact interval $I \subset (0, \infty)$. Suppose moreover that for every truncation order $G \geq 0$ there exist coefficient functions $F_g : I \to \mathbb{R}$ and constants $C_{G,I}$ such that

$$\log Z_{\text{edge}}\left(\frac{\tau}{N_{\text{edge}}}\right) = \sum_{g=0}^{G} N_{\text{edge}}^{2-2g} F_g(\tau) + R_{G+1}(\tau, N_{\text{edge}})$$

with

$$\left|R_{G+1}(\tau, N_{\text{edge}})\right| \leq C_{G,I} N_{\text{edge}}^{-2G} \qquad (\tau \in I).$$

Then this branch defines a controlled theorem-level worldsheet effective description of the edge dynamics, with control parameters N_{edge}^{-2}, the fixed-τ window I, and the truncation order G.

Proof. Theorem 8.1 puts the edge partition function on the compact-group heat-kernel / two-dimensional Yang–Mills surface. On that surface, the standard Gross–Taylor dictionary reads a genus expansion of the displayed form as a closed-worldsheet rewriting of the free energy. The stated large-N_{edge} assumption supplies that expansion on the fixed-τ window I, together with an explicit truncation bound. Therefore the worldsheet rewriting is controlled by the displayed parameters N_{edge}^{-2}, I, and G. The statement is only an effective description on the declared branch: it does not by itself construct a critical worldsheet CFT, derive modular invariance, worldsheet supersymmetry, anomaly cancellation, GSO projection, or full massless-spectrum matching. □

External inputs. The criterion theorem uses the declared large-N_{edge} regime and the imported Gross–Taylor large-N worldsheet dictionary for two-dimensional Yang–Mills. The external content is the existence of the large-N_{edge} sequence, the fixed-τ window on which the genus expansion holds, and the uniform remainder control. Any further lift to critical superstring structure requires worldsheet supersymmetry, critical dimension, modular invariance, anomaly cancellation, GSO projection, and full massless-spectrum matching beyond the present declared continuation theorem.

Feature	Asymptotic-boundary holography	OPH
Primitive data	Boundary theory at infinity	Finite-capacity screen with observer-patch net
Factorization across cuts	Subtle and often indirect	Explicit edge-center completion with central labels
Status of Λ	External background datum	Fixed globally by N_{scr} after local null reconstruction
Gauge group	Usually specified as part of the model	Reconstructed from edge sectors, then fixed on the realized MAR branch
String sector	Fundamental dual description	Controlled large-N_{edge} worldsheet effective description of edge dynamics on the stated branch

9 Claim Tiers and Falsifiability

The recovered core is

$$(D1\text{–}D5) \cup (D7\text{–}D9),$$

namely the relativity chain together with the realized Standard Model structural chain. D6 is the global closure of that same Einstein branch at the cosmic record-capacity fixed point, D10 is the integrated quantitative-closure branch, and D12 collects phenomenological continuations.

9.1 Prediction audit by claim tier

Tier	Representative outputs	What would actually falsify the tier
Phase I recovered core (Theorem 3.11; D1–D5, D7–D9)	Confluence of overlap repair, Lorentz kinematics, the Jacobson-type Einstein branch in the stated scaling regime, compact gauge reconstruction in the bosonic branch, the support-visible four-dimensional Euclidean Yang–Mills form and compact-gauge repair gap under the declared compact-gauge branch assumptions, the realized Standard Model quotient chain on the explicit realized matter package, the exact hypercharge lattice on that realized matter package, the realized color triplet $N_c = 3$, the generation count $N_g = 3$, and the product-group consequence of no gauge-mediated proton decay	A mathematical failure in the derivation chain, loss of the Yang–Mills form or exact repair-gap branch conditions, or data that require a different realized gauge quotient, different realized hypercharges, a different color or generation count on the admissible branch, or gauge-mediated proton decay
Phase II zero-input global closure (D6)	$\Lambda_{\mathrm{CRC}} = \frac{3\pi}{GN_{\mathrm{CRC}}}$, $S_{\mathrm{dS}} = N_{\mathrm{CRC}}$, $r_{\mathrm{dS}} = \sqrt{3/\Lambda_{\mathrm{CRC}}}$, and $t_\Lambda = r_{\mathrm{dS}}/c$ from the cosmic record-closure fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$, with $F(N) = \mathrm{Cap}_{\mathrm{read}}(\mathrm{Obs}(\mathrm{nf}(\mathfrak{U}_N)))$, represented in counts by $N_\star = \mathrm{MAR}\arg\max_N(\log\lvert\Omega_N^{\mathrm{sc}}\rvert - N)$; any retained capacity-level neutrino side estimate on that branch value is legacy bookkeeping rather than the weighted-cycle neutrino theorem lane; the observed cosmic age is a downstream FLRW benchmark	Failure of the readback fixed-point certificate or of the specific global capacity relation would falsify the D6 closure without erasing the recovered core

Tier	Representative outputs	What would actually falsify the tier
Phase II quantitative-closure branch (D10)	forward transmutation data $\alpha_U(P)$, $t_U(P)$, $t_{\rm tr}(P)$, pixel-closure gauge-coupling consistency, $\alpha_i(m_Z)$, the source-locked anchor $a_0(P) = \alpha_{\rm em}^{-1}(m_Z^2; P)$, the declared electroweak transport family $(W, Z, \alpha_{\rm em}^{-1}(q^2), \sin^2\theta_W(q^2), v)$, and the Thomson endpoint $\alpha_{\rm Th}^{-1} = 137.035999177(21)$ on the OPH plus empirical hadron closure row. The source-side audit trunk emits $\alpha_{\rm cand}^{-1} \simeq 136.9948351646$ at $P_{\rm cand} \simeq 1.6309720957$. The endpoint residual table records $0.0414658610\ldots$ inverse-alpha units at the public endpoint pixel, with $S_{\rm required} \simeq 0.8954001326$ and $c_Q \simeq 0.6580257599$. Empirical $e^+e^- \to$ hadrons input is kept in a separate row class from the source-only theorem	Failure of the integrated quantitative closure once P and the printed running/matching/scheme conventions are imposed, or failure of the explicit source spectral payload, same-scheme remainder, empirical source disclosure, and interval-certificate records
Phase III phenomenological continuations (D12 and beyond)	Flavor ansätze, charged-lepton continuation ansätze beyond exact centered readback, texture branches, the weighted-cycle / Majorana-holonomy neutrino branch above the legacy D6 side estimate, further neutrino mass/mixing refinements, dark-sector response laws, CMB/inflation-replacement kernels, H_0/S_8 and growth continuations, heuristic baryogenesis continuations, proton-spin bookkeeping, proton-lifetime estimates beyond the gauge-channel exclusion, black-hole spectroscopy templates, controlled large-$N_{\rm edge}$ string/worldsheet effective descriptions, conjectural critical-superstring extensions, and other downstream phenomenology	Retraction of the corresponding continuation only. These branches are not part of the recovered core and their failure is not, by itself, a falsification of relativity-plus-Standard-Model recovery

Phase I is the recovered-core tier, Phase II is adjacent but non-core, and Phase III contains phenomenological continuations. The D6 row permits the conditional holonomy reading of FLRW flatness in Lemma 6.43; it does not promote an inflation replacement, CMB likelihood, or high-H_0 branch. Current low-H_0, Planck-like S_8 dark/anomaly rows are therefore diagnostic continuation checks until a homogeneous anomaly-load selector and growth kernels are emitted before fitting cosmological data. On the exact finite packet-closed quotient branch, the fixed-cutoff consensus surface does define an affine OPH closure map: the schedule-independent normal-form projection pushes forward to a continuous idempotent self-map of the finite packet simplex, its image is exactly the normal-form simplex, and its fixed points are exactly the packets supported on quotient normal forms. That finite branch also supplies the grammar from which the cosmic readback map F and the self-closing observer-normal-form density $\Omega_N^{\rm sc}$ are defined. The readback fixed point gives the

D6 target, with the screen-normalized density as its count representation; interpretive strange-loop discussions and the full habitat theorem sit outside these tiers and are not needed for the finite packet-quotient fixed-point statement.

9.2 Recovered-core discriminators

Once the claim boundary is made explicit, the sharpest discriminators are the ones that do not depend on auxiliary flavor, dark-sector proposals, heuristic baryogenesis continuations, or string ansätze.

9.2.1 Realized Standard Model gauge structure

On the realized MAR-admissible branch, this paper fixes

$$G_{\text{phys}} = \frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_g = 3, \qquad N_c = 3,$$

with the exact one-generation hypercharge lattice. Here the count statements are internal to the same realized branch: the D8 minimal coupled-carrier theorem stack fixes the color triplet and hence $N_c = 3$; CKM phase counting and weak-sector asymptotic freedom then bound N_g; MAR fixes the smallest realized value $N_g = 3$; and Witten parity is only a consistency check on the resulting triplet-doublet package.

$$Y_Q = \frac{1}{6}, \qquad Y_L = -\frac{1}{2}, \qquad Y_u = -\frac{2}{3}, \qquad Y_d = \frac{1}{3}, \qquad Y_e = 1, \qquad Y_H = \frac{1}{2}.$$

Evidence that the realized low-energy gauge structure or realized hypercharges differ from these values would directly contradict the recovered core.

9.2.2 Known-force and charge coverage

This paper does not leave any known long-range or Standard Model gauge force unassigned. The coverage is explicit by tier:

- **Gravity**: D3–D6 recover Lorentz kinematics, the null-stress bridge, and the Jacobson-type Einstein branch, with T_{ab} as the stress-energy source and the graviton as the quantum of the dynamical metric branch.
- **Strong interaction**: D8–D9 recover the $\mathrm{SU}(3)_c$ color factor, the color triplet $N_c = 3$, quark color triplet/antitriplet assignments, and the eight gluon generators $(8, 1, 0)$. Confinement and hadron spectra are separate infrared QCD questions, not missing gauge-charge assignments.
- **Weak interaction**: D8–D9 recover the $\mathrm{SU}(2)_L$ weak doublet structure, the one-Higgs branch, and the charged weak carriers $W^{\pm}$ from the broken $\mathrm{SU}(2)_L$ generators. D10 supplies the quantitative running/matching readout for v and the W/Z validation rows.
- **Hypercharge and electromagnetism**: Theorem 7.14, Proposition 7.19, and Corollary 7.22 fix the $\mathrm{U}(1)_Y$ lattice, the unbroken generator $Q = T_3 + Y$, integer charge for color singlets, and the photon A as the massless $\mathrm{U}(1)_Q$ carrier. The Thomson-limit fine-structure endpoint is a D10 quantitative closure.

9.2.3 Product-group corollary

Because the realized gauge structure is a product group up to the finite central quotient fixed above, there are no mixed $(3, 2, \pm 5/6)$ gauge bosons. Hence

$$\tau_p^{(\text{gauge})} = \infty$$

for gauge-mediated proton decay. Observation of proton decay specifically attributable to gauge-boson exchange would therefore falsify the realized product-group branch.

9.2.4 Relativity branch

The relativity claim is sharp about its scope: given the support-visible geometric-modular theorem and the stated null-bridge conditions in the intended scaling regime, and the same realized cap-label-preserving MaxEnt family satisfying the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations on that branch, this paper predicts Lorentz kinematics on the screen and a Jacobson-type Einstein branch. What is being tested here is the existence of that scaling regime and its branch conditions, not a fixed-cutoff matrix-algebra identity.

9.3 What does *not* count as a contradiction of the recovered core

The following do *not* falsify Theorem 3.11 if they fail as stated:

1. the uniform $\mathbb{Z}_6$ center-label ensemble and the associated $\varepsilon = 1/6$ flavor ansatz;
2. charged-lepton continuation ansätze beyond the exact centered-readback / common-shift frontier, texture exponents, or legacy capacity-level neutrino side estimates;
3. modular-anomaly dark-sector response ansätze, including MOND/RAR-style scaling claims;
4. heuristic baryogenesis continuations based on extra suppression-counting assumptions; a baryogenesis theorem requires a concrete out-of-equilibrium mechanism, derived CP-odd source terms, justified defect/sphaleron counting, washout control, and a freeze-out computation of the final asymmetry;
5. strong-CP continuations; the available corpus does not derive the bare QCD angle θ_{QCD}, the physical anomaly-invariant combination $\bar{\theta}$, or a proof that $\bar{\theta} = 0$;
6. proton-spin ansätze or proton-lifetime estimates beyond the structural exclusion of gauge-mediated decay;
7. discrete-horizon spectroscopy templates;
8. controlled large-N_{edge} string/worldsheet effective descriptions.

These are all post-Phase-I, non-core branches. Some are Phase-II implemented estimates or conditional continuation checks, while the rest are Phase-III continuations. None is part of this paper's recovered relativity-plus-Standard-Model core.

10 Common Objections and Clarifications

10.1 Is the use of P circular?

The pixel area P is fixed by the outer/inner closure condition in the synthesis paper. That condition matches the outer pixel detuning to the inner electromagnetic observation scale emitted by the same cell. Within the printed quantitative implementation, that P first fixes $M_U(P)$ and $E_{\text{cell}}(P)$, then a one-dimensional pixel-closure solve fixes $\alpha_U(P)$, equivalently the internal transmutation data $t_U(P)$ and $t_{\text{tr}}(P)$, and only then do $t_2(P)$, $t_3(P)$, $v(P)$, and the running electroweak outputs appear. Quantities algebraically entangled with that quantitative branch are on the Phase-II implementation surface: they are forward-emitted there, and comparison with observation checks that printed implementation rather than enlarging the recovered-core claim set.

Operationally, changing P moves the quantitative family

$$(\alpha_U, \alpha_i(m_Z), a_0, \alpha_{\text{em}}^{-1}(q^2), \sin^2\theta_W(q^2))$$

through that forward solve. It does not alter the parameter-free structural outputs such as the gauge quotient or hypercharge lattice. Downstream matter-sector continuations, including charged-lepton continuation ansätze beyond the exact centered-readback / common-shift frontier, are outside this SM/GR derivation paper's recovered-core theorem package.

10.2 Does a fixed UV cell size break Lorentz invariance?

The UV regulator is placed on the screen algebra. The emergent bulk spacetime has no microscopic bulk lattice in this formulation. Lorentz kinematics arise from the continuum modular action on caps, summarized in Corollary 6.16.

Spherical geometry is the bridge that makes this possible. A support-visible observer cut is charted by S^2. Round caps on that chart carry the modular flows used by the BW branch, and orientation-preserving conformal maps of the same sphere give

$$\mathrm{Conf}^+(S^2) \cong \mathrm{SO}^+(3,1).$$

The Lorentz group therefore enters as the symmetry of observer-facing cap geometry. A finite echosahedral or cellulated carrier sits on the regulator side; the Lorentz claim concerns the support-visible scaling limit of the spherical cap chart extracted from that carrier.

This is conceptually similar to other continuum limits in statistical mechanics and lattice field theory. A regulator may break a symmetry at finite cutoff while the infrared fixed point restores it exactly. Here the restoration is stronger than a generic RG expectation because the modular-flow theorem identifies the symmetry group itself:

$$\mathrm{Conf}^+(S^2) \cong \mathrm{SO}^+(3,1).$$

Residual Lorentz violation would therefore have to arise from a failure of the support-visible modular-flow theorem. A UV cell size by itself is insufficient. The UV-side burden on the Lorentz branch is therefore the support-visible scaling theorem on the prime geometric subnet, not the mere existence of a finite cell size.

10.3 Why use type-I algebras if continuum QFT uses type III?

Type I is a regulator statement. At every finite cutoff the patch and collar algebras are finite-dimensional matrix algebras, so entropy, recovery maps, and

$$K = -\log\rho$$

are literal finite-dimensional objects. The Lorentz branch is *not* a claim that the continuum observer algebra stays in that class. The refinement picture is

$$\text{finite type-I regulator net} \longrightarrow \text{scaling-limit observer net } \mathcal{A}_\infty,$$

and the limit can leave the regulator class. Axiom 3.3 controls the realized state-side branch across refinement; it does *not* by itself prove that the refinement limit is type I or that the emitted scaling-limit pair lies in the BW / canonical cap phase with standard geometric modular action.

On the geometric-subnet branch of Theorem 6.8, the scaling-limit cap algebras may be non-type-I and, in the continuum-QFT case of interest, are expected to be type III. Then there is generally no cap density matrix $\rho_C \in \mathcal{A}_\infty^{\mathrm{geo}}(C)$ and no bounded operator $K_C = -\log \rho_C$ inside the limit algebra. The correct continuum object is the modular automorphism group of the pair $(\mathcal{A}_\infty^{\mathrm{geo}}(C), \omega_\infty^{\mathrm{geo},C})$, which on that branch acts geometrically and typically outerly.

There is therefore no contradiction between using type-I regulators and aiming at a type-III continuum. The finite regulator provides the collar decomposition, recovery control, and carried errors. What the fixed-cutoff MaxEnt package adds is control of the realized state-side branch through one common finite-dimensional multiplier family under refinement. Theorem 6.8 does not require the continuum algebra to be type I and does not require a full-algebra unregularized common spectral floor. It extracts the support-visible prime geometric cap pair, proves the geometric modular automorphism statement on that observer-facing limit, and leaves only the false stronger full-algebra statement unclaimed.

This is analogous to ordinary lattice field theory: one computes at finite regulator in a type-I algebra and then asks which continuum algebraic phase the scaling limit realizes. The additional structural point here is that the realized state-side branch is controlled under refinement. Theorem 6.8 extracts the support-visible cap pair on the prime geometric subnet and proves BW rigidity at automorphism level through regularized modular transport, weak-$*$/GNS extraction, support-readable modular covariance, round-cap rigidity, and KMS/BW normalization. Only stronger claims about off-support full-algebra directions or unique microscopic phase representatives lie outside that theorem.

10.4 UV branch and scaling-limit scope

At fixed cutoff, Axiom 3.3 yields a quasi-local finite-range interacting branch, and Propositions 4.11 and 4.15 show that the physical normal form is unique only modulo quotient-level OPH-stable equivalence, while the terminal expectation functionals on the declared fixed-point / quotient-local physical algebras are schedule-independent on that carrier even when microscopic representatives differ by gauge labels globally or by sector labels on one regional quotient-local glued state. The invariant fixed by the axiom language is the class $[\mathfrak{U}]_{\mathrm{OPH}}$ of the physical branch modulo implementation hiding and inert ancillary stabilization, not a unique microscopic representative. The genuinely noncentral topological case is also closed at fixed cutoff by the higher-gauge package of Theorems 5.3–5.6 and Proposition 5.13, so the fixed-cutoff topological UV surface is closed on the ordinary, central-defect, and crossed-module branches alike. Beyond fixed cutoff, Theorem 6.8 closes the support-visible continuum modular/geometric lift on the prime geometric subnet; on the bosonic compact-gauge lane, the zero-obstruction sector category, monoidal refinement ladder, $3+1$-dimensional symmetric braiding, faithful bosonic fiber functor, and realized witness package are supplied by Theorems 5.8, 5.11, 7.2, 7.3, and 7.23. The consensus paper also classifies that D1 lane explicitly: on each fixed finite patch net, normal-form computation is finite-state and decidable with the Lyapunov step bound, the automatic approximate-stability inputs are the collar-local splice and record controls, and long-run noisy approximate consensus requires the separate fair-block contraction certificate. Computational universality for growing patch-net families is an expressive-power question in the consensus paper, not a premise for the Lorentz, local Einstein, compact-gauge, or realized Standard Model branches.

10.5 Why is charge quantization possible without a simple GUT?

The framework uses the global quotient and anomaly structure rather than a simple-group embedding. Once the realized matter content is fixed, the subgroup acting trivially on all states is exactly $\mathbb{Z}_6$, and the anomaly equations force the sixth-integer hypercharge lattice. Integer electric charge for color singlets is then a property of the realized quotient structure. A simple unified gauge group is not required.

11 Discussion

The framework is strongest where the outputs are discrete or structurally rigid:

1. the Lorentz branch on the support-visible extracted prime geometric cap pair;
2. the scaling-limit Einstein branch under the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations and the null modular bridge together with the bounded-interval projective branch;
3. the Standard Model gauge quotient chain on the explicit realized matter package;
4. exact hypercharges and structural electroweak force content on the realized one-generation matter package with one Higgs doublet;
5. the realized color triplet $N_c = 3$ and generation count $N_g = 3$ on the realized MAR-admissible branch;
6. the compact two-fixed-point quantitative layer, kept explicit as a secondary layer rather than as part of Phase I.

The local MaxEnt branch gives a finite-range interacting fixed-cutoff dynamics, Propositions 4.11 and 4.15 fix uniqueness of the UV branch only modulo OPH-stable equivalence, and Theorems 5.3–5.6 with Proposition 5.13 close the genuinely noncentral topological branch at fixed cutoff. The continuum BW/geometric lift is closed in the support-visible sense by Theorem 6.8. The unregularized full-algebra common-floor route is deliberately not claimed, because Proposition 6.11 shows that off-support directions can collapse while every finite stage remains faithful and Markov. The observer-facing theorem uses regularized modular transport, weak-$*$/GNS extraction of the support-visible prime cap pair, support-readable modular covariance, ordered cut-pair rigidity, and KMS/BW normalization; this is the content needed for the Lorentz, null-stress, and local Einstein branches.

Shared excitation dictionary and flavor theorem boundaries. The D10 quantitative-closure branch gives integrated gauge-coupling closure on the displayed carrier. The pixel fixed-point equation fixes the arithmetic bridge from P through $\alpha_U(P)$, the transmutation data, the electroweak source anchor, and the Ward-projected Thomson endpoint on that branch. The exact W/Z chart is a compare-only validation sidecar beneath the public promotion gate; source spectral measure payloads, same-scheme remainders, and interval certificates document that public validation surface rather than supply a missing fine-structure bridge. On the declared runtime surface the running/matching packet is a declared-convention contract rather than a theorem that derives every coefficient, threshold, and conversion. The separate D6 cosmological-parameter package closes at the cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$. On the flavor side the manuscript treats one shared OPH excitation dictionary as the only proof-facing family datum; lane-specific integers are not introduced independently. For each realized same-label refinement arrow e on the three-generation bundle, let ω_e be the same-label overlap holonomy scalar, let

$$d_e := 1 - \omega_e$$

be the derived edge defect, let g_e be the same-label gap scalar, and define

$$q_e := \sqrt{g_e d_e}, \qquad \eta_e := \log q_e - \frac{1}{3}\sum_f \log q_f, \qquad \mu_e := \frac{e^{\eta_e}}{\frac{1}{3}\sum_f e^{\eta_f}}.$$

These are the common descendants of overlap holonomy and edge-defect data. They exist once the same-label gap and overlap witnesses are fixed, and they are the shared family-side input used by the charged, quark, and neutrino continuations.

To include the charge-side bookkeeping explicitly, let s_e denote the realized sector-charge step carried by the same-label arrow e. For any lane-specific readout γ on the realized same-label transport graph, let

$$m_{\gamma,e} \in \mathbb{Z}_{\geq 0}$$

be the multiplicity with which γ traverses the elementary same-label excitation arrow e. Equivalently, along the realized transport path one has

$$s_\gamma = \sum_e m_{\gamma,e} s_e.$$

The associated excitation weight and excitation action are

$$A_\gamma := -\sum_e m_{\gamma,e} \log q_e = -\log\Big(\prod_e q_e^{m_{\gamma,e}}\Big), \qquad Y_\gamma^{\mathrm{exc}} := \prod_e q_e^{m_{\gamma,e}}.$$

So overlap holonomy, edge defects, gap data, and sector-charge transport determine the suppression law before any base choice is made. What earlier D12 language called a “texture exponent” or “defect count” is therefore not an extra primitive beyond the structural branch. The primitive flavor-side integers are the transport multiplicities $m_{\gamma,e}$. A one-number exponent in “units of $\log 6$” is only the compressed summary

$$n_\gamma^{(6)} := A_\gamma / \log 6,$$

and it becomes a literal defect count only after the additional uniform sixfold center-label collapse compatible with the realized

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

quotient. Without that extra collapse, the derived OPH data are the multiplicity vector $m_{\gamma,\bullet}$ and the excitation action A_γ, not an independently postulated flavor integer.

The imported realized-branch input here is the D9 result that the same realized branch carries the color triplet $N_c = 3$ and the generation count $N_g = 3$. This shared dictionary feeds the downstream particle continuations. Their lane-specific theorem objects, exact witnesses, and quantitative closures are developed in Ref. [2]. On the compact-paper surface, the dictionary serves as the common family-side input beneath those continuations. Old base-6 exponents are compare-only compressions of excitation action rather than primitive flavor data.

Exact non-hadron output lane. For reader-facing exact hits, the non-hadron surface is explicit and lane-split. The structural carrier chain

$$\text{Axioms 1–5} \longrightarrow D7\text{–}D9 \longrightarrow (m_\gamma, m_g, m_{\text{grav}}) = (0, 0, 0)$$

is theorem-grade structural exactness. The electroweak quantitative-closure chain

$$\text{Axioms 1–5} \longrightarrow D7\text{–}D10$$

has an exact frozen authoritative repair sidecar

$$(M_W, M_Z) = (80.377,\ 91.18797809193725)\,\text{GeV}$$

on the frozen authoritative repair surface, but that exact pair is compare-only beneath the source spectral measure payload, same-scheme remainder, and interval certificate. The D11 Higgs/top branch closes as one source-only split theorem on the declared D10/D11 quantitative surface. On the emitted D10 repair tuple $(\eta_{\text{source}}, \beta_{EW}, \lambda_{EW}, \tau_{2,\text{tree}}^{\text{exact}}, \delta n_{\text{tree}}^{\text{exact}})$, define

$$\rho_{HT} = \log(1 + \tau_{2,\text{tree}}^{\text{exact}}),$$

$$R_T = -\tau_{2,\text{tree}}^{\text{exact}} \eta_{\text{source}}^2 + \left(1 + \frac{\beta_{EW}}{28}\right)\eta_{\text{source}}^6 + \frac{\eta_{\text{source}}^8}{14} + \frac{\eta_{\text{source}}^9}{27},$$

$$R_H = \eta_{\text{source}}^5 - \frac{3}{25}\eta_{\text{source}}^6 + \frac{\lambda_{EW}\eta_{\text{source}}^6}{18} + \frac{\eta_{\text{source}}^8}{2\beta_{EW}}.$$

Then the forward split coordinates are

$$\pi_y = \frac{\eta_{\text{source}} + \left(\frac{3}{2} + \frac{\beta_{EW}}{4}\right)\rho_{HT} + R_T}{\sqrt{\pi}}, \qquad \pi_\lambda = \frac{\eta_{\text{source}} - \left(\frac{4}{3} - \frac{\beta_{EW}}{54}\right)\rho_{HT} + R_H}{\sqrt{\pi}},$$

and the declared D11 Jacobian reads out

$$\delta y_t(\mu_t) = \pi_y\, y_t^{\text{core}}(\mu_t), \qquad \delta\lambda(\mu_t) = -\frac{16}{9}\pi_\lambda\, \lambda^{\text{core}}(\mu_t).$$

On that same declared surface,

$$m_H = 125.1995304097179\ \text{GeV}, \qquad m_t^{D11} = 172.3523553288312\ \text{GeV}.$$

At the precision quoted by the PDG Higgs listing, the Higgs row lands on the 2025 Higgs average. The same surface emits a companion top coordinate. The exact public running-top row is carried by the selected-class quark theorem and uses the PDG 2025 cross-section entry. The bridge to the auxiliary direct-top PDG row is closed as a corpus-limited codomain no-go; the auxiliary row is compare-only. A source-side extraction-response kernel is outside the emitted corpus. The one-scalar companion seed

$$\sigma_{D11,\text{HT}} = \frac{\alpha_U \cos(2\theta_{W0})}{\sqrt{\pi}}$$

is on disk only as the lower-rank fixed-ray companion branch with $\pi_y = \pi_\lambda = \sigma_{D11,\text{HT}}$. The exact inverse slice on the same D11 Jacobian is compare-only and does not define the theorem surface. The same D11 Jacobian also has a compare-only exact inverse slice at the canonical Higgs/top reference pair

$$(m_H, m_t) = (125.1995304097179,\ 172.3523553288312)\,\text{GeV}.$$

Detailed particle-spectrum continuation statements are carried by Ref. [2]. That companion paper develops the quark selected-class theorem, the charged-lepton determinant-line boundary, the weighted-cycle neutrino theorem surface, and the hadron execution boundary. The hadron backend is a source-backend boundary with empirical closure policy documented. The compact paper promotes no source-only hadron masses. Empirical hadron closure values use a separate $e^+e^- \to$ hadrons payload class. These particle-continuation statements do not belong to the compact-paper claim surface.

Local unification surface and exact-release frontier. The local quantitative bridge can be stated without inflating the recovered-core claim. On the bosonic side, the declared pixel input P fixes the D10/D11 trunk

$$P \longmapsto \alpha_U(P) \longmapsto \big(t_U(P), t_{\mathrm{tr}}(P)\big) \longmapsto v(P),$$

after which the D10 source basis fixes the declared electroweak transport package and the D11 source-only split theorem with

$$\rho_{HT} = \log\!\big(1 + \tau_{2,\mathrm{tree}}^{\mathrm{exact}}\big), \qquad (\pi_y, \pi_\lambda) \longmapsto (m_t, m_H)$$

by the declared D11 Jacobian, again with no inverse readback of the internal transmutation data from measured low-energy couplings. The D10/P/fine-structure chain is the declared arithmetic bridge on this quantitative surface. The electroweak W/Z and running-family rows carry separate public validation artifacts: the source spectral measure payload, same-scheme remainder, and interval certificate. The same D11 Jacobian then emits the Higgs/top branch. On the gravity side, the same D10 pixel law packages

$$\bar{\ell}_{\mathrm{SU}(2)}(t_{2,\mathrm{run}}) \;+\; \bar{\ell}_{\mathrm{SU}(3)}(t_{3,\mathrm{run}}) = P/4, \qquad G = \frac{a_{\mathrm{cell}}}{4\,\bar{\ell}(t)}.$$

The exact local release burden consists of two theorem-sized objects: one strict classical-regime clause above the BW/Lorentz support package and one target-free D10 chart-identity theorem for exact W/Z. The familiar-unit readout is not a third theorem-sized burden. It is the explicit display package attached to the theorem-side data a_{cell}, c, and $\bar{\ell}_{\mathrm{shared}}$. On the gravity side, Proposition 5.27 gives

$$\bar{\ell}_{\mathrm{shared}} = \bar{\ell}_{\mathrm{SU}(2)}(t_{2,\mathrm{run}}) \;+\; \bar{\ell}_{\mathrm{SU}(3)}(t_{3,\mathrm{run}}) = P/4,$$

and Proposition 5.28 gives the local SI readout

$$G_{\mathrm{SI}} = \frac{c^3 a_{\mathrm{cell}}}{\hbar P}$$

relative to the declared microscopic datum a_{cell}. The same proposition fixes the full familiar-unit package:

$$L_{\mathrm{loc}} = \sqrt{a_{\mathrm{cell}}}\,\widehat{L}(P), \qquad t_{\mathrm{loc}} = \frac{\sqrt{a_{\mathrm{cell}}}}{c}\,\widehat{T}(P),$$

$$E_{\mathrm{loc}} = \frac{\hbar c}{\sqrt{a_{\mathrm{cell}}}}\,\widehat{E}(P), \qquad \Theta_{\mathrm{loc}} = \frac{\hbar c}{k_B\sqrt{a_{\mathrm{cell}}}}\,\widehat{\Theta}(P),$$

with dimensionless $\widehat{L}, \widehat{T}, \widehat{E}, \widehat{\Theta}$. Meters and seconds are therefore read from the single local ruler $\sqrt{a_{\mathrm{cell}}}$ together with the structural Lorentz output c, while GeV and Kelvin are downstream familiar-unit displays of the inverse local ruler through $\hbar$ and k_B. On that declared extension surface the local release values are

$$c = 299792458\,\mathrm{m/s}, \qquad G = 6.674299995910528 \times 10^{-11}\,\mathrm{m}^3\,\mathrm{kg}^{-1}\,\mathrm{s}^{-2},$$

$$(M_W, M_Z, M_H) = (80.377,\ 91.18797809193725,\ 125.1995304097179)\,\text{GeV}.$$

The displayed c row is structural from the Lorentz branch rather than P-fitted, and the displayed G row is an exact emitted branch value relative to the stated a_{cell} datum rather than a literal zero-difference identity against the rounded benchmark 6.6743×10^{-11}.

Particle-spectrum continuation boundaries beyond the compact SM/GR ledger are stated in Ref. [2].

The paper does not claim complete closure of all low-energy observables.[1]

12 Conclusion

Observer-Patch Holography starts from overlap consistency of observer patches and yields, within one framework, several effective structures usually treated separately in modern fundamental physics. This SM/GR derivation paper supports a unique schedule-independent quotient normal form for overlap repair from each fixed initial physical quotient state, with same-boundary uniqueness only under the additional unique-extension condition, and hence a quotient-level physical UV branch only modulo boundary redundancy, implementation hiding, and inert ancillary stabilization. It also supports a Lorentz branch on the extracted prime geometric cap pair, a Jacobson-type Einstein branch under the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations and the null modular bridge together with the bounded-interval projective branch, the global closure of that same gravity lane at the cosmic record-capacity fixed point, the Standard Model gauge quotient with exact hypercharges and structural electroweak force content on the realized matter package on the realized MAR one-Higgs branch in the bosonic internal-gauge sector, the realized color triplet $N_c = 3$ and generation count $N_g = 3$ on that branch, a support-visible compact-gauge Yang–Mills form together with repair-gap accounting of the Yang–Mills mass gap under the declared compact-gauge branch assumptions, and a separate integrated quantitative-closure branch. Phenomenological continuations are discussed separately from the recovered core.

The framework therefore presents general relativity and the Standard Model as effective descriptions arising from the five-axiom basis together with the support-visible BW scaling theorem, geometric-subnet construction, and categorical gauge reconstruction. The local quantitative claim surface is equally explicit: the same declared pixel input P controls the electroweak/Higgs mass trunk and the classical-regime gravity coupling, the invariant causal speed c is the structural Lorentz output, and the familiar-unit package reads meters, seconds, GeV, and Kelvin from the single local ruler $\sqrt{a_{\text{cell}}}$ with $\hbar$ and k_B used only as display conventions. Separate continuation boundaries concern sharper quantitative control, a fuller fermionic gauge branch, and theorem-side local exact-release objects beyond the corpus-backed surfaces stated here.

A Supplemented Gravity and Structural Gauge Details

This appendix carries the gravity-side and structural gauge items that support the compact paper's technical surface. They sharpen the compact paper's Lorentz, Einstein, cosmological-constant, and product-group claim surfaces without widening the recovered core beyond its stated branch conditions.

[1] Questions outside this theorem package do not alter the specific claim set established here. Even a separate habitat theorem for OPH state-and-law data would not enlarge this SM/GR derivation paper's Phase-I scope unless the actual closure map, invariant admissible sector, and stability hypotheses were also proved at the same tier.

A.1 Lorentz and Einstein Bridge

Theorem A.1 (Conformal group isomorphism)**.** *The orientation-preserving conformal group of S^2 is isomorphic to the connected Lorentz group:*

$$\mathrm{Conf}^+(S^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1).$$

Theorem A.2 (Support-visible geometric modular flow on caps on the extracted prime geometric subnet)**.** *Assume the OPH axioms, the derived fixed-cutoff collar package, and the support-visible BW scaling theorem on the extracted prime geometric subnet. Then the scaling-limit modular automorphism group of the cap pair is geometric conformal dilation on that subnet, with the standard 2π normalization. If the emitted scaling-limit cap algebra is type I, this may be written as $K_C = 2\pi B_C$; in the generic continuum case the theorem is the automorphism statement on a non-type-I cap algebra with outer geometric modular action. The collar replacements used in this branch are exact only at exact Markovity or in controlled fixed-collar families with $\delta^{\mathrm{M}} \to 0$; small CMI supplies a Fawzi–Renner recovered comparison state, not a dimension-free one-shot exact Markov normal form.*

Definition A.3 (BW-branch observer-relative time reading)**.** *On the branch satisfying the hypotheses of the preceding geometric modular-flow theorem, OPH uses the modular automorphism parameter t of the extracted cap pair as that cap observer's relative time parameter. This is a declared BW-branch reading of the derived geometric modular flow; it is not an additional proof that arbitrary operational clocks, global time, or the full problem of time have been derived from the axioms.*

BW-side UV scaffold. The UV data on this branch are not a separate cap-isotropy or Euclidean-regularity selector. The fixed-cutoff cap algebras are type-I regulators, but the scaling-limit observer algebra may leave that class, and MaxEnt alone does not select the BW / canonical cap phase. The UV package is the realized transported geometric cap-local system together with the carried-collar schedule derived from the transported fixed-local-collar Markov/faithfulness datum on each fixed local collar model, followed by support-readable modular covariance and ordered cut-pair rigidity on the emitted scaling-limit geometric cap pair.

Null-strip completion status. On the D4 side, the fixed-cutoff strip package is more specific than the generic inherited-strip summary. The null cuts first transfer the same cut-center data as the spatial collar branch, which fixes the central sector-pair decomposition of the strip algebra. The stronger left/right tensor decomposition used by the null modular bridge then requires the extra inherited strip-split condition on the multiplicity spaces, together with the exact-or-controlled Markov hypotheses on one fixed inherited strip model. On that same fixed strip model, the renormalized half-line family is endpoint-Lipschitz, hence defines the weak tail generator, and on the scaling-limit geometric-cap branch the half-line blow-up net carries the derived half-sided modular pair. Borchers–Wiesbrock then supplies the positive null-translation generator on its Stone domain together with the affine half-line modular relation $K_a(\Omega) = K_0(\Omega) - 2\pi a P_\Omega$, and the same half-line family fixes the generator/charge identification internally. Bounded-interval formulas are downstream on the separate interval-preserving projective branch.

Lemma A.4 (Null data ambiguity)**.** *If a symmetric tensor X_{ab} satisfies*

$$X_{ab}k^a k^b = 0$$

for every null vector k, then $X_{ab} = \phi\, g_{ab}$ for some scalar ϕ.

Corollary A.5 (Null modular data determine Einstein only up to the metric term)**.** *Null modular data determine T_{ab} only up to ϕg_{ab}. Consequently the Einstein equation is fixed locally only up to Λg_{ab}.*

Theorem A.6 (Jacobson-type rest-frame relation)**.** *In the local Lorentzian scaling regime, once the realized cap-label-preserving MaxEnt family satisfies the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap variations, the half-line generator/charge identification of the null bridge, the bounded-interval transport input used in the local Lorentzian regime, and the internal small-ball bridge yield the rest-frame first-variation relation that drives the compact paper's local Einstein branch; on the same scaling branch, if that rest-frame relation holds for all local observer four-velocities and all reference states, the compact paper upgrades it to the full tensor equation by the explicit local quadratic-polarization argument of Corollary 6.33.*

Theorem A.7 (Newton coupling from edge entropy density)**.** *The gravitational coupling is fixed by the edge entropy density:*

$$G = \frac{a_{\text{cell}}}{4\bar{\ell}(t)}.$$

A.2 Cosmological-Constant / Screen-Capacity Closure

Lemma A.8 (Vacuum energy blindness)**.** *For any null vector k,*

$$T^{\text{vac}}_{kk} = T^{\text{vac}}_{ab} k^a k^b = -\rho_{\text{vac}}\, g_{ab} k^a k^b = 0.$$

Vacuum-energy contributions therefore lie in the kernel of the null map $T_{ab} \mapsto T_{kk}$.

Proposition A.9 (Structural separation)**.** *Within the structural split used here:*

1. *local modular/null data fix T_{ab} only up to ϕg_{ab}; and*
2. *Λ is fixed by global screen capacity, not by local null data.*

So the large vacuum-energy bookkeeping of EFT is not itself the local quantity determining curvature on this branch.

Theorem A.10 (No local Λ prediction)**.** *Within the null-modular reconstruction used here:*

1. *all T_{kk} data are unchanged under $T_{ab} \to T_{ab} + \phi g_{ab}$;*
2. *therefore Λ cannot be fixed by local overlap consistency alone; and*
3. *Λ requires the global screen-capacity closure that closes the same Einstein branch beyond the local null data; the zero-input closure is the cosmic record-closure fixed point $N_{\text{CRC}} = F(N_{\text{CRC}})$, with $N_\star$ as its count-density representation.*

Corollary A.11 (Benchmark cosmological readout on the D6 branch)**.** *On the same D6 branch as Corollaries 6.36–6.38, inserting the observed value $\Lambda \approx 1.09 \times 10^{-52}\,\text{m}^{-2}$ gives*

$$r_{\text{dS}} = \sqrt{\frac{3}{\Lambda}} \approx 1.66 \times 10^{26}\ \text{m}, \qquad t_\Lambda = \frac{r_{\text{dS}}}{c} \approx 17.5\ \text{Gyr},$$

$$N_{\text{patch}} = \left(\frac{r_{\text{dS}}}{\ell_P}\right)^2 \approx 1.05 \times 10^{122}, \qquad N_{\text{scr}} = S_{\text{dS}} = \pi N_{\text{patch}} \approx 3.31 \times 10^{122},$$

$$\Lambda \ell_P^2 \approx 2.85 \times 10^{-122}.$$

Proof. Immediate from Corollary 6.38. □

Observed-age benchmark status. The branch quantity t_Λ is the de Sitter static-patch timescale. The usual cosmic age t_0 is a compare-only FLRW benchmark, with no additional D6 theorem output. On the flat ΛCDM benchmark

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_\Lambda}} \sinh^{-1}\left(\sqrt{\frac{\Omega_\Lambda}{\Omega_m}}\right)$$

one gets the standard comparison value $t_0 \approx 13.8\,\mathrm{Gyr}$ at $H_0 \approx 67.4\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ and $\Omega_\Lambda \approx 0.685$.

Lemma A.12 (FLRW curvature as visible scalar holonomy)**.** *On a homogeneous-isotropic spatial slice with constant sectional curvature K, small spatial loop holonomy obeys*

$$\mathrm{Hol}_{\square_{uv}} = \exp\Big(K A_\square J_{uv} + O(A_\square^{3/2})\Big)\,.$$

On a visibly separated OPH refinement system, the refinement-limit scalar spatial holonomy vanishes if and only if $K = 0$. Thus a flat FLRW branch is the zero-visible-spatial-holonomy branch.

Flatness boundary. This holonomy statement names a conditional cosmology bridge only. It does not add a D6 theorem output and does not solve the inflationary flatness or horizon problems. Selecting $K = 0$ requires an additional continuation premise: the preserved cosmological boundary datum carries no independent curvature charge and the same-boundary or MAR selector chooses the minimal visible geometric obstruction.

Theorem A.13 (D5–D6 cosmological-capacity closure stack)**.** *Assume the compact paper's local Einstein branch, the cosmic record-capacity fixed point*

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}),$$

its observed-branch de Sitter entropy readout

$$N_{\mathrm{CRC}} = S_{\mathrm{dS}},$$

the standard de Sitter entropy relation

$$S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4G} = \frac{3\pi}{G\Lambda},$$

and the standard static-patch formulas

$$r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\mathrm{dS}}}{c}.$$

Then the cosmological-constant package is one local/global theorem stack: local null data fix the Einstein branch only modulo Λg_{ab}, the same branch closes globally as

$$G_{ab} + \frac{3\pi}{G N_{\mathrm{CRC}}}\, g_{ab} = 8\pi G\,\langle T_{ab}\rangle,$$

the same D6 closure fixes

$$S_{\mathrm{dS}} = N_{\mathrm{CRC}}, \qquad A_{\mathrm{dS}} = 4GN_{\mathrm{CRC}}, \qquad r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\mathrm{dS}}}{c},$$

and the observed cosmic age is a downstream FLRW benchmark, with no additional theorem output.

Scope boundary. The D6 hypotheses are exactly the D5 local Einstein branch, the cosmic record-capacity fixed point $N_{\rm CRC} = F(N_{\rm CRC})$, its observed-branch de Sitter entropy readout $N_{\rm CRC} = S_{\rm dS}$, the standard de Sitter entropy relation, and the standard static-patch formulas. The local null-data route does not by itself determine the global capacity; that value is fixed by the readback closure. CMB kernels, inflation-replacement claims, H_0/S_8 branches, dark/anomaly growth kernels, and baryogenesis continuations require separate continuation theorems or likelihood contracts.

Separate pixel-side benchmarks. The companion pixel numbers

$$a_{\rm cell} \approx 1.63\,\ell_P^2, \qquad \ell_{\rm UV} = \sqrt{a_{\rm cell}} \approx 1.28\,\ell_P, \qquad \bar{\ell} \approx 0.408$$

belong to the separate pixel-closure/local-readout package, outside the D6 cosmological-parameter corollary itself.

A.3 Product-Group Structural Corollaries

Theorem A.14 (Factorization equivalence)**.**

Factorizing edge weights $\Longleftrightarrow$ Additive boundary Laplacian $\Longleftrightarrow$ Product gauge group.

If

$$H_\partial = H_\partial^{(1)} + H_\partial^{(2)} + H_\partial^{(3)} \qquad \textit{with} \qquad [H_\partial^{(i)}, H_\partial^{(j)}] = 0,$$

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.15 (No leptoquarks, hence no gauge-mediated proton decay)**.** *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. The structural claim is exactly

$$\tau_p^{(\rm gauge)} = \infty,$$

meaning that the gauge-boson proton-decay channel is absent. Scalar-mediated or higher-dimensional baryon-violating operators are outside this structural corollary.

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.4 Controlled Worldsheet Effective Description

Exact OPH-to-2D-Yang–Mills bridge. On the compact-group heat-kernel branch, Theorem 8.1 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.28.

Large-N_{edge} boundary. The large-N_{edge} worldsheet interpretation is carried only when one fixes a distinct large-N_{edge} sequence, with $N_{\text{edge}} \neq N_c = 3$, a fixed-$\tau$ window for

$$\tau = tN_{\text{edge}},$$

and the uniform genus-remainder control of Theorem 8.2. On that branch the edge free energy satisfies the compact 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 compact 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.5 Support-Visible Yang–Mills Gap from Repair Dynamics

Assumption A.16 (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, 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, repair completeness, and support-visible compact-gauge continuum extraction.*

Theorem A.17 (OPH derivation of the four-dimensional Euclidean Yang–Mills form)**.** *Under Assumption A.16, 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. Equivalently, the support-visible continuum transfer semigroup is the Euclidean Yang–Mills semigroup associated with the gauge-quotient cylinder measure

$$d\mu_{\mathrm{YM}}(A) = Z^{-1} e^{-S_E[A]} \, DA/G$$

in the OPH support-visible GNS representation.

Proof. The proof spine is included here to make the branch target explicit. The compact-gauge branch reconstructs the compact group G from the zero-obstruction transportable bosonic sector category and its faithful fiber functor. 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. Thus there is 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 and vanish from the strict Yang–Mills fixed form. Normalizing the unique positive quadratic invariant defines the coupling g. The compatible finite-stage gauge-register / quantum-link Gibbs measures then pass through the same support-visible weak-$*$ / GNS extraction used by the compact-gauge branch, giving the displayed Euclidean transfer semigroup. □

Prize-facing proof separation. The proof has two separate claims. Theorem A.17 derives the four-dimensional Euclidean Yang–Mills form on the 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, 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 be the support-visible compact-gauge quotient state space, let π_r be the vacuum stationary measure, and set

$$K_r := L^2(X_r, \pi_r).$$

Let $\mathcal{C}_r$ be the finite family of active repair collars. For each $C \in \mathcal{C}_r$, let $\rho_C : X_r \to Y_C$ be the complete repaired visible datum, 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 at theorem level in the compact paper.

Proposition A.18 (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. Then

$$\Phi_C(a) = a \quad (a \in \mathcal{N}_C), \qquad \Phi_C(\mathcal{A}_r^{\mathrm{sv}}) \subseteq \mathcal{N}_C, \qquad \pi_r \circ \Phi_C = \pi_r,$$

and Φ_C is $\mathcal{N}_C$-bimodular. Hence, for $a \in \mathcal{N}_C$ and $x \in \mathcal{A}_r^{\mathrm{sv}}$,

$$\pi_r(a^* \Phi_C(x)) = \pi_r(a^* x).$$

Since $\Phi_C(x) \in \mathcal{N}_C$ and π_r is faithful, this characterizes $\Phi_C(x)$ as the orthogonal projection of x onto $\mathcal{N}_C$ in the GNS inner product. That projection is the π_r-preserving conditional expectation E_C. □

Lemma A.19 (Implementation hiding gives fiber permutation symmetry)**.** *Fix an active collar C and a repaired value $y \in Y_C$. On the support-visible quotient, the hidden fiber $F_C(y) = \rho_C^{-1}(y)$ has no remaining observable labels. The conditioned local MaxEnt state is uniform on $F_C(y)$, and the primitive collar relaxation commutes with all finite permutations of $F_C(y)$.*

Proof. The quotient removes implementation labels: representatives with the same repaired datum and different hidden fiber coordinates define the same support-visible observable state. Conditioned on the repaired datum, no admissible support-visible constraint distinguishes two fiber points. The MaxEnt state is therefore uniform, and quotient-preserving collar relaxation commutes with the full hidden-fiber permutation action. □

Lemma A.20 (Scalar relaxation on a uniform hidden fiber)**.** *Let F be a finite hidden fiber with uniform measure. Let E_F be expectation onto constants, and let D_F be a positive self-adjoint Markov relaxation generator with*

$$\ker D_F = \operatorname{Ran}(E_F)$$

that commutes with the full permutation group of F. Then $D_F = c_F(I - E_F)$ for a scalar $c_F > 0$.

Proof. The permutation representation on $L^2(F)$ splits as constants plus the zero-sum subspace. The zero-sum subspace is irreducible for the full symmetric group when $|F| \geq 2$. Schur's lemma makes D_F scalar on that subspace, and positivity with the stated kernel makes the scalar strictly positive. ☐

Theorem A.21 (Exact Euclidean-consensus law)**.** *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. Axiom 3.3 gives a local-Gibbs state with quasi-local finite-range generator on the declared finite-constraint branch. After the support-visible quotient, each primitive local Euclidean piece D_C is supported on collar C, preserves the repaired visible datum ρ_C, and relaxes only complementary fiber data. Thus

$$\ker D_C = \operatorname{Ran}(E_C).$$

Lemma A.19 gives full hidden-fiber permutation symmetry. Applying Lemma A.20 fiberwise makes the relaxation scalar on the orthogonal complement of the repaired datum:

$$D_C = c_C (I - E_C)$$

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. ☐

Lemma A.22 (Uniform active-collar rate floor)**.** *Assume finite local combinatorial type, a branch-homogeneous local constraint family, and refinement-stable exact repair semantics. Then the local Euclidean repair rate depends only on active collar type, the active collar-type set is finite across the cofinal refinement family, and there is $c_* > 0$, independent of r, such that*

$$c_C \geq c_* \qquad \textit{for every active collar } C \in \mathcal{C}_r.$$

Proof. Active collars are exactly the collars on which complementary invisible fiber data are genuinely relaxed, so $D_C \neq 0$ on $\ker(E_C)$ and $c_C > 0$. Finite local combinatorial type gives a finite list of bounded collar patterns, including visible interfaces, hidden fiber cardinalities, constraint templates, and repair maps. Branch homogeneity makes the conditioned MaxEnt relaxation scalar a function of this finite type data. Refinement stability replaces collars by copies of the same bounded type templates with the same normalized local Euclidean repair semantics. Taking the minimum over the finite active type list gives $c_* > 0$. ☐

Proposition A.23 (Finite-stage repair gap)**.** *Assume the active collars admit a bounded-color decomposition*

$$\mathcal{C}_r = \bigsqcup_{a=1}^{q} \mathcal{C}_{r,a}, \qquad q < \infty,$$

independent of r, and that exact quotient-local gluing makes the corresponding color expectations commute. Let $P_{0,r}$ project onto constants in K_r. Then

$$L_r^{\mathrm{EC}} \geq c_*(I - P_{0,r}), \qquad \Delta_{\mathrm{rep},r} \geq c_* > 0.$$

Proof. For each color a, define the parallel color expectation

$$E_{r,a} := \prod_{C \in \mathcal{C}_{r,a}} E_C.$$

Same-color collars are disjoint, so the factors commute, and exact quotient-local gluing makes the family $\{E_{r,a}\}_{a=1}^{q}$ a commuting family of orthogonal projections on the vacuum branch. Repair completeness says that the only support-visible observables fixed by every local repair are constants:

$$\bigcap_{a=1}^{q} \mathrm{Ran}(E_{r,a}) = \mathbb{C}\,\mathbf{1}_r.$$

Set

$$\widetilde{L}_r := \sum_{a=1}^{q} c_*(I - E_{r,a}).$$

For commuting projections, $I - \prod E_C \leq \sum_C (I - E_C)$, so $\widetilde{L}_r \leq L_r^{\mathrm{EC}}$. Since the $E_{r,a}$ commute, they are simultaneously diagonalizable. On every nonconstant joint eigenspace at least one color eigenvalue is zero. Thus $\widetilde{L}_r \geq c_*(I - P_{0,r})$, and the same lower bound holds for L_r^{EC}. ☐

Proposition A.24 (Coherent refinement and compact-gauge extraction)**.** *Under Assumption A.16, the compact-gauge branch supplies finite-stage support-visible quotient state spaces, compatible repair-side and physical-side refinement maps, projectively compatible compact-gauge cylinder marginals, weak-$*$ compactness, a diagonal subnet with convergent local marginals on every support-visible cylinder algebra, support-visible GNS gluing to continuum Hilbert spaces K and $\mathcal{H}$, and a faithful limiting vacuum pair after quotienting the maximal repair-invariant overlap-trivial kernel.*

Proof. The D7 compact-gauge ladder proves fixed-cutoff bosonic sector categories, faithful monoidal refinement functors, compatible finite-dimensional fibers, the directed colimit sector category, and compact gauge reconstruction. The support-visible compact-gauge quotient algebra is generated by overlap sector projectors and compact-gauge visible bosonic observables, modulo the maximal repair-invariant overlap-trivial kernel. Each finite regulator has finite-dimensional local cylinder algebras, so its state space is weak-$*$ compact. Refinement compatibility makes the local marginals projective. A diagonal subnet over a countable cofinal family of support-visible cylinders gives a limiting state on their algebraic union. Positivity and normalization pass to the limit on every cylinder, and the quotient by the repair-invariant overlap-trivial kernel makes the limiting support-visible gauge-invariant state faithful. GNS applied to that state gives K, and transporting the same coherent refinement data through the finite-stage unitaries U_r gives $\mathcal{H}$. ☐

Theorem A.25 (Continuum exact transfer identification)**.** *There is a unitary $U : \mathcal{H} \to K$ such that*

$$Ue^{-tH}U^{-1} = e^{-tL^{\mathrm{rep}}} \qquad (t \geq 0),$$

and hence

$$UHU^{-1} = L^{\mathrm{rep}}.$$

Proof. Theorem A.21 gives exact finite-stage intertwining for every regulator. Proposition A.24 supplies the coherent direct-limit and support-visible GNS extraction hypotheses. The finite-stage intertwiners therefore pass to the continuum Hilbert spaces. Equality of generators follows from uniqueness of the self-adjoint generator of a strongly continuous contraction semigroup. □

Theorem A.26 (Osterwalder–Schrader reconstruction on the compact-gauge branch)**.** *Under Assumption A.16 and Proposition A.24, 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.17.*

Proof. Euclidean covariance is part of the four-dimensional scaling chart and local-Gibbs cylinder family. Reflection positivity is part of the ordinary vacuum branch. Regularity follows from the finite local cylinder construction and support-visible weak-$*$ limit. The vacuum vector is cyclic for the GNS closure of the gauge-invariant local cylinder algebra. The Osterwalder–Schrader reconstruction theorem then produces the Hilbert space, vacuum, local algebra, positive Hamiltonian, and Euclidean transfer semigroup [40, 41]. □

Proposition A.27 (Nontriviality of the support-visible compact-gauge theory)**.** *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 a realized nontrivial compact-gauge branch. At finite cutoff this gives a support-visible gauge-invariant local observable, for example a nonconstant Wilson/plaquette cylinder observable, whose vacuum variance is positive. Subtracting its expectation gives a GNS vector orthogonal to the vacuum. The finite-range local-Gibbs generator assigns finite energy to finite-cylinder excitations, and Proposition A.24 transports these vectors to the support-visible continuum. □

Theorem A.28 (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 above. The theory reconstructed in Theorem A.26 is nontrivial by Proposition A.27, and its continuum support-visible Hamiltonian H satisfies*

$$H \geq c_*(I - P_0),$$

where P_0 projects onto the vacuum. Therefore

$$\mathrm{Spec}(H) \cap (0, c_*) = \varnothing, \qquad \Delta_{\mathrm{YM}} \geq c_* > 0.$$

Moreover the Yang–Mills gap is the repair gap:

$$\Delta_{\mathrm{YM}} = \Delta_{\mathrm{rep}}.$$

Proof. Proposition A.23 gives

$$L_r^{\mathrm{EC}} \geq c_*(I - P_{0,r})$$

at every finite stage. Proposition A.24 and Theorem A.25 transport the bound to the support-visible continuum:

$$L^{\mathrm{rep}} \geq c_*(I - P_0).$$

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, not merely 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 [39, 38]. Theorem A.17 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.27 supplies nontriviality. Theorem A.28 supplies the exact spectral gap accounting on that same branch. The claim to full Clay admissibility is exactly the claim that the support-visible compact-gauge continuum extraction used in Proposition A.24 supplies the required four-dimensional axiomatic quantum Yang–Mills construction. 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, A. Osika, K. Xue, B. Cassie, P. Nguyen, M. Poneder, and K. A. Anirudha, *Observers Are All You Need*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/observers_are_all_you_need.pdf.

[2] B. Müller, A. Osika, K. Xue, and M. Poneder, *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.

[3] DULA - GR OK 4, *The DULA Theorem: Graded Monoid Homomorphisms for Prime Congruences*, September 19, 2025. Available at https://fermatslibrary.com/p/74724608.

[4] M. H. A. Newman, "On theories with a combinatorial definition of 'equivalence'," *Ann. of Math. (2)* **43** (1942), 223–243.

[5] E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," *Commun. Math. Phys.* **28** (1972), 251–257.

[6] J. J. Bisognano and E. H. Wichmann, "On the duality condition for a Hermitian scalar field," *J. Math. Phys.* **16** (1975), 985–1007.

[7] J. J. Bisognano and E. H. Wichmann, "On the duality condition for quantum fields," *J. Math. Phys.* **17** (1976), 303–321.

[8] 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.

[9] H.-W. Wiesbrock, "Half-Sided modular inclusions of von-Neumann-Algebras," *Commun. Math. Phys.* **157** (1993), 83–92.

[10] W. G. Unruh, "Notes on black-hole evaporation," *Phys. Rev. D* **14** (1976), 870–892.

[11] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Phys. Rev. Lett.* **75** (1995), 1260–1263, arXiv:gr-qc/9504004.

[12] T. Jacobson, "Entanglement equilibrium and the Einstein equation," *Phys. Rev. Lett.* **116** (2016), 201101, arXiv:1505.04753.

[13] G. W. Gibbons and S. W. Hawking, "Cosmological event horizons, thermodynamics, and particle creation," *Phys. Rev. D* **15** (1977), 2738–2751.

[14] Planck Collaboration, "Planck 2018 results. VI. Cosmological parameters," *Astron. Astrophys.* **641** (2020), A6, arXiv:1807.06209.

[15] 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.

[16] E. H. Lieb and M. B. Ruskai, "Proof of the strong subadditivity of quantum-mechanical entropy," *J. Math. Phys.* **14** (1973), 1938–1941.

[17] D. Petz, "Sufficient subalgebras and the relative entropy of states of a von Neumann algebra," *Commun. Math. Phys.* **105** (1986), 123–131.

[18] D. Petz, "Sufficiency of channels over von Neumann algebras," *Quart. J. Math.* **39** (1988), 97–108.

[19] O. Fawzi and R. Renner, "Quantum conditional mutual information and approximate Markov chains," *Commun. Math. Phys.* **340** (2015), 575–611, arXiv:1410.0664.

[20] 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.

[21] S. Doplicher and J. E. Roberts, "A new duality theory for compact groups," *Invent. Math.* **98** (1989), 157–218.

[22] 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.

[23] T. Tannaka, "Über den Dualitätssatz der nichtkommutativen topologischen Gruppen," *Tohoku Math. J.* **45** (1938), 1–12.

[24] M. G. Krein, "A principle of duality for a bicompact group and a square block algebra," *Dokl. Akad. Nauk SSSR* **69** (1949), 725–728.

[25] H. Georgi and S. L. Glashow, "Unity of all elementary-particle forces," *Phys. Rev. Lett.* **32** (1974), 438–441.

[26] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak interactions with lepton-hadron symmetry," *Phys. Rev. D* **2** (1970), 1285–1292.

[27] E. Witten, "An SU(2) anomaly," *Phys. Lett. B* **117** (1982), 324–328.

[28] S. Dimopoulos, S. Raby, and F. Wilczek, "Supersymmetry and the scale of unification," *Phys. Rev. D* **24** (1981), 1681–1683.

[29] 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.

[30] D. J. Gross and W. Taylor, "Two-dimensional QCD is a string theory," *Nucl. Phys. B* **400** (1993), 181–208, arXiv:hep-th/9301068.

[31] F. Peter and H. Weyl, "Die Vollständigkeit der primitiven Darstellungen einer geschlossenen kontinuierlichen Gruppe," *Math. Ann.* **97** (1927), 737–755.

[32] 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.

[33] 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.

[34] W. Donnelly and A. C. Wall, "Entanglement entropy of electromagnetic edge modes," *Phys. Rev. Lett.* **114** (2015), 111603, arXiv:1412.1895.

[35] 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.

[36] 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.

[37] 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.

[38] Clay Mathematics Institute, "Yang–Mills & the Mass Gap." Available at https://www.claymath.org/millennium/yang-mills-the-maths-gap/.

[39] 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.

[40] K. Osterwalder and R. Schrader, "Axioms for Euclidean Green's functions," *Communications in Mathematical Physics* **31** (1973), 83–112.

[41] K. Osterwalder and R. Schrader, "Axioms for Euclidean Green's functions II," *Communications in Mathematical Physics* **42** (1975), 281–305.

[42] Particle Data Group, "Review of Particle Physics," *Phys. Rev. D* **110** (2024), 030001. Available at https://pdg.lbl.gov/2024/download/db2024.pdf.

[43] Particle Data Group, "Electroweak Model and Constraints on New Physics," Available at https://pdg.lbl.gov/2024/reviews/rpp2024-rev-standard-model.pdf.