Observers Are All You Need

B. Müller Alexander Osika Kai Xue Ben Cassie Peter Nguyen Mario Poneder
Kale Arnav Anirudha

June 8, 2026

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

Abstract

Defined from the inside, Observer-Patch Holography (OPH) builds physics from finite records on horizon-screen cuts. Overlapping observer patches must agree on every observable they can both compare. Near each observer, a patch carries only the finite records available for that comparison. The consistency rule, local entropy selection, recoverable records, and minimal admissible realization turn those comparisons into structure. Public records become the route by which space, time, matter, and measurement appear inside the network.

Assembled here are the main OPH results: Lorentzian and null structure, the Einstein branch, the realized Standard Model quotient $SU(3) \times SU(2) \times U(1)/\mathbb{Z}_6$, exact hypercharge, $N_c = 3$, $N_g = 3$, record/Born-rule structure, and the cosmological capacity closure. No physical constants are inserted as free inputs. In their place, the local pixel ratio $P \simeq 1.630968$ is the unique outer/inner fixed point linking screen-cell area to electromagnetic observation strength, and the cosmic record capacity $N_{\mathrm{CRC}} \simeq 3.31 \times 10^{122}$ is the unique fixed point where the universe reads back its own horizon capacity as an internal public record. Consequently, the fine-structure and cosmological constant readouts, together with the remaining particle, gravity, and observer-facing quantities, are downstream of those two closures and the recovered structural branch.

A reader who follows the mathematical proofs presented should conclude that, contrary to popular belief, space, time, and matter are emergent appearances, while information and computation form the fundamental substrate of reality.

Contents

1 Introduction

Observer-Patch Holography asks how much of physics is forced once the world is described from the inside. No observer has access to a view from nowhere. In the idealized description there is a shared screen net, often charted by S^2, but each finite observer has access only to a local patch P_O of that net. The observer carries records inside that patch and can compare only the observables shared with neighboring patches. OPH takes that comparison rule literally: where descriptions overlap, they must agree. Thus an "observer screen" is an operational access cut, not a private physical sphere assigned to the observer.

The formalism used here is deliberately quantum-algebraic. Patches carry algebras and states, record events use the trace/Born rule on their declared operator surface, and the gravity branch uses generalized entropy. OPH's theory-of-everything claim is that the observed effective universe can be recovered from this axiom set and its stated branch conditions if the framework is correct. It is not a demand that every mathematical ingredient be rebuilt from a blank starting point before the reconstruction can proceed.

That one demand is strong enough to organize the familiar world. The effective bulk is reconstructed from compatible screen data. Locality comes from the way collars control recovery across patch boundaries. Gauge freedom is the freedom to change the local description without changing any shared observable. Records are the stable parts of a patch that can be checked by other patches. Matter is the stable charge and excitation content that survives transport, fusion, and minimal admissible realization.

The first major achievement is spacetime. Modular flow on screen caps becomes geometric, giving the Lorentzian structure of local physics. Generalized entropy at fixed cap then gives the Einstein relation in the large-scale regime. The cosmological constant is not treated as a vacuum-energy guess. It is read from the cosmic record-capacity fixed point N_{CRC}, with

$$\Lambda_{\mathrm{CRC}} = \frac{3\pi}{G N_{\mathrm{CRC}}}.$$

The Gibbons–Hawking entropy and the Planck-2018 late-time de Sitter benchmark give a concrete normalization [14, 15]. With $R_{\mathrm{dS}} \simeq 1.66 \times 10^{26}$ m and $\ell_P \simeq 1.616 \times 10^{-35}$ m, the bare horizon ratio is

$$N_{\mathrm{patch}} = \left(\frac{R_{\mathrm{dS}}}{\ell_P}\right)^2 \simeq 1.05 \times 10^{122}.$$

The entropy capacity is

$$N_{\mathrm{scr}} = S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4\ell_P^2} = \pi N_{\mathrm{patch}} \simeq 3.31 \times 10^{122},$$

so $\Lambda \ell_P^2 \simeq 2.85 \times 10^{-122}$. OPH uses N_{scr} for the Gibbons–Hawking entropy capacity and N_{patch} for the bare horizon area ratio. The input-free closure is sharper: define the OPH readback map $F(N)$ to be the active cosmic record capacity reconstructed by stable observers inside the universe supplied with capacity N. The cosmic record-closure capacity satisfies

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}).$$

The fixed-point proof gives uniqueness and stability. The observed branch read-off is $N_{\mathrm{obs}} \simeq 3.31 \times 10^{122}$, so the cosmological constant is the downstream readout $\Lambda_{\mathrm{CRC}} = 3\pi/(G N_{\mathrm{CRC}})$.

The second achievement is the Standard Model structure. Under the stated coherent transportable-sector refinement ladder, 3 + 1D symmetry, and compatible bosonic finite-fiber conditions, edge sectors reconstruct the bosonic compact-gauge branch. On the realized low-energy branch with the explicit one-generation/one-Higgs matter package, minimal admissible realization then selects the compact gauge quotient

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

The realized low-energy branch then fixes exact hypercharge, three generations, three colors, the masslessness of the photon, gluons, and graviton, integer charge for color singlets, and the absence of gauge-mediated proton decay.

The third achievement is quantitative and sits outside the recovered core. OPH's zero-input constants program has two closure values: the local pixel fixed point and the global cosmic record-closure fixed point. On the observed branch, the global scale is $N_{\mathrm{CRC}} \simeq N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}$. The local scale is

$$P := \frac{a_{\mathrm{cell}}}{\ell_P^2}.$$

The value of P is not inserted as a separate measured constant. OPH computes it by giving the same screen cell two readings. From outside, the cell sits a little above the exact self-similar entropy balance point $\varphi = (1+\sqrt{5})/2$. From inside, the same cell appears as the smallest electromagnetic observation step available in the encoded world. The local scale is the value of P where those two readings coincide.

With the P fixed point in place, the same scale organizes the weak-boson compare-only validation pair, the low-energy electromagnetic endpoint, the Higgs/top quantitative surface, the running quark package, a weighted-cycle neutrino theorem branch, and the local gravity-facing readout. The status table records the scope of the weak compare-only validation row, charged-lepton target-anchored witnesses, and hadron backend.

The fourth achievement is implementation. The screen-microphysics construction makes patches, overlaps, edge observables, records, repair maps, and observer checkpoints explicit at finite cutoff through a federated patch-carrier architecture. OPH has a regulated observer-and-record surface on which the basic consistency machinery can be written directly.

Axioms in words. The theory uses five axioms. A screen net assigns algebras to connected screen patches. Overlap consistency requires agreement on shared observables. Local MaxEnt selects the least biased local state compatible with the finite constraint data and keeps the realized construction stable under refinement. Recoverable generalized entropy supplies the collar-recovery and focusing structure used by the gravity lane. Minimal admissible realization selects the simplest low-energy sector compatible with the consistency constraints. These axioms define a quantum-algebraic starting point. Quantum mechanics and quantum field theory are treated as effective descriptions carried by the deeper observer-patch architecture, not as structures whose use here depends on first deriving every mathematical ingredient from operational records alone.

This paper is the synthesis surface. It gives the operational picture first, then places the technical proof machinery after the reader has the construction in hand.

2 Main Achievements

The OPH papers support the following unified picture.

1. **Observers and records become physical objects.** Finite patches, shared overlaps, record algebras, repair maps, and observer checkpoints are represented on the screen rather than assumed from outside.

2. **Spacetime is reconstructed.** On the certified support-visible BW scaling branch, cap-pair extraction, regularized modular transport, support-readable modular covariance, round-cap rigidity, and KMS/BW normalization make the cap modular automorphism geometric and supply Lorentz kinematics. Recoverable generalized entropy together with the null modular bridge and bounded-interval transport then supplies the Einstein relation. The screen-capacity branch ties the cosmological constant to the zero-input cosmic record-closure fixed point:

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}), \qquad \Lambda_{\mathrm{CRC}} = \frac{3\pi}{G N_{\mathrm{CRC}}}.$$

   Here $F(N)$ is the active cosmic record capacity read back by observers inside the universe supplied with capacity N; uniqueness and stability follow from the OPH readback fixed-point certificate.

3. **The Standard Model quotient is selected by MAR.** Under the stated coherent transportable-sector refinement conditions, edge sectors reconstruct a bosonic compact-gauge branch. The overlap/holonomy calculus classifies transportable sectors; it does not by itself choose the Standard Model. The realized low-energy branch with the explicit one-generation/one-Higgs matter package and MAR then selects the compact gauge quotient, exact hypercharge, three generations, three colors, and the structural absence of gauge-mediated proton decay.

4. **The local particle scale is computed by fixed-point closure.** The pixel ratio P is fixed by matching a screen cell's outer entropy detuning to its inner electromagnetic observation scale. The public readout is $P \simeq 1.6309682094$, $\alpha^{-1}(0) = 137.035999177(21)$, and $\alpha(0) \simeq 0.00729735256433$. The calculation uses the source anchor and Ward-projected electromagnetic transport to enforce the outer/inner pixel fixed point. A different fine-structure value would move the pixel away from that consistency point. The status table records the source-side diagnostic trunk and endpoint residual.

5. **The framework has a regulated implementation surface.** The microphysics paper turns the observer language into explicit finite patch carriers, edge observables, records, and repair cycles.

6. **Edge sectors also admit two-dimensional Yang–Mills and worldsheet-style effective descriptions.** The heat-kernel edge partition function reorganizes the same boundary data into the two-dimensional Yang–Mills form, and on the declared large-edge branch this yields a controlled worldsheet effective description. This 2D bridge sits beside the compact paper's support-visible compact-gauge repair-gap theorem.

7. **The compact-gauge branch carries the 4D Euclidean Yang–Mills form and isolates the gap theorem.** On the declared support-visible compact-gauge branch, with the four-dimensional scaling chart, reflection-positive ordinary vacuum, no additional gauge-invariant relevant dimension-four pure-gauge operator on that branch besides curvature squared, repair completeness, and support-visible continuum extraction in force, OPH obtains the Euclidean Yang–Mills action from compact-gauge holonomy data and the local MaxEnt/Gibbs continuum limit. Exact local repair acts as conditional expectation, the Euclidean transfer generator is the sum of active collar relaxations, and a uniform finite-stage repair gap passes through continuum extraction to the compact-gauge Yang–Mills Hamiltonian. The accounting identity is $\Delta_{\mathrm{YM}} = \Delta_{\mathrm{rep}}$. The Clay-facing note isolates this same branch theorem and treats its support-visible continuum extraction as the claimed axiomatic construction surface.

3 Physical Implementation and A_5/E_8 Significance

The implementation question is what must exist below the observer-level axioms so they have a physical carrier. OPH's answer is a fixed-cutoff federation of finite patch carriers. Each carrier exposes overlap ports, record registers, repair maps, and checkpoint interfaces. The screen is the observer-facing geometry chart. It is not the carrier itself in the sense of one literal smooth ball. The implementation is the finite algebraic machinery that makes compare, write, repair, and re-read operations available.

Spherical geometry carries several pieces of the construction at once. An observer-accessible cut has a closed two-dimensional angular chart. Caps and collars on that chart supply the cut data used by modular flow, entropy variation, and overlap comparison. The orientation-preserving conformal group of the same S^2 chart is the celestial-sphere realization of $\mathrm{SO}^+(3,1)$, so the sphere is the kinematic bridge between finite screen cuts and the emergent $3+1$-dimensional Lorentz branch once the support-visible cap-pair theorem is satisfied. Finite cellulations of the chart provide the regulator surface on which ports, edge sectors, and local comparison data can be made explicit; they are not by themselves a Lorentz invariant continuum.

The echosahedral carrier is the concrete local realization of this chart data. Its multi-port interface supplies discrete boundary directions, face and edge incidence, exposed readout channels, record slots, and repair channels. Recurrent toroidal subchannels supply local winding and memory structure. The smooth spherical chart describes the observer-facing cut; the echosahedral carrier supplies the fixed-cutoff local machine whose quotient data can present that cut.

The A_5 and E_8 labels mark different roles in that machinery. A_5-icosahedral symmetry is the local finite-screen language used when a patch carrier is made concrete: it organizes ports, faces, and overlap data with enough discrete symmetry to stabilize support-visible cuts. E_8-type language is the exceptional closure language used by the high-symmetry branch. It names the root-lattice and affine exceptional structure in which the icosahedral data can sit. The term is representation-closure language for the branches that call for it.

This distinction matters for the synthesis paper. The physical claim lives at the observer-visible quotient: finite patches must expose the same records and shared observables after allowed implementation-hiding changes. The microphysics paper gives the concrete reference architecture for that quotient surface, including the federated patch-carrier model, fixed-cutoff edge heat-kernel/Casimir law, central-record/Born–Lüders measurement interface, Bell/CHSH event surface, and checkpoint/restoration package. See Ref. [3], available as the numbered GitHub reference in the bibliography.

Particle Prediction Surface

The particle paper treats the completion pipeline as an explicit closure matrix. The non-hadron bundle reports the following status:

Companion Papers

The detailed depth surfaces used throughout this synthesis are:

1. *Recovering Relativity and the Standard Model from Observer Overlap Consistency* [1], which gives the compact technical core for relativity, gravity, zero-obstruction compact-gauge reconstruction, and MAR-selected Standard Model structure.

2. *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics* [2], which carries the finite patch-net fixed-point, defect, quotient, and operator-record consensus

Lane	Output	Role	Status note
Massless carriers	$m_\gamma = m_g = m_{\mathrm{grav}} = 0$	structural theorem	symmetry-protected zeros
W/Z	80.377 GeV, 91.18797809193725 GeV	weak-sector validation pair	declared compare-only validation surface; ledger records exact scope
Fine structure	$\alpha^{-1}(0) = 137.035999177(21)$, $P = 1.630968209403959324879279847782648941\ldots$	fixed-point derivation	source audit and empirical hadron closure records are separate bookkeeping surfaces
Higgs/top	125.1995304097179 GeV, 172.35235532883115 GeV	declared D10/D11 quantitative surface	auxiliary direct-top PDG codomain is compare-only in the available corpus
Charged leptons	0.0005109989499999994, 0.10565837550000004, 1.7769324651340912 GeV	target-anchored same-family exact witness	source landing from P to physical charged data lacks a theorem-grade uncentered trace lift in the available corpus
Quarks	$u, d, s, c, b, t = (0.00216, 0.00470, 0.0935, 1.273, 4.183, 172.35235532883115)$ GeV	selected-class exact theorem	exact only on the selected public quark frame class; global public-frame classification has a no-go boundary in the available corpus; strong CP is work in progress
Neutrinos	0.017454720257976796, 0.019481987935919015, 0.05307522145074924 eV	weighted-cycle theorem branch	absolute masses are not directly measured; PMNS comparison tension is visible
Hadrons	no source-only prediction emitted; empirical closure rows separate	source backend absent; empirical closure policy emitted	source-only masses require a working OPH hadron backend, such as GLORB/Echosahedron; empirical closure uses a separate $e^+e^- \to$ hadrons payload class

package, with the default overlap code treated as a finite constraint code, termination separated from the local-diamond and repair-completeness conditions needed for confluence, and QECC/min-cut, spectral, BFT, and hardware-speedup claims gated by separate certificates.

3. *Federated Echosahedral Screen Microphysics: Patch Hardware, Records, and Observer Synchronization in OPH* [3], which carries the regulated federated patch-carrier architecture, the fixed-cutoff edge heat-kernel / Casimir theorem, and the measurement and observer checkpoint/restoration packages.

4. *Deriving the Particle Zoo from Observer Consistency* [4], which carries the particle derivations, masses, couplings, and sector calculations.

The sections that follow give the synthesis view. Detailed proof burdens remain with the companion papers and with the late dependency checklist in this paper.

4 Recovered Core: Foundations and Structural Branches

Observer-Patch Holography is a reconstruction program for fundamental physics. Physical data live on a horizon screen S^2, and the effective world is encoded by the mutual consistency of overlapping patch descriptions. This recovered-core section develops the structural branch: fixed-cutoff overlap repair, the separated cofinal refinement-limit consensus bridge, collar recovery, the Lorentz and Einstein route, compact gauge reconstruction in the bosonic branch, and the realized-branch Standard Model quotient.

4.1 Overview

The recovered core has a simple conceptual order. First, overlap consistency turns local patch states into a constrained gluing problem. The fixed-cutoff branch gives collar decompositions, edge centers, and finite repair dynamics on the declared physical quotient. Second, the continuum geometric branch turns modular flow on caps into Lorentz kinematics and then into the Einstein relation through fixed-cap generalized-entropy stationarity. Third, the transportable edge-sector branch first classifies zero-obstruction transportable sectors and reconstructs some compact internal gauge group in the bosonic branch; minimal admissible realization, anomaly constraints, and the explicit CKM/weak-sector clauses of MAR then select the realized Standard Model quotient with exact hypercharge, structural electroweak force content, the realized color triplet $N_c = 3$, and the generation count $N_g = 3$ on the explicit realized one-generation/one-Higgs package. This branch does not by itself supply a super-Tannakian internalization of fermions/chirality.

The local pixel ratio $P \equiv a_{\mathrm{cell}}/\ell_P^2$ and the cosmic record capacity N_{CRC} enter only on their declared quantitative branches. P is fixed by the outer/inner closure construction in Section 6. The D6 branch uses the cosmic record-closure fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$, where $F(N)$ is the active capacity read back by stable observers inside the universe supplied with capacity N, to close the metric-term ambiguity of the Einstein branch on the cosmological-capacity surface. The finite-count representation selects the MAR maximizer of $\log|\Omega_N^{\mathrm{sc}}| - N$, where the subtraction is division by the full screen Hilbert dimension e^N. The late classification section records the formal dependency checklist and separates recovered-core claims from quantitative closures and continuations.

4.2 Model and Axioms

4.2.1 Observers and access model

An observer O is a tuple $(P_O, \mathcal{A}(P_O), \rho_O, R_O)$ where:

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

Observers are internal patterns in the global state. Different observers correspond to different patches and their compatible marginals.

4.2.2 Screen, patches, and algebra net

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

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

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

4.2.3 Five OPH axioms

This paper uses five OPH axioms. The support-visible BW scaling theorem below handles the scaling/BW step. Transportability, the fixed-cutoff bosonic sector category, refinement/fiber descent, and the realized MAR-admissible compact-gauge witness are theorem-level results cited directly where they are used. Other branch-local hypotheses are written out in prose where needed instead of being assigned a separate label family.

1. **Screen Net.** A horizon screen S^2 carries a net of algebras $P \mapsto \mathcal{A}(P)$.
2. **Overlap Consistency.** Local states agree on shared observables for any overlap.
3. **Local MaxEnt and Refinement Stability.** At the UV scale the realized state maximizes entropy subject to a finite family of gauge-invariant local constraints, and the resulting family of states is stable under coarse-graining so symmetry-allowed relevant operators are not held at zero by unexplained fine tuning.
4. **Recoverable Generalized Entropy.** A generalized entropy functional exists on caps, obeys quantum focusing on null generators, and comes with the recoverability structure used throughout the manuscript: collar tripartitions have small CMI with controlled recovery maps, with stronger collar/null-strip hypotheses stated explicitly where needed.
5. **Minimal Admissible Realization (MAR).** On the admissible low-energy branch, the realized sector package is the lexicographically minimal one under the complexity vector $C(\mathfrak{S})$. MAR is an explicit structural-economy axiom on admissible realized branches, not a theorem derived from the preceding four axioms.

4.2.4 Closure values and theorem-local technical data

The quantitative branches discussed here use two fixed-point closure values: the Phase-II pixel fixed point and the global cosmic record-capacity fixed point:

$$P \equiv a_{\mathrm{cell}}/\ell_P^2, \qquad N_{\mathrm{CRC}} \equiv \log \dim \mathcal{H}_{\mathrm{tot}}.$$

Here $P = P_\star$ is the local UV-area ratio fixed on the outer/inner closure branch summarized in Section 6, while N_{CRC} is the unique stable fixed point of the OPH-derived cosmic record-closure readback map. On the observed branch this is the de Sitter entropy capacity usually denoted $N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}$. Capacity-level neutrino side estimates tied to that branch value are legacy bookkeeping terms separate from the weighted-cycle theorem lane. These are quantitative-branch quantities, not extra axioms.

The capacity normalization uses the de Sitter static-patch entropy [14]. 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 late-time de Sitter scale [15] gives $N_{\rm patch} \simeq 1.05 \times 10^{122}$, $N_{\rm scr} \simeq 3.31 \times 10^{122}$, and $\Lambda \ell_P^2 \simeq 2.85 \times 10^{-122}$.

For the recovered core, the controlled scaling/BW step for Lorentz/null-modular/Einstein statements is Theorem 4.2, the support-visible BW scaling theorem. Transportability, the fixed-cutoff bosonic category, refinement/fiber descent, the realized compact-gauge witness, fixed-cutoff regulator bookkeeping, quasi-local propagation, endpoint-Lipschitz control, collar decomposition, and related branch-internal statements are cited directly from axioms and earlier theorems.

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

1. A finite set $\{O_a\}$ of gauge-invariant local operators, each supported on a ball of radius $\leq r_0 = O(\ell_{\rm UV})$.
2. Constraint equations $\langle O_a(x) \rangle = c_a$ for each cell x in the UV lattice.
3. Optionally, a finite number of global constraints (total energy, charge).

This is the minimal specification needed to derive the local Gibbs form of Lemma 2.6 from Axiom 3.

Clarification (MaxEnt $\neq$ thermal equilibrium). MaxEnt here is **local state selection** given constraints, not "the universe is in thermal equilibrium." The Lagrange multipliers (inverse temperatures) may vary slowly in space and time. Non-equilibrium physics appears as gradients in these multipliers and as controlled violations of exact Markov additivity under the explicit mixing hypotheses stated later in Section 2.3. Equilibrium is an approximation regime with explicit error terms.

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

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

Central-defect subbranch. On triple overlaps, when the failure of strict coherence is central, one has

$$\varphi_{ij}\varphi_{jk}\varphi_{ki} = \mathrm{Ad}(z_{ijk}), \qquad z_{ijk} \in Z(\mathcal{A}_{ijk}).$$

Collar double-scaling hypothesis. There exists a UV length $\ell_{\rm UV}$ such that for any cap C and collar width δ, in the refinement limit $\delta \to 0$ and $\ell_{\rm UV} \to 0$ with $\delta/\ell_{\rm UV} \to \infty$, the Markov error satisfies

$$I(A_\delta : D_\delta \mid B_\delta)_\omega \leq \varepsilon(\delta/\ell_{\rm UV}), \qquad \varepsilon(x) \to 0 \text{ as } x \to \infty.$$

See Section 2.3 for the collar tripartition definitions and the regulated EC proof that yields this limit from the fixed-cutoff realized presentation.

Geometric-subnet BW lift. The continuum Lorentz statement is asked only on the support-visible extracted prime geometric subnet for caps, together with the controlled collar/refinement limit that carries the Markov and recovery remainders explicitly. The finite cap regulators are type-I algebras, but the scaling-limit observer algebra may leave that class, so the theorem is automorphism-level on the emitted geometric cap pair and $K_C = 2\pi B_C$ is only the special type-I limit form. The proof of the support-visible BW scaling theorem combines regularized modular transport, weak-$*$/GNS extraction, support-readable modular covariance, ordered cut-pair rigidity, and KMS/BW normalization.

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

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

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

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

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

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

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

Heuristic motivation. The modular Hamiltonian $K_C = -\log \rho_C$ is quasi-local by Lemma 4.1a-b (modular additivity localizes it to the collar). The derived Lieb-Robinson control then bounds support spreading under $e^{iK_C t}$. In the double-scaling collar limit ($\delta/\ell_{\mathrm{UV}} \to \infty$), the thickening vanishes in macroscopic units. This is heuristic support for the tangent-limit package used later.

Continuum-limit heuristic. Define the induced region flow

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

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

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

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

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

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

for an induced map $R_{\ell \to L}$ on one common finite-dimensional multiplier family. Because the same finite constraint family is preserved, the regulator states lie in that shared multiplier family rather than in unrelated state spaces at different cutoffs, and the realized branch is the persistent trajectory or invariant subset selected inside it under the induced refinement maps. This is a state-side persistence statement; it does not by itself upgrade fixed-cutoff edge labels to a transportable refinement-persistent sector colimit.

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

External mathematical inputs: SSA and recovery theorems (Petz 1986, 1988; Fawzi and Renner 2015), the fixed-volume area-variation identity for small geodesic balls used in the local Einstein rest-frame relation, and the local quadratic-polarization argument used in the tensor upgrade. The null modular bridge and the small-ball kernel are carried internally on the compact surface rather than imported from a separate EFT small-ball law. For SM contact we also use the Doplicher-Roberts reconstruction (Doplicher and Roberts 1989, 1990) once localized zero-obstruction sectors are assembled in the small-region limit. Full citations appear in the References.

Recovered-core theorem status. The Lorentz/null-modular/Einstein recovered core runs through the support-visible automorphism theorem stated in Section 4.2. Fixed-cutoff statements remain regulator statements.

4.2.5 Notation

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

4.3 Information-Theoretic Tools

4.3.1 Strong subadditivity and Markov states

For any tripartite state ρ_{ABC},

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

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

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

4.3.2 Approximate recovery

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

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

4.3.3 Collar refinement and sufficient mechanisms

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

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

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

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

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

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

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

This is the fixed-cutoff realized presentation used throughout the collar analysis. The EC theorem works on the invariant-state realization and uses the sector-preserving induced collar algebra there; the entire fixed-point algebra on the unreduced tensor product is outside that theorem surface. When the consensus paper speaks of quotient repair, the physical repair law is defined on this fixed-point / overlap-invariant data and any representative-level map is only a lift of that quotient update. Descent to the gauge quotient is therefore built into the physical-algebra formulation as quotient-level repair. On the declared fixed-cutoff collar branch, the local repair step itself is read from exact Markov splice or a declared Petz/Fawzi–Renner recovery channel. The declared fixed-cutoff branch adds repair completeness, together with control on the declared Petz domain where that branch is used. The consensus paper proves those clauses on a nontrivial rooted-tree packet-net export: finite packet labels live on a rooted tree, hidden labels are quotient-only, weighted parent-copy repair strictly lowers Φ, normal forms are exactly consistent packet assignments, and the classical full-support Petz channel is CPTP and trace-norm contractive on its positive-support-gap domain. Gauge-invariant observables factor through the quotient normal form on that verified branch. The declared touched-overlap acceptance contract makes Φ the finite-patch Lyapunov functional for accepted moves, and the support-local disjoint-commutation plus restriction-compatible union-collar glue that turn overlap associativity into the local diamond are part of the declared repair law itself. On the same fixed-cutoff quantum lift, the terminal expectation functional on each declared physical observable algebra is therefore unique even when microscopic representative lifts differ by gauge labels globally or by sector labels inside one quotient-local glued state. The unqualified coding statement at this stage is only a finite constraint-code statement: the overlap-net codewords are the globally consistent states $C = \Phi^{-1}(0)$. A bare graph does not determine code distance–the same graph can carry constant readouts with distance 1 or repetition constraints with distance $|V|$. Topological-code distance/min-cut, Knill–Laflamme resilience, spectral-gap convergence, BFT wall-clock liveness, and hardware search-work reduction therefore require their own certificates and are not imported by the core overlap graph. The logic is not “termination implies uniqueness”: Lyapunov descent gives termination, while the local diamond on the physical quotient plus repair completeness gives confluence. The theorem therefore gives one schedule-independent terminal quotient normal form from a fixed initial quotient state. A stronger same-boundary conclusion requires a preserved boundary/sector map and at most one consistent quotient extension in that boundary fiber. If accepted schedules terminate at different observer-facing quotient normal forms without a declared holonomy or higher-gauge obstruction, that repair law is outside the

OPH consensus theorem. For separated cofinal refinement systems, the consensus paper adds an inverse-limit bridge: if finite-stage restriction maps commute with the quotient normal-form maps and holonomy maps, and if compatible families are visibly separated on cofinal finite stages, then finite normal forms and holonomy obstructions assemble into unique refinement-limit classes with finite-stage witnesses for nonzero holonomy. It also adds the RG-facing comparison: when a chosen coarse-graining channel shadows the finite-stage normal forms and obstruction maps with declared defects ε^n and ε^h, reconciling first and then coarse-graining gives the same macroscopic law readout as coarse-graining first and then reconciling, up to $\max\{\varepsilon^n, \varepsilon^h\}$; exact naturality is the zero-defect case. The same consensus surface is classified explicitly as a finite-state exact theorem package plus this controlled inverse-limit bridge: decidable normal-form computation with the Lyapunov step bound, automatic approximate stability only through the collar-local splice and record controls, a conditional fair-block contraction branch for long-run noisy approximate consensus, and a verified rooted-tree packet-net subdomain. Computational universality for growing patch-net families belongs to the consensus paper's expressive-power boundary and is not a dependency for the recovered-core branches.

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

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

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

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

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

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

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

and the sector-preserving collar algebra

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

with

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

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

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

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

Then

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

By Schur's lemma,

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

Therefore the invariant subspace is

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

as claimed. The induced sector-preserving collar algebra is

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

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

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

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

Exact Markov route. If the reference state on $A_\delta B_\delta D_\delta$ is exact Markov, or in an explicitly stated idealized limit that reduces to exact Markovity, then

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

This is the HJPW normal form applied to the EC blocks, and it is the exact identity used for literal Markov-modular equalities.

Interpreting collar refinement as the inductive limit of these regulators with $\delta/\ell_{\mathrm{UV}} \to \infty$, Theorem 2.3 supplies the kinematic edge-center decomposition. Exact Markovity is one idealized route. The following lemma and axiom provide the quantitative decay route used when the manuscript keeps the approximation explicit.

Lemma 2.6 (MaxEnt with local constraints implies local Gibbs form). Under the local MaxEnt branch of Axiom 3 and the finite-dimensional regulator realization above, the MaxEnt state has the Gibbs form

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

where the sum runs over UV cells x and constraint operators O_a. The effective Hamiltonian H_{eff} is quasi-local with range $O(\ell_{\mathrm{UV}})$.

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

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

Strict concavity of von Neumann entropy ensures uniqueness. When the constraints are "translated local" (the same O_a at each cell x), the exponent is a sum of local terms. QED.

Exponential mixing hypothesis. There exist constants c and correlation length $\xi = O(\ell_{\mathrm{UV}})$ such that

$$I_\omega(A_\delta : D_\delta \mid B_\delta) \leq c\, |\partial C|_{\mathrm{UV}}\, e^{-\delta/\xi}, \qquad |\partial C|_{\mathrm{UV}} \sim \frac{\operatorname{length}(\partial C)}{\ell_{\mathrm{UV}}}.$$

This is the standard fixed-cutoff clustering/mixing condition used for local Gibbs states. One sufficient realization on a chosen fixed-cutoff collar family is a Dobrushin-type uniqueness estimate for the corresponding local Gibbs specifications, or a spectral-gap estimate uniform across that family. Used here, it is a collar-local recoverability condition on the selected branch; it does not imply infinite-volume uniqueness and it does not decide whether the separately constructed refinement-limit gauge-sector colimit is trivial or nontrivial.

Theorem 2.5 (Local Gibbs + mixing implies collar refinement). Under Lemma 2.6 and the exponential mixing hypothesis above, the collar double-scaling hypothesis holds in the limit $\delta \to 0$, $\ell_{\mathrm{UV}} \to 0$ with $\delta/\ell_{\mathrm{UV}} \to \infty$.

Proof. The mixing bound above has polynomial growth in $|\partial C|_{\mathrm{UV}}$ and exponential decay in $\delta/\ell_{\mathrm{UV}}$. In the double-scaling limit the exponential dominates, so $I_\omega(A_\delta : D_\delta \mid B_\delta) \to 0$. QED.

This bound is the quantitative hinge for constructive gluing.

4.3.4 Concrete UV realization: quantum link models

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

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

Degrees of freedom. Attach to every oriented link ℓ a **finite-dimensional** Hilbert space $\mathcal{H}_\ell$. In ordinary Wilson lattice gauge theory, the continuum/refinement-limit edge description is modeled by $L^2(G)$ (infinite-dimensional for continuous G), but the microscopic OPH regulator is instead a **quantum link model** with finite-dimensional link Hilbert spaces that preserve gauge symmetry in operator form [34]. Optionally attach matter Hilbert spaces $\mathcal{H}_v$ at vertices. Then:

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

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

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

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

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

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

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

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

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

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

with center generated by the projectors P_α. This is precisely the Schur-lemma mechanism of Theorem 2.3. The labels α are the familiar edge-mode / electric-flux labels appearing whenever one factorizes gauge theories across an entangling cut [35]. Exact Markovity depends on the state and is not forced by the decomposition alone.

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

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

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

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

Scope of this example. The quantum-link realization makes the fixed-cutoff matrix/fixed-point bookkeeping concrete. Modular flow on caps becomes geometric conformal dilation with the 2π KMS normalization through the support-visible BW scaling theorem on the prime geometric subnet. Viable architectures for this include holographic quantum error-correcting codes [36] and quantum double / string-net Hamiltonians [37], but any use of QECC distance, min-cut resilience, or Knill–Laflamme correction requires the corresponding code subspace, logical-operator, error-family, and recovery certificate.

4.3.5 Conformal-modular fixed-point microphysics

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

Local finite-constraint MaxEnt branch. The constraint family $\mathcal{C}$ is generated by finitely many gauge-invariant local densities $\{O_a(x)\}$ of UV range $O(\ell_{\mathrm{UV}})$, with the same finite label set retained under refinement. This is not an extra postulate beyond the third OPH axiom; it is that axiom unpacked at regulator level.

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

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

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

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

Derived propagation control on the same branch. Because H_{MaxEnt} is a finite-range or quasi-local sum on a finite type-I regulator net, standard Lieb–Robinson estimates [6] apply to the automorphism group it generates, or to any branch generator lying in the same bounded-support algebraic closure. Thus there is a finite propagation velocity v_{LR} and constants C, ξ such that for local observables A_X, B_Y,

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

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

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

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

Refinement-stable multiplier branch. Because the same finite constraint family is preserved under coarse-graining, the regulator states lie in one common finite-dimensional multiplier family rather than in unrelated state spaces at different cutoffs. The "refinement-stable" language used later therefore means persistence along the stable or fixed branch of this multiplier family. This state-side notion is enough to compare realized states across cutoffs. The fixed-cutoff bosonic sector category is constructed in FixedCutoffBosonicSectorCategory, and the refinement functors plus finite bosonic fiber descent are constructed in RefinementFunctorAndFiberDescent. The realized compact-gauge witness theorem supplies nonempty realized MAR-admissible witness data, but the Standard Model selection step is the later MAR application to sector packages. Accordingly, whenever the gauge derivation speaks of a refinement-stable directed colimit of zero-obstruction sectors, the third OPH axiom supplies the realized state branch along which those theorem-produced sector objects persist.

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

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

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

Proposition 2.6 (Geometric modular action on caps on the extracted prime geometric subnet). For any support-visible extracted scaling-limit geometric cap pair $(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C})$ supplied by Theorem 4.2, let $\alpha_{\lambda_C(s)}$ be the automorphism induced by the standard cap-preserving conformal subgroup $\lambda_C(s)$. Then

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

If the limit cap algebra happens to be type I, this may be written as $K_C = 2\pi B_C$. In the generic continuum case, the same statement is an outer modular action on a non-type-I algebra.

Proof. Ordered cut-pair rigidity on the extracted prime geometric cap pair identifies the scaling-limit cap modular group with the standard geometric cap subgroup up to normalization, and the modular KMS condition fixes the normalization to 2π. If the limit algebra is type I, the automorphism statement may be represented by a modular Hamiltonian K_C; otherwise the automorphism statement is the full content. QED.

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

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

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

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

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

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

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

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

- **Derived half-sided modular inclusion.** Nested null half-line algebras satisfy half-sided modular inclusion on the declared geometric scaling branch by Corollary 5.2e; Borchers–Wiesbrock then identifies the positive null-translation generator on its Stone domain together with the affine half-line modular relation [10].

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

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

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

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

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

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

QNEC internalization. QNEC has rigorous QFT proofs in broad settings [16]. On the support-visible scaling branch with the extracted geometric cap pair in place, the Recoverable Generalized Entropy axiom can be supported internally by:

- EC + MaxEnt derive $S_{\text{gen}} = S_{\text{bulk}} + \langle L_C \rangle$ (Section 5.4).

- On the null-stress and Einstein theorem branch, focusing becomes a semiclassical consequence in the same scaling regime.

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

- **Internal branch consequences:** the local-Gibbs form, quasi-local propagation, bounded-interval endpoint-Lipschitz control, and the realized state-side refinement branch along which later transport questions are asked, all from the local finite-constraint MaxEnt branch.

- **Support-visible scaling statement:** the scaling limit emits the support-visible prime geometric cap pair and ordered cut-pair rigidity holds on it. On that branch the limit algebra may be non-type-I, the modular action may be outer, and the 2π normalization and later null half-sided-inclusion bridge follow.

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

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

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

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

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

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

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

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

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

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

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

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

BW-side closure status. Proposition 2.6a blocks the stronger unregularized full-algebra conclusion from the proved finite-cutoff package. The regularized theorem supplies the constructive replacement used by the paper: support-visible modular transport is controlled with an explicit cutoff schedule, and Theorem 4.2 uses exactly that observer-facing content to close the BW/geometric cap-pair statement. Uniform-on-compact-time convergence of the regularized modular automorphism groups would require an additional equicontinuity, strong-resolvent, or Trotter-style lemma. A black-box AQFT/BW reconstruction route would require separately verifying conformal-net properties such as locality, additivity, duality, standardness, positive energy, and suitable modular inclusions.

4.4 Overlap Consistency and Gluing

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

4.4.1 Constructive gluing on tree covers

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

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

with

$$\delta_k = 2\sqrt{\ln 2\, \varepsilon_k}.$$

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

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

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

4.4.2 Gauge-as-gluing and loops

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

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

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

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

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

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

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

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

with

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

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

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

4.4.3 Loop obstruction class (central defect)

On the central-defect subbranch, define central defects z_{ijk} by

$$\varphi_{ij}\varphi_{jk}\varphi_{ki} = \mathrm{Ad}(z_{ijk}) \quad \text{on } \mathcal{A}_{ijk}.$$

Then $\{z_{ijk}\}$ is a Čech 2-cocycle, and its cohomology class $[z]$ is gauge invariant. Central defects do not obstruct the collar block decomposition: they only replace K_Σ by its central extension $\widehat{K}_\Sigma$. Ordinary loop-coherent gluing exists iff $[z] = 0$. (A full proof appears in Section 6.4 below, in the algebra-net language.)

4.4.4 Non-central obstruction (2-group cocycle)

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

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

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

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

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

obeying the 2-cocycle conditions

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

and on quadruple overlaps,

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

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

Theorem 3.4 (non-central obstruction). Loop-coherent gluing exists iff the 2-cocycle (g_{ij}, h_{ijk}) is equivalent to the trivial cocycle in nonabelian Čech H^2 with values in the crossed module $(H \to G)$.

Proof sketch. Strict gluing corresponds to $h_{ijk} = 1$ and $g_{ij}g_{jk} = g_{ik}$. Gauge changes are exactly the crossed-module coboundaries, so strictification exists iff the 2-class is trivial. QED.

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

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

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

The physical higher-gauge collar is

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

with

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

and the defect class

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

is invariant under local rechartings, classifies fixed-cutoff genuinely noncentral sectors, and vanishes iff the defect is removable.

Proof sketch. Pair-overlap rechartings are inner, while triple-overlap associators strictify to compact crossed-module data. Finite-dimensional unitary $\mathcal{T}_\Sigma$-modules decompose semisimply, and Schur matching leaves only diagonal left/right blocks. The transport statement is the crossed-module Čech analogue of the central-defect case. QED.

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

On the ordinary branch, this transport is strictly path-independent iff every closed overlap loop has trivial holonomy on the collar-sector block. On the central branch, elementary triangle moves accumulate the central cocycle z_{ijk}; strict path-independent transport exists iff the central loop-coherence class $[z]_\Sigma \in \check{H}^2(N_\Sigma, Z)$ vanishes. When $[z]_\Sigma = 0$, a central 1-cochain strictifies the lifts U_{ij}, and when $[z]_\Sigma \neq 0$, the residual path dependence is exactly the represented central multiplier.

On the genuinely noncentral branch, path comparisons are H_Σ-valued 2-morphisms in the crossed module $H_\Sigma \to G_\Sigma$. Strict ordinary transport exists iff $q_\Sigma = [(g,h)]$ vanishes. If $q_\Sigma = 0$, the crossed-module data strictify to a genuine G_Σ-valued 1-cocycle and the ordinary path-composite construction applies. If $q_\Sigma \neq 0$, the sector is a higher-gauge fixed-cutoff sector with higher transport rather than an ordinary compact-group DR sector. Thus transportability is a theorem-level zero-obstruction classification, not an independent input. QED.

4.5 Modular Flow and Lorentz Kinematics

4.5.1 Modular additivity in the Markov collar limit

Consider a collar tripartition $A : B : D$ around a cap boundary, with the EC decomposition of Section 2.3 understood. Define, for a faithful reference state ω,

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

Exact modular additivity is not a consequence of EC alone. It is available only on the exact Markov set, or along a controlled fixed-cutoff family that approaches that set on one fixed collar model. Three quantities must therefore be kept separate:

1. the raw collar conditional mutual information $I(A : D \mid B)$;
2. the constructive Fawzi–Renner comparison error

$$r_{\mathrm{FR}}(\varepsilon) := 2\sqrt{1 - e^{-\varepsilon}} \leq 2\sqrt{\varepsilon};$$

3. the fixed-collar exact-Markov replacement modulus

$$\delta^{\mathrm{M}}_{A:B:D}(\varepsilon) := \sup\left\{ \inf_{\sigma \in \mathfrak{M}_{A:B:D}} \|\rho - \sigma\|_1 : I(A : D \mid B)_\rho \leq \varepsilon \right\},$$

   where

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

Lemma 4.1a (Exact Markov implies exact additivity up to a central term). If $I(A : D \mid B)_\omega = 0$, then the EC decomposition of Section 2.3 puts the state in HJPW block form

$$\omega_{ABD} = \bigoplus_\alpha p_\alpha\, \omega^{(\alpha)}_{Ab^\alpha_L} \otimes \omega^{(\alpha)}_{b^\alpha_R D},$$

and $\Delta K(\omega)$ is central. On the canonical HJPW block model one may take

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

Equivalently, there exists a central operator $K_{\partial,ABD}(\omega) \in Z(\mathcal{A}(B))$ such that

$$K_{ABD}(\omega) = K_{AB}(\omega) + K_{BD}(\omega) - K_B(\omega) + K_{\partial,ABD}(\omega).$$

Proof. On each HJPW block the modular Hamiltonians of ABD, AB, BD, and B are the logarithms of tensor-product states with the same classical block weight p_α. The blockwise logarithms therefore cancel exactly in the combination $K_{ABD} - K_{AB} - K_{BD} + K_B$. If one keeps the blockwise endpoint-label bookkeeping explicit rather than absorbing it into the canonical block identification, the remainder is a direct sum of block constants, hence central. QED.

Proposition 4.1b (Controlled exact-Markov replacement on a fixed collar). Because the state space is compact and conditional mutual information is continuous in finite dimension,

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

Thus small collar CMI converges to the exact HJPW normal form only in this controlled fixed-cutoff sense; the manuscript does not claim a dimension-free one-shot trace-norm bound from small $I(A:D \mid B)$ directly to an exact Markov state.

Theorem 4.1b' (Finite-collar Markov replacement stability). On one fixed faithful finite-dimensional collar model, the qualitative modulus above can be made collar-local and quantitative. If all relevant marginals stay above a chosen floor $\lambda_* > 0$, then there are constants $C_{A:B:D,\lambda_*} > 0$ and $\theta_{A:B:D,\lambda_*} > 0$, depending only on that collar model and floor, such that

$$d_{\mathrm{M}}(\rho) \le C_{A:B:D,\lambda_*}\, I(A:D \mid B)_\rho^{\theta_{A:B:D,\lambda_*}}.$$

This is a fixed-collar Lojasiewicz-type rate for the analytic function $I(A:D \mid B)$ on the faithful state manifold. It is not a dimension-free stability theorem for arbitrary tripartite quantum systems.

Corollary 4.1c (Carried collar defect operator). Let ω_ε satisfy $I(A:D \mid B)_{\omega_\varepsilon} \le \varepsilon$, and choose $\sigma_\varepsilon \in \mathfrak{M}_{A:B:D}$ with

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

Define the carried defect operator

$$\mathfrak{D}_{A:B:D}(\omega_\varepsilon, \sigma_\varepsilon) := \Delta K(\omega_\varepsilon) - \Delta K(\sigma_\varepsilon).$$

Then

$$K_{ABD}(\omega_\varepsilon) = K_{AB}(\omega_\varepsilon) + K_{BD}(\omega_\varepsilon) - K_B(\omega_\varepsilon) + K_{\partial,ABD}(\sigma_\varepsilon) + \mathfrak{D}_{A:B:D}(\omega_\varepsilon, \sigma_\varepsilon),$$

where $K_{\partial,ABD}(\sigma_\varepsilon) = \Delta K(\sigma_\varepsilon)$ is central. For every bounded observable X supported on $A \cup B$,

$$|\mathrm{Tr}[X(\omega_\varepsilon - \sigma_\varepsilon)]| \le \|X\|_\infty\, \delta^{\mathrm{M}}_{A:B:D}(\varepsilon).$$

If the relevant marginals of ω_ε and σ_ε are uniformly faithful with lower spectral bound $\lambda_* > 0$, then

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

Independently, Fawzi–Renner recovery supplies a constructive comparison state $\omega_\varepsilon^{\mathrm{rec}}$ with

$$\|\omega_\varepsilon - \omega_\varepsilon^{\mathrm{rec}}\|_1 \le r_{\mathrm{FR}}(\varepsilon), \qquad r_{\mathrm{FR}}(\varepsilon) := 2\sqrt{1 - e^{-\varepsilon}} \le 2\sqrt{\varepsilon}.$$

This $O(\varepsilon^{1/2})$ term is the finite-stage observable error. The modulus $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon)$ is instead the fixed-collar quantity that justifies replacing the physical state by an exact Markov normal form in the later geometric modular arguments.

Accordingly, the manuscript's later exact collar formulas take one of two forms:

1. literal exactness when the reference state is exact Markov on the relevant collar; or

2. a controlled collar family for which $\delta^{\mathrm{M}}_{A:B:D}(\varepsilon_\delta) \to 0$, with the constructive Fawzi–Renner remainder and the carried operator defect $\mathfrak{D}_{A:B:D}$ kept explicit at finite stage.

Theorem 4.1d (Finite-stage modular-defect propagation and dimension quarantine). For any fixed downstream branch calculation that uses finitely many collar or strip modular-additivity identities, replacing each exact identity by its controlled finite-stage form changes every bounded downstream modular observable O by at most

$$\mathcal{P}_O\Big(\{r_{\mathrm{FR}}(\varepsilon_j)\}, \{\delta^{\mathrm{M}}_j(\varepsilon_j)\}, \{\eta^{\mathrm{reg}}_j\}\Big),$$

where $\mathcal{P}_O \to 0$ as all listed errors vanish. The polynomial-continuity modulus depends only on the declared fixed collar models, support-visible dimensions, bounded observable class, faithful floors or regularization schedules, and bounded modular-time intervals. No BW/Lorentz, null-modular, or Einstein scaling step uses a dimension-free trace-norm stability theorem from small CMI directly to exact Markov normal form.

4.5.2 Theorem: BW_{S^2} on the extracted prime geometric subnet

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 fixed-cutoff matrix-algebra identity and not the slogan "finite cells imply Lorentz invariance." Its target is the support-visible refinement-limit geometric cap pair $(\mathcal{A}^{\mathrm{geo,sv}}_\infty(C), \omega^{\mathrm{geo},C}_\infty)$. The branch theorem is the conditional implication

cap-pair extraction+ regularized modular transport+ support-readable modular covariance+ round-cap rigidity+

The support-visible theorem below uses regularized modular transport, projectively compatible exact-Markov replacement on fixed collars, weak-∗/GNS extraction, support-readable modular covariance, and ordered cut-pair rigidity; it does not require a type-I continuum algebra or a full-algebra unregularized common spectral floor.

Fix a cap $C \subset S^2$ and a shrinking collar family $(A_\delta, B_\delta, D_\delta)$ around ∂C. 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,

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

where $\lambda_* > 0$ is the lower spectral bound used to compare modular Hamiltonians. All exact collar formulas in this subsection are therefore to be read either literally at exact Markovity or along a controlled collar family satisfying

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

with $r_{\mathrm{FR}}(\varepsilon_\delta)$ and η^{M}_δ 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.

Theorem 4.2 (Support-visible BW_{S^2} scaling theorem on the extracted prime geometric subnet). Let $C \subset S^2$ be a round cap. For every OPH-realized observer-supporting refinement branch satisfying the five OPH axioms and the derived fixed-cutoff regulator/collar/consensus package, the support-visible prime geometric cap net admits a weak-$*$/GNS scaling-limit cap pair $(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C})$. Let $\lambda_C(s)$ denote the standard cap-preserving conformal one-parameter subgroup, normalized so that the null blow-up near a smooth cut acts by $v \mapsto e^{-s}v$. Then the scaling-limit modular automorphism group is

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

At finite regulator stage the nonadditive part of

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

is confined to the shrinking collar up to the carried errors $r_{\mathrm{FR}}(\varepsilon_\delta)$, the fixed-collar replacement modulus δ^{M}, and the regularized support-visible modular transport bound. The support-visible modular limit is independent of the exact-Markov replacement sequence after projective compatibility and GNS quotienting, and support-readable modular covariance turns the limiting modular group into a cap-local support map before round-cap rigidity is applied. No separate cap-isotropy/SO(2)-equivariance selector, finite-cell Lorentz-invariance premise, or full-algebra unregularized common floor is used; the conformal support-map clause is the explicit scaling-limit branch condition. Record/pointer and interface-inert auxiliary registers are outside the extracted geometric subnet. If the scaling-limit cap algebra is type I, the same automorphism identity may be represented by

$$K_C = 2\pi B_C;$$

otherwise the automorphism identity is the full statement and the action is outer. 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.

Proof. Markov locality localizes the modular defect to the shrinking collar, with carried finite-stage errors $r_{\mathrm{FR}}(\varepsilon_\delta)$ and δ^{M}. Proposition 2.6a shows why a full-algebra unregularized floor is unavailable in general. Theorem 2.6b supplies the replacement: regularized support-visible modular matrix elements converge on every fixed local collar model under an explicit cutoff schedule. Projective replacement compatibility and the support-visible convergence audit ensure that the emitted modular limit is independent of the chosen exact-Markov comparison family. Support-visible cap-pair extraction on the local GNS support quotient uses weak-$*$ compactness of finite-stage state spaces and the consensus refinement mechanism to supply local limit states, then applies GNS construction to obtain the scaling-limit cap pair. Support-readable modular covariance reads the modular group as a cap-local support map, and round-cap rigidity from surviving cut data identifies the scaling-limit modular automorphism group with the standard cap-preserving conformal flow up to normalization:

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

for some $\kappa_C > 0$. Because s is normalized by the null blow-up $v \mapsto e^{-s}v$, the modular KMS condition fixes $\kappa_C = 2\pi$. Thus

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

If the limit algebra is type I, this automorphism identity is represented by $K_C = 2\pi B_C$; otherwise the automorphism identity is the full statement and the action is outer. QED.

Definition 4.2a (BW-branch observer-relative time reading). On the branch satisfying the hypotheses of Theorem 4.2, OPH uses the modular automorphism parameter t of the extracted cap pair $(\mathcal{A}_\infty^{\mathrm{geo,sv}}(C), \omega_\infty^{\mathrm{geo},C})$ as that cap observer's relative time parameter. This is a declared BW-branch reading of the geometric modular flow derived in Theorem 4.2; it is not an additional proof that arbitrary operational clocks, global time, or the full problem of time have been derived from the axioms.

Scope note. The theorem fixes the 2π normalization and the carried refinement limit on the support-visible extracted geometric cap pair. The UV-side scaffold is the realized transported geometric cap-local system, regularized support-visible modular transport on fixed local collars, support-visible cap-pair extraction on the local GNS support quotient, support-readable modular covariance, and round-cap rigidity from surviving cut data. No separate cap-isotropy input is used, and the conformal support-map clause is the explicit scaling-limit branch condition rather than a finite-regulator Lorentz premise. Stronger compact-time group convergence or a black-box conformal-net BW reconstruction would be extra analytic certificates, not hidden steps in this theorem.

4.5.3 Theorem: BW_{S^2} implies Lorentz kinematics

Theorem 4.3 (Lorentz kinematics on the screen). Under the hypotheses of Theorem 4.5.2,

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

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 4.5.2, 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. QED.

4.6 Gravity from Fixed-Cap Generalized-Entropy Stationarity

4.6.1 Cap first law

Section 4.2 fixes the cap modular statement first at the automorphism level. On the extracted prime geometric subnet, the scaling-limit cap modular flow is geometric with the standard 2π normalization,

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

If the realized scaling-limit cap algebra happens to be type I, one may choose a density matrix $\rho_C^{\omega_\infty^{\mathrm{geo},C}}$ and modular Hamiltonian

$$K_C := -\log \rho_C^{\omega_\infty^{\mathrm{geo},C}},$$

and then for a perturbation $\rho(\varepsilon)$ with $\rho(0) = \omega_\infty^{\mathrm{geo},C}$ the first-law identity reads

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

In the generic continuum/non-type-I case, Section 4.2 does not supply an inner operator identity $K_C = 2\pi B_C$; the theorem surface there is the geometric modular automorphism statement, so the cap first-law formula is restricted to the special type-I realization.

4.6.2 Null-surface modular bridge

This subsection extends the fixed-cutoff weak-tail-generator boundary to the derived positive null-translation stage and the exact half-line generator/charge identification. At fixed cutoff one first transfers cut-center data to regulated null strips, then imposes the extra inherited split condition needed for the spatial-collar-type tensor decomposition, obtains exact or controlled strip additivity on one inherited strip model, and derives a weak tail generator for the renormalized half-line family. On the scaling-limit geometric-cap branch of Theorem 4.2, the null half-line blow-up net then inherits geometric dilation and therefore half-sided modular inclusion, so Borchers–Wiesbrock supplies an explicit positive null-translation generator on its Stone domain. The downstream boundary is narrower: bounded-interval formulas require the separate interval-preserving projective branch, and the tensor upgrade carries the null-invisible metric ambiguity. Accordingly the null modular bridge consists of:

- null-cut center transfer yielding the central cut-label decomposition, together with the extra inherited split condition needed for the spatial-collar-type tensor decomposition;
- exact or controlled four-term strip additivity on one fixed inherited strip model;
- endpoint-Lipschitz control for the renormalized half-line family and the resulting weak tail generator;
- derived half-sided modular inclusion on the null half-line blow-up net, and the resulting Borchers positive translation generator.

The later density-upgrade template is kept below as an explicit downstream template rather than as a fixed-cutoff theorem proved in this subsection. Lemma 5.2f makes the positive null-translation generator itself explicit, and Theorem 5.2g closes the half-line generator/charge identification on that same family: they record the Borchers unitary group, positivity and self-adjointness on the Stone domain, the corrected affine commutator relation, the half-line modular-Hamiltonian identity $K_a = K_0 - 2\pi a P_\Omega$, and the exact identification of P_Ω with the local null-stress charge. Bounded-interval formulas are downstream.

Proposition 5.2a (Null-cut center transfer and inherited split). Fix a regulated null tripartition

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

with cuts $\Gamma_- := \{v = v_2\}$ and $\Gamma_+ := \{v = v_3\}$. Under the fixed-cutoff realized presentation, assume that the two null cuts inherit the ordinary or central-defect boundary-redundancy data used in the spatial collar branch, so that in a compatible type-I regulator presentation one has

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

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

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

where opposite sides of each cut carry inverse transport. If the strip algebra is the commutant of the transported cut actions, then

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

If, in addition, each multiplicity space factors as

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

with $\mathcal{A}(I_- \cup J)$ acting blockwise only on $\mathcal{H}_{i_-^{\alpha_-}} \otimes \mathcal{H}_{j_L^{\alpha_-,\alpha_+}}$ and $\mathcal{A}(J \cup I_+)$ acting blockwise only on $\mathcal{H}_{j_R^{\alpha_-,\alpha_+}} \otimes \mathcal{H}_{i_+^{\alpha_+}}$, then

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

and

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

Proof. Complete reducibility gives the displayed decompositions. Commuting with the two transported cut actions forces strip operators to preserve the sector pair (α_-, α_+) and act only on the multiplicity space; Schur's lemma kills the off-diagonal intertwiners and leaves the direct-sum algebra above. Taking invariants across both cuts in the glued tripartition again uses Schur's lemma and leaves exactly the matching sector pairs. The left/right split of the multiplicity spaces is additional structure; it is not forced by the center transfer alone. QED.

Corollary 5.2b (Exact or controlled four-term null modular relation on an inherited strip model). On one fixed finite-dimensional strip model satisfying Proposition 5.2a, define

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

and

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

For any strip state η, let

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

If the strip reference state ω is exact Markov, then $\Delta K_J(\omega)$ is central, and on the canonical inherited HJPW model one may take

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

Equivalently, 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 $I(I_- : I_+ \mid J)_\omega \le \varepsilon$, choose $\widetilde{\omega}_J \in \mathfrak{M}_{I_-:J:I_+}$ with

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

Define

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

Then

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

with $K_{\partial,J}(\widetilde{\omega}_J) = \Delta K_J(\widetilde{\omega}_J) \in Z(\mathcal{A}(J))$. Every bounded observable on $I_- \cup J$ or $J \cup I_+$ then differs from the exact-Markov reference by at most

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

If the relevant strip marginals 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, Fawzi–Renner gives a recovered comparison state with trace-norm error

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

Thus the exact four-term strip relation is available only at exact Markovity or in a controlled strip family on one fixed inherited strip model with $\delta^{\mathrm{M}}_{I_-:J:I_+}(\varepsilon_J) \to 0$, while $\mathfrak{D}_J$ is carried explicitly at finite stage.

Proof. On the inherited split of Proposition 5.2a, the exact-Markov case is the same HJPW block calculation as in the spatial collar theorem: the four logarithms cancel blockwise on the canonical model, and any residual bookkeeping term is central. The controlled replacement is the same fixed-model compactness argument used for ordinary collars, after relabeling $A, B, D \mapsto I_-, J, I_+$; the operator-norm bound on $\mathfrak{D}_J$ is the corresponding relabeling of the modular-transport estimate. QED.

Definition (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 5.2b on the inherited strip model, or by its controlled exact-Markov replacement when that model is used as reference. Define

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

Proposition 5.2c (Endpoint-Lipschitz null modular families and weak tail generator). For renormalized modular Hamiltonians on one null generator,

$$\widetilde{K}[I,\Omega] := K[I,\Omega] - K_\partial(I,\Omega), \qquad \widetilde{K}_a(\Omega) := \widetilde{K}[(a,\infty),\Omega],$$

the branch-internal endpoint-control coming from the local finite-constraint MaxEnt branch gives:

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

for bounded intervals in a compact endpoint window, and

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

for half-lines. Hence

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

is locally Lipschitz and therefore absolutely continuous, and its distributional derivative defines the weak tail generator

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

For $a < b$,

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

$q(a,\Omega)$ is a weak tail generator rather than a local operator-valued density; its identification with the positive self-adjoint Borchers generator occurs only after Corollary 5.2e and Lemma 5.2f.

Proof. The derived endpoint-control estimate applies to the renormalized family after removal of the central endpoint term. For bounded intervals, the symmetric-difference length is bounded by $|a'-a| + |b'-b|$; for half-lines it is exactly $|a'-a|$. Therefore the matrix-element functions are

Lipschitz in the endpoints. A Lipschitz function on a compact interval is absolutely continuous, so $f_{\psi,\phi}(a)$ has an L^∞_{loc} derivative and obeys the fundamental theorem of calculus. Defining q as the rescaled negative derivative gives the displayed integral relation. QED.

Downstream c.12–c.13 boundary. The fixed-cutoff bridge established above is the end of this subsection's proved content.

Lemma 5.2d (Downstream density-upgrade template). If, in addition, the weak tail generator $q(a,\Omega)$ is weakly differentiable in a and both $f_{\psi,\phi}(a)$ and $q_{\psi,\phi}(a)$ vanish at $+\infty$, then one may define a density

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

and obtain the half-line formula

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

This is a downstream density-upgrade step and is not proved from the fixed-cutoff strip arguments above.

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

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

By isotony of the null interval net,

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

Corollary 5.2e (Derived half-sided modular inclusion on null half-lines). On the scaling-limit geometric-cap branch of Theorem 4.2, 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. After the harmless convention change $t\mapsto -t$, this is the standard positive-time half-sided-inclusion form. The half-sided modular inclusion is therefore derived from the null-interval structure, isotony, and the scaling-limit geometric action rather than imported separately.

Proof. Blow up the cap modular flow of Theorem 4.2 near a smooth cut and restrict to the chosen null generator. In the tangent limit the cap-preserving flow becomes the null dilation $v\mapsto e^{-2\pi t}v$. For every bounded interval $(c,d)\subset H_a$, this sends $\mathcal{A}((c,d),\Omega)$ to $\mathcal{A}((e^{-2\pi t}c, e^{-2\pi t}d),\Omega)$, and taking the von Neumann closure of the interval net yields the displayed identity for $\mathcal{M}_a(\Omega)$. If $t\leq 0$, then $e^{-2\pi t}a\geq a$, so $H_{e^{-2\pi t}a}\subseteq H_a$, and isotony gives $\mathcal{M}_{e^{-2\pi t}a}(\Omega)\subseteq\mathcal{M}_a(\Omega)$. QED.

Lemma 5.2f (Positive null-translation generator). For the derived half-sided modular inclusion

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

let $\Delta_0(\Omega)$ be the modular operator of the standard pair $(\mathcal{M}_0(\Omega),\omega)$ and define $K_0(\Omega) := -\log\Delta_0(\Omega)$. Borchers–Wiesbrock then 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

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

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,

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

If $\Delta_a(\Omega)$ denotes the modular operator of $(\mathcal{M}_a(\Omega),\omega)$ and $K_a(\Omega) := -\log\Delta_a(\Omega)$, then

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

so

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

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 normalization of Corollary 5.2b, the weak endpoint derivative from Proposition 5.2c is therefore exactly the Borchers generator. Bounded-interval modular-Hamiltonian formulas are downstream and require the separate interval-preserving branch recorded in Theorem 5.2g.

Proof. Corollary 5.2e gives the required half-sided modular inclusion. Borchers–Wiesbrock then yields the unique strongly continuous unitary group $U_\Omega(a)$, its positive self-adjoint generator P_Ω, and the affine covariance relation with the modular group; Stone's theorem gives the displayed domain formula. Differentiating

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

first at $a = 0$ and then at $t = 0$ on the common analytic core gives

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

Because $U_\Omega(a)\omega = \omega$ and $U_\Omega(a)\mathcal{M}_0(\Omega)U_\Omega(a)^* = \mathcal{M}_a(\Omega)$, uniqueness of modular data for the translated standard pair gives the displayed formulas for $\Delta_a(\Omega)$ and $K_a(\Omega)$. Differentiating $K_a(\Omega) = U_\Omega(a)K_0(\Omega)U_\Omega(a)^*$ in a yields $dK_a/da = -2\pi P_\Omega$ as a quadratic-form identity, and integrating from 0 to a gives the stated half-line modular-Hamiltonian relation. QED.

Theorem 5.2g (The half-line generator is the local null-stress charge). In the canonical normalization of the preceding half-line construction,

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

for all $\psi, \phi \in D(K_0(\Omega)) \cap D(P_\Omega)$. In the effective spacetime description of that same renormalized half-line family, the local null-stress charge is defined by the identical endpoint derivative,

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

so

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

as a quadratic-form identity on the common domain. The only continuum input used here 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. Bounded-interval transport through the affine-covariant kernel $g_I(v)$ is a separate projective-branch step, and reconstruction of a full tensor from the directional charges is subject to the null-invisible metric ambiguity.

Boundary of the null bridge. The null bridge is an explicit theorem branch rather than an automatic fixed-cutoff identity. Its required data are:

- transferred cut-center data alone give the central sector decomposition of the null strip rather than the left/right HJPW split;
- the extra decomposition-inheritance condition of Proposition 5.2a is exactly what upgrades those sectors to the same collar-type block structure used in the spatial branch;
- the fixed-cutoff null-strip bridge is exact for exact Markov strip states on that inherited decomposition, and otherwise uses a controlled limit with the carried defect operator $\mathfrak{D}_J$;
- the local finite-constraint MaxEnt branch gives endpoint-Lipschitz control strong enough to define a weak tail generator for renormalized half-lines, and the scaling-limit geometric-cap branch of Theorem 4.2 then derives half-sided modular inclusion on the half-line blow-up net;
- Borchers–Wiesbrock therefore supplies a genuine positive null-translation generator P_Ω inside the bridge itself, together with its Stone domain and the half-line modular-Hamiltonian relation $K_a = K_0 - 2\pi a P_\Omega$;
- the half-line generator/charge identification is fixed inside the bridge itself, while bounded-interval transport and the later tensor upgrade are downstream.

This is the route from null-strip factorization to the later D5 Einstein branch, with the half-line generator/charge identification proved inside the null modular bridge itself.

4.6.3 Modular energy as stress-tensor charge on null half-lines

The effective-theory input here is the standard stress-tensor representation used to name the operator fixed by the null bridge. In the canonical normalization of the preceding null-half-line construction, the positive null generator is the endpoint derivative of the renormalized half-line modular family:

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

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

In the effective spacetime description of that same half-line family, the local null-stress charge is defined by the identical endpoint derivative,

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

so on the common quadratic-form domain

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

Thus the null generator is exactly the local null-stress charge in the effective spacetime description; this step is not a separate scaling-limit input. When the effective description admits a local stress tensor, for example in a UV CFT regime or any local Lorentzian regime with local modular Hamiltonians, the same operator is the standard null-stress charge associated with the chosen null generator.

For comparison, in a CFT vacuum on a ball one has the familiar local formula

$$H_\zeta = \int_\Sigma T_{ab}\zeta^b \, d\Sigma^a,$$

so the present identification is the null-half-line version of that same modular-energy locality rather than an additional EFT postulate.

The downstream boundary is narrower. The separate interval-preserving projective branch is needed to transport this half-line statement to bounded null intervals with affine-covariant kernel $g_I(v)$, and the null-to-tensor upgrade determines T_{ab} only up to the familiar ambiguity

$$T_{ab} \mapsto T_{ab} + \phi g_{ab}.$$

But the generator/charge identification itself is fixed inside the null bridge.

4.6.4 Localized generalized entropy from Markov + MaxEnt

Using the collar decomposition and the collar double-scaling hypothesis, supplied either as a stated collar-limit condition or derived via Theorem 2.5 from the local-Gibbs form together with exponential mixing, the state takes the Markov normal form. MaxEnt selection maximizes entropy within each edge sector, producing

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

The entropy splits as

$$S(\rho_C) = H(p_\alpha) + \sum_\alpha p_\alpha S(\rho^\alpha_{\mathrm{bulk},C}) + \sum_\alpha p_\alpha \log d_\alpha.$$

Convention: Throughout this paper, "log" denotes the natural logarithm (ln), so entropies are measured in **nats** (1 nat = 1/ln 2 ≈ 1.443 bits). This is standard in thermodynamics and QFT; the Bekenstein-Hawking formula S = A/4G uses nats. When clarity requires it, we write $\log_2$ explicitly for bits.

Define

$$S_{\mathrm{bulk}}(C) := H(p_\alpha) + \sum_\alpha p_\alpha S(\rho^\alpha_{\mathrm{bulk},C}),$$

and the central area operator

$$L_C := \sum_\alpha (\log d_\alpha) P_\alpha.$$

Then

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

Deriving Newton's constant from edge entropy density.

Rather than normalize L_C by fiat, we *derive* the relation to G from the UV edge structure. In the collar double-scaling limit, the edge contribution becomes extensive along the entangling surface $\Sigma = \partial C$:

$$\text{Tr}(\rho L_C) \approx N_\Sigma \cdot \bar{\ell}(t), \qquad \bar{\ell}(t) := \sum_\alpha p_\alpha \log d_\alpha,$$

where N_Σ is the number of UV cut elements covering Σ and $\bar{\ell}(t)$ is the **single-cell edge entropy** from the heat-kernel distribution (Theorem 6.20). Similarly, the geometric area is extensive:

$$A(C) \approx N_\Sigma \cdot a_{\text{cell}},$$

where a_{cell} is the area per UV cut element in the emergent metric.

Matching these expressions gives the **derived formula for Newton's constant**:

$$G = \frac{a_{\text{cell}}}{4\,\bar{\ell}(t)}$$

where:

- a_{cell} is fixed operationally from the UV correlation/mixing length ξ via $a_{\text{cell}} \sim \xi^2$ (from the exponential mixing hypothesis),
- $\bar{\ell}(t) = \sum_R p_R(t) \log d_R$ is computed from the heat-kernel edge distribution with $p_R \propto d_R e^{-t\lambda_R}$.

Explicitly:

$$\bar{\ell}(t) = \frac{\sum_R d_R e^{-t\lambda_R} \log d_R}{\sum_R d_R e^{-t\lambda_R}}.$$

Thus G is the inverse edge-entropy density per geometric area, computable from the UV regulator and the reference-state Gibbs parameter t, rather than a normalization convention.

4.6.5 Fixed-cap generalized-entropy stationarity from MaxEnt

MaxEnt selection implies that for variations preserving cap labels (fixed size and charges),

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

Using the split above and the first law for the bulk term,

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

4.6.6 Einstein equation from fixed-cap generalized-entropy stationarity

In the d=4 scaling regime, the small-ball bridge is internal to the derived gravity chain once the separate bounded-interval projective branch is included. The only extra standard geometric input is the fixed-volume area-variation identity for a small geodesic ball. The modular kernel itself is not imported from a separate EFT law: on the extracted prime geometric subnet, a sufficiently small cap is the tangent causal diamond D_ℓ, whose preserving conformal Killing field is

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

On the t=0 slice its lapse is

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

The bounded-interval kernel comes from combining the D3 geometric cap generator with the separate interval-preserving projective branch used downstream of the D4 half-line stress bridge, so the D3 cap-modular theorem and the D4 stress bridge together yield

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

If $\delta\langle T_{00}\rangle$ is approximately constant across the ball and the carried remainder is $o(\ell^4)$, then

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

Stationarity of generalized entropy therefore gives the observer-covariant scalar relation

$$\delta\Big[(G_{ab} + \Lambda g_{ab})u^a u^b\Big] = 8\pi G\, u^a u^b\,\delta\langle T_{ab}\rangle$$

for the local diamond rest-frame four-velocity u^a. In the adapted rest frame $u^a = (1,0,0,0)$, this is

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

4.6.7 Overlaps supply all timelike directions

This scalar-to-tensor upgrade is internal once the Lorentz branch is in place. Overlapping observers through the same bulk point realize every local timelike four-velocity u. Define

$$Y_{ab} := G_{ab} + \Lambda g_{ab} - 8\pi G\,\langle T_{ab}\rangle.$$

The rest-frame relation of the previous subsection says

$$u^a u^b\,\delta Y_{ab} = 0$$

for every such u. In local inertial coordinates,

$$u^a = \gamma(1, v^i), \qquad |\vec{v}| < 1,$$

so

$$0 = u^a u^b\,\delta Y_{ab} = \gamma^2\Big(\delta Y_{00} + 2v^i\delta Y_{0i} + v^i v^j\,\delta Y_{ij}\Big)$$

for all $|\vec{v}| < 1$. The polynomial therefore vanishes coefficientwise, hence

$$\delta Y_{00} = 0, \qquad \delta Y_{0i} = 0, \qquad \delta Y_{ij} = 0.$$

So the full tensor variation vanishes:

$$\delta Y_{ab} = 0.$$

Along any connected scaling branch of reference states, Y_{ab} is therefore constant. The only purely local freedom left by the null reconstruction is the metric term, and fixing that term on one maximally symmetric reference state gives

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

Thus the tensor upgrade does not require a separate EFT handoff once the geometric-cap theorem, the null bridge, and the overlap-complete timelike family are in place.

4.6.8 Non-tunable numerical constants

The gravity chain yields specific numerical constants as rigid outputs of the axiom chain.

The 2π KMS normalization. From Theorem 4.2, cap modular flow on the support-visible scaling-limit geometric cap pair is fixed by the KMS condition to the standard 2π normalization. This is the same rigidity that fixes Unruh/Hawking temperature normalization.

The geometric coefficient $\Omega_{d-2}/(d^2-1)$. This coefficient appears in both (a) the CFT-ball modular Hamiltonian weight integral and (b) the geometric area-variation identity. It is an exact integral identity:

$$\int_{B_\ell^{d-1}} \frac{\ell^2 - r^2}{2\ell}\, d^{d-1}x = \frac{\Omega_{d-2}\, \ell^d}{d^2-1}.$$

In $d = 4$:

$$\frac{\Omega_2}{4^2-1} = \frac{4\pi}{15} \approx 0.8377580409572781.$$

This is the reason prefactors cancel cleanly when going from $\delta S_{\mathrm{gen}} = 0$ to the Einstein equation (leaving $8\pi G$ with the 2π fixed by Theorem 4.2).

What is predicted. The framework cleanly separates:

- **Non-tunable constants**: 2π (KMS period), $\Omega_{d-2}/(d^2-1)$ (geometric coefficient), the existence of the Einstein form.
- **Global closure constants**: G is the conversion between edge entropy and geometric area on the emitted local gravity branch. Λ is not fixed by the local reference-state variation data alone; it closes through the cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$.

4.6.9 Quantitative Markov error and controlled corrections

The role of this subsection is to keep three distinct quantities separate:

1. the raw collar conditional mutual information

$$\varepsilon_\delta := I(A_\delta : D_\delta \mid B_\delta)_\omega;$$

2. the constructive recovery error

$$r_{\rm FR}(\varepsilon_\delta) := 2\sqrt{1-e^{-\varepsilon_\delta}} \le 2\sqrt{\varepsilon_\delta};$$

3. the exact-Markov replacement error on one fixed faithful collar model,

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

   where $\lambda_* > 0$ is the lower spectral bound needed to compare modular Hamiltonians.

The exact identity for the modular defect expectation is unchanged:

$$\langle \Delta K_\delta\rangle_\omega = -I(A_\delta : D_\delta \mid B_\delta)_\omega = -\varepsilon_\delta,$$

with

$$\Delta K_\delta := K_{A_\delta B_\delta D_\delta} - K_{A_\delta B_\delta} - K_{B_\delta D_\delta} + K_{B_\delta}.$$

What changes is the interpretation of later "exact" collar formulas. For each fixed finite-dimensional collar one may choose an exact Markov reference state σ_δ with

$$\|\omega_\delta - \sigma_\delta\|_1 \le \delta^{\rm M}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta),$$

and a constructive recovered comparison state $\omega_\delta^{\rm rec}$ with

$$\|\omega_\delta - \omega_\delta^{\rm rec}\|_1 \le r_{\rm FR}(\varepsilon_\delta).$$

These two comparisons do different jobs:

Bounded-observable control. For every bounded observable X supported on $A_\delta \cup B_\delta$,

$$|{\rm Tr}[X(\omega_\delta - \omega_\delta^{\rm rec})]| \le \|X\|_\infty\, r_{\rm FR}(\varepsilon_\delta),$$

and

$$|{\rm Tr}[X(\omega_\delta - \sigma_\delta)]| \le \|X\|_\infty\, \delta^{\rm M}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta).$$

The first bound is constructive and dimension-free; the second is the fixed-collar route to the exact Markov set.

Modular-additivity control. If the relevant marginals are uniformly faithful with lower spectral bound $\lambda_* > 0$, then

$$\|\Delta K_\delta(\omega) - \Delta K_\delta(\sigma_\delta)\|_\infty \le \eta_\delta^{\rm M}.$$

Since $\Delta K_\delta(\sigma_\delta)$ is central by exact Markovity, the finite-stage modular defect is within $\eta_\delta^{\rm M}$ of the exact splice or additivity value.

Therefore the manuscript's exact collar identities arise as limits with a carried remainder

$$\eta_\delta := r_{\rm FR}(\varepsilon_\delta) + \eta_\delta^{\rm M},$$

together with the separate long-wavelength derivative remainders present in the small-ball expansion. Exact identities at finite collar width require exact Markovity; otherwise one works with η_δ-controlled approximations and then lets $\eta_\delta \to 0$ in the controlled collar limit.

Finite-cutoff Einstein remainder bound. In the small-ball rest-frame calculation, write

$$\delta S_C = \frac{8\pi^2\ell^4}{15}\,\delta\langle T_{00}\rangle + \delta\langle E_C^{(\delta)}\rangle + O(\ell^5\partial T),$$

where $E_C^{(\delta)}$ packages the carried collar remainder. Then the derived fixed-cap generalized-entropy stationarity theorem gives

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

with

$$\mathcal{E}_\delta := \frac{15G}{\pi\ell^4}\,\delta\langle E_C^{(\delta)}\rangle + O(\ell\,\partial T),$$

and, by the bounds above,

$$|\delta\langle E_C^{(\delta)}\rangle| \le C_C\,\eta_\delta$$

for some collar-dependent constant C_C on the fixed faithful model. Equivalently, before the controlled collar and small-ball limits, one may write

$$|\mathcal{E}_{\ell,\delta}| \le C_1 r_{\mathrm{FR}}(\varepsilon_\delta) + C_2\delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta) + C_3\eta_\delta^{\mathrm{reg}} + C_4\ell\,\|\partial T\| + o_\delta(1) + o_\ell(1),$$

with constants depending only on the fixed collar model, the bounded observable class, and the small-ball chart. Thus the exact Einstein relation is the controlled limit, not a finite-cutoff identity inherited from approximate Markovity.

If one wishes to package this carried remainder as an effective anomalous energy density, the natural definition is

$$\delta\langle T_{00}^{\mathrm{anom}}\rangle := \frac{15}{8\pi^2\ell^4}\,\delta\langle E_C^{(\delta)}\rangle.$$

Then

$$\delta(G_{00} + \Lambda g_{00}) = 8\pi G\,\delta(\langle T_{00}\rangle + \langle T_{00}^{\mathrm{anom}}\rangle) + O(\ell\,\partial T).$$

This is a bookkeeping definition of the finite-stage remainder, not an upgrade of the recovered-core theorem. Its significance is exactly that it is controlled:

$$|\delta\langle T_{00}^{\mathrm{anom}}\rangle| \le \frac{15\,C_C}{8\pi^2\ell^4}\,\eta_\delta.$$

The framework carries the Markov error through three distinct controls:

1. $r_{\mathrm{FR}}(\varepsilon_\delta)$ controls bounded observables at one stage;
2. $\delta^{\mathrm{M}}_{A_\delta:B_\delta:D_\delta}(\varepsilon_\delta)$ controls convergence to the exact Markov normal form on a fixed collar;
3. η_δ^{M} controls modular-Hamiltonian replacements once faithfulness is assumed;
4. the Einstein-branch remainder is a carried $O(\eta_\delta)$ term, not a silent exact identity at finite cutoff.

This is the concrete bridge from "axioms about screens" to "precision GR predictions plus a disciplined carried remainder".

4.6.10 Focusing/QNEC internalization via relative entropy

Once the fixed-cutoff null modular structure of Section 5.2 is combined with the exact half-line generator/charge identification and the bounded-interval projective branch, focusing constraints follow from information-theoretic principles in the same scaling-limit branch rather than from a separate fixed-cutoff postulate.

Derivation chain. QNEC and focusing are supported within the same branch:

- Fixed-cutoff null modular bridge (§5.2) $\Longrightarrow$ exact-or-controlled strip additivity and a weak tail generator
- Derived half-sided modular pair + exact half-line generator/charge identification $\Longrightarrow$ local null-stress charge on half-lines and [K, P] = i 2π P
- Relative entropy monotonicity $\Longrightarrow$ QNEC
- Einstein (Thm 5.1) + Raychaudhuri $\Longrightarrow$ QFC for S_gen

Relative entropy monotonicity argument. The key input is the monotonicity of relative entropy under partial trace, which is pure information theory:

$$S(\rho_{AB}\|\sigma_{AB}) \geq S(\rho_A\|\sigma_A).$$

For null deformations parameterized by λ, consider nested null regions $R(\lambda) \subset R(\lambda')$ obtained by varying the entangling cut along v. The modular Hamiltonian K_λ generates the modular flow, and relative entropy satisfies convexity:

$$\frac{d^2}{d\lambda^2} S(\rho_\lambda\|\sigma_\lambda) \geq 0.$$

Proposition 5.10a (Internal QNEC). Under the fixed-cutoff null bridge of §5.2 together with the downstream density-upgrade and bounded-interval projective conditions, the second null variation of von Neumann entropy satisfies

$$\frac{d^2 S_{\text{bulk}}}{d\lambda^2} \leq 2\pi\langle T_{kk}(\lambda)\rangle,$$

with the 2π normalization fixed by Theorem 4.2.

Proof. Use the derived half-sided modular-inclusion statement of Corollary 5.2e and its Borchers-Wiesbrock consequence from Lemma 5.2f. This gives the translation structure with $[K, P] = -i2\pi P$.

Consider the relative entropy $S(\rho_\lambda\|\omega_\lambda)$ between the state ρ restricted to $R(\lambda)$ and the reference state ω. Monotonicity under restriction to smaller regions ($\lambda' > \lambda$) gives:

$$S(\rho_{R(\lambda)}\|\omega_{R(\lambda)}) \leq S(\rho_{R(\lambda')}\|\omega_{R(\lambda')}).$$

Using the first law $\delta S = \delta\langle K\rangle$ and the geometric half-line modular form supplied by Section 5.2, expand to second order in the deformation. The convexity of relative entropy yields the QNEC inequality. The bound saturates for coherent states in the standard way. QED.

Corollary 5.10b (QFC for generalized entropy). With the central area operator L_C from EC/MaxEnt (Section 5.4), define

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

Given the Einstein branch (Theorem 5.1) and the classical Raychaudhuri identity for null congruences, the Quantum Focusing Conjecture (QFC) follows within the same scaling regime: the generalized expansion Θ_{gen} is non-increasing along null generators.

Proof sketch. The Raychaudhuri equation relates expansion evolution to R_{kk}. Einstein's equation gives $R_{kk} = 8\pi G(T_{kk} - \frac{1}{2}g_{kk}T)$. For null k, this simplifies to $R_{kk} = 8\pi G T_{kk}$. The QNEC (Prop 5.10a) then bounds the bulk entropy production, ensuring $d\Theta_{\text{gen}}/d\lambda \le 0$. QED.

Significance. This section shows how the Recoverable Generalized Entropy focusing package is internally supported inside the same null-modular branch once the stress-tensor identification and Einstein relation are in place. That package is part of the stated axiom set.

Theorem 5.1 (Observer-consistency implies a scaling-limit Einstein branch). Under the five OPH axioms, the collar and null-strip hypotheses of Sections 2 and 5, the hypotheses of Theorem 4.2, the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations on the realized cap-label-preserving MaxEnt family, the half-line generator/charge identification of Section 5.2 together with the bounded-interval projective branch used there, and admissible fixed-cap MaxEnt variations about a maximally symmetric reference state, the fixed-cap generalized-entropy stationarity condition implies the rest-frame first-variation relation

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

If this relation holds for all local directions and reference states in the locally Lorentzian scaling regime, overlap consistency upgrades it to the semiclassical Einstein equation modulo the expected Λg_{ab} ambiguity.

Proof sketch. The cap first law identifies δS_C with $\delta\langle K_C\rangle$, while the null bridge together with the bounded-interval projective branch relates the bulk modular variation to the required null stress-tensor charge on the diamond. Since the reference state is maximally symmetric, the fixed-volume small-ball area identity enters at first order in the perturbation, and the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap MaxEnt variations on the realized cap-label-preserving MaxEnt family yields the displayed rest-frame relation. Overlap consistency then supplies all local directions, and the null-to-tensor lemma upgrades the relation to the tensor equation modulo Λg_{ab}. QED.

4.6.11 Black-hole structural statements and continuation-level spectroscopy

This subsection mixes two claim tiers and is organized accordingly. The retained structural black-hole package is a four-step chain: fixed-cutoff edge-center collar decomposition on the exact-Markov or idealized exact-recoverability horizon carrier, the small-CMI recoverability reading of interior encoding on that same carrier, Hawking/KMS normalization on the geometric modular branch, and the discrete area spectrum together with the Schwarzschild transition identity $\Delta E = k_B T_H \ln(d'/d)$ for an actual sector change $d \to d'$. Any clean comb structure, QNM selector, Page-type linewidth estimate, PBH burst template, Kerr/LIGO horizon spectroscopy template, or Page-curve/island closure is continuation-level and uses continuation inputs stated below.

Derived structural statements.

Area eigenvalues from edge sectors. The central area operator (Section 5.4) is

$$L_C = \sum_{\alpha} (\log d_\alpha) P_\alpha,$$

where $d_\alpha \in \mathbb{N}$ is the dimension of the edge Hilbert space in sector α. With the normalization $\mathrm{Tr}(\rho L_C) = \langle A\rangle/4G$, the area eigenvalues are

$$A_\alpha = 4G \log d_\alpha = 4\ell_p^2 \ln d_\alpha,$$

where $\ell_p^2 = \hbar G/c^3$ is the Planck area. Since d_α is a positive integer, **areas are discretely spaced** with logarithmic gaps.

Hawking emission energy quantization. For a Schwarzschild black hole with $A(M) = 16\pi G^2 M^2/c^4$, a transition between sectors $d \to d'$ changes the area by

$$\Delta A = 4\ell_p^2 \ln(d'/d).$$

The corresponding ADM energy change of the hole is $\Delta(Mc^2) = c^2 \Delta M$, with $\Delta M = \Delta A/(dA/dM)$. This gives

$$\Delta(Mc^2) = \frac{\hbar c^3}{8\pi GM} \ln(d'/d).$$

Using the Hawking temperature $T_H = \hbar c^3/(8\pi G k_B M)$, whose 2π normalization is fixed by the BW_{S^2} tangent-limit normalization of Theorem 4.2:

$$\Delta(Mc^2) = k_B T_H \ln(d'/d).$$

Continuation-level spectroscopy templates. The remainder of this subsection is not part of the recovered core. It records what follows only if one adds extra discrete-horizon selection rules and, where stated, standard semiclassical inputs such as evaporation-power models, QNM/transition identifications, or greybody matching.

Integer transitions. If dominant emission steps reduce the edge dimension by an integer factor k (i.e., $d' = d/k$), the emitted spectrum becomes a discrete comb:

$$\Delta E_k = k_B T_H \ln k, \qquad \Delta f_k = \frac{c^3}{16\pi^2 GM} \ln k.$$

Structural condition: comb vs. generic discreteness. The log-integer *comb* structure requires the additional dynamical assumption that integer-ratio emission steps ($d \to d/k$) dominate. If generic transitions between arbitrary integers dominate instead, the set of $|\ln(d'/d)|$ values becomes a dense log-rational set that may appear quasi-continuous after folding in linewidths and astrophysical effects. What follows directly from the axioms is the *discrete area spectrum*; the clean comb pattern follows only with that selection rule.

Continuation-level template (Discrete Hawking spectrum). If the additional integer-transition selection rule is realized, the Hawking emission spectrum consists of discrete lines with spacing $\Delta E_k = k_B T_H \ln k$, where k is an integer characterizing the dominant sector transitions, instead of a continuous thermal profile.

Mass-independent fractional linewidth (continuation-level estimate). Using Page's semiclassical calculation for emission power $P(M) = p_0 \hbar c^6/(G^2 M^2)$ with $p_0 \approx 2 \times 10^{-4}$, the emission rate is $\dot{N} \approx P/\langle E \rangle$ where $\langle E \rangle = a\, k_B T_H$ with $a \sim \mathcal{O}(1-10)$. The natural linewidth $\Gamma \sim \hbar \dot{N}$ divided by the level spacing gives:

$$\frac{\Gamma}{\Delta E_k} \approx \frac{64\pi^2 p_0}{a \ln k} \approx 3 - 5\%$$

Within this continuation branch, the emission lines are narrow (few-percent fractional width) and the estimated fraction is mass-independent.

Connection to quasinormal modes (interpretive continuation). The highly-damped Schwarzschild quasinormal modes have asymptotic real part (Motl, 2002):

$$\mathrm{Re}\,\omega \to \frac{c^3}{8\pi GM}\ln 3.$$

This matches exactly the $k = 3$ transition frequency $\Delta E_3/\hbar$. If one adopts a Bohr-type identification between quantum transition frequencies and asymptotic QNM frequencies, this selects

$$\Delta A = 4\ell_p^2 \ln 3 \approx 4.39\,\ell_p^2$$

as the fundamental area quantum.

Scope statement. The area quantization follows from the edge-sector structure (derived). The $k = 3$ selection requires the additional interpretive identification with QNM frequencies (not derived from axioms). The linewidth prediction uses standard semiclassical inputs.

Numerical examples. For Δf_k = ($c^3/16\pi^2$GM) ln k:

- **M = 30 M⊙**: k=2 at 29.7 Hz, k=3 at 47.1 Hz
- **M = 1 M⊙**: k=2 at 891 Hz, k=3 at 1412 Hz
- **M = 10^{12} kg (primordial)**: k=2 at 7.3 MeV, k=3 at 11.6 MeV

Within the integer-transition continuation, these frequencies track $k_B T_H \ln k$ exactly and are in principle distinguishable from a continuous thermal spectrum.

Continuation-level PBH burst search template.

Within the discrete-horizon continuation, the Hawking comb would provide a distinctive gamma-ray template. A template-level discriminant is **log-integer energy ratios**: if two emission lines are observed at energies E_2 and E_3, their ratio must satisfy

$$\frac{E_3}{E_2} = \frac{\ln 3}{\ln 2} \approx 1.585$$

exactly, independent of black hole mass. Within that continuation template, the ratio is fixed once the log-integer rule is assumed.

Available instruments and energy coverage. The $k = 2$ line energy $E_2 = k_B T_H \ln 2$ determines which instruments can see a given BH mass:

Instrument	Energy band	BH mass range (k=2 in band)
Fermi GBM (BGO)	0.15–40 MeV	2×10^{11}–5×10^{13} kg
Fermi LAT	0.1–300 GeV	2×10^{7}–7×10^{10} kg
H.E.S.S.	0.1–100 TeV	7×10^{4}–7×10^{7} kg
LHAASO-WCDA	1–15 TeV	5×10^{5}–7×10^{6} kg

Detector resolution vs. intrinsic linewidth. The continuation-level linewidth estimate is 3–5% (mass-independent). Current detector energy resolutions:

- Fermi GBM: $< 10\%$ (0.1–1 MeV), $\sim 4\%$ at 10 MeV (BGO)
- Fermi LAT: $< 10\%$ (1–100 GeV)
- H.E.S.S.: $\sim 15\%$ (TeV)
- LHAASO-WCDA: $\sim 33\%$ (TeV)

Within this template, the comb could be resolvable with GBM/LAT; at TeV energies it would appear as moderately broad bumps rather than sharp lines.

Search protocol. A dedicated OPH-comb search would:

1. Select burst-like candidates (10–120 s time windows, matching existing PBH burst search protocols).
2. Fit each candidate with null model (smooth continuum) vs. OPH comb model (peaks at $E_k = E_0 \ln k$ convolved with detector response).
3. Scan over the single scale parameter $E_0 = k_B T_H$ (equivalently, BH mass).
4. Require at least two lines satisfying log-integer ratio to claim detection.
5. Correct significance for trials (time windows $\times$ sky positions $\times$ E_0 scan).

Observational context. Dedicated PBH burst searches (H.E.S.S., LHAASO) report **no significant bursts**. An OPH-specific comb-template analysis of archival data would:

- Set upper limits on OPH-comb PBH burst rates
- Demonstrate direct testability of the discrete spectrum prediction
- Provide constraints comparable to or stronger than generic PBH burst limits

Data availability. Fermi GBM provides public Time-Tagged Event (TTE) burst data; Fermi LAT provides public photon event lists with documented analysis workflows. H.E.S.S. has a small public test data release.

Continuation-level GW horizon spectroscopy template for Kerr remnants.

The same continuation-level discrete-horizon branch extends to gravitational wave observables. For Kerr black holes, the thermodynamic first law is $\delta M = T_H \delta S + \Omega_H \delta J$, so the entropy change for absorbing a quantum with frequency ω and azimuthal number m is:

$$\delta S = \frac{\hbar(\omega - m\Omega_H)}{k_B T_H}.$$

In the edge-sector framework, $\delta S = \ln(d'/d)$, so the discreteness condition becomes:

$$\hbar(\omega - m\Omega_H) = k_B T_H \ln k, \qquad k \in \{2, 3, 4, \ldots\}$$

Under those additional discrete-horizon assumptions, this gives the **GW horizon spectroscopy comb**: a continuation-level set of discrete resonant frequencies where the horizon can efficiently absorb or emit energy.

Kerr line frequencies. For a remnant with mass M and dimensionless spin $\chi = a_*/M$, define the spin correction factor:

$$g(\chi) = \frac{2\sqrt{1-\chi^2}}{1+\sqrt{1-\chi^2}}, \qquad \Omega_H(M,\chi) = \frac{c^3}{2GM} \cdot \frac{\chi}{1+\sqrt{1-\chi^2}}.$$

The line frequencies are:

$$f_{k,m}(M,\chi) = \frac{m\,\Omega_H(M,\chi)}{2\pi} + \frac{c^3}{16\pi^2 GM}\, g(\chi)\, \ln k$$

Within this continuation template, once LIGO/Virgo infers (M, χ) for a remnant, the line pattern is fixed by the inferred remnant parameters and the assumed discrete-horizon rule. This is not a recovered-core theorem.

Line weights from GR envelope + discretization. In this continuation template, the line strengths are modeled by matching to the known GR greybody absorption spectrum in the semiclassical limit. The discretization rule gives bin width $\Delta\omega_k \approx \omega_T \ln(1 + 1/k)$ where $\omega_T = k_B T_H/\hbar$. The net line weight (absorption minus stimulated emission) is:

$$W^{\text{net}}_{k,\ell m} = \Gamma^{\text{GR}}_{\ell m}(\omega_{k,m}) \cdot \Delta\omega_k \cdot \frac{k-1}{k}$$

where $\Gamma^{\text{GR}}_{\ell m}$ is the standard GR greybody factor and the $(k-1)/k$ factor arises from KMS detailed balance with $e^{(\omega - m\Omega_H)/T_H} = k$.

Universal stacking coordinate. Define the dimensionless rescaled frequency:

$$x := \frac{GM}{c^3 g(\chi)}(\omega - m\Omega_H).$$

Then the predicted line locations collapse to universal constants:

$$x_k = \frac{\ln k}{8\pi} \qquad (k = 2, 3, 4, \ldots)$$

Numerically: $x_2 = 0.02758$, $x_3 = 0.04371$, $x_4 = 0.05516$, $x_5 = 0.06404$.

Stacking test. Multiple BBH events can be mapped to this universal x coordinate and stacked. If the comb is real, peaks align across events with different (M, χ); detector noise does not stack coherently.

Comparison to existing work. Prior area-quantization searches [38] used parameterized models with one free spacing constant. The OPH prediction is more constrained: multiple lines with exact $\ln k$ ratios, plus the $(k-1)/k$ weight hierarchy from detailed balance.

Numerical example (GW170608). Remnant parameters: $M_f \approx 18.0 M_\odot$, $\chi_f \approx 0.69$. For $m = 2$, the horizon rotation frequency is $m\Omega_H/(2\pi) \approx 719$ Hz. The **thermal comb spacing** (the part that encodes the area quantization) is:

k	$\Delta f_k := \frac{c^3 g(\chi)}{16\pi^2 GM} \ln k$ (Hz)	Relative weight $(k-1)/k$
2	41.6	0.500
3	65.9	0.667
4	83.2	0.750
5	96.5	0.800
6	107.5	0.833

The full physical frequencies are $f_{k,2} = 719 + \Delta f_k$ Hz (i.e., 760–827 Hz), outside LIGO's most sensitive band for this remnant. However, the **stacking analysis** uses the rescaled coordinate $x = GM(\omega - m\Omega_H)/(c^3 g(\chi))$, which maps the thermal spacing to universal constants $x_k = \ln k/8\pi$ regardless of the rotation offset.

Template-matching criterion. After rescaling by (M, χ), spectral features in this continuation template must satisfy $f_k/f_2 = \ln k/\ln 2$ exactly, independent of remnant parameters. Absence of coherent stacking at the predicted x_k values would challenge the discrete-horizon continuation template, not by itself the derived area-spectrum statement.

4.6.12 Classical mechanics from emergent GR

Once the Einstein equation is established, the framework inherits standard GR consequences. This section makes explicit how classical mechanics emerges.

Stress-energy conservation is automatic. The contracted Bianchi identity is geometric:

$$\nabla^a G_{ab} = 0.$$

Combined with the Einstein equation, this implies:

$$\nabla^a \langle T_{ab} \rangle = 0.$$

Geodesic motion from dust limit. For pressureless classical matter ("dust"), $T^{ab} = \rho\, u^a u^b$. Conservation yields:

$$\nabla_a(\rho u^a u^b) = 0 \quad \Rightarrow \quad u^b \nabla_a(\rho u^a) + \rho\, u^a \nabla_a u^b = 0.$$

Projecting orthogonally to u^b using $h^b{}_c = \delta^b{}_c + u^b u_c$ kills the first term, giving:

$$\rho\, u^a \nabla_a u^b = 0 \quad \Rightarrow \quad u^a \nabla_a u^b = 0.$$

This is the geodesic equation: free classical bodies follow spacetime geodesics.

Newtonian limit from weak-field GR. Take the weak-field, slow-motion limit with metric:

$$g_{00} \approx -(1 + 2\Phi/c^2), \qquad g_{0i} \approx 0, \qquad g_{ij} \approx \delta_{ij}(1 - 2\Phi/c^2),$$

and velocities $|\mathbf{v}| \ll c$. Then $G_{00} \approx 2\nabla^2\Phi/c^2$ (leading order), and $T_{00} \approx \rho c^2$. The Einstein equation reduces to:

$$\nabla^2 \Phi = 4\pi G \rho.$$

Geodesic motion reduces to:

$$\ddot{\mathbf{x}} = -\nabla\Phi.$$

These are Newton's gravitational law and Newton's second law. Classical mechanics is recovered as a controlled limit of the emergent GR dynamics.

Precision classical predictions. Once the field equation is fixed to Einstein form, the framework inherits the standard GR precision toolbox (post-Newtonian expansion, lensing, time delay, etc.), with no free "shape" parameters beyond G and Λ.

Selected precision predictions (in the regime where the GR derivation applies):

Light bending by mass M: For impact parameter b,

$$\Delta\theta = \frac{4GM}{c^2 b}.$$

For the Sun with $b \approx R_\odot$: $\Delta\theta \approx 1.751$ arcsec.

Mercury perihelion advance: Per orbit,

$$\Delta\varpi = \frac{6\pi GM}{a(1 - e^2)c^2}.$$

Using Mercury's orbital parameters: $\Delta\varpi \approx 42.98$ arcsec/century.

Gravitational redshift: Between two radii in a static potential,

$$\frac{\Delta\nu}{\nu} \approx \frac{\Delta\Phi}{c^2}.$$

For the Sun (surface to infinity): $z \approx 2.12 \times 10^{-6}$.

These predictions are fixed functions of G and known source parameters, and are confirmed observationally to high precision. The framework contains them automatically on the derived Einstein branch.

4.6.13 Precision gravity predictions and experimental bounds

The gravity sector makes symmetry-protected exact-zero predictions that can be confronted with the tightest available experimental bounds. This section translates the theoretical predictions into the specific observables that experiments actually constrain.

Speed of gravitational waves. The derived GR regime implies massless gravitons propagating on the same null cones as photons:

$$\frac{c_{\rm GW} - c}{c} = 0 \text{ exactly.}$$

Current bound (GW170817 + GRB 170817A multi-messenger):

$$-3 \times 10^{-15} < \frac{c_{\rm GW} - c}{c} < +7 \times 10^{-16} \quad \text{(90\% credibility)}.$$

For a source at $\sim$ 40 Mpc, this fractional difference corresponds to only a few seconds of propagation-time mismatch across $\sim 10^8$ years of travel.

Graviton mass. The gauge redundancy (diffeomorphism invariance) forbids a hard mass term:

$$m_g = 0 \text{ exactly.}$$

Current bound (GW dispersion analysis, PDG 2025):

$$m_g \leq 1.76 \times 10^{-23}\ \text{eV}/c^2 \quad \text{(90\% credibility)}.$$

This corresponds to a reduced Compton wavelength $\bar{\lambda}_C \gtrsim 1.6 \times 10^{16}$ m, i.e., order ~ 1.6 light-years.

No dipole radiation. Many modified gravity theories predict extra channels (scalar/vector) producing dipolar radiation at (-1)PN order. The derived GR limit predicts no such channel.

Current bound (GW170817 inspiral phasing, PDG 2025):

$$-4 \times 10^{-6} < \delta\hat{p}_{-2} < 2 \times 10^{-5} \quad \text{(90\% credibility)}.$$

Only tensor polarizations. The GR outcome means only the two tensor (helicity-2) modes propagate. Pure non-tensor hypotheses are disfavored by observational constraints, and mixed tensor-scalar/vector models are tightly constrained.

Equivalence principle tests. Additional null checks from the derived GR structure:

- Universality of free fall (space tests): precision $\sim 10^{-15}$
- Nordtvedt parameter ($\eta = 4\beta - \gamma$): $(0.47 \pm 0.55) \times 10^{-4}$
- Binary pulsar radiative damping (PSR J0737-3039): 0.999963 ± 0.000063

4.6.14 Theory-side error propagation from Markov bounds

The framework provides exact-zero predictions and quantitative control over how well those predictions hold. The Markov/recovery machinery can be propagated through the entire GR emergence chain.

The key quantitative hook. From Theorem 3.1, if the target state satisfies

$$I(A_k : C_k \mid B_k) \leq \varepsilon_k,$$

then recovery maps exist with trace-distance error

$$\delta_k = 2\sqrt{\ln 2 \cdot \varepsilon_k}.$$

Trace distance gives immediate bounds on observable errors. Using the standard dual norm inequality:

$$|\langle O\rangle_\rho - \langle O\rangle_\sigma| \leq \|O\|_\infty \|\rho - \sigma\|_1 = 2\|O\|_\infty D(\rho, \sigma),$$

where $D(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$ is the trace distance.

Exponential decay from the mixing hypothesis. The mixing assumption (Section 2.3) provides:

$$I_\omega(A_\delta : D_\delta \mid B_\delta) \leq c \cdot |\partial C|_{\mathrm{UV}} \cdot e^{-\delta/\xi}.$$

Combining these gives an explicit precision dial:

$$\delta_{\mathrm{step}} \lesssim 2\sqrt{\ln 2 \cdot c \cdot |\partial C|_{\mathrm{UV}}} \cdot e^{-\delta/(2\xi)}.$$

What precision requires. To match the GW speed bound ($\sim 10^{-15}$ fractional accuracy), the recovery-map error must satisfy:

$$\delta \lesssim 10^{-15} \quad \Rightarrow \quad \varepsilon \lesssim \frac{(\delta/2)^2}{\ln 2} \approx 3.6 \times 10^{-31}.$$

This is extremely small, but achievable: with a macroscopic boundary ($|\partial C|_{\mathrm{UV}} \sim 10^{35}$ for a meter-scale boundary at Planck UV scale), the exponential decay $e^{-\delta/\xi}$ with $\delta/\xi \sim$ a few hundred easily pushes below 10^{-31} once the prefactor is included.

Precision summary. The framework provides:

1. Exact-zero predictions ($m_g = 0$, $c_{\mathrm{GW}} = c$) from symmetry protection.
2. Translation of those zeros into the specific observables experiments constrain.
3. Explicit bounds on how far derived geometric statements can drift, using the conditional mutual information $\to$ trace distance $\to$ observable error chain.

This is the concrete path from "axioms about screens" to "precision GR predictions with quantitative error control."

4.6.15 Dark-sector response from the modular anomaly

This subsection records a D12 continuation benchmark obtained by adding a specific deep-IR response ansatz on top of the modular-anomaly term. It is not a theorem-level closure of dark-matter phenomenology: neither the observational dark-matter identification nor MOND/RAR-like dynamics are derived here.

The modular anomaly term T_{ab}^{anom} derived in Section 5.9 supplies one structural ingredient for a possible dark-sector continuation, without introducing new particle species.

The identification. The anomalous stress-energy contribution

$$\langle T_{00}^{\text{anom}} \rangle = \frac{15}{8\pi^2} \cdot \frac{\delta \langle K_C^{(\text{anom})} \rangle}{\ell^4}$$

is "dark" by construction: it arises from information-theoretic/gravitational structure (modular Markov imperfections), rather than Standard Model fields. Its coupling is gravitational, with no electromagnetic coupling. These properties motivate the dark-sector interpretation; observational identification requires separate phenomenology.

Connection to the cosmological parameter package. The framework makes Λ a global capacity parameter and fixes the same D6 static-patch radius:

$$\Lambda = \frac{3\pi}{G N_{\text{scr}}}, \quad r_{\text{dS}} = \sqrt{\frac{3}{\Lambda}}.$$

This imports the D6 static-patch scale into the continuation surface. Any galaxy-scale continuation built from the anomaly term would therefore be an IR phenomenon, appearing only when accelerations are small and distances are large. The imported scale does not by itself determine the source/response law.

Acceleration benchmark. If one assumes the relevant deep-IR response is controlled only by r_{dS}, c, and the fixed anomaly prefactor, then the benchmark scale must:

1. Vanish if $r_{\text{dS}} \to \infty$ (infinite capacity, no de Sitter static-patch scale)
2. Be controlled by r_{dS} as the only new IR scale
3. Carry non-tunable coefficients from the derivation

The anomaly enters with prefactor $\frac{15}{8\pi^2}$. The corresponding benchmark acceleration scale constructible from (Λ, c) is:

$$\boxed{a_0^{(\text{OPH})} := \frac{15}{8\pi^2} \cdot c^2 \sqrt{\frac{\Lambda}{3}} = \frac{15}{8\pi^2} \cdot \frac{c^2}{r_{\text{dS}}}}$$

Normalization estimate. Using Planck 2018 ΛCDM parameters ($H_0 \approx 67.4$ km/s/Mpc, $\Omega_\Lambda \approx 0.685$):

- $\Lambda \approx 1.09 \times 10^{-52}\ \text{m}^{-2}$
- $r_{\text{dS}} \approx 1.66 \times 10^{26}$ m
- Therefore:

$$\boxed{a_0^{(\text{OPH})} \approx 1.03 \times 10^{-10}\ \text{m/s}^2}$$

For comparison, observational fits to galaxy regularities (RAR/MDAR/MOND phenomenology) quote $a_0 \sim 1.2 \times 10^{-10}$ m/s^2. This numerical proximity is only a benchmark coincidence unless a separate response law and a controlled galaxy-scale limit are derived.

One illustrative response ansatz. If one further assumes that T_{00}^{anom} is the dominant deep-IR source and that the response organizes into a MOND/RAR-like law, the Newtonian limit could be written as:

$$\nabla^2\Phi = 4\pi G(\rho_b + \rho_{\text{anom}}),$$

i.e., baryons plus an effective extra density. Under the same extra ansatz the radial acceleration relation (RAR) could be written as:

$$g_{\text{obs}} \approx g_b + \sqrt{a_0 \cdot g_b}, \quad g_{\text{DM}} := g_{\text{obs}} - g_b \approx \sqrt{a_0 \cdot g_b}.$$

With $a_0 = a_0^{(\text{OPH})}$ fixed, that same ansatz would imply:

(i) Baryonic Tully-Fisher relation.

$$V^4 \approx G \cdot M_b \cdot a_0^{(\text{OPH})}$$

where V is the asymptotic rotation velocity and M_b is baryonic mass.

(ii) Flat rotation curves. For a point mass M_b:

$$g_{\text{DM}}(r) = \frac{\sqrt{GM_b\, a_0^{(\text{OPH})}}}{r} \quad \Rightarrow \quad M_{\text{DM}}(r) \propto r$$

i.e., inferred dark mass grows linearly with radius, producing flat rotation curves.

(iii) Characteristic surface density.

$$\Sigma_0^{(\text{OPH})} = \frac{a_0^{(\text{OPH})}}{2\pi G} \approx 0.25 \text{ kg/m}^2 \approx 120\, M_\odot/\text{pc}^2.$$

This lies in the range of observed central halo surface densities, but it is not a derived halo theorem.

Scope. What is grounded in the framework developed here:

- The modular anomaly term exists with fixed coefficient $\frac{15}{8\pi^2}$
- Λ and r_{dS} are determined by screen capacity
- The benchmark scale $a_0^{(\text{OPH})}$ follows only after importing the D6 static-patch scale and assuming that no additional IR scale enters the continuation

Scope boundary for theorem-level closure:

- A controlled nonrelativistic limit from the anomaly term to galaxy observables
- A derived sign and closure analysis for the effective anomaly contribution on the relevant galaxy-scale states
- A derived response law selecting the MOND/RAR functional rather than alternative IR behavior
- A proof that the anomaly term supplies the relevant dominant source on galaxy scales rather than only a benchmark normalization
- Lensing, cluster, and Bullet-Cluster phenomenology
- Cosmological abundance and structure-formation checks

- Environment-dependence and stability control

Cosmology/Boltzmann contract. The static galaxy ansatz above is not an FLRW perturbation kernel. A cosmological dark/anomaly claim must instead expose the variables needed by an Einstein–Boltzmann implementation:

$$\bar{\rho}_A(a), \quad \bar{\rho}_{A,\mathrm{eq}}(a), \quad w_A(a), \quad c^2_{s,A}(k,a), \quad \sigma_A(k,a), \quad Q^\mu_A, \quad B_A(k,a), \quad \Gamma_{\mathrm{rec}}(k,a).$$

The cold transported limit is the check case: when exchange, pressure, sound-speed, and anisotropic stress corrections are turned off, the anomaly slot must reduce to a CDM-like component before recombination. Any nonzero-field response, late-time growth suppression, or S_8-relief branch has to be emitted by the finite-collar parent evaluator through $B_A(k,a)$ and $\Gamma_{\mathrm{rec}}(k,a)$, not fitted as a free environmental kernel to CMB, weak-lensing, SPARC, or cluster data. Current compressed H_0, Ω_m, σ_8, and S_8 rows are therefore plumbing diagnostics for a low-H_0, Planck-like branch, not theorem-grade cosmology.

This branch is a phenomenological continuation above D6 and D12 inputs. It contains structural ingredients and an IR benchmark, but it does not derive the precise galaxy-scale response from the recovered core.

Falsifiability. A completed continuation must supply the derivation steps above and remain compatible with galaxy, lensing, cluster, Bullet-Cluster, and cosmological data. Failure of that continuation retracts the modular-anomaly dark-sector continuation. It does not retract the recovered core theorem package.

4.6.16 De Sitter holography: static patch vs boundary-at-infinity

A natural question arises: how does this framework relate to the "unsolved problem" of de Sitter holography?

What the usual dS holography problem is. When people say "dS holography is unsolved," they typically mean that we do not have anything as sharp as AdS/CFT, where the bulk has a timelike asymptotic boundary supporting a well-defined dual CFT with a precise dictionary. For de Sitter, there is no asymptotic timelike boundary in the static patch where one can simply place the dual theory. The classic dS/CFT proposal at future infinity has familiar difficulties, including non-unitarity worries and complex conformal weights.

Static-patch/horizon-screen setup. The framework begins with an observer's static patch and its horizon screen S^2, building a net of subregion algebras on that screen. At finite cutoff those algebras are type-I regulators. The Lorentz statement, when invoked, is a support-visible scaling-limit statement about the refinement-limit observer net, and that limit may leave the regulator class. By Theorem 4.5.2, the realized scaling-limit cap modular action on the extracted prime geometric cap pair is geometric and, in the non-type-I case of interest, generally outer.

This is therefore a fundamental fork away from AdS/CFT-style holography:

AdS/CFT	This framework
Codimension-1 boundary at infinity	Codimension-2 horizon screen (S^2)
Single global boundary theory	Observer-dependent patches that overlap
Dual CFT required	Only algebras + consistency conditions
Negative Λ	Positive Λ natural

This aligns with the static-patch/complementarity intuition in the dS literature, where the fundamental description is patch-based and different static patches are related by consistency rules, not by a single global boundary theory.

The mechanism: Λ as global capacity, not local physics. A key structural result is that null modular data reconstruct the stress tensor only up to an additive metric term. This is the statement that vacuum-energy or cosmological-constant shifts are invisible to the local null-data route. The Einstein equation derived from the fixed-cap generalized-entropy stationarity theorem is therefore fixed only up to Λg_{ab}.

Theorem boundary: cosmological-constant / screen-capacity closure stack. On the D5→D6 branch, assume the local Einstein equation is fixed only modulo Λg_{ab}, 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:

1. local null data determine the Einstein branch only modulo Λg_{ab};
2. the same branch closes globally as
$$G_{ab} + \frac{3\pi}{G N_{\mathrm{CRC}}}\, g_{ab} = 8\pi G \, \langle T_{ab} \rangle;$$
3. the same D6 closure fixes
$$S_{\mathrm{dS}} = N_{\mathrm{CRC}}, \qquad A_{\mathrm{dS}} = 4G N_{\mathrm{CRC}}, \qquad r_{\mathrm{dS}} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\mathrm{dS}}}{c};$$
4. the observed cosmic age is a downstream FLRW benchmark rather than an additional theorem output.

Thus the cosmological-constant package is one local/global theorem stack rather than a split local-plus-global story.

Reason. The local null-data route is blind to metric-term shifts, so the Einstein branch is fixed locally only modulo Λg_{ab}. The cosmic record-capacity fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$, together with its observed-branch de Sitter entropy readout, then yields

$$\Lambda_{\mathrm{CRC}} = \frac{3\pi}{G N_{\mathrm{CRC}}},$$

which closes the same Einstein branch globally and fixes the displayed static-patch package. The observed cosmic age is a later FLRW comparison quantity rather than part of this theorem stack.

The D6 hypotheses are the cosmic record-capacity fixed point, its observed-branch de Sitter entropy readout, the de Sitter entropy relation, and the standard static-patch formulas. The local null-data route does not by itself determine the global capacity; that global value is fixed by the cosmic record-closure readback fixed point. This logic is compatible with the BW scaling branch but does not depend on the scaling-limit algebra being type I.

Input-free capacity closure. For capacity N, let $F(N)$ be the active cosmic record capacity read back by stable observers inside the OPH universe supplied with capacity N:

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

The cosmic record-closure capacity is the fixed point

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}), \qquad \Lambda_{\mathrm{CRC}} = \frac{3\pi}{G N_{\mathrm{CRC}}}.$$

If F is an OPH-derived contraction on the admissible capacity interval, then this fixed point is unique and stable. The count-density representation of the same target is obtained by letting Ω_N^{sc} be the terminal OPH normal forms that are repair-closed, observer/checkpoint-supporting, locally recovered-core closed, and whose own horizon record surface reads back capacity N. Since $\log\dim\mathcal{H}_{\partial,N} = N$, define the screen-normalized self-closure density

$$\Pi(N) = \frac{|\Omega_N^{\mathrm{sc}}|}{\dim\mathcal{H}_{\partial,N}} = |\Omega_N^{\mathrm{sc}}|e^{-N}.$$

The corresponding selector 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 map

$$T_\eta(N) = N + \eta\ell'(N)$$

has a unique stable fixed point under the derivative-sign certificate for $H_N = \ell'$. Informally, this is the single screen size where the universe reads back its own boundary without deficit or slack. On the observed branch this fixed point is the de Sitter entropy capacity.

Capacity readout. The branch uses N_{scr} as the entropy capacity. The bare radius-squared ratio is $N_{\mathrm{patch}} = (r_{\mathrm{dS}}/\ell_P)^2$, and

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

The observed late-time scale gives $N_{\mathrm{patch}} \simeq 1.05\times10^{122}$, $N_{\mathrm{scr}} \simeq 3.31\times10^{122}$, and $\Lambda\ell_P^2 \simeq 2.85\times10^{-122}$.

What this solves vs. what it assumes. The model does **not** solve the classic "give me a unitary CFT at future infinity" problem. It does not aim there. It also does **not** prove that every refinement-stable MaxEnt branch lands in the static-patch geometric modular branch with emitted cap pair and standard modular action. What it does provide is a coherent route to **patch holography** in which de Sitter static patches are natural:

1. the fundamental object is a horizon screen in a static-patch description;

2. Λ is a capacity parameter tied to finite Hilbert-space dimension, not a locally reconstructible vacuum-energy term;

3. Einstein-like dynamics emerge up to Λg_{ab};
4. on the support-visible BW scaling branch, the realized scaling-limit cap modular action is geometric and may be outer on a non-type-I observer algebra.

BW-side status. The Lorentz side is closed at the support-visible scaling level by Theorem 4.2. The theorem intentionally avoids the false stronger route through a full-algebra unregularized common floor; the automorphism statement on the observer-facing geometric cap pair is the required static-patch content.

Many observers, one Λ. In this framework, each timelike observer is associated with a horizon patch rather than a single global description. The de Sitter parameter Λ is the shared global capacity constraint across overlap-consistent descriptions.

Summary. The model gets de Sitter by moving the holographic screen from "infinity" to an observer's horizon and by treating de Sitter entropy, i.e. finite screen capacity, as the global N_{CRC} closure datum. The input-free closure is the cosmic record-closure readback fixed point stated above, with the screen-normalized self-closure density as its count representation. The usual dS-holography obstacles are precisely the ones avoided by refusing the boundary-at-infinity viewpoint. This is not a claim of a solved dS/CFT dual. It is a static-patch holography program with an explicit local/global split and an explicit extracted-geometric-subnet scaling boundary.

4.7 Gauge Reconstruction and Standard Model Structure

4.7.1 Edge sector category and gauge group reconstruction

At any fixed UV cutoff, edge-center completion provides finite collar-sector packages of edge charges, intertwiners, fusion, and duals. 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 the question whether zero-obstruction edge sectors survive refinement, so the compact gauge argument uses the theorem-produced fixed-cutoff category, refinement/fiber ladder, and realized MAR-admissible compact-gauge witness theorem rather than treating the mixing estimate as a gauge-sector existence proof.

On the ordinary or central-defect branch, path-independent movement of collar charges is supplied by overlap gluing. TransportabilityFromOverlapGluing constructs transport from overlap paths and proves the exact zero-obstruction criterion: ordinary/central strict transport exists 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 remains a higher-gauge sector rather than an ordinary DR field-algebra sector. Thus the overlap obstruction calculus classifies and routes sectors; it does not by itself select the Standard Model.

Classification is not realization.

Proposition (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

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

DR/Tannaka reconstruction returns $G = \mathrm{Aut}_\otimes(\mathcal{F})$ for the constructed sector category and fiber functor. If that category is trivial, G is trivial; in general G is whatever compact group the tensor-fiber data reconstructs. The Standard Model quotient enters only after MAR is applied to a nonempty realized one-Higgs chiral sector package. QED.

Theorem (FixedCutoffBosonicSectorCategory). At every fixed regulator cutoff, on the ordinary or central-defect zero-obstruction branch and in the bosonic internal-gauge sector of the $3+1$-dimensional EFT regime, the localized collar charges form a semisimple rigid symmetric C^*-tensor category $\mathsf{Sect}_r^{\mathrm{bos}}$. Edge-center completion identifies the simple objects as minimal central summands P_α with irreducible boundary carriers W_α. Zero-obstruction overlap transport identifies localizations. Collar concatenation gives tensor product; overlap-gluing associativity gives the associator; orientation reversal gives duals; the finite-dimensional lifted collar algebra gives the *-operation and C^*-norm; and $3+1$-dimensional spacelike exchange gives bosonic symmetry. 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.

Theorem (RefinementFunctorAndFiberDescent). On a cofinal refinement tail of the ordinary or central zero-obstruction bosonic EFT branch, the fixed-cutoff categories $\mathsf{Sect}_r^{\mathrm{bos}}$ of FixedCutoffBosonicSectorCategory assemble into a directed refinement system. For $s \succeq r$, refinement induces faithful monoidal *-functors

$$U_{rs} : \mathsf{Sect}_r^{\mathrm{bos}} \to \mathsf{Sect}_s^{\mathrm{bos}}$$

that preserve the tensor unit, tensor product, associator, duals, *-structure, bosonic symmetry, and zero-obstruction transport classes on cofinal tails, with $U_{rt} \simeq U_{st} \circ U_{rs}$. The lifted finite collar Hilbert spaces also give compatible finite-dimensional bosonic multiplicity fibers $F_r : \mathsf{Sect}_r^{\mathrm{bos}} \to \mathsf{Hilb}_{\mathrm{fd}}$, with natural unitary identifications $F_s(U_{rs}X) \cong F_r(X)$.

Proof. Edge-center completion realizes each fixed-cutoff bosonic sector as a minimal central summand with a finite boundary carrier. A refinement embeds the coarse collar presentation into the refined collar presentation and sends each overlap-visible zero-obstruction summand to the direct sum of its refined descendants with the same transported charge class. Intertwiners refine by conjugating with the collar inclusion, so the induced map on sectors is faithful and *-preserving. Collar concatenation commutes with refinement up to the overlap-gluing associator, while orientation reversal and $3+1$-dimensional spacelike exchange commute with the same inclusion; hence the induced functors are monoidal, preserve duals, and preserve bosonic symmetry. Compatibility of successive inclusions gives $U_{rt} \simeq U_{st} \circ U_{rs}$. Finally, the lifted collar carrier of a sector is finite-dimensional at every fixed cutoff, and refinement only changes its presentation by a unitary identification inside the refined lifted collar algebra. These carriers define finite multiplicity fibers F_r, and the same inclusions give the required natural unitary compatibilities. QED.

Consequently the EFT branch has a directed ladder $(\mathsf{Sect}_r^{\mathrm{bos}}, U_{rs}, F_r)$ of theorem-produced bosonic edge-sector categories satisfying that zero-obstruction criterion, with monoidal refinement functors and compatible finite multiplicity fibers. Write

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

for the directed colimit retaining the sectors and intertwiners that persist in that system. The theorem below is neutral about whether Sect_∞ is trivial or nontrivial: if the realized branch furnishes only the tensor unit, the reconstructed compact group is the trivial group.

Persistence lemma. Write $U_{rs} : \mathsf{Sect}_r^{\mathrm{bos}} \to \mathsf{Sect}_s^{\mathrm{bos}}$ for the refinement functor from a regulator scale r to a finer scale $s \succeq r$. If a realized zero-obstruction edge-sector class α_{r_0} appears on a sufficiently fine collar, has overlap-visible edge-center support, and has representatives $\alpha_s \in \mathsf{Sect}_s^{\mathrm{bos}}$ on a cofinal refinement tail such that $U_{st}\alpha_s \simeq \alpha_t$ for all later $t \succeq s$ on that tail, then it determines a unique object $[\alpha] \in \mathsf{Sect}_\infty$. The reason is just the directed-colimit equivalence relation: two representatives define the same object once their images agree at a common finer stage. Therefore a later stage cannot erase, split, or change the zero-obstruction status of that class without either breaking the theorem-produced refinement-system hypotheses or introducing new overlap-visible support. This is a persistence result in the constructed ladder, not a proof that a nontrivial colimit exists.

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

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

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

Theorem 6.1 (Constructed bosonic sector category and Tannaka/DR reconstruction). On the ordinary or central zero-obstruction bosonic EFT branch constructed by FixedCutoffBosonicSectorCategory and RefinementFunctorAndFiberDescent, the colimit Sect_∞ is a rigid symmetric C^* tensor category with faithful bosonic fiber functor $\mathcal{F}$, and therefore

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

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

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

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

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

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

Hence $G = \mathrm{Aut}_\otimes(\mathcal{F})$ is a closed subgroup of the compact product $\prod_X U(\mathcal{F}(X))$, so G is compact. The DR/Tannaka hypotheses are therefore satisfied on the constructed pair $(\mathsf{Sect}_\infty, \mathcal{F})$, and the Doplicher–Roberts/Tannaka reconstruction theorem gives a symmetric C^*-tensor equivalence $\mathsf{Sect}_\infty \simeq \mathrm{Rep}(G)$ [22, 23]. If another compact group G' gave the same category through a faithful bosonic fiber functor, the induced symmetric tensor equivalence $\mathrm{Rep}(G) \simeq \mathrm{Rep}(G')$ compatible with the forgetful functors would identify both groups as the automorphism group of the same fiber functor. Thus $G \cong G'$, so the reconstruction is unique up to isomorphism. QED.

Corollary 6.1 (field algebra reconstruction on zero-obstruction sectors). If in the small-region limit the edge sectors are localized and satisfy the zero-obstruction transport criterion of Theorem 3.4b, then there exists a field algebra $\mathcal{F}$ and a compact group G such that $\mathcal{A} = \mathcal{F}^G$. This is the Doplicher–Roberts reconstruction of local gauge symmetry from the constructed transportable sector category. QED.

Corollary 6.1a (The compact-gauge setup and realized witness). On the central branch, the strict DHR-transportability condition used in Corollary 6.1 is exactly the theorem-level loop-coherence condition $[z]_\Sigma = 0$ of TransportabilityFromOverlapGluing. On the genuinely noncentral branch, strict ordinary transport is exactly $q_\Sigma = 0$, while $q_\Sigma \neq 0$ remains a higher-gauge fixed-cutoff sector handled by crossed-module data rather than by the ordinary DR field-algebra corollary. The fixed-cutoff bosonic category is constructed by FixedCutoffBosonicSectorCategory, and the faithful monoidal refinement functors plus finite bosonic fibers are constructed by RefinementFunctorAndFiberDescent. Thus the compact-gauge route uses the zero-obstruction sectors classified by gluing and theorem-produced category/refinement/fiber data. DR/Tannaka reconstruction then yields a compact group from that category; MAR, not the cocycle calculus alone, selects the realized Standard Model branch.

Theorem (gauge-sector classification-selection factorization). On the OPH compact-gauge lane, the realized Standard Model claim factors as

$$\text{overlap/gluing data} \longrightarrow \text{obstruction class} \longrightarrow \mathsf{Sect}_\infty^{\text{bos}} \longrightarrow G = \mathrm{Aut}_\otimes(\mathcal{F}) \longrightarrow \mathfrak{S}_{\mathrm{MAR}}.$$

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. QED.

4.7.2 Selecting the SM factors (derived from MAR)

Theorem 6.1 yields *some* compact G. Axiom 5 (MAR, Minimal Admissible Realization) acts next on admissible sector packages; on the explicit one-Higgs chiral matter branch below, it selects the realized Standard Model quotient.

Axiom 5 (MAR): Minimal Admissible Realization. Among all OPH-realizable sector packages $\mathfrak{S}$ consisting of the connected Lie gauge-sector image relevant in the low-energy EFT, its admissible light chiral matter content, and one Higgs doublet, and which are (i) loop-coherent / transportable (vanishing relevant obstruction: $[z] = 0$ on the central branch or $q_\Sigma = 0$ on the higher-gauge branch), (ii) anomaly-free, (iii) refinement-stable with light chiral matter, (iv) single-Higgs Yukawa-completable with one connected abelian charge factor acting nontrivially on the coupled carrier, (v) intrinsically quark-sector CP-capable, (vi) weak-sector UV-completable on that same one-Higgs branch, the realized low-energy package is the lexicographically minimal one under

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

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

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

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

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

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

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

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

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

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

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

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

Proposition 6.1e (meaning of MAR uniqueness). MAR uniqueness is uniqueness of the observer-visible low-energy package modulo the physical equivalence relation in Definition 6.1b. It is not uniqueness of a microscopic regulator representative. The subsequent 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, the same exact hypercharge lattice after normalization, the same structural electroweak force content, $N_c = 3$, and $N_g = 3$. Any remaining differences are gauge-center quotient choices, relabelings, or inert implementation data and therefore do not change the physical predictions recorded in D8–D9.

Note on transport obstruction. The gluing obstruction is tracked explicitly so the gauge lane states its branch conditions. TransportabilityFromOverlapGluing identifies $[z]_\Sigma = 0$ as the central-branch strict transport criterion and $q_\Sigma = 0$ as the genuinely noncentral strict ordinary transport criterion. When $q_\Sigma \neq 0$, the sector is handled by crossed-module data rather than by the ordinary DR field-algebra corollary. FixedCutoffBosonicSectorCategory constructs the fixed-cutoff category, RefinementFunctorAndFiberDescent constructs the refinement/fiber descent, and the realized witness theorem supplies nonempty ordinary/central compact-gauge witness data. The realized Standard Model branch enters only after MAR is applied to the admissible one-Higgs sector packages.

What MAR derives. Product structure, minimal sector content, and coupled edge-capacity minimality are consequences of MAR applied to the admissible class:

- **Product structure**: follows from the minimal coupled carrier $\mathbb{C}^3 \otimes \mathbb{C}^2$, which enforces commuting color and weak actions.
- **Minimal sector content**: the pseudoreal doublet and complex triplet are the minimal nonabelian representations satisfying the admissibility conditions, while the connected abelian charge factor is an explicit admissibility input.
- **Coupled edge-capacity minimality**: this is the first component of MAR's complexity vector.

With MAR stated, the SM derivation proceeds via standard lemmas:

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

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

QED.

Lemma 6.3 (SU(2) from a pseudoreal doublet). Let H be a positive-dimensional compact connected Lie group with a faithful 2D pseudoreal unitary representation V. Then the semisimple part of H contains an $SU(2)$ factor acting as the fundamental doublet. Finite or disconnected counterexamples are not part of the theorem package, which only applies the lemma to the identity component on the relevant nonabelian image. QED.

Lemma 6.4 (SU(3) from an irreducible triplet). Let H be a positive-dimensional compact connected Lie group with a faithful irreducible complex 3D unitary representation W. Then the semisimple image contains an $SU(3)$ factor acting as the fundamental triplet. Finite or disconnected counterexamples are not part of the theorem package, which only applies the lemma to the identity component on the relevant nonabelian image. QED.

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

Proposition 6.6 (physical group quotient). If the realized matter spectrum has hypercharges quantized in sixths, then the kernel acting trivially on all realized sectors is Z_6, so

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

QED.

Proposition 6.6a (SM from MAR). Under Axiom 5 (MAR):

- By the definition of χ_{cpl}, the realized MAR package contains a light chiral pseudoreal weak-type role and a light chiral complex color-type role on a common coupled carrier; clause (iii) imposes refinement stability on that branch, and clause (iv) adds one connected abelian charge factor acting nontrivially on the same carrier together with one-Higgs Yukawa completeness.
- The minimal faithful pseudoreal representation is the doublet ($\chi = 2$), giving $SU(2)$. No intrinsically complex 2D irrep qualifies in the connected compact Lie class, because the connected derived image of any irreducible 2D unitary representation lies in $SU(2)$, so the nonabelian action is pseudoreal up to abelian twist. The first intrinsically complex case is therefore the triplet ($\chi = 3$), giving $SU(3)$.
- The minimal coupled carrier for both is $\mathbb{C}^3 \otimes \mathbb{C}^2$, giving coupled edge capacity $\chi_{\text{cpl}} = 6$.
- The block-diagonal faithful representation $\mathbb{C}^3 \oplus \mathbb{C}^2$ of $S(U(3) \times U(2))$ has dimension 5, but it is not coupled and therefore is not the MAR minimizer.
- The maximal compact subgroup of $U(6)$ acting on $\mathbb{C}^3 \otimes \mathbb{C}^2$ with commuting actions is $(SU(3) \times SU(2) \times U(1))/(\text{finite center})$.
- The commutant of $SU(3) \times SU(2)$ inside $U(6)$ is exactly $U(1)$, so any connected abelian factor acting nontrivially on the coupled carrier is necessarily a single $U(1)$, and no additional continuous factors appear without increasing χ_{cpl}.
- Product structure is not separately assumed: it follows from the tensor product structure of the minimal coupled carrier.

Combined with Proposition 6.6 (hypercharges quantized in sixths from the realized spectrum), this yields:

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

The full proof route is given in the remainder of Section 6.2.

Theorem 6.6aa (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

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

with

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

The fixed-cutoff edge heat-kernel law gives $p_R(t) \propto d_R e^{-tC_2(R)}$, hence positive support for every finite-dimensional witness sector at finite t. The refinement functors carry those zero-obstruction labels cofinally. The witness is anomaly-free and one-Higgs Yukawa-complete by the hypercharge theorem, contains the weak-type and color-type roles used by MAR, and satisfies the CP-capability and weak-sector UV clauses that force $N_g = 3$ once $N_c = 3$. Therefore the MAR-admissible class is nonempty, and every OPH-admissible microscopic UV completion of the realized branch projects, modulo physical equivalence, to

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

with the Standard Model hypercharge lattice. Microscopic regulator representatives remain underdetermined up to gauge presentation, implementation hiding, and inert ancillary stabilization. QED.

Theorem (Realized Standard Model branch). Assume the ordinary or central zero-obstruction bosonic branch; the fixed-cutoff bosonic sector category, refinement/fiber descent, and compact-group reconstruction above; a nonempty realized one-Higgs chiral MAR-admissible class witnessed by $\mathcal{W}_{\mathrm{SM}}$; weak-type and intrinsically complex color-type roles on a common coupled carrier; one connected abelian charge factor acting nontrivially on that carrier; and the anomaly-free, one-Higgs Yukawa-complete, CP-capable, and weak-sector UV clauses of Axiom 5. Then MAR minima exist, and every MAR-minimal observer-visible package in that class is physically equivalent to

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

with

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

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

The proof is the earlier well-founded MAR-minima proposition plus the weak/color classification lemmas, the minimal coupled-carrier result $\mathbb{C}^3\otimes\mathbb{C}^2$, the $U(1)$ commutant, hypercharge theorem, three-color corollary, MAR-branch generation theorem, and the $\mathbb{Z}_6$ kernel proof. QED.

Corollary 6.6b (Three colors on the realized branch). On the realized one-generation matter package carried by the minimal coupled block of Proposition 6.6a, the color multiplicity is fixed by the same MAR derivation: the color factor is the fundamental triplet of SU(3), so

$$N_c=3.$$

This is not a later selector layered on top of the theorem stack. QED.

Corollary 6.6c (structural electroweak force and charges). On the realized one-Higgs branch, the electroweak gauge factor has Lie algebra

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

The Higgs doublet $H=(1,2)_{1/2}$ selects a 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 numerical W/Z masses belong to the D10 running/matching surface rather than to this recovered-core structural corollary. QED.

4.7.3 Refinement stability and unprotected relevant operators

The refinement-stability used later in the gauge branch is contained in the third OPH axiom. Because the same finite local constraint family is preserved across cutoffs, the realized UV states lie in one common finite-dimensional MaxEnt family rather than in a new coupling space at every refinement step. The selected stable branch of this family is therefore not an extra axiom beyond the MaxEnt branch itself.

Concretely, if

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

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

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

The resulting refinement-stable branch is simply a trajectory or invariant subset of this finite-dimensional multiplier map. Accordingly, when sections speak of a refinement-stable directed colimit of sectors, the refinement-stable qualifier means persistence along this MaxEnt branch itself; the extra ingredients are transportability, symmetry, and the bosonic fiber-functor clauses, and the MaxEnt/mixing package does not decide whether the resulting colimit is trivial or nontrivial.

The consensus paper makes the state-side coarse-graining connection explicit. Given quotient normal-form maps n_r, obstruction maps h_r, and coarse-graining maps (ρ_{sr}, χ_{sr}), reconciliation commutes with coarse-graining at the macroscopic readout scale whenever the two square defects

$$d_r^Q(\rho_{sr} n_s(x), n_r \rho_{sr}(x)), \qquad d_r^{\mathcal{H}}(\chi_{sr} h_s(x), h_r \rho_{sr}(x))$$

are controlled. Exact refinement naturality gives zero defect; approximate RG matching is theorem-grade only to the extent that these defects are bounded on the selected branch. This keeps the MaxEnt refinement map $R_{\ell \to L}$ connected to the reconciliation law without promoting arbitrary coarse-graining channels to OPH laws.

Relevant operators that are neither symmetry-forbidden nor retained in the selected constraint family are precisely the directions that try to push the flow off that stable branch.

Lemma 6.7 (refinement-stable MaxEnt branch forbids unprotected relevant operators). Assume the local finite-constraint MaxEnt/refinement branch of the third OPH axiom. Let $\mathcal{O}$ be a gauge-invariant Lorentz-scalar relevant deformation in the emergent EFT sense ($\Delta < 4$ in $3+1$D), allowed by symmetry and absent from the retained constraint family. Then the selected branch can keep the coupling of $\mathcal{O}$ at zero only if symmetry forbids that direction or the constraint family explicitly retains it. Otherwise generic refinement induces a nonzero component along that direction and the flow leaves the selected branch. This is a branch-persistence statement, not a universal entropy-ordering theorem for arbitrary off-branch phases.

Proof sketch. Linearize the induced refinement map $R_{\ell \to L}$ on the finite-dimensional multiplier space around the selected branch. A relevant operator gives an unstable direction with scaling exponent $y > 0$. If $\mathcal{O}$ is not fixed by symmetry and is not part of the retained constraint family, generic UV mismatch produces a nonzero component along that direction. Because $y > 0$, repeated coarse-graining amplifies the component and pushes the flow off the selected branch unless one fine-tunes it away at every scale or protects it by symmetry or by the declared constraint family. QED.

Corollary 6.8 (chirality selector). A gauge-invariant Dirac mass term is a relevant scalar. If both chiralities exist in conjugate representations, the mass term is allowed and will be generated under refinement unless symmetry-forbidden or explicitly retained as a protected constraint.

Therefore keeping light fermions on the selected refinement-stable branch requires chiral matter content or an explicit protecting mechanism. QED.

4.7.4 Generation number from CKM CP capability and weak-sector completability (derived from MAR)

Anomaly cancellation is generation-by-generation, so it does not fix the number of generations. On the realized one-Higgs branch, the lower and upper bounds come from the CP-capability and weak-sector UV clauses declared in MAR, and MAR then selects the minimum.

Proposition 6.9 (The number of generations is N_g = 3). On the realized one-Higgs quark branch of Proposition 6.6a, with the derived $N_c = 3$ from Corollary 6.6b and with the CP-capability and weak-sector UV clauses contained in Axiom 5, the generation number is

$$N_g = 3.$$

Inputs.

1. **Intrinsic CKM CP capability** is part of the realized quark branch through clause (v) of Axiom 5.
2. **Weak-sector UV completability** is part of the same realized one-Higgs branch through clause (vi) of Axiom 5.
3. **MAR minimality** acts on the same realized branch once the first three complexity entries are fixed.
4. Use the derived $N_c = 3$ from Corollary 6.6b.

Step 1: CKM CP-capability lower bound. The number of physical CP-violating phases in an N_g × N_g CKM matrix is:

$$\#(\text{CP phases}) = \frac{(N_g - 1)(N_g - 2)}{2}.$$

- For N_g = 1, 2: this is 0 → **no intrinsic CKM CP capability**.
- For N_g = 3: this is 1 → **intrinsic CKM CP capability is available**.

So intrinsic CKM CP capability requires:

$$N_g \geq 3.$$

Step 2: SU(2) asymptotic freedom upper bound. The one-loop coefficient is:

$$b_{\text{SU}(2)} = \frac{22}{3} - \frac{1}{3}N_g(N_c + 1) - \frac{1}{6},$$

where the final $-1/6$ is the contribution of one complex Higgs doublet. Asymptotic freedom means $b_{\text{SU}(2)} > 0$, i.e.,

$$N_g(N_c + 1) < \frac{43}{2}.$$

With N_c = 3, we have N_c + 1 = 4, so:

$$4N_g < \frac{43}{2} \quad \Rightarrow \quad N_g \leq 5.$$

Combining: 3 ≤ N_g ≤ 5.

Step 3: MAR selection. These are not extra post-MAR selectors: they are the numerical content of clauses (v) and (vi) of Axiom 5 on the realized one-Higgs branch. Given the allowed window {3, 4, 5}, MAR (fourth component of the complexity vector $C(\mathfrak{S})$) selects the smallest viable choice:

$$N_g = 3.$$

QED.

Why this is convincing.

- It predicts a **single integer**.
- It uses **two explicit MAR clauses** (intrinsic CKM CP capability and weak-sector UV completability) plus MAR's lexicographic minimality, rather than a separate post hoc admissibility menu.
- It is not a fit to a continuous number.
- Under the stated gauge-selection hypotheses, this is a derived result.

Corollary 6.9a (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 6.1 through Proposition 6.9. Then every MAR-minimal representative in that branch 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 Proposition 6.6 and the finite quotient fixed by Proposition 6.6. The only residual freedom is physical equivalence in the sense of Definition 6.1b.

Proof. Proposition 6.1d supplies the MAR minima. Lemmas 6.3–6.5 and Proposition 6.6a fix the minimal weak, color, coupled-carrier, and connected abelian roles at the least first three complexity entries. Corollary 6.6b fixes $N_c = 3$, Proposition 6.9 fixes $N_g = 3$, and the hypercharge and quotient propositions fix the visible charge lattice and $\mathbb{Z}_6$ kernel. The equivalence relation removes only relabelings, gauge-center conventions, and inert implementation data. QED.

4.7.5 Hilbert-space formulation of gluing data

Let {P_i} be a good cover of the screen. For each patch, fix a representation

$$\pi_i : \mathcal{A}_i \to \mathcal{B}(\mathcal{H}_i).$$

For each overlap, choose a unitary intertwiner

$$U_{ij} : \mathcal{H}_j \to \mathcal{H}_i$$

such that for all $O \in \mathcal{A}_{ij}$,

$$\pi_i(O) = U_{ij} \pi_j(O) U_{ij}^{\dagger}.$$

Normalize U_ii = 1 and U_ji = U_ij†.

Lemma 6.10 (centrality on triple overlaps). On a triple overlap define

$$\Omega_{ijk} := U_{ij}U_{jk}U_{ki}.$$

For all $O \in \mathcal{A}_{ijk}$,

$$\Omega_{ijk}\pi_i(O) = \pi_i(O)\Omega_{ijk}.$$

Proof. Conjugation by U_ki sends π_i(O) to π_k(O), by U_jk to π_j(O), by U_ij back to π_i(O). Thus conjugation by Ω_ijk fixes π_i(O), so Ω_ijk commutes with π_i(O). QED.

Lemma 6.11 (gauge behavior). If Ũ_ij = V_i U_ij V_j† with V_i acting trivially on overlap observables, then

$$\tilde{\Omega}_{ijk} = V_i\Omega_{ijk}V_i^{\dagger}.$$

In particular, if Ω_ijk is central, its class is gauge invariant. QED.

4.7.6 Loop obstruction class (central defect)

Assume the defect is central and write

$$\varphi_{ij} := \mathrm{Ad}(U_{ij}) \quad \text{on } \mathcal{A}_{ij}.$$

This is the abelian truncation of the full 2-group obstruction in Section 3.4.

Then there exist central unitaries z_{ijk} such that

$$\varphi_{ij}\varphi_{jk}\varphi_{ki} = \mathrm{Ad}(z_{ijk}).$$

Theorem 6.12 (loop-coherent gluing iff vanishing obstruction). The family $\{z_{ijk}\}$ is a Čech 2-cocycle, and its class $[z]$ is gauge invariant. On any quadruple overlap P_{ijkl},

$$z_{jkl}z_{ikl}^{-1}z_{ijl}z_{ijk}^{-1} = 1.$$

A loop-coherent global gluing exists iff [z] = 0.

Proof. Compare two parenthesizations of ϕ_ij ϕ_jk ϕ_kl ϕ_li on a quadruple overlap to obtain the cocycle condition above. Gauge changes shift z by a coboundary. If [z]=0, rephase by a 1-cochain to eliminate defects and obtain path-independent transport. Conversely, loop-coherent gluing implies z_ijk = 1. QED.

4.7.7 EFT reduction to anomaly cancellation

Assume ExtEFT: a low-energy 3+1D chiral gauge theory exists with group G. Then the obstruction class $[z]$ is structurally analogous to the EFT 't Hooft anomaly class and is expected to map to it after a separate anomaly-descent construction. That map is not constructed in this manuscript, so $[z] = 0$ is used here as an internal gluing/transportability condition rather than a proved equivalence to EFT anomaly cancellation.

4.7.8 Hypercharge from anomaly freedom and Yukawas

Theorem 6.13 (Hypercharge from anomaly freedom and Yukawas). Assume gauge group SU(N_c) × SU(2) × U(1)_Y and one generation of left-handed Weyl fermions (Q, u^c, d^c, L, e^c), with a Higgs doublet H and Yukawa terms

$$QHu^c, \qquad QH^\dagger d^c, \qquad LH^\dagger e^c.$$

Then anomaly freedom and Yukawa invariance fix the hypercharges up to an overall normalization, yielding the Standard Model pattern for N_c = 3.

Proof. Yukawa invariance gives

$$Y_u = -(Y_Q + Y_H), \quad Y_d = -Y_Q + Y_H, \quad Y_e = -Y_L + Y_H.$$

Anomaly cancellation yields

$$\begin{aligned} SU(2)^2U(1): &\quad N_cY_Q + Y_L = 0, \\ \text{grav}^2U(1): &\quad 2N_cY_Q + N_cY_u + N_cY_d + 2Y_L + Y_e = 0. \end{aligned}$$

Solving gives

$$Y_L = -N_cY_Q, \quad Y_H = N_cY_Q, \quad Y_u = -(N_c+1)Y_Q, \quad Y_d = (N_c-1)Y_Q, \quad Y_e = 2N_cY_Q.$$

With these relations, SU(N_c)^{2}U(1) and U(1)^3 anomalies vanish automatically. Fixing the normalization by Q = T_3 + Y and Q(ν_L)=0 gives

$$Y_Q = \frac{1}{2N_c}.$$

For N_c = 3,

$$Y_Q = \frac{1}{6}, \quad Y_L = -\frac{1}{2}, \quad Y_e = 1, \quad Y_u = -\frac{2}{3}, \quad Y_d = \frac{1}{3}, \quad Y_H = \frac{1}{2}.$$

Without Yukawas, the cubic anomaly leaves two discrete branches (Y_u, Y_d exchange). Yukawa invariance selects the branch with a single Higgs doublet. QED.

Corollary 6.13a (Exact rational hypercharges). With the derived N_c = 3, the hypercharge assignments are uniquely fixed to exact rational values:

$$Y_Q = \tfrac{1}{6}, \quad Y_L = -\tfrac{1}{2}, \quad Y_u = -\tfrac{2}{3}, \quad Y_d = \tfrac{1}{3}, \quad Y_e = 1, \quad Y_H = \tfrac{1}{2}.$$

Why this is convincing.

- These are **exact rationals**, not approximate numbers.
- Their ratios are fixed by anomaly freedom + Yukawa invariance, and the absolute lattice is fixed by the standard normalization, with no continuous parameters to adjust.
- This high-precision set of numbers strongly constrains the realized matter package and matches observation exactly.

4.7.9 Witten anomaly on the realized color branch

Theorem 6.14 (Witten anomaly consistency on the realized color branch). On the realized one-generation package emitted by Proposition 6.6a and Corollary 6.6b, the global SU(2) anomaly is absent generation by generation.

Inputs.

1. The realized color factor is the SU(3) triplet from Corollary 6.6b.
2. The matter content per generation includes:
   - one left-handed quark doublet Q which is an SU(2) doublet and carries color,
   - one left-handed lepton doublet L which is an SU(2) doublet and color singlet.
3. **Witten's global SU(2) anomaly constraint** (Witten, 1982): the number of left-handed SU(2) doublets must be even.

Proof. Count SU(2) doublets per generation:

- Quark doublets: N_c copies (one per color),
- Lepton doublets: 1 copy.

Total doublets per generation:

$$N_c + 1.$$

Substituting the realized value $N_c = 3$ gives

$$N_c + 1 = 4.$$

Witten anomaly cancellation therefore requires an even number of doublets, and the realized branch satisfies it exactly:

$$4 \equiv 0 \pmod 2.$$

So the realized triplet-doublet package is globally consistent generation by generation. QED.

Theorem-stack role. The D8 minimal-coupled-carrier derivation emits the specific value $N_c = 3$ and carries it into D9. Witten's anomaly is retained as a nontrivial consistency check on that emitted triplet-doublet package, not as a residual selector that upgrades oddness to 3.

Why this is convincing.

- It checks a **global anomaly constraint** on the emitted branch without introducing a new selector.
- The parity constraint is independent of continuous parameters, RG running, masses, or Yukawa values.
- It confirms that the realized SU(3) triplet plus lepton doublet package is globally consistent generation by generation.

4.7.10 Bond-dimension gatekeeping

In tensor-network or code realizations, gauge actions act on edge factors of size χ, so emergent compact gauge groups embed in $U(\chi)$. This shows a capacity constraint: accommodating SU(3) color and SU(2) weak factors shows $\chi \geq 6$ in the minimal case, consistent with the MAR-derived gauge group.

4.7.11 Inevitability of photon and graviton

The model requires photons and gravitons.

Photon inevitability chain:

1. Gauge-as-gluing states that overlap identifications have redundancy forming a local groupoid.
2. Theorem 2.3 (edge-center completion) decomposes collar Hilbert spaces into sectors labeled by boundary gauge representations.
3. Theorem 6.1 constructs the refinement-limit bosonic edge-sector category from theorem-produced fixed-cutoff category, refinement, and fiber data, then recovers a compact gauge group G from it.
4. Corollary 6.1 (on the zero-obstruction sector system of TransportabilityFromOverlapGluing) reconstructs a field algebra with G as a local gauge symmetry.
5. For the Standard Model, the realized one-Higgs branch leaves the $Q = T_3 + Y$ stabilizer $\mathrm{U}(1)_Q$ after electroweak symmetry breaking.
6. A gauge boson is the quantum of the gauge field. Once $\mathrm{U}(1)_Q$ emerges from overlap redundancy, its gauge field exists, and its quantum (the photon) must exist.

The photon is not postulated. It is forced by the axioms through the chain above. The photon mediates the correlations between charged excitations in different patches; it is how the $\mathrm{U}(1)_Q$ redundancy structure propagates through the algebra net.

Graviton inevitability chain:

1. Theorem 4.2 (BW_{S^2}) shows that on the support-visible prime geometric cap net, with collar Markov locality and controlled refinement-limit transport, modular flow on caps becomes geometric conformal dilation.
2. Theorem 4.3 identifies the induced kinematic group as $\mathrm{Conf}^+(\mathrm{S}^2) \cong \mathrm{PSL}(2,\mathbb{C}) \cong \mathrm{SO}^+(3,1)$, the Lorentz group.
3. Theorem 5.1 shows that, under the stated scaling-limit null-stress and reference-state conditions, the condition δS_gen = 0 implies the rest-frame first-variation Einstein relation, with overlap consistency upgrading it to the semiclassical Einstein equations in the EFT regime.
4. The metric tensor emerges as the compression of modular flow data, and its dynamics are fixed by the derived fixed-cap generalized-entropy stationarity theorem.
5. A dynamical metric in a quantum theory requires a spin-2 quantum field. Its quantum (the graviton) must exist.

The graviton is not postulated. It is forced by the axioms through the chain above. Diffeomorphism invariance emerges because the bulk spacetime description is a compression of screen data; different coordinate descriptions are redundancies in how that compression is presented.

4.7.12 Quotient-protected charge quantization

Theorem 6.19 (Charge quantization and no fractional color singlets). If the global gauge group is

$$G_{\text{phys}} = \frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{Z_6},$$

as derived in Proposition 6.6, then

All color-singlet states have integer electric charge.

Equivalently: no stable isolated particles with charges like ±1/3 can exist as color singlets.

Proof. The Z_6 quotient identifies the center elements (e^{2πi/3}, -1, e^{iπ/3}) ∈ SU(3) × SU(2) × U(1) with the identity. For a color-singlet state ($\tau = 0$), the SU(3) factor acts trivially. The resulting identification requires the SU(2) × U(1) quantum numbers to satisfy

$$(-1)^{2j} \cdot e^{i\pi n/3} = 1,$$

where j is the SU(2) spin and n = 6Y is the integer hypercharge label. This gives n ≡ -6j (mod 6), i.e., n ≡ 0 (mod 6) for integer j and n ≡ 3 (mod 6) for half-integer j. Equivalently: Y is integer when j is integer, and Y is half-integer when j is half-integer.

After electroweak breaking, Q = T_3 + Y. For integer j, T_3 ∈ ℤ and Y ∈ ℤ, so Q ∈ ℤ. For half-integer j, T_3 ∈ ℤ + 1/2 and Y ∈ ℤ + 1/2, so Q = (half-integer) + (half-integer) ∈ ℤ. In both cases, Q ∈ ℤ. QED.

Experimental status: No fractionally charged color-singlet particles have been observed. Three independent high-precision bounds confirm this:

1. **Neutrality of matter** (PDG 2024): The proton-electron charge sum satisfies

$$|q_p + q_e|/e < 1 \times 10^{-21},$$

confirming charge quantization to 21 decimal places.

2. **Fractional charge searches in bulk matter**: Silicone oil drop experiments limit fractionally charged particle abundance to

$$\text{(fractionally charged particles)/nucleon} \lesssim 10^{-22}.$$

3. **Collider searches** (CMS, PRL 134, 2025): Exclusions for stable particles with $q \in [e/3, 0.9e]$ up to masses ∼ 640 GeV (95% CL).

Existing searches are fully consistent with these structural consequences.

These are exact symmetry- or quotient-protected outputs: the empirical content is that violations are excluded up to available experimental sensitivity.

4.7.13 Coupling extraction from edge-sector probabilities

The edge-center completion (Theorem 2.3) yields sector probabilities p_α on collar boundaries. These probabilities encode the renormalized gauge coupling through a heat-kernel/Laplacian weighting law.

Abelian case (Z_n). For a Z_n gauge theory, the edge sectors are labeled by charge q ∈ {0, 1, ..., n-1}. The correct "Casimir" eigenvalue is the Laplacian eigenvalue of the boundary random walk:

$$\lambda_q = 4 \sin^2 \left(\frac{\pi q}{n} \right).$$

Note: only in the limit n → ∞ and q ≪ n does λ_q ≈ $(2\pi q/n)^2$ ∝ q^2. For finite n, the exact form is essential.

The sector probabilities follow a heat-kernel law:

$$p_q \propto e^{-t(\mu)\lambda_q},$$

where t(μ) is the "modular time" parameter encoding the scale. The extraction formula is:

$$t(\mu) = -\frac{\log(p_q/p_0)}{\lambda_q}, \qquad g^2_{\text{ent}}(\mu) = \frac{t(\mu)}{2\pi}.$$

Consistency requires that t extracted from different charges q agrees; this has been verified numerically (see Section 6.14).

Electric-center measurement. The edge sectors are measured using the *electric-center* prescription. For a region A and boundary vertex v ∈ ∂A, define the restricted star operator:

$$Q_v^{(A)} = \prod_{\ell \in \text{star}(v) \cap A} X_\ell^{\pm 1},$$

where X_ℓ is the shift operator on link ℓ. The sector projectors are:

$$P_{v,q} = \frac{1}{n} \sum_{m=0}^{n-1} \omega^{-mq} \left(Q_v^{(A)}\right)^m, \qquad \omega = e^{2\pi i/n},$$

and the probabilities are p_{v,q} = ⟨P_{v,q}⟩. This electric-center operator, built from X's rather than Z's, correctly captures the boundary gauge charge/flux that labels entanglement edge sectors.

Non-abelian generalization. For SU(N) gauge theories, the edge sectors are labeled by irreducible representations with probabilities:

$$p_j \propto d_j \, e^{-t(\mu) C_2(j)},$$

where d_j is the dimension and C_2(j) the quadratic Casimir. Extraction:

$$t(\mu) = -\frac{\log(p_j/p_0)}{C_2(j)}, \qquad g^2_{\text{ent}}(\mu) = \frac{t(\mu)}{2\pi}.$$

Theoretical derivation. The fixed-cutoff same-overlap edge law is carried by the regulated microphysics package once the declared overlap-sector projectors, reversible thermal branch, and quasi-local local-Gibbs generator of Theorem 2.6 are in hand. On this synthesis surface, the point is to summarize that fixed-cutoff source package and to state the refinement/Peter–Weyl lift boundary explicitly rather than to relocate the leaf-level source surface.

Theorem 6.20 (Fixed-cutoff edge-sector law and refinement lift boundary). Under the OPH axioms, the fixed-cutoff overlap-gauge realization of Sections 2.3 and 3.2, and the quasi-local local-Gibbs form supplied by Theorem 2.6, the regulated same-overlap edge-sector probability distribution satisfies:

$$p_R = \frac{d_R \, e^{-t\lambda_R}}{\sum_{R'} d_{R'} \, e^{-t\lambda_{R'}}}$$

where λ_R is the Laplacian eigenvalue on the R-isotypic component and t is determined by the collar Gibbs parameter.

Proof.

Step 1 (Edge Hilbert space). From the fixed-cutoff overlap-gauge realization, the edge degrees of freedom at a boundary circle Σ = ∂C live in a Hilbert space transforming under the gauge group G. Microscopically this is a finite-dimensional quantum-link edge space on the declared overlap interface; the effective representation-theoretic description used only for the refinement lift is modeled by $L^2(G)$. By the Peter–Weyl theorem [32]:

$$L^2(G) \cong \bigoplus_R V_R \otimes V_R^*$$

where V_R is the carrier space of irrep R.

Step 2 (Gauge invariance). The Gauss law constrains physical states. For an entanglement cut at Σ, the physical edge Hilbert space decomposes as $\mathcal{H}$_{edge}^{phys} = ⊕_R W_R where W_R contains states with flux in representation R.

Step 3 (Natural Hamiltonian). From the local-Gibbs form, the MaxEnt generator restricted to edge modes takes the form H_{edge} = Σ_R h_R P_R where P_R is the projector onto the R-sector. The key claim is that h_R = λ_R.

Justification: For a compact simple factor, the group Laplacian Δ_G = -Σ_a (T^a)^2 is the unique bi-invariant second-order differential operator up to scale. For a product group $G = \prod_i G_i$, the most general bi-invariant second-order operator is a positive linear combination $\sum_i c_i \Delta_{G_i}$, with one coefficient per factor. Any other gauge-invariant local choice would require higher derivatives, violating locality. For finite groups, the Cayley graph Laplacian plays the same role: λ_R = |S| - (1/d_R) Σ_{s∈S} χ_R(s).

Step 4 (MaxEnt selection). MaxEnt selects the Gibbs state:

$$\rho_{\text{edge}} = \frac{1}{Z} e^{-tH_{\text{edge}}} = \frac{1}{Z} \sum_R e^{-t\lambda_R} P_R.$$

Step 5 (Sector probabilities). The probability of sector R is p_R = Tr(ρ_{edge} P_R). The effective dimension for entanglement is d_R (not d_R^2) because we trace over one side of the cut. This gives:

$$p_R = \frac{d_R \, e^{-t\lambda_R}}{Z}.$$

QED.

Why the entropy rank is d_R (instead of d_R^2). The full edge space has dimension d_R^2 in sector R (from $V_R \otimes V_R^*$). Entanglement entropy, however, measures correlations *across* the cut. After tracing over one side, the reduced density matrix has effective rank d_R. Mathematically: in the Markov normal form, the edge factor on one side contributes $\log d_R$ to the entropy.

Scope. The fixed-cutoff same-overlap derivation is the theorem-bearing part of the edge-law package. The quasi-local local-Gibbs generator used here is derived from Theorem 2.6: if MaxEnt constraints are expectations of finitely many quasi-local operators, the entropy maximizer is automatically a Gibbs state with a quasi-local generator. The additional representation-theoretic step summarized here is the compact-group / Peter–Weyl refinement lift, together with its factorwise Laplacian identification. The fixed-cutoff same-overlap Casimir law is supplied on the microphysics surface; model-family uniqueness and large-edge continuations have their own status boundaries.

Normalization anchor: 2D Yang-Mills. The parameter t can be exactly matched to a conventional coupling in 2D Yang-Mills, where the physical Hamiltonian is literally the group Laplacian:

$$H = \frac{g^2}{2} \Delta_G, \qquad \Delta_G \chi_R = -C_2(R) \chi_R \quad \Rightarrow \quad E_R = \frac{g^2}{2} C_2(R).$$

Euclidean evolution for "time" A (the area of a cylinder in 2D YM) gives weight(R) ∝ exp(-A E_R) = exp(-g^2 A C$_2$(R)/2). Comparing with the heat-kernel expansion K_t(U) = Σ_R d_R χ_R(U) e^{-t C$_2$(R)} yields the exact identification:

$$t_{\text{phys}} = \frac{g^2 A}{2} \quad \text{(in 2D YM, no ambiguity).}$$

This shows that the Laplacian + MaxEnt → heat-kernel structure is exact in a solvable two-dimensional Yang-Mills case. The coefficient in front of C_2 is fixed. This normalization anchor is a two-dimensional edge-theory check, separate from the compact paper's support-visible compact-gauge repair-gap theorem. In any regime where the edge theory reduces to an effective 2D YM with known "Euclidean thickness" A_eff:

$$g^2(\mu) = \frac{2}{A_{\text{eff}}(\mu)} \cdot \frac{\Delta_R(\mu)}{C_2(R)},$$

and the RHS must be R-independent. This R-independence is an internal precision consistency test; the formula itself is the normalization map that connects t to the conventional gauge coupling.

4.7.14 Numerical validation of the heat-kernel law

The heat-kernel/Laplacian weighting of edge sectors has been validated in explicit 2D Z_n gauge models on closed geometries.

Model. A 2×2 periodic lattice gauge theory (8 links) with Z_n link Hilbert spaces and Hamiltonian:

$$H = -K \sum_p \text{Re}(B_p) - h \sum_\ell \text{Re}(X_\ell) - \Gamma \sum_v \text{Re}(A_v),$$

where X_ℓ is the Z_n shift on link ℓ, B_p is the oriented plaquette operator (product of Z's around plaquette p), and A_v is the oriented star/Gauss operator (outgoing X, incoming X†). With K = 1 and Γ = 5, the ground state satisfies ⟨A_v⟩ = 1 at all vertices to numerical precision.

Region and edge operator. Region A consists of links whose tail has x = 0 ("half-lattice" cut). At each boundary vertex v, the electric-center edge charge is the restricted star Q_v^{(A)} = ∏_{ℓ ∈ star(v) ∩ A} X_ℓ^{±1}.

Results for $\mathbf{Z_2}$. With $\lambda_1 = 4\sin^2(\pi/2) = 4$:

h	p_0	p_1	t	g_ent
0.5	0.8266	0.1734	0.391	0.249
1.0	0.9612	0.0388	0.803	0.357
2.0	0.9917	0.0083	1.194	0.436

Results for $\mathbf{Z_3}$ (overconstrained test). With $\lambda_1 = \lambda_2 = 4\sin^2(\pi/3) = 3$:

h	p_0	p_1	p_2	t(q=1)	t(q=2)	g_ent	m_plaq
0.2	0.4395	0.2803	0.2803	0.1500	0.1500	0.154	2.22
0.5	0.7509	0.1245	0.1245	0.5989	0.5989	0.309	1.75
1.0	0.9606	0.0197	0.0197	1.2956	1.2956	0.454	4.07
1.5	0.9851	0.0074	0.0074	1.6288	1.6288	0.509	7.06
2.0	0.9921	0.0039	0.0039	1.8440	1.8440	0.542	10.10

The equality $p_1 = p_2$ is exact (charge conjugation symmetry in Z_3). The equality t_{q=1} = t_{q=2} is the crucial **overconstrained** check: at h = 1.0, extracting t from q = 1 and q = 2 independently gives t_{q=1} ≈ 1.2956389318579 and t_{q=2} ≈ 1.2956389318521. The agreement to ~10^{-14} (machine precision) confirms that the edge distribution genuinely follows the heat-kernel/Laplacian form.

Region-choice robustness. At h = 1, the extracted g_ent is nearly independent of region size:

- 2 links (one vertex's outgoing links): g_ent ≈ 0.453
- 4 links (half-lattice): g_ent ≈ 0.454
- 6 links (three vertices): g_ent ≈ 0.453

This locality confirms that the coupling is dominated by physics near the cut, not global bookkeeping, exactly what is expected if this behaves like a local QFT observable.

Results for Z_5 (golden ratio test). The Z_5 case provides a stringent test because the Laplacian eigenvalues have a distinctive ratio involving the golden ratio $\phi = (1+\sqrt{5})/2$:

$$\lambda_q = 4\sin^2\left(\frac{\pi q}{5}\right), \qquad \frac{\lambda_2}{\lambda_1} = \frac{\sin^2(72°)}{\sin^2(36°)} = \phi^2 \approx 2.618.$$

This ratio distinguishes the Laplacian law from naive alternatives: a linear model (λ_q ∝ q) would predict ratio 2, while a quadratic model (λ_q ∝ q^2) would predict ratio 4.

Simulations on a 2×2 torus in the dual/flux basis (125 states in the zero-winding sector) give:

h	Measured ratio $\ln(p_2/p_0)/\ln(p_1/p_0)$	Deviation from ϕ^2
0.5	2.25	14%
1.0	2.51	4%
2.0	2.619	< 0.1%

In the weak-field limit (h → 0, strong magnetic coupling), the simulation converges to the golden ratio squared. This confirms that the vacuum entanglement spectrum encodes the precise geometric structure of the gauge group Laplacian.

Significance. This validates the mathematical law (sector probabilities weighted by Laplacian eigenvalues) in explicit 2D gauge-invariant models with non-flat sector distributions. The Z_3 and Z_5 tests are structurally identical to SU(2)/SU(3): multiple irreps overconstrain the slope, and agreement confirms the mechanism works before jumping to nonabelian groups.

Results for S_3 (first nonabelian test). The abelian tests above use charge-sector projectors that reduce to Fourier modes. For nonabelian groups, the edge-sector projector must be generalized to character projectors:

$$P_{v,R} = \frac{d_R}{|G|}\sum_{h\in G}\chi_R(h^{-1})Q_v^{(A)}(h),$$

where d_R is the dimension of irrep R, χ_R is its character, and Q_v^{(A)}(h) is the restricted gauge action at boundary vertex v acting only on links in region A.

For S_3 (the smallest nonabelian group, order 6), there are three irreps: trivial (d=1), sign (d=1), and standard (d=2). The Cayley-graph Laplacian eigenvalues for the transposition generating set are:

$$\lambda_{\text{triv}} = 0, \qquad \lambda_{\text{sign}} = 6, \qquad \lambda_{\text{std}} = 3.$$

Exact reduction on one plaquette. For the single-plaquette model (4 links), imposing Gauss's law at all vertices means the physical wavefunction depends only on the plaquette holonomy's conjugacy class. Since S_3 has exactly 3 conjugacy classes, the gauge-invariant Hilbert space is 3-dimensional, spanned by the character states {|χ_R⟩}. In this basis, the edge-sector probabilities are exactly p_R = $|c\_R|^2$ where $|\psi_0\rangle$ = Σ_R c_R |χ_R⟩. This is not an approximation; it is an exact identity for the one-plaquette gauge-invariant sector.

The heat-kernel ansatz predicts p_R ∝ d_R exp(-t λ_R). Extracting t independently from the sign and standard irreps provides an overconstrained test: the ratio λ_sign/λ_std = 6/3 = 2 is a parameter-free prediction. Results from a single-plaquette S_3 lattice gauge model (K=1, Γ=5):

h	p_triv	p_sign	p_std	t (sign)	t (std)	Δt/t	log-ratio
0.5	0.909	0.0013	0.089	1.09	1.01	8.4%	2.17
1.0	0.980	7.5×10^{-5}	0.020	1.58	1.54	2.8%	2.06
2.0	0.996	4.3×10^{-6}	0.004	2.06	2.04	1.0%	2.02
5.0	0.9993	1.0×10^{-7}	0.00066	2.68	2.67	0.3%	2.006
12	0.9999	3.0×10^{-9}	0.00011	3.27	3.27	0.1%	2.002
100	1.0000	6.1×10^{-13}	2.0×10^{-6}	4.69	4.69	0.009%	2.0002

The "Δt/t" column shows the fractional difference (t_sign - t_std) / $\bar{t}$. The "log-ratio" column shows log(p_sign/p_0) / log(p_std/(2 p_0)), which should equal λ_sign/λ_std = 2 if the heat-kernel form holds exactly.

As h increases, both diagnostics converge: Δt/t drops below 10^{-4} and the log-ratio approaches 2.000. This is exactly the expected behavior: finite-size corrections are largest at strong coupling; the heat-kernel form becomes exact as the perturbative regime is approached.

This is the first nonabelian validation of the edge-sector extraction mechanism. The structure (character projectors, Laplacian eigenvalues from the group's Cayley graph, overconstrained t extraction) is identical to what will be used for SU(2) and SU(3).

Parameter-free predictions for SU(2) and SU(3). The heat-kernel law yields exact, parameter-free ratio predictions that require no scheme matching. Define the "Casimir log-gap":

$$\Delta_R \equiv \ln\left(\frac{p_0}{d_0}\right) - \ln\left(\frac{p_R}{d_R}\right) = t\, C_2(R).$$

Ratios of Δ_R cancel all unknowns (t, partition function):

$$\frac{\Delta_{R_1}}{\Delta_{R_2}} = \frac{C_2(R_1)}{C_2(R_2)} \quad \text{(exact, parameter-free).}$$

SU(2) predictions. Irreps labeled by spin j have d_j = 2j+1 and C_2(j) = j(j+1). The framework predicts:

- $\Delta_1/\Delta_1/_2$ = 2/(3/4) = **8/3** ≈ **2.667**
- $\Delta_3/_2/\Delta_1/_2$ = (15/4)/(3/4) = **5**
- $\Delta_3/_2/\Delta_1$ = (15/4)/2 = **15/8 = 1.875**

SU(3) predictions. Irreps labeled by Dynkin indices (p,q) have C_2(p,q) = (p^2 + q^2 + pq + 3p + 3q)/3. Using the fundamental **3** = (1,0) with C_2 = 4/3 as the reference:

- $\Delta_8/\Delta_3 = 3/(4/3) = \mathbf{9/4 = 2.25}$
- $\Delta_6/\Delta_3 = (10/3)/(4/3) = \mathbf{5/2 = 2.5}$
- $\Delta_{10}/\Delta_3 = 6/(4/3) = \mathbf{9/2 = 4.5}$
- $\Delta_{15}/\Delta_3 = (16/3)/(4/3) = \mathbf{4}$
- $\Delta_{27}/\Delta_3 = 8/(4/3) = \mathbf{6}$

These are the SU(2)/SU(3) analogs of the Z_5 golden-ratio test: exact rational numbers fixed entirely by group theory, with no adjustable parameters.

Preliminary SU(3) results. A one-plaquette SU(3) "quantum link" model (finite truncated irrep basis, n_max = 12, $\kappa = 2$) has been used to extract t from 14 different irreps simultaneously. The results show internal consistency at the 1-3% level:

bare g^2	extracted t (mean±std)	g_ent	gap
0.3	0.314±0.0005	0.224	1.92
0.5	0.539±0.0025	0.293	1.83
0.8	0.896±0.012	0.378	1.72
1.0	1.144±0.025	0.427	1.64

The standard deviation across irreps provides a built-in error estimate. This is a QCD proton-physics surrogate rather than a full proton calculation; it lacks dynamical quarks and operates on a single plaquette. It nevertheless demonstrates that the nonabelian extraction machinery produces self-consistent outputs without tuning.

Extracting the normalization factor A_eff. The 2D YM anchor (Section 6.13) gives t = g^2 A / 2, so the "effective Euclidean thickness" is

$$A_{\text{eff}} = \frac{2t}{g^2}.$$

Computing this from the SU(3) table:

bare g^2	extracted t	A_eff
0.3	0.314	2.093
0.5	0.539	2.156
0.8	0.896	2.240
1.0	1.144	2.288

Mean: A_eff ≈ 2.19 with point-to-point scatter ~4%.

Extrapolation to weak coupling. The systematic drift in A_eff shows fitting A_eff(g^2) = A_0 + a · g^2. A weighted linear fit gives:

$$A_0 = 2.004 \pm 0.012$$

with χ^2/dof ≈ 0.09, indicating excellent consistency. This strongly shows that, in this toy UV completion, the "missing normalization" converges to A_eff → 2 as $g^2 \to 0$.

The normalization factor behaves like a quasi-constant and extrapolates to a simple value (≈ 2) in the weak-coupling limit. This provides a concrete path to absolute coupling predictions: once

A_eff is determined from microphysics, the conversion $g^2 = 2t/A\_eff$ fixes the gauge coupling without additional free parameters.

Internal validation summary. The heat-kernel law has been validated with increasing precision across multiple gauge groups:

- $\mathbf{Z_3}$: Overconstrained t extraction (q=1 vs q=2), precision ~10^{-14}
- $\mathbf{Z_5}$: Golden ratio squared ($\lambda_2/\lambda_1 = \phi^2$), precision 0.04%
- $\mathbf{S_3}$: Casimir log-ratio ($\lambda\_sign/\lambda\_std = 2$), precision 0.01%
- **SU(3)**: 14-irrep simultaneous extraction, precision 1-3%

The Z_3 test achieves machine precision because it is exactly overconstrained. The Z_5 and S_3 tests converge to their predicted ratios as coupling decreases. This provides strong internal validation of the mechanism "MaxEnt + Laplacian => heat-kernel sector weights" before applying it to physical gauge groups.

4.7.15 IBM Quantum Cloud hardware benchmarks

The lattice calculations above are internal numerical validations. The IBM Quantum Cloud benchmarks add a hardware check of whether OPH-motivated reduced-sector structures survive preparation and readout on real superconducting-qubit devices. These runs are small-sector consistency benchmarks rather than a standalone confirmation of the full framework. Their value is that they probe the same recoverability and heat-kernel observables on physical hardware and in simulation.

Detailed circuits, representative raw outputs, and rerun instructions are public in the project write-up `extra/IBM_QUANTUM_CLOUD.md` and the associated code/data bundle `code/ibm_quantum_cloud/`.

- **Stage 1 (Markov/recoverability benchmark).** On `ibm_marrakesh` and `ibm_fez`, the structured state reconstructs below both controls in conditional mutual information and above both controls in Petz fidelity. On `ibm_marrakesh`, the observed ordering is $0.2309 < 0.3890 < 0.9474$ for CMI and $0.9297 > 0.8649 > 0.6066$ for Petz fidelity; on `ibm_fez`, it is $0.1498 < 0.4992 < 0.9166$ and $0.9479 > 0.8117 > 0.6449$. This matches the predicted ordering on two independent real backends.
- $\mathbf{Z_3}$ **hardware sanity check.** The overconstrained Z_3 extraction passes cleanly on hardware: across prepared $t = 0.30, 0.60, 0.90$, the mean extracted t is within about 0.02 of the target, the two independent extractions agree to within 0.0142, 0.0034, and 0.0043, and leakage stays below 0.1%. This shows that the reduced-sector preparation and readout path is internally coherent on-device before sharper ratio claims are interpreted.
- $\mathbf{Z_5}$ **exact-ratio test.** The parameter-free OPH target is $\Delta_2/\Delta_1 = \varphi^2 \approx 2.618$. Across repeated sweeps on `ibm_marrakesh` and one replication on `ibm_fez`, the best `ibm_marrakesh` points land within 0.8%, 1.2%, and 1.5% of the target, while a focused high-shot rerun at $t = 0.90$ lands within 2.6%. The `ibm_fez` replication is noisier but is in the same neighborhood rather than producing a clean contradictory ratio.
- $\mathbf{S_3}$ **nonabelian ratio test.** The first nonabelian hardware run revealed a real layout-dependent bias. After reversing the qubit layout, the mitigated ratio moves from 1.8724 to 2.0299 against the OPH target 2.0000, and a direct confirmation run returns 2.0657. The nonabelian target reappears once the hardware mapping is corrected, indicating an identifiable device/layout effect instead of a structural failure of the OPH prediction.

These IBM runs are best read as hardware benchmarks for the local edge-sector picture. They show that the predicted recoverability ordering survives on two real backends, the abelian exact-ratio tests land near the parameter-free targets, and the first nonabelian test returns to the expected ratio once a diagnosed layout bias is corrected. They add real-device evidence to the simulation-based checks, but they do not validate the full theorem stack by themselves.

4.8 Status, Tests, and Scope Boundaries

4.8.1 Classification of results and dependencies

The documentary node labels D1–D12 are used only to keep the dependency boundary compact. The phase split is:

$$\text{Phase I} = (D1\text{–}D5) \cup (D7\text{–}D9), \qquad \text{Phase II} = D6 \cup D10, \qquad \text{Phase III} = D12.$$

Phase	Nodes	Meaning
I	D1–D5, D7–D9	Recovered core: fixed-cutoff overlap repair and collar structure; the Lorentz/null-modular/Einstein branch; bosonic compact gauge reconstruction; the support-visible four-dimensional Euclidean Yang–Mills form and repair-gap theorem on the compact-gauge branch; realized-branch Standard Model quotient, exact hypercharge, structural electroweak force content, $N_c = 3$, and $N_g = 3$ under the explicit matter-package and admissibility inputs.
II	D6, D10	Quantitative closures: screen-capacity closure of the same Einstein branch and the P-driven electroweak/gauge-coupling branch. These depend on declared external or branch-specific inputs.
III	D12	Continuations: flavor details beyond the stated theorem surfaces, dark-sector proposals, CMB/inflation-replacement kernels, H_0/S_8 and growth branches, heuristic baryogenesis branches, spectroscopy, hadrons, and string/worldsheet effective-description branches.

Phase I is the theorem-bearing core using the stated scaling, transport, and gauge theorem stack. Phase II supplies quantitative closures on top of that core. Phase III contains program branches whose extra ansätze are not part of the recovered-core theorem package. In particular, low-ℓ CMB kernels, parity envelopes, inflation-replacement claims, and dark/anomaly H_0/S_8 growth modifications remain continuation gates until their finite-collar source functions and likelihood contracts are supplied.

The effective low-energy summary is

$$\mathcal{L}_{\text{eff}}^{\text{OPH}} \approx \sqrt{-g}\left[\frac{1}{16\pi G}(R - 2\Lambda) + \mathcal{L}_{\text{SM}}^{\text{realized branch}}\right] + \sum_i \frac{c_i}{M_*^{\Delta_i - 4}}\mathcal{O}_i.$$

Here $\mathcal{L}_{\text{SM}}^{\text{realized branch}}$ denotes the realized $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)/\mathbb{Z}_6$ sector. The higher-dimension terms absorb UV-sensitive and continuation-level corrections outside the recovered core.

4.8.2 Structural assessment

- **Dynamics**: the GR chain runs through geometric modular covariance, the null bridge, fixed-cap generalized-entropy stationarity, and the tensor upgrade. The scaling branch uses the support-visible BW scaling theorem on the prime geometric subnet.

- **Gauge structure**: compact gauge reconstruction uses the transportable bosonic sector package. MAR plus anomaly algebra and the explicit CKM/weak-sector clauses on the same one-Higgs branch select the realized Standard Model quotient, $N_c = 3$, $N_g = 3$, and the structural electroweak force content $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \to \mathrm{U}(1)_Q$.
- **Microscopic theory**: finite quantum-link style presentations give concrete fixed-cutoff examples. They do not select a unique microscopic representative.
- **Noncentral gluing**: the crossed-module higher-gauge package closes the fixed-cutoff topological branch. The realized zero-obstruction bosonic compact-gauge branch is carried by a separate theorem stack on the ordinary or central branch.

4.8.3 Known-force and charge coverage

The recovered stack does not leave any known long-range or Standard Model gauge force unassigned:

- **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**: The hypercharge theorem, the $\mathbb{Z}_6$ quotient, and Corollary 6.6c 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.

4.8.4 Testable Items and Phenomenological Branches

The detailed prediction surfaces live in the companion papers. The synthesis-level list is:

- **Exact structural tests**: massless photon/gluons/graviton, exact color and electroweak charge assignments, exact hypercharge quantization, $N_c = 3$, $N_g = 3$, and no gauge-mediated proton decay on the product-group branch.
- **Information-theoretic gravity bound**: modular-additivity defects give explicit upper bounds on GR deviations wherever the Markov/mixing hypotheses apply.
- **Edge-sector tests**: Casimir-ratio predictions, including $\Delta_8/\Delta_3 = 9/4$ for the SU(3) edge-sector benchmark, test the heat-kernel mechanism.
- **Quantitative particle checks**: the P-driven electroweak, Higgs/top, quark, charged, and neutrino lanes are compared on their declared surfaces in Ref. [4], with W/Z compare-only, Higgs/top exact on the declared D10/D11 surface, quarks selected-class exact, charged absolute masses closed as a current-corpus no-go, and neutrinos theorem-grade on the weighted-cycle branch.
- **Continuation templates**: black-hole combs, PBH burst templates, dark-sector response laws, heuristic baryogenesis continuations, and critical-string lifts are not recovered-core tests.

4.8.5 Scope Boundaries

The recovered core does not derive the following:

- charged-lepton absolute masses on the available corpus, full flavor-labeled neutrino closure, CKM/PMNS closure, global public quark-frame classification beyond the selected class, or general Yukawa hierarchy;
- hadron masses and resonances, which require nonperturbative production computation;
- dark-sector response laws, heuristic baryogenesis continuations, strong-CP proposals, proton-spin fractions, and late-stage spectroscopy templates;
- inflation-replacement, CMB low-ℓ/parity kernels, H_0/S_8 growth modifications, and cosmological dark/anomaly Boltzmann kernels beyond the contract stated in Section 4.6.15;
- Page-curve/island closure, PBH/LIGO comb claims, or critical-superstring completion;
- a unique microscopic representative or a full fermionic/super-Tannakian gauge reconstruction.

N_{CRC} is the global capacity fixed point for the cosmological branch. P is selected on the Phase-II outer/inner closure branch and is not a Phase-I axiom. A contradiction in a Phase-III continuation retracts that continuation. A contradiction in D10 challenges the quantitative-closure branch. A contradiction in Phase I challenges the recovered-core claim set.

4.8.6 Comparison with other unification approaches

Unified models attempting to tie together QFT, gravity, and SM structure tend to encounter a repeatable set of conceptual difficulties. This subsection examines how the observer-patch holography framework addresses these common pitfalls.

1. Subsystem factorization in gauge theory and gravity.

In gauge theories and gravity, the Hilbert space does not cleanly split as "inside $\otimes$ outside" across a cut. This infects entanglement entropy definitions, area terms, edge modes, and observable identification. Many unification attempts handwave this or patch it with conventions.

Resolution: The framework builds from a net of von Neumann algebras on patches plus overlap consistency. It does not start from naïve tensor factorization. The gauge-as-gluing + regulator package yields edge-center completion: a canonical block decomposition on collars where the center captures superselection data at the cut, and the state becomes (exactly or approximately) Markov across the collar. The entropy split S(ρ_C) = S_bulk + $\langle$L_C$\rangle$ follows from having a center with sector labels. This replaces the ad hoc "add an area term" move.

2. Modular Hamiltonian nonlocality.

Many entanglement-based gravity derivations depend on modular Hamiltonians that look like local stress-tensor charges (true only in special states/regions). In generic QFT states, modular Hamiltonians are nonlocal, making "first law of entanglement $\Rightarrow$ Einstein equation" arguments fragile.

Resolution: The Markov collar condition does heavy lifting: approximate Markov implies approximate modular additivity, with the defect controlled by conditional mutual information. This makes "modular locality" a controlled approximation. It is not treated as an assumption. On the extracted prime geometric subnet of Theorem 4.2, the controlled tangent-half-space comparison then locks modular flow to geometric dilations with rigid 2π normalization.

3. Lorentz invariance as a derived output.

Discrete microscopic models generally break Lorentz symmetry, and many unified proposals simply postulate Lorentz invariance in the IR.

Resolution: Lorentz kinematics are tied to geometric modular flow on caps. On the extracted prime geometric subnet, the support-visible BW scaling theorem gives cap-pair extraction, support-readable modular covariance, ordered cut-pair rigidity, and modular flow as conformal transformations on S^2. Hence $Conf^+(S^2) \cong PSL(2,\mathbb{C}) \cong SO^+(3,1)$. No external spacetime symmetry axiom is added.

4. Dynamics beyond kinematics.

Many approaches produce emergent geometry/kinematics but stall at dynamics: why Einstein's equations (with the right coefficient) rather than some other geometric PDE?

Resolution: The framework combines the MaxEnt-selected fixed-cap generalized-entropy stationarity theorem, the derived K_C = 2πB_C structure on the extracted geometric cap branch, the internal null modular bridge identifying the half-line generator with the local null-stress charge, and the internal small-ball bridge from the geometric cap generator. The explicit scaling-limit regularity, bounded-interval transport through the separate projective branch, and tensor-upgrade conditions are stated in the Einstein branch. It does not rely only on "assume a UV CFT."

5. Gauge symmetry origin and compactness.

Most unification stories pick a gauge group and work out consequences. Emergent-gauge approaches sometimes produce noncompact groups or uncontrolled redundancies.

Resolution: Gauge symmetry is recast as redundancy in overlap identifications (gauge-as-gluing). From fixed-cutoff edge sectors, fusion, coherent refinement transport, and compatible finite multiplicity fibers, a refinement-limit bosonic tensor category is constructed; Tannaka-Krein / Doplicher-Roberts reconstruction then yields a compact group G on that bosonic branch. "Gauge symmetry" names the gluing redundancy at the conceptual level. "Compact group" is the mathematical form compatible with finite-dimensional sector/fiber-functor structure.

6. Massless photon and graviton usually hand-imposed.

Getting massless gauge bosons is easy if exact gauge invariance is assumed, but that restates the problem. Massless graviton is more delicate (mass terms, vDVZ discontinuity, strong coupling scales).

Resolution: Once gauge and diffeomorphism invariance are emergent redundancies of description (from gluing consistency / emergent geometry), hard mass terms are forbidden: "a coordinate system's Jacobian can't show up as a physical mass." These symmetry-protected zeros emerge from the same consistency machinery that gives the symmetries.

7. Global consistency, anomalies, and loop patching.

Building physics from local patches hits loop/holonomy problems: consistent gluing on a tree but obstructions around loops. These obstructions are often anomalies or global topological constraints.

Resolution: This is elevated to a first-class organizing principle: gluing data on overlaps define cocycles; central defects define a Čech obstruction class [z] (and more generally a 2-group/crossed-module cocycle for noncentral defects). "Global consistency exists iff the obstruction class vanishes" becomes the universal statement. Anomalies become "failure to glue" rather than a mysterious quantum pathology.

8. Charge quantization without a GUT.

Without embedding into a simple GUT group, explaining charge quantization (why all isolated color singlets are integer charged) is awkward. Standard lore requires grand unification or monopoles.

Resolution: The framework leans on global group structure (the Z_6 quotient) and derives congruence/selection rules for allowed representations/hypercharges. This gives a structural explanation for integer-charged color singlets without introducing the proton-decay channel of simple-GUT models.

9. Coupling unification usually forces proton decay.

Traditional simple-group unification introduces leptoquark gauge bosons (X, Y) mediating proton decay. Experiment keeps pushing limits up, pressuring minimal GUTs.

Resolution: The retained D10 discussion is geometric/entropic only at the calibration level: shared edge diffusion data, heat-kernel weights, printed running/matching/threshold/scheme conventions, and extra calibration assumptions can mimic unification-style running without embedding into a simple Lie group. If the reconstructed gauge group genuinely factorizes as a product (sector factorization selector), there are no mixed generators playing the X/Y role. "Unification-like couplings" and "group unification" therefore come apart.

10. Cosmological constant locality.

The cosmological constant problem is a graveyard of unified theories: local QFT estimates are enormous, and tiny observed Λ seems to demand absurd fine tuning.

Resolution: From null modular data, T_ab is reconstructed only up to ϕg_ab. Local consistency conditions and null focusing are blind to vacuum-energy shifts, so the Einstein equation is fixed only up to Λg_ab. Λ becomes a global "capacity" parameter of the static patch (tied to log dim H_tot). It is not a locally computable quantity. This resolves the conceptual tension: local microphysics *cannot* fix Λ by structural information-theoretic reasons.

11. UV infinities and nonrenormalizability.

Unified programs struggle to give sharp, finite microscopic definitions. Formal continuum structures, infinite entropies, and regularization dependence abound.

Resolution: The regulator construction uses local patch algebras that are type-I and finite-dimensional, with a MaxEnt branch whose generator is quasi-local and obeys a Lieb-Robinson bound. So the fixed-cutoff UV branch is interacting in the ordinary finite-range sense, and the fundamental degrees of freedom are finite and live on the screen. The framework does *not* claim a unique microscopic UV completion: physical uniqueness is only modulo gauge or implementation hiding together with inert ancillary stabilization. The genuinely noncentral topological branch is also closed at fixed cutoff by the higher-gauge crossed-module collar theorem, while the support-visible BW scaling branch is closed by theorem and the realized zero-obstruction bosonic compact-gauge branch is carried by its declared theorem stack.

12. Predictivity vs. parameter explosion.

Unified models often explode in parameters, sectors, or vacua, becoming hard to test directly because everything depends on choices.

Resolution: The framework compresses freedom into a "pixel area" (resolution) parameter and a total Hilbert space capacity (size) parameter, then derives structure from consistency (Lorentz, Einstein form, compact gauge group reconstruction, exact zeros, quantization patterns). The explicit MAR selector then picks the SM factors and sector factorization on the admissible class discussed in Section 6.2.

Structural pattern. The framework treats locality, Lorentz invariance, gauge symmetry, and gravity as consequences of consistency conditions among overlapping descriptions together with information-theoretic properties of states. Modular rigidity then supplies the familiar symmetry and dynamical structures.

Engineering deliverables. Certain problems can be stated as explicit closure tasks:

- Λ is structurally explained as a global capacity parameter; the input-free numerical prediction is the self-closure fixed-point statement
- A full microphysical derivation of geometric modular action is required

These are shared challenges across unification approaches. The framework provides an explicit map of where they live and what would resolve them.

5 Consensus, Defects, and Implementation Hiding

The fixed-cutoff consensus package admits an equivalent finite patch-net formulation. Local patch descriptions agree on overlaps, local recovery-derived repair rules remove mismatches, and the physical output is the schedule-independent quotient normal form together with the induced terminal expectation functionals on the declared physical observable algebras. This section records the fixed-point, defect, and gauge-quotient statements in that language.

The consensus lane also carries a simple finite-candidate law-selection model: once one equips candidate repair rules with a fitness functional, replicator dynamics gives a clean toy picture in which more successful reconciliation laws dominate a competing pool. That selection picture is useful for the overall synthesis picture, but in the consensus paper it is a finite-candidate monotonicity result, not part of the recovered-core gravity/gauge theorem surface, not a universality claim, and not a literal cosmological dynamics claim.

On its declared fixed-cutoff branch, *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics* [2] imports four core D1 surfaces into this synthesis surface:

1. **Constraint-code firewall.** A bare finite overlap net is a finite constraint code: its codewords are exactly the globally consistent states $C = \Phi^{-1}(0)$. It is not automatically a QECC/topological code, and its graph min-cut does not determine code distance. Distance/min-cut, Knill–Laflamme resilience, spectral convergence, BFT liveness, and hardware speedup enter only through separate certificates.

2. **Asynchronous confluence.** For the declared accepted repair law, the local-fit contract makes Φ a Lyapunov functional, hence gives termination on the finite patch net. The fixed-cutoff union-collar gluing package supplies the local diamond on the physical quotient, and only that confluence condition together with repair completeness yields a unique schedule-independent normal form from a fixed initial quotient state. Same-boundary uniqueness requires the additional unique consistent extension condition in the preserved boundary/sector fiber.

3. **Cycle obstruction and higher-gauge defects.** On the abelian branch, global consistency holds exactly when cycle holonomy vanishes; on the genuinely noncentral branch, the crossed-module defect class
$$q_\Sigma \in \check{H}^2(N_\Sigma, H_\Sigma \to G_\Sigma)$$
labels the fixed-cutoff higher-gauge obstruction.

4. **Gauge quotient and observable-level confluence.** The repair law descends to the overlap-invariant quotient, and the induced terminal state on the declared physical observable algebras is unique there even when microscopic representatives differ by gauge or sector relabelings inside one quotient-local glued state.

5. **Record algebra and stability.** On the declared fixed-cutoff observer-accessible operator surface, central record projectors carry the Born/Lüders rules directly, while approximate record projectors inherit explicit $(\varepsilon, \delta_{\mathrm{rec}})$ stability bounds on that same event surface. This is a theorem about the stated algebraic record interface, not a requirement that OPH first rebuild every mathematical ingredient from operational records alone.

The consensus paper is equally explicit about the imported repair-law data and what sits outside the fixed-cutoff theorem stack. The declared repair step includes the touched-overlap local-fit contract and the union-collar gluing package on the physical quotient; the branch conditions above that step are repair completeness and, on the Petz branch, the stated support/CPTP clause. On that same fixed-cutoff surface, exact normal-form computation is finite-state and decidable, automatic approximate stability is only collar-local through the splice and record estimates, and long-run noisy approximate consensus is theorem-grade only after a fair-block contraction certificate is supplied for the chosen exported patch-net family. That certificate boundary is not a dependency for the support-visible BW theorem, local Einstein branch, compact-gauge reconstruction, or realized Standard Model branch.

6 Particle-Spectrum Branch

The particle branch follows the same logical order as *Deriving the Particle Zoo from Observer Consistency* [4]. One first fixes the compact-gauge classification basis, applies MAR to the realized one-Higgs branch to identify the Standard Model quotient, and only then asks how transport-stable overlap data organize the observed particle families.

Structural particle theorem package. On the realized compact-gauge branch, the structural carrier skeleton is fixed by the compact gauge and MAR-admissibility chain:

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

This determines the exact structural zero rows carried by the gauge and gravity sectors: the photon, gluons, and graviton are massless structural carriers on the realized branch.

Proof route. The proof route is the one developed in *Recovering Relativity and the Standard Model from Observer Overlap Consistency* [1] and *Deriving the Particle Zoo from Observer Consistency* [4]. Under the explicit coherent bosonic refinement ladder, symmetry, and finite-fiber conditions of the compact paper, the surviving transportable edge-sector category reconstructs a compact internal symmetry group. MAR then selects the realized admissible branch, anomaly freedom fixes the hypercharge lattice, the minimal coupled carrier fixes $N_c = 3$, and the same one-Higgs quark branch fixes $N_g = 3$. The particle branch builds on that structural skeleton and does not bypass it.

Sector	Claim tier	What is fixed here	Boundary to stronger claim
Structural carriers	realized-branch structural theorem	exact $m_\gamma = m_g = m_{\mathrm{grav}} = 0$; realized quotient; $N_c = 3$; $N_g = 3$	theorem-grade structural exactness on the particle side
Electroweak bosons	validation branch + exact sidecar	target-free continuation rows plus the exact frozen-repair pair $M_W = 80.377$ GeV, $M_Z =$ 91.18797809193725 GeV	status ledger records the validation scope
Fine-structure endpoint	fixed-point derivation + source audit trunk	$\alpha^{-1}(0) =$ 137.035999177(21) with $P =$ 1.63096820940395932487927984778264894 1...; source audit value $\alpha^{-1}_{\mathrm{cand}} =$ 136.99483516462164945794999458578719326 2029	status ledger records the endpoint residual package

Sector	Claim tier	What is fixed here	Boundary to stronger claim
Higgs/top stage	exact source-only D11 split theorem on the declared D10/D11 running, matching, and threshold surface + compare-only exact sidecar	shared scalar ρ_{HT}, source-only split selectors, and the declared D11 Jacobian together with an exact Higgs/top inverse pair	the split theorem is closed on the declared quantitative surface; the inverse pair is compare-only; the auxiliary direct-top PDG codomain is compare-only and is closed as a corpus-limited no-go
Quark family	selected-class theorem + exact supporting surfaces	exact running quark sextet together with explicit exact forward Yukawas on the public physical quark frame class f_P	selected-class closure on f_P; supporting exact surfaces and the separate mass bridge are carried by Ref. [4]; global classification of all public quark frame classes is closed as a corpus-limited no-go; the physical strong-CP invariant and its vanishing are work in progress
Charged leptons	continuation branch + exact witness	exact same-family charged triple together with exact same-carrier centered readback on a target-anchored witness	same-family-only on the exact witness; theorem lane carries the closed common-shift no-go, and the determinant-line lift applies only on theorem-grade physical charged data; the theorem-grade uncentered trace lift is not emitted
Neutrinos	weighted-cycle continuation + emitted bridge invariant + diagnostic sidecars	emitted absolute masses/splittings on the weighted-cycle branch through the bridge-rigidity invariant $C_\nu = G^2 Q S^{-1/2}$ above the emitted proxy P_ν, with induced paper-facing amplitude B_ν	theorem-grade on the declared weighted-cycle branch; the exact positive-segment adapter, bridge corridor, and correction audit are diagnostic-only and do not feed back into theorem state
Hadrons	backend-gated nonperturbative continuation	stable-channel and readout architecture are defined	source-only masses require one production backend export bundle, then executed unquenching, runtime receipt, and production systematics

Claim-tier boundary. Structural carrier content and the realized Standard Model branch are theorem-bearing outputs. The electroweak D10 branch, the Higgs/top stage, quark rows, charged-lepton rows, neutrino rows, and hadron rows sit on their stated quantitative-closure, continuation, or nonperturbative claim tiers. On the target-free electroweak D10 lane, the logical order is forward: the outer/inner closure relation fixes P, and then P fixes $\alpha_U(P)$, equivalently $t_U(P)$ and $t_{\mathrm{tr}}(P)$, then the downstream scale data and public electroweak rows. Measured electroweak quantities are compare-only validation targets and do not feed back into the transmutation solve. Source payloads, same-scheme remainders, interval certificates, and empirical hadron records document public validation and endpoint bookkeeping; they do not supply the arithmetic bridge from P to the fine-structure endpoint.

Screen-normalized self-consistency for total capacity. The global counterpart to the local P-closure is the cosmic record-closure capacity,

$$N_{\mathrm{CRC}},$$

the active horizon capacity at which the universe can contain observers who read back the same boundary capacity that supports them. This synthesis paper records the derivation here, next to the pixel derivation, because the two quantitative scales play parallel roles. The local equation for P says that one screen cell has the same outside geometric reading and inside electromagnetic reading. The global equation for N_{CRC} says that the whole horizon has the same outside supplied capacity and inside observer-readable cosmic record capacity.

For a trial entropy capacity N, the outside global readout is the de Sitter curvature term

$$\Lambda_{\text{out}}(N) = \frac{3\pi}{GN}.$$

That trial value defines a candidate OPH universe: horizon size, accessible records, dilution rate, checkpoint continuation, observer-supporting capacity, and late-time de Sitter boundary. Observers inside that candidate universe do not see an external parameter N directly. They infer boundary capacity from internal geometry, causal accessibility, stable records, and self-read closure. The corresponding capacity readback map is

$$F(N) := N_{\text{read}}(N),$$

where $F(N)$ is the active cosmic record capacity read back by observers inside the universe supplied with capacity N. The global closure equation is

$$\boxed{N_{\text{CRC}} = F(N_{\text{CRC}}).}$$

In words, N_{CRC} is the only global capacity at which the universe can read back its own boundary without deficit or slack.

At fixed cutoff r, the finite OPH machinery makes this map explicit. Let $\mathfrak{U}_{r,N}$ be the finite OPH universe candidate with supplied active boundary capacity N. The consensus branch supplies a quotient normal-form map $\text{nf}_{r,N}$ on the declared terminating/confluent repair surface. Let Obs select the stable self-reading observer sector, and let Cap_{read} return the capacity reconstructed by that sector. Then the finite readback map has the schematic form

$$F_r(N) = \text{Cap}_{\text{read}}\left(\text{Obs}\left(\text{nf}_{r,N}(\mathfrak{U}_{r,N})\right)\right).$$

The refinement-limit map F is the cofinal limit of this finite readback construction when that limit exists on the admissible capacity interval.

The uniqueness proof is the direct global analogue of the local P-fixed-point proof. Let I be the admissible interval of active horizon capacities. If the OPH-derived readback map $F : I \to I$ satisfies

$$0 \leq F'(N) \leq \kappa < 1 \qquad (N \in I)$$

and the observed branch satisfies $F(N_{\text{obs}}) = N_{\text{obs}}$, then N_{obs} is the unique stable fixed point. Indeed, if N_1 and N_2 are two fixed points, then

$$|N_1 - N_2| = |F(N_1) - F(N_2)| \leq \kappa |N_1 - N_2|.$$

Since $\kappa < 1$, this implies $N_1 = N_2$. Banach iteration $N_{k+1} = F(N_k)$ also converges to the same fixed point from any starting value in I.

The physical meaning of the contraction condition is that changing supplied raw capacity changes observer-readable active capacity by less than one-for-one. Extra Hilbert-space room does not automatically become public record capacity. It must be reachable, recordable, self-readable, stable under repair, coupled to observer continuation, and visible on the quotient. Capacity is physical capacity only insofar as the universe can read it.

The existing self-closing normal-form count is the finite-count representation of the same closure target. For a candidate entropy capacity N, let $\mathcal{H}_{\partial,N}$ be the observer-facing horizon record Hilbert space with $\log\dim \mathcal{H}_{\partial,N} = N$, and let $\Omega^{\text{sc}}_{r,N}$ be the terminal normal forms that are repair-closed, observer/checkpoint-supporting, locally recovered-core closed, and whose own horizon record surface reads back capacity N. Define

$$A_r(N) := |\Omega^{\mathrm{sc}}_{r,N}|, \qquad \Pi_r(N) := \frac{A_r(N)}{\dim \mathcal{H}_{\partial,N}} = A_r(N)e^{-N}.$$

The subtraction by N in

$$\log \Pi_r(N) = \log |\Omega^{\mathrm{sc}}_{r,N}| - N$$

is not a tunable penalty; it is exactly division by the full screen Hilbert-space size e^N. Thus the count-density selector is

$$N_{r,\star} = \mathrm{MAR}\arg\max_N \left[\log |\Omega^{\mathrm{sc}}_{r,N}| - N\right].$$

In the refinement limit this becomes

$$N_\star = \mathrm{MAR}\arg\max_N \left[\log |\Omega^{\mathrm{sc}}_{N}| - N\right].$$

When the limiting pressure

$$\ell(N) := \log |\Omega^{\mathrm{sc}}_{N}| - N$$

is differentiable, the count-density proof can also be written as the stationarity fixed point

$$T_\eta(N) := N + \eta \ell'(N), \qquad T_\eta(N) = N \Longleftrightarrow \ell'(N) = 0.$$

The proof certificate is the derivative-sign certificate for

$$H_N(N) := \ell'(N),$$

namely $H_N(N_0) = 0$ and $H_N'(N) < 0$ on the admissible interval, or the stronger concavity bound

$$-M \le \ell''(N) \le -m < 0.$$

This certificate is another way to express uniqueness of the global self-read closure. A numerical solver may locate a candidate root, but uniqueness and stability come from the OPH-derived readback or pressure certificate. Software checks, when used, check the finite grammar or transfer-function certificate; they are not a brute-force enumeration of cosmological normal forms and are not the proof by themselves.

The observed de Sitter read-off is

$$N_{\mathrm{obs}} = S_{\mathrm{dS}} = \frac{A_{\mathrm{dS}}}{4\ell_P^2} = \frac{3\pi}{\Lambda \ell_P^2} \simeq 3.31 \times 10^{122}.$$

On the observed branch,

$$N_{\mathrm{CRC}} = N_\star = N_{\mathrm{obs}}, \qquad \Lambda_{\mathrm{CRC}} = \frac{3\pi}{G N_{\mathrm{CRC}}}.$$

The fixed-point certificate makes this value the unique stable capacity of the self-reading solution we inhabit, not a freely tunable creator-set parameter. Informally, N_{CRC} is the single horizon capacity at which the outside total horizon capacity and the inside observer-accessible public record agree.

Outer/inner self-consistency for the pixel ratio. This synthesis paper is the canonical place where the local pixel ratio P is defined, so the compact and particle papers can cite one derivation rather than duplicate it. The construction has six steps.

1. **Outer pixel ratio.** The outer-side local screen datum is the dimensionless area ratio

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

2. **Exact 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 self-similar balance condition

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

gives $x^2 - x - 1 = 0$ and hence $x = \varphi = (1+\sqrt{5})/2$. The realized branch must sit near, but not exactly at, this point because exact equilibrium is too symmetric to support durable records, structure, and dynamics.

3. **Outer detuning variable.** On the declared Phase-II closure surface we parametrize the outer-side detuning by

$$\alpha_{\text{ext}}(P) := \frac{P - \varphi}{\sqrt{\pi}}, \qquad \text{equivalently} \qquad P = \varphi + \alpha_{\text{ext}}(P)\sqrt{\pi}.$$

4. **Source-only D10 anchor from the same pixel.** Let $I \subset \mathbb{R}_+$ be the physical pixel interval and let

$$F_{D10} : I \longrightarrow \mathcal{D}_{D10}$$

denote the declared D10 source map. For any trial $P \in I$, the forward D10 lane emits

$$P \longmapsto \alpha_U(P) \longmapsto (t_U(P), t_{\text{tr}}(P)) \longmapsto (t_2(P), t_3(P), v(P)) \longmapsto A_Z(P) = \alpha_{\text{em}}^{-1}(m_Z^2; P).$$

5. **Ward-projected Thomson transport.** Let T_Q be the declared Ward-projected $\mathrm{U}(1)_Q$ transport operator from the m_Z-anchor to the Thomson limit. This is the endpoint functional that reads the fixed point as $\alpha^{-1}(0) = 137.035999177(21)$. Define

$$A_T(P) := T_Q\big(A_Z(P), F_{D10}(P)\big) = \alpha_{\text{Th}}^{-1}(P).$$

The inner coupling used in the outer/inner closure is the coupling, not the inverse coupling:

$$\alpha_{\text{in}}(P) := \frac{1}{A_T(P)}.$$

6. **Closure.** The realized pixel ratio is the fixed point for which the outer detuning equals the inner observation coupling:

$$\alpha_{\text{ext}}(P) = \alpha_{\text{in}}(P).$$

Equivalently,

$$H(P) := P - \varphi - \frac{\sqrt{\pi}}{A_T(P)} = 0, \qquad \text{or} \qquad P = \varphi + \frac{\sqrt{\pi}}{A_T(P)} = \varphi + \alpha_{\text{in}}(P)\sqrt{\pi}.$$

Writing

$$G(P) := \varphi + \frac{\sqrt{\pi}}{A_T(P)},$$

the closure problem is the fixed-point condition $G(P) = P$. On any physical interval where this induced map is a self-map and a contraction, the closure root is locally unique.

The closure interpretation assigns one screen cell two coupled roles. On the outer side it is one cell of the screen whose displacement above the exact self-similar equilibrium φ sets the geometric detuning. On the inner side the same cell emits the smallest electromagnetic observation step available on the particle branch. The pixel ratio P is the size of that cell in Planck areas, and the realized P is the value at which those outer and inner readings agree. The fixed-point equation is the mathematical form of that consistency condition. The public root is

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

The value is forced on the declared branch because the same screen cell must be one object with two compatible descriptions. The geometric detuning fixes the outer side, the Ward-projected Thomson endpoint fixes the inner electromagnetic side, and the fixed point is the sole shared value. The status table records the source-audit trunk, endpoint residual, and interval-certificate disclosures. Once the fixed point $P_\star$ is selected, every downstream quantity is propagated forward from that same $P_\star$. A separate hardware note reports an optical-cavity check of the same fixed-point geometry; this is treated as corroborating engineering evidence. The same closure surface also admits an inverse observational use case in which observers infer P and total screen capacity from inside the world.

One clean refinement studies the same closure directly at zero momentum. That route would close the electromagnetic row without transporting from the m_Z anchor. The present paper keeps the fixed-point idea itself simple and treats that zero-momentum promotion as a separate quantitative task.

Exact non-hadron hit surface. For quick orientation, the non-hadron exact-output surface consists of structural zero rows, the electroweak exact sidecar, the Higgs/top exact sidecar, a same-family charged witness, the selected-class quark theorem with exact forward Yukawas on f_P, and the weighted-cycle neutrino branch. This synthesis paper uses those facts as lane markers. Ref. [4] carries the exact values, supporting surfaces, diagnostic sidecars, and per-sector caveats.

Extended derivation. The structural-to-family route and the full per-sector claim-tier audit are developed in *Deriving the Particle Zoo from Observer Consistency* [4].

7 Screen Microphysics, Records, and Observer Continuation

The screen-microphysics lane is the engineering side of the paper corpus. It asks whether some microscopic model might exist in principle and gives one explicit regulated carrier architecture: a federation of finite observer patches with echosahedral multi-port interfaces, recurrent toroidal subchannels, exposed overlap data, records, repair instruments, and observer-facing interfaces all made concrete at fixed cutoff. A spherical screen is used there only as an observer-facing regulator chart for support-visible cuts, not as a claim that the universe is a literal spherical quantum computer.

This lane has five theorem-bearing components. First, the reference architecture itself is explicit. Second, the regulated patch-net embedding theorem is closed on its fixed-cutoff object/local-interface surface, and the edge-sector thermal/Casimir branch carries an explicit fixed-cutoff law with one compact lift above that theorem surface. Third, the fixed-cutoff measurement/Born-rule package is closed on its declared operator and record-event surface. Fourth, the fixed-cutoff Bell/CHSH package is closed on the declared two-wing surface. Fifth, the checkpoint/restoration observer package is closed. The downstream boundaries on this lane are packet-level quotient closure where one wants an autonomous exported dynamics, compact-group/Peter–Weyl bookkeeping on the D10/compact surfaces, fair-block contraction certificates for long-run noisy approximate consensus on exported packet nets, and model-family universality. The support-visible BW / geometric-modular / local-Einstein closure lives on the companion recovered-core papers rather than on this fixed-cutoff engineering lane. A reconstruction of Hilbert, C^*- or von Neumann algebra structure, Born probabilities, trace structure, and entropy from operational records alone remains a separate question outside this synthesis lane.

Reference-architecture takeaway. The synthesis-level point is that the regulated screen-side architecture supplies a concrete habitat in which local observables, overlap observables, record registers, and synchronization maps can all be written without pretending to identify a unique final UV completion. Hardware evidence is not imported by this synthesis theorem surface unless it appears in a public, hash-stable OPH evidence bundle.

Imported theorem from Ref. [3] (Fixed-cutoff record algebra and Born-Lüders package). For each completed compare/write/verify slice of the regulated microphysics, the declared pointer and overlap-sector projectors generate a finite commutative central record algebra

$$\mathcal{Z}_{\mathrm{rec}}.$$

Every observer-accessible event E in that algebra has probability

$$\mathbb{P}(E) = \mathrm{Tr}(\rho P_E),$$

and conditioning on E gives the operational post-measurement state

$$\rho|_E = \frac{P_E \rho P_E}{\mathrm{Tr}(\rho P_E)}.$$

If a practical readout instead uses projectors Q_a with commuting central reference projectors $\widehat{Q}_a$ on the same declared slots and

$$\delta_{\mathrm{rec}} = \max_a \|Q_a - \widehat{Q}_a\|,$$

then

$$\|[Q_a, Q_b]\| \le 4\,\delta_{\mathrm{rec}},$$

and any accessible-state perturbation $\|\widetilde{\rho} - \rho\|_1 \le \varepsilon$ changes each declared elementary record-event probability by at most $\varepsilon + \delta_{\mathrm{rec}}$.

Imported theorem from Ref. [3] (Fixed-cutoff Bell / CHSH package). On the declared fixed-cutoff physical observable algebra, fix commuting left/right wing subalgebras, central binary setting registers, binary projective readouts on each wing, and a physical source state on the joint wing algebra. Then the compare slice carries the joint law

$$p(a, b \mid x, y) = \mathrm{Tr}\big(\rho_{LR}\, P^{(L)}_{a|x} P^{(R)}_{b|y}\big),$$

the local marginals are independent of the remote setting, the Bell correlators satisfy

$$E(x, y) = \mathrm{Tr}(\rho_{LR}\, A_x B_y),$$

and the CHSH combination obeys

$$\big|E(0,0) + E(0,1) + E(1,0) - E(1,1)\big| \leq 2\sqrt{2}.$$

If the source family contains an explicit two-qubit branch with the stated Pauli readouts, the same fixed-cutoff surface saturates $2\sqrt{2}$ exactly on that branch. The source-state input is explicit: the theorem derives the Bell law and the Tsirelson bound from the stated fixed-cutoff operator surface, while Bell-pair preparation and two-qubit factor structure are stated branch conditions.

Imported theorem from Ref. [3] (Checkpoint/restoration and observer backup). For an observer patch P_O, let the checkpoint data consist of the observer-facing record algebra, the accessible local state, the future update schedule, and the externally visible overlap interface data. Exact restoration reproduces the full future law of observer-accessible events, while an ε-accurate restoration changes that future law by at most ε in total variation.

Observer-facing conclusion. The theorem-level content is modest but sharp: the paper corpus supplies a fixed-cutoff measurement interface, a fixed-cutoff checkpoint/restoration package, and an operational observer-identity criterion on observer-accessible event algebras. Stronger substrate-selection or strange-loop closure claims are outside the recovered theorem package.

Extended derivation. The fixed-cutoff edge-law, measurement, Born-rule, Bell/CHSH, and checkpoint/restoration packages are developed in *Federated Echosahedral Screen Microphysics: Patch Hardware, Records, and Observer Synchronization in OPH* [3].

8 Worldsheet/String Branch

The synthesis paper carries the string/worldsheet branch on the same theorem boundary used across the paper corpus: a genuine theorem-level bridge from the edge-sector theorem surface to the two-dimensional Yang–Mills partition function, followed by a controlled large-N_{edge} effective-description theorem on the stated branch. This two-dimensional heat-kernel partition statement and declared-branch large-edge effective description sit beside the compact paper's support-visible compact-gauge four-dimensional Yang–Mills form theorem and repair-gap theorem.

Imported theorem (Heat-kernel edge-sector bridge). Import the fixed-cutoff same-overlap edge-law package from *Federated Echosahedral Screen Microphysics* [3]. Under the same overlap-gauge realization and local-Gibbs/MaxEnt edge branch, the edge-sector weights satisfy

$$p_R(t) \propto d_R e^{-tC_2(R)}.$$

Consequently the edge partition function matches the standard two-dimensional Yang-Mills heat-kernel form. The compact paper carries OS/Wightman-strength continuum reconstruction for the four-dimensional Euclidean Yang–Mills form and the OPH mass-gap claim separately on the support-visible compact-gauge branch.

Imported proof route. The microphysics source surface carries the fixed-cutoff thermal/Casimir law on the declared overlap interface. Peter–Weyl decomposition [32] then identifies the resulting sum with the standard heat-kernel expression for two-dimensional Yang-Mills on the compact-group refinement lift.

Compact-carrier theorem. Writing the closed edge partition function as

$$Z_{\text{edge}}(t) = \sum_R d_R^2 e^{-tC_2(R)},$$

the compact carrier states the exact bridge itself as a named theorem: Peter–Weyl identifies it with the compact-group heat kernel at the identity,

$$Z_{\text{edge}}(t) = K_t(1).$$

The Chapman–Kolmogorov gluing law for K_t is the precise sense in which collar sewing matches the two-dimensional Yang–Mills heat-kernel surface before any large-N_{edge} continuation is invoked. The compact carrier is a partition-identity theorem on its stated branch. The four-dimensional OPH claim identifies the compact-gauge branch with the Euclidean Yang–Mills form and applies the repair-dynamics theorem on that same support-visible branch.

Controlled large-N_{edge} branch. Fix a distinct large-N_{edge} realization, with $N_{\text{edge}} \neq N_c = 3$, and the fixed-$\tau$ variable

$$\tau = tN_{\text{edge}}$$

on a compact window I. If the resulting edge free energy satisfies the compact paper's large-N_{edge} criterion, with remainder control on I, then the Gross-Taylor rewriting [31] is a controlled theorem-level worldsheet effective description of the edge dynamics on that branch. Critical-superstring claims, worldsheet CFT closure, anomaly cancellation, and full massless-spectrum matching are outside the recovered core theorem package.

Higher-gauge continuation slot. The genuinely noncentral crossed-module gluing data on cuts provide the natural continuation-level slot for B-field / gerbe data in any string-style reorganization of the OPH edge sector. This is a structural analogy only; it is not an open/closed-string boundary formalism or a critical-string theorem.

9 Cross-Lane Theorem Boundary and Continuation Directions

The paper corpus has the following theorem boundary:

1. The fixed-cutoff collar, higher-gauge, consensus, record, Bell/CHSH, checkpoint, and restoration packages are theorem-bearing on their stated finite-regulator surfaces.
2. Physical UV uniqueness is quotient-level and stable under inert ancillas. OPH fixes terminal values on declared physical observables, not a unique microscopic representative.
3. The Lorentz/null-modular/Einstein chain is closed on the support-visible refinement/scaling branch by the BW scaling theorem on the prime geometric subnet: cap-pair extraction, fixed-collar Markov replacement with carried r_{FR}, δ^{M}, and regularized modular remainders, regularized modular transport, support-readable modular covariance, round-cap rigidity, and KMS/BW normalization make the cap modular automorphism geometric. This is not a finite-cell Lorentz-invariance claim or a one-shot upgrade from small CMI to exact Markov geometry.

4. The cosmological constant branch uses the cosmic record-closure fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$ and fixes $\Lambda_{\mathrm{CRC}} = 3\pi/(GN_{\mathrm{CRC}})$ on the same Einstein branch. Its count-density representation is $N_\star = \mathrm{MAR}\arg\max_N(\log|\Omega_N^{\mathrm{sc}}| - N)$. On the observed branch this readback fixed point is the de Sitter entropy capacity.

5. Compact gauge reconstruction uses the theorem-produced transportability criterion, fixed-cutoff bosonic sector category, and bosonic refinement ladder to reconstruct a compact gauge group from zero-obstruction sectors. That is the classification stage, not the Standard Model selector. The realized MAR-admissible witness data plus the explicit one-Higgs matter package then select the realized Standard Model quotient, exact hypercharge, $N_c = 3$, and $N_g = 3$.

6. The particle branch is sector-split: electroweak W/Z is a validation pair, Higgs/top is exact on its declared quantitative surface, selected-class quarks close on their declared public frame class, weighted-cycle neutrinos sit on their declared theorem surface, strong CP is work in progress, charged-lepton source landing from P to physical charged data is a corpus-limited no-go boundary, the auxiliary direct-top PDG row is compare-only, and source-only hadron masses require a production OPH backend.

7. The string/worldsheet branch is a controlled heat-kernel continuation branch.

Local unification status. At the synthesis level, one local unification surface is explicit and one uniform theorem tier for c, G, W, Z, and H is absent. The same declared local input P governs the electroweak/Higgs trunk and the gravity-side edge-entropy bridge, while the invariant causal speed c is structural from the Lorentz branch and gains its SI label only on the local readout package.

Quantity	Best value on the paper surface	Common paper-surface chain	Caveat
c	299792458 m/s	structural Lorentz/BW branch $\to$ invariant causal speed $\to$ local SI readout	structural rather than P-fitted; the SI presentation sits on the stated BW/Lorentz support package
G	$6.674299995910528 \times 10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$ on the stated local extension surface	$P \to \bar{\ell}_{\mathrm{SU}(2)} + \bar{\ell}_{\mathrm{SU}(3)} = \bar{\ell}_{\mathrm{shared}} \to G_{\mathrm{SI}} = c^3 a_{\mathrm{cell}}/(\hbar P)$	stated relative to the declared a_{cell} datum; on that surface the D10 pixel law fixes $\bar{\ell}_{\mathrm{shared}} = P/4$, and the strict classical-regime clause is explicit
M_W	target-free continuation row 80.37700001539531 GeV; exact compare-only sidecar 80.377 GeV	$P \to \alpha_U(P) \to (t_U(P), t_{\mathrm{tr}}(P)) \to v(P) \to$ target-free electroweak repair chart $\to M_W$	compare-only validation sidecar; public validation records are the certified P root, source spectral measure payload, same-scheme remainder, and interval certificate
M_Z	target-free continuation row 91.18797807794321 GeV; exact compare-only sidecar 91.18797809193725 GeV	same target-free electroweak trunk	same compare-only validation caveat
M_H	exact D11 split-theorem pair $M_H = 125.1995304097179$ GeV, $m_t = 172.3523553288312$ GeV	same P-driven electroweak core $\to$ D10 repair chart $\to$ ρ_{HT} and source-only split selectors $\to$ D11 Jacobian	Higgs row is exact on the declared running, matching, and threshold surface; the inverse sidecar is compare-only; the auxiliary direct-top PDG row is a compare-only codomain with a corpus-limited no-go boundary

The exact-release package on that surface is local and finite: the strict classical-regime clause and the target-free electroweak chart identity for W/Z. The familiar-unit display split on that same surface is

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

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

with dimensionless $\widehat{L}, \widehat{T}, \widehat{E}, \widehat{\Theta}$. The gravity-side shared-edge identity and the local SI readout are therefore fixed relative to the declared microscopic datum a_{cell}, while meters, seconds, GeV, and Kelvin are downstream display conventions built from the structural c row and the familiar constants $\hbar$, k_B.

Companion papers. Detailed derivations appear in Refs. [1, 2, 3, 4].

10 Conclusion

This paper presents the full framework on one synthesis surface. The five-axiom basis, the derived gravity and gauge branches, the consensus formulation, the particle-spectrum branch, the regulated screen-microphysics lane, the fixed-cutoff measurement and observer package, and the controlled large-N_{edge} worldsheet effective-description branch appear on one common synthesis surface.

The corpus is uneven but concrete. It has a sharp fixed-cutoff local picture, a derived recovered core for relativity and Standard-Model structure, real particle-side quantitative outputs, and a usable microscopic engineering lane. The scope boundaries are equally explicit: bounded-interval projective inputs in the null-stress branch, the cosmic record-capacity fixed point for the global branch constant, microscopic representative uniqueness, public validation artifacts for the D10 electroweak rows, strong CP, full flavor/CKM closure, charged-lepton source landing from $P_\star$, and source-only hadron masses without a production OPH backend. Refs. [1, 2, 3, 4] are the depth surfaces for sector-by-sector proofs, quantitative details, and continuation-level material. The local cross-lane numerical surface is explicit enough to summarize concisely: the fixed point $P_\star$ organizes the bosonic mass trunk and the gravity-side shared edge-entropy identity beneath the emitted local G row, while the same familiar-unit package reads meters and seconds from the single ruler $\sqrt{a_{\text{cell}}}$ together with the structural invariant speed c, and reads GeV and Kelvin from the inverse local ruler through $\hbar$ and k_B.

A Candidate Microphysics Reference Architecture

The synthesis paper keeps the observer-facing theorem package in the main text, but the concrete federated patch-carrier architecture belongs in an appendix. The theorem-bearing fixed-cutoff source surface for the edge heat-kernel / Casimir leaf is the microphysics companion paper; the appendix records only the synthesis-level restatement imported from that surface. The reference architecture used here is a federation of finite observer patches with echosahedral multi-port interfaces, boundary-fixed patch algebras, boundary-fixed overlap algebras on collars, declared shared readout packets, local repair instruments, and record registers exposed on the observer-facing interface. A spherical screen is a regulator chart for support-visible cuts, not a literal substrate claim.

Theorem (Regulated patch-net embedding). On a fixed finite patch federation and declared finite overlap cover, the regulated microphysics construction realizes an explicit finite patch-net layer: patch vertices are genuine finite observer patches, overlap edges are genuine collars or port interfaces, packet fields are actual readouts of declared overlap observables, and allowed repair branches are concrete local channels on the same finite regulator.

Proof sketch. All patch and overlap algebras are finite-dimensional at fixed cutoff, and the declared packet fields are readouts of explicit overlap observables in those algebras. The candidate repair menu is implemented by local finite-support channels on the same neighborhoods. Therefore the abstract finite patch-net syntax is embedded into an explicit regulator object rather than introduced as a separate formal layer.

Theorem (Fixed-cutoff edge heat-kernel branch). On the thermalized overlap branch of the same regulated architecture, the stationary sector law is

$$\pi_\beta(\alpha) \propto d_\alpha e^{-\beta C_2(\alpha)},$$

and along refinement ladders converging to a compact-group Peter–Weyl decomposition [32] the same weights converge cylinderwise to the heat-kernel/Casimir law used in the OPH edge-sector branch.

Proof sketch. The regulated proposal kernel is chosen so that the weighted proposal symmetry $d_\alpha q_{\alpha\beta} = d_\beta q_{\beta\alpha}$ holds. Detailed balance with the Casimir acceptance rule then fixes the stationary distribution in finite dimension, and the compact-group heat-kernel law is the corresponding refinement lift.

Extended derivation. The full fixed-cutoff architecture, validation package, and theorem stacks are developed in *Federated Echosahedral Screen Microphysics* [3].

B Interpretive Epilogue: State-and-Law Habitat

This appendix records the OPH state-and-law habitat available inside the framework. A strange-loop closure theorem is outside this appendix.

The OPH inputs used here are:

1. the patch net $P \mapsto A(P)$ of von Neumann algebras together with isotony and overlap restriction maps;
2. the global state on the inductive-limit algebra, whose restrictions witness a nonempty overlap-consistent local sector;
3. the finite Axiom-3 MaxEnt/refinement branch, which supplies a common finite family of gauge-invariant local constraint observables across cutoffs;
4. on the fixed-cutoff branch, the derived finite type-I presentations and compact boundary gluing groups used elsewhere in the OPH package.

The law slot is encoded by the finite expectation-value coordinates of the retained Axiom-3 constraint family. That choice furnishes a compact convex state-and-law habitat for the framework.

Theorem (Internal state-and-law habitat theorem). Fix finitely many screen patches $P_1, \dots, P_N$ in the standard OPH setup. For each i, let

$$M_i := A(P_i), \qquad M_{ij} := A(P_i \cap P_j),$$

where M_i is the patch von Neumann algebra and $M_{ij} \subset M_i, M_j$ is the overlap algebra. Let

$$S_i := S(M_i) \subset M_i^*$$

be the full state space of M_i, endowed with the weak* topology $\sigma(M_i^*, M_i)$.

Choose a finite family of self-adjoint gauge-invariant local constraint observables

$$\mathcal{O} = \{O_1, \ldots, O_m\}$$

from the Axiom-3 MaxEnt branch, with each O_a supported in one of the chosen patches; write $i(a)$ for an index such that $O_a \in M_{i(a)}$.

Define the overlap-consistent state sector

$$X_{\text{ov}} := \left\{ (\omega_1, \ldots, \omega_N) \in \prod_{i=1}^{N} S_i : \omega_i|_{M_{ij}} = \omega_j|_{M_{ij}} \text{ for all } i, j \right\}.$$

Define the OPH law-coordinate map

$$c : X_{\text{ov}} \to \mathbb{R}^m, \qquad c(\omega_1, \ldots, \omega_N) := \big(\omega_{i(1)}(O_1), \ldots, \omega_{i(m)}(O_m)\big).$$

Let

$$X_{\mathcal{O}} := \left\{ (\omega_1, \ldots, \omega_N, \ell) \in \left(\prod_{i=1}^{N} S_i \right) \times \mathbb{R}^m : (\omega_1, \ldots, \omega_N) \in X_{\text{ov}},\ \ell = c(\omega_1, \ldots, \omega_N) \right\}.$$

Then:

1. the ambient product
$$E := \left(\prod_{i=1}^{N} M_i^* \right) \times \mathbb{R}^m$$
is a Banach space for any product norm;

2. $X_{\mathcal{O}}$ is nonempty and convex;

3. $X_{\mathcal{O}}$ is norm-closed in E;

4. $X_{\mathcal{O}}$ is compact in the product topology
$$\tau := \left(\prod_{i=1}^{N} \sigma(M_i^*, M_i) \right) \times \text{Euclidean topology on } \mathbb{R}^m;$$

Proof. Each M_i is a von Neumann algebra in the OPH patch net, hence a Banach space, so each dual M_i^* is Banach. A finite product of Banach spaces is Banach, hence so is E.

For each i, the state space

$$S_i = \{\omega \in M_i^* : \omega \geq 0,\ \omega(1) = 1\}$$

is convex and weak* compact by Banach–Alaoglu, because it is a weak* closed subset of the dual unit ball. For every pair (i, j), restriction along $M_{ij} \subset M_i$ and $M_{ij} \subset M_j$ defines affine weak* continuous maps

$$r_{ij} : S_i \to S(M_{ij}), \qquad r_{ji} : S_j \to S(M_{ij}).$$

Hence X_{ov} is an intersection of affine equalizer sets, so it is convex and τ-closed inside $\prod_i S_i$. Since $\prod_i S_i$ is compact and convex in the product weak* topology, X_{ov} is compact and convex as well.

Nonemptiness comes from the OPH global state on the inductive-limit algebra: its restrictions

$$(\omega|_{M_1}, \dots, \omega|_{M_N})$$

belong to X_{ov} because the restrictions agree on every overlap by construction.

Each coordinate $\omega_{i(a)} \mapsto \omega_{i(a)}(O_a)$ is affine and weak* continuous, since evaluation at a fixed algebra element is weak* continuous. Thus c is affine and τ-continuous. The graph map

$$G : X_{\mathrm{ov}} \to E, \qquad G(x) := (x, c(x)),$$

is affine and continuous, so its image $X_{\mathcal{O}} = G(X_{\mathrm{ov}})$ is convex and τ-compact.

To see norm-closedness in E, let $(x_n, c(x_n)) \to (x, \ell)$ in norm. Norm convergence implies weak* convergence on each coordinate, so $x \in X_{\mathrm{ov}}$ because X_{ov} is weak* closed. Since c is weak* continuous, $c(x_n) \to c(x)$, while norm convergence in $\mathbb{R}^m$ also gives $c(x_n) \to \ell$. Hence $\ell = c(x)$, so $(x, \ell) \in X_{\mathcal{O}}$.

Finally, $X_{\mathcal{O}}$ is a compact convex subset of the locally convex topological vector space E equipped with τ. □

The additional inputs needed for a strange-loop closure theorem are:

1. a closure map T built from internal operations such as overlap repair, Axiom-3 MaxEnt reprojection, collar recoverability, and record-sector coarse-graining;
2. a proof that the stronger "observer-supporting" or "selected-world" criteria carve out a nonempty T-invariant observer-supporting subset of $X_{\mathcal{O}}$;
3. any uniqueness, contraction, or Lyapunov-type stability estimate for that closure map.

The habitat theorem supplies only the ambient Banach/compact-convex setting. It does not define a canonical internal self-map, it does not identify any nonempty invariant observer-supporting subset, and it does not prove uniqueness or stability. In particular, nonemptiness and compact-convexity of the ambient habitat do not by themselves imply the existence of a nonempty invariant observer-supporting sector. This does not erase the exact finite packet branch proved on the consensus surface: there the quotient normal-form map pushes forward to a continuous affine idempotent self-map of the finite packet simplex, with fixed points exactly the packets supported on quotient normal forms. Any further strange-loop reading of the ambient habitat as a closure story for existence is interpretive only and is not part of the recovered theorem package.

B.1 Additional Problem Closures

Using the phase split of Section 4.8.1, the broader closure claims read as follows. The table also records appendix-tier closure stories so they cannot float outside the official status table. In this table, "Proved corollary of earlier theorems" means proved under the dependencies named in the row, including declared branch conditions and cited companion-paper theorem surfaces.

Claim	Status	Boundary
Quantum gravity consistency	Proved corollary of earlier theorems	Lorentz and Jacobson-type Einstein branches follow from the support-visible BW scaling theorem together with the stated null-bridge and fixed-cap conditions.
UV completion / microscopic uniqueness	Separate scope boundary	Fixed-cutoff theorem packages exist, but a unique microscopic representative is not identified. The support-visible BW scaling theorem and the realized zero-obstruction bosonic compact-gauge branch are closed on their declared companion-paper theorem surfaces; the separate boundary is microscopic representative uniqueness rather than a missing recovered-core lift.
Measurement problem	Separate scope boundary	Ref. [3] closes the fixed-cutoff central-record / Born-Lüders interface, but a full philosophical or continuum-level measurement closure is not derived here.
Cosmological principle and horizon homogeneity	Separate scope boundary	MaxEnt symmetry inheritance is suggestive only; an independent derivation of cosmological isotropy and homogeneity sits outside the recovered-core theorems.
Observer-relative modular time	Proved corollary of earlier theorems	On the BW_{S^2} geometric branch, observer-relative time is read from the cap modular automorphism parameter of the observer's accessible algebra/state pair. This is the D3 reading only.
Global / operational problem of time	Separate scope boundary	OPH does not prove a separate operational-clock theorem or a global solution of the problem of time.
Black-hole information paradox	Separate scope boundary	Edge-center sectorization and recoverability motivate an internal resolution template, but the full black-hole-information branch sits outside Phase I.
D6 cosmological-capacity closure at the cosmic record fixed point	Proved corollary of earlier theorems	At the fixed point $N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}})$, with observed-branch readout $N_{\mathrm{CRC}} = S_{\mathrm{dS}}$, the same Einstein branch fixes $\Lambda_{\mathrm{CRC}} = 3\pi/(GN_{\mathrm{CRC}})$ and the static-patch readout; $N_\star$ is the count-density representation of the same closure.
Bare cosmological-constant problem from local null data	Separate scope boundary	The local null-data route determines the Einstein branch only modulo Λg_{ab} and does not derive the screen-capacity identification from bare OPH axioms.
Magnetic monopole expectation from simple GUT groups	Proved corollary of earlier theorems	The realized product-group branch $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)/\mathbb{Z}_6$ removes the standard simple-GUT symmetry-breaking monopole channel.
Proton stability against gauge-mediated decay	Proved corollary of earlier theorems	The realized product-group branch removes the standard simple-GUT gauge-mediated proton-decay channel.
Proton spin fraction	Separate scope boundary	The proton-spin claim requires nonperturbative QCD completion beyond Phase I.
Dark matter phenomenology	Separate scope boundary	Modular-anomaly response laws are conjectural beyond the recovered gravity chain.
Baryon asymmetry scale	Separate scope boundary	Suppression-counting estimates are not a substitute for a derived out-of-equilibrium mechanism.
Three generations	Proved corollary of earlier theorems	$N_g = 3$ is recovered on the realized MAR-admissible branch.
Downstream flavor structure, Koide, and charged-lepton fits	Separate scope boundary	Companion flavor constructions require extra ansätze and sit outside the recovered core.
Edge-to-2D-YM / controlled large-N_{edge} worldsheet branch	Proved corollary of earlier theorems	On the stated overlap-gauge, local-Gibbs/MaxEnt, compact-group, and large-N_{edge} conditions, the edge-sector heat-kernel bridge and controlled Gross–Taylor worldsheet effective description are theorem-level. This is a two-dimensional partition and worldsheet-effective branch, separate from the four-dimensional compact-gauge theorem.
Support-visible compact-gauge Yang–Mills form and mass gap	Proved corollary of earlier theorems	On the ordinary or central zero-obstruction compact-gauge branch with a four-dimensional Euclidean scaling chart, reflection-positive vacuum, active exact-Markov repair collars, bounded-color collar covers, repair completeness, and support-visible continuum extraction, OPH derives the Euclidean Yang–Mills action, local repair is conditional expectation, Euclidean transfer is the repair generator, and the exact accounting identity is $\Delta_{\mathrm{YM}} = \Delta_{\mathrm{rep}}$.
Critical superstring / worldsheet CFT lift	Separate scope boundary	Critical-superstring claims, worldsheet CFT closure, anomaly cancellation, and full massless-spectrum matching require extra ingredients outside the recovered core.
Why anything exists / strange-loop closure	Separate scope boundary	Appendix B provides only the habitat theorem and fixed-point setting. The OPH closure map, invariant observer-supporting sector, and uniqueness/stability proofs are not part of this theorem package.
Fixed-cutoff check-point/restoration/backup	Proved corollary of earlier theorems	Ref. [3] proves same-interface checkpoint/restoration and backup for observer-accessible event algebras with explicit future-law error control.
Substrate-transfer / stronger observer continuation	Separate scope boundary	Stronger substrate-selection, redesigned-environment continuation, strange-loop closure, uniqueness, and stability

C Observer Continuation and Backup

This appendix uses the fixed-cutoff checkpoint/restoration/error/identity theorem package proved in Ref. [3] and summarizes its algebraic interface in observer language. It does not define a strange-loop closure map on an invariant observer-supporting sector or establish uniqueness and stability for such a map. Nor does it upgrade the fixed-cutoff theorem stack to continuation across redesigned environments: the proved package is same-interface restoration on observer-accessible event algebras with explicit error control.

C.1 Observer as Algebraic Pattern

Here, an observer is represented by

$$O = (P, \mathcal{A}(P), \rho, R),$$

where P is a screen patch, $\mathcal{A}(P)$ is its local algebra, ρ is the local state, and R is the record algebra. Records are carried exactly by central record projectors and, on practical readout surfaces, by approximately commuting projectors in overlap centers, so they are shareable without violating no-cloning constraints for generic quantum states.

Ref. [3] proves a fixed-cutoff observer-facing measurement interface: a finite central record algebra, Born probabilities for its event projectors, and Lüders conditioning on that same commuting record algebra. This appendix uses that fixed-cutoff measurement package in observer language. If a practical readout uses projectors that are $\delta_{\rm rec}$-close in operator norm to that central reference algebra and the accessible state is perturbed by at most ε in trace norm, each declared elementary record-event probability shifts by at most $\varepsilon + \delta_{\rm rec}$.

C.2 Markov Collar Factorization

For a collar tripartition A-B-D, the exact Markov normal form is used only when $I(A:D|B) = 0$ holds literally, or along a controlled fixed-collar family for which the exact-Markov replacement modulus $\delta^{\rm M}_{A:B:D}(\varepsilon) \to 0$. In that exact or controlled limit one obtains

$$\rho_{ABD} = \bigoplus_\alpha p_\alpha \, \rho^{(\alpha)}_{Ab^\alpha_L} \otimes \rho^{(\alpha)}_{b^\alpha_R D}.$$

The sector label α is classical center data. This decomposition is the mathematical basis for extracting an interior observer state with a controlled boundary interface. Small CMI by itself supplies a recovered comparison state; it is not silently promoted to an exact Markov state.

C.3 Checkpoint and Restoration Map

A checkpoint is the tuple

$$\mathcal{C} = (R, \alpha, \rho^{(\alpha)}_{\rm int}),$$

with $\rho^{(\alpha)}_{\rm int}$ the interior state after fixing α. Given a compatible target environment state $\sigma^{(\alpha)}_{\rm env}$, a restored state is

$$\rho^{(\alpha)}_{\rm new} = \rho^{(\alpha)}_{\rm int} \otimes \sigma^{(\alpha)}_{\rm env}.$$

For approximate Markov collars, recovery is controlled by the standard bound

$$\|\rho_{ABD} - ({\rm id}_A \otimes \mathcal{R}_{B\to BD})(\rho_{AB})\|_1 \le 2\sqrt{\ln 2\, \varepsilon}.$$

This gives quantitative trace-distance control on the recovered comparison state under finite CMI (finite collar error). The exact HJPW splice form instead requires literal exact Markovity or the separate fixed-collar replacement modulus tending to zero. Interpreting either control as observer continuation requires additional modeling assumptions beyond the bound itself.

The fixed-cutoff microphysics claim is stronger than the bare recoverability estimate written above. In Ref. [3], exact checkpoint restoration preserves the full future law of observer-accessible events, while an ε-accurate restoration changes that future law by at most ε in total variation. The stronger strange-loop / substrate-transfer closure story is outside this theorem.

C.4 Physical Meaning

At fixed cutoff, Ref. [3] proves a checkpoint/restoration theorem and backup corollary for observer-accessible event algebras. The proved package stops at same-interface backup/restoration with explicit future-law error control; stronger continuation and substrate-selection claims are outside it. This appendix states the observer-facing meaning of that theorem and its algebraic prerequisites. The stronger substrate-selection, strange-loop closure, uniqueness, and stability package is separate from the fixed-cutoff theorem.

D Cosmology, Horizons, and Modular-Anomaly Continuations

This appendix carries the cosmology- and continuation-facing technical statements for the synthesis paper. They do not enlarge the recovered core; they make the structural continuation boundary explicit on the synthesis surface.

D.1 Cosmological Principle

Lemma 4.1 (MaxEnt symmetry inheritance). Let G act on the screen by automorphisms α_g. If the constraint set is G-invariant and von Neumann entropy is G-invariant, then the MaxEnt optimizer ω can be taken G-invariant. If the optimizer is unique, it is automatically G-invariant.

Theorem 4.2 (Isotropy from SO(3)-invariant constraints). Under SO(3)-invariant constraints and MaxEnt uniqueness, the reference state satisfies

$$\omega \circ \alpha_g = \omega \qquad \forall g \in \mathrm{SO}(3).$$

The corresponding stress tensor has perfect-fluid form:

$$\langle T_{ab} \rangle = \rho\, u_a u_b + p(g_{ab} + u_a u_b).$$

Theorem 4.3 (Schur-type homogeneity). Let (Σ, h_{ij}) be a three-dimensional Riemannian manifold. If at every point $p \in \Sigma$, the curvature tensor is SO(3)-invariant, then

$$R_{ijkl}(p) = K(p)(h_{ik}h_{jl} - h_{il}h_{jk}),$$

and the second Bianchi identity forces $\nabla_m K = 0$, so K is constant.

Corollary 4.4. If OPH supplies isotropy for all observers, spatial geometry is a constant-curvature space form: S^3, $\mathbb{R}^3$, or H^3.

Theorem 4.5 (FLRW emergence). With the semiclassical Einstein equation from recoverable generalized entropy / entanglement equilibrium, a perfect-fluid stress tensor from SO(3)-invariant MaxEnt, and positive Λ from finite screen capacity, the emergent geometry is FLRW:

$$ds^2 = -dt^2 + a(t)^2 \left(\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right),$$

with Friedmann equations

$$H^2 + \frac{k}{a^2} = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3}, \qquad \dot{H} - \frac{k}{a^2} = -4\pi G(\rho + p).$$

D.2 Horizon-Problem Control

Theorem 6.1 (Anisotropy bound from Markov control). For a tripartition A-B-D with collar width δ, the Markov error satisfies

$$I(A : D \mid B) \lesssim c \cdot |\partial C|_{\text{UV}} \cdot e^{-\delta/\xi},$$

where ξ is the correlation length.

Corollary 6.2. For bounded observable O with $\|O\| \sim 1$,

$$|\Delta\langle O \rangle| \leq 2\sqrt{\ln 2 \cdot \varepsilon}.$$

On the CMB anisotropy target $\delta T/T \lesssim 10^{-5}$, this gives the familiar ε-scale bound used in the horizon-problem bookkeeping:

$$\varepsilon \lesssim 3.61 \times 10^{-11} \text{ bits}, \qquad \frac{\delta}{\xi} \gtrsim \ln\!\left(\frac{c \cdot |\partial C|_{\text{UV}}}{\varepsilon_{\max}}\right) \approx 24.$$

With $\xi \approx \sqrt{a_{\text{cell}}}\,\ell_P \approx 1.28\,\ell_P$, the corresponding collar benchmark is

$$\delta_{\text{CMB}} \approx 31\,\ell_P.$$

Theorem 6.3 (Homogeneity as the MaxEnt default). If MaxEnt constraints are uniform across UV cells with no marked points, the MaxEnt state is invariant under triangulation automorphisms. In the continuum limit this gives SO(3) invariance.

Reading rule. The cosmological-principle and horizon-homogeneity package is a branch-local theorem package on the realized symmetric MaxEnt branch. The exact theorem-level content is:

1. SO(3)-invariant constraint data plus MaxEnt uniqueness force an isotropic reference state and a perfect-fluid stress tensor.
2. If the same isotropy condition holds for all observers, the Schur/Bianchi argument forces constant-curvature spatial slices.
3. Combined with the semiclassical Einstein branch and positive Λ from finite screen capacity, the metric takes FLRW form.
4. Markov control supplies an explicit patch-overlap anisotropy bound and the collar benchmark used in the CMB homogeneity bookkeeping.

The package does not prove from the bare OPH axioms alone that the realized cosmological branch must satisfy SO(3)-invariant constraint data, that MaxEnt uniqueness holds in every cosmological sector, or that the observed universe saturates the displayed CMB benchmark.

D.3 Cosmological-Constant / Screen-Capacity Closure

Theorem 6.4 (D5–D6 local/global closure stack). On the gravity lane, assume the local Einstein branch has been recovered only modulo Λg_{ab}, 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

$$S_{\rm dS} = \frac{A_{\rm dS}}{4G} = \frac{3\pi}{G\Lambda},$$

and the standard de Sitter static-patch formulas

$$r_{\rm dS} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\rm dS}}{c}.$$

Then:

1. local null data determine the Einstein branch only modulo Λg_{ab};
2. the same branch closes globally as

$$G_{ab} + \frac{3\pi}{GN_{\rm CRC}}\, g_{ab} = 8\pi G\, \langle T_{ab} \rangle;$$

3. the same D6 closure fixes

$$S_{\rm dS} = N_{\rm CRC}, \qquad A_{\rm dS} = 4GN_{\rm CRC}, \qquad r_{\rm dS} = \sqrt{\frac{3}{\Lambda}}, \qquad t_\Lambda = \frac{r_{\rm dS}}{c};$$

4. the observed cosmic age is a downstream FLRW benchmark rather than an additional theorem output.

So the cosmological-constant package is one local/global theorem stack rather than a split local argument plus an unrelated global readout.

Scope boundary. The D6 hypotheses are the local Einstein branch, the cosmic record-capacity fixed point $N_{\rm CRC} = F(N_{\rm CRC})$, its observed-branch de Sitter entropy readout, the 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.

Input-free capacity closure. The zero-input global capacity is the cosmic record-closure fixed point. For candidate capacity N, let $F(N)$ be the active horizon capacity read back by stable observers inside the OPH universe supplied with capacity N:

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

The target is

$$N_{\rm CRC} = F(N_{\rm CRC}), \qquad \Lambda_{\rm CRC} = \frac{3\pi}{GN_{\rm CRC}}.$$

The screen-normalized density of terminal self-closing observer normal forms is the count representation of the same target. Let Ω_N^{sc} be the terminal OPH normal forms that are repair-closed, observer- and checkpoint-supporting, carry the recovered local package, and whose own horizon record surface reads back capacity N. Since $\log\dim\mathcal{H}_{\partial,N} = N$, define

$$\Pi(N) := \frac{|\Omega_N^{\mathrm{sc}}|}{\dim\mathcal{H}_{\partial,N}} = |\Omega_N^{\mathrm{sc}}|e^{-N}.$$

The corresponding selector 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 in the synthesis section. Informally, this is the single balance point where the universe reads back its own boundary without deficit or slack. On the observed branch,

$$N_{\mathrm{CRC}} = N_\star \simeq 3.31 \times 10^{122}.$$

Capacity normalization. The branch uses N_{scr} as the de Sitter entropy capacity. The bare horizon ratio is

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

so

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

For the observed late-time scale, $N_{\mathrm{patch}} \simeq 1.05 \times 10^{122}$ and $N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}$.

D.4 Black-Hole Structural Package and Continuation Boundary

The AMPS-style information trilemma assumes:

1. an outgoing mode B entangled with an interior partner A (smooth horizon),
2. late radiation B entangled with early radiation R (unitarity), and
3. monogamy of entanglement.

The false assumption in this setting is the naive tensor factorization

$$\mathcal{H} \stackrel{?}{=} \mathcal{H}_{\mathrm{inside}} \otimes \mathcal{H}_{\mathrm{outside}} \otimes \mathcal{H}_{\mathrm{radiation}}.$$

Theorem 7.1 (Edge-center black-hole decomposition). On the fixed-cutoff edge-center collar branch, assume the black-hole collar state lies in the exact Markov class or in the idealized exact-recoverability limit used elsewhere in the paper stack. Then the collar decomposition gives

$$\rho_{A_\delta B_\delta D_\delta} = \bigoplus_\alpha p_\alpha \left(\rho_{A_\delta b_L^\alpha} \otimes \rho_{b_R^\alpha D_\delta}\right).$$

The glue between inside and outside is the edge-sector label α in the center. Given α, inside and outside factorize.

Theorem 7.2 (Hawking/KMS normalization). On the geometric modular branch where the horizon generator carries the standard 2π KMS normalization, the outside algebra is in a KMS state at inverse temperature

$$\beta = \frac{2\pi}{\kappa}$$

with respect to the horizon Killing generator.

Corollary 7.3 (Recoverability-style interior encoding). On the same fixed-cutoff collar carrier, when $I(A_\delta : D_\delta \mid B_\delta) \le \varepsilon$, there exists a recovery channel $\mathcal{R}_{B_\delta \to A_\delta B_\delta}$ such that

$$\|\rho_{A_\delta B_\delta D_\delta} - (\mathcal{R}_{B_\delta \to A_\delta B_\delta} \otimes \mathrm{id}_{D_\delta})(\rho_{B_\delta D_\delta})\|_1 \le 2\sqrt{\ln 2 \cdot \varepsilon}.$$

So the interior collar data are encoded in the exterior collar together with the shared cut data, rather than supplied by an independent tensor factor.

Corollary 7.4 (Discrete area spectrum and Schwarzschild transition identity). The same edge-center package carries a discrete area operator

$$L_C = \sum_\alpha (\log d_\alpha) P_\alpha, \qquad A_\alpha = 4G \log d_\alpha = 4\ell_P^2 \ln d_\alpha,$$

and for a Schwarzschild sector transition $d \to d'$ one has

$$\Delta(Mc^2) = k_B T_H \ln(d'/d).$$

On the emission branch $d' = d/k$ with $k > 1$, the emitted line energy magnitude is

$$E_{\text{emit}} = k_B T_H \ln k.$$

Scope boundary. The retained structural theorem package is exactly Theorem 7.1, Theorem 7.2, Corollary 7.3, and Corollary 7.4. Integer-transition combs, QNM selectors, Page-type linewidth estimates, PBH burst templates, Kerr/LIGO horizon spectroscopy templates, and any Page-curve or island closure are outside that structural core. Those items require continuation inputs such as an integer-transition selection rule, a semiclassical evaporation-power model, a QNM/transition identification, greybody matching, or an explicit island prescription.

D.5 Modular-Anomaly Continuation

Theorem 9.1 (Modular additivity defect). Define the collar modular additivity defect by

$$\Delta K_\delta := K_{ABD} - K_{AB} - K_{BD} + K_B.$$

Then

$$\langle \Delta K_\delta \rangle_\omega = -I(A : D \mid B)_\omega.$$

This is the precise structural identity underneath the modular-anomaly dark-sector continuation. By itself it does not close a dark-matter theorem.

The corresponding continuation packages the anomalous contribution into the modified Einstein equation

$$G_{ab} + \Lambda g_{ab} = 8\pi G \left(\langle T_{ab} \rangle + \langle T_{ab}^{\text{anom}} \rangle \right),$$

with rest-frame normalization

$$\langle T_{00}^{\text{anom}} \rangle := \frac{15}{8\pi^2} \frac{\delta \langle K_C^{\text{anom}} \rangle}{\ell^4}.$$

On that continuation surface the anomaly gravitates, is dark by construction, is covariantly conserved, and is effectively classical because it is organized by central/recoverability data.

If one tentatively identifies the imported D6 static-patch radius r_{dS} as the unique IR scale available to the continuation, the natural benchmark acceleration is

$$a_0^{(\text{OPH})} := \frac{15}{8\pi^2} c^2 \sqrt{\frac{\Lambda}{3}} = \frac{15}{8\pi^2} \frac{c^2}{r_{\text{dS}}} \approx 1.03 \times 10^{-10} \text{ m/s}^2,$$

close to the empirical MOND scale. That numerical proximity is only a benchmark unless a separate source/response law and a controlled galaxy-scale limit are derived. An illustrative MOND/RAR-style response would read

$$g_{\text{obs}} \approx g_b + \sqrt{a_0 g_b},$$

but that functional form is not derived from the OPH axioms.

Reading rule. The anomaly term gravitates, is dark by construction, and is covariantly conserved, but the synthesis surface keeps the entire dark-sector package on the D12 continuation tier rather than promoting any MOND/RAR-style law, galaxy-rotation template, or cluster phenomenology to theorem status. The imported D6 static-patch scale fixes only the benchmark normalization. The missing inputs are the controlled nonrelativistic limit, derived sign and closure control for the effective anomaly contribution, a derived source/response law, proof of the relevant dominant galaxy-scale source, lensing and cluster phenomenology, cosmological abundance and structure-formation checks, and environment/stability bounds. Those are continuation-level branches above the identity written here.

References

[1] B. Müller, A. Osika, K. Xue, and P. Nguyen, *Recovering Relativity and the Standard Model from Observer Overlap Consistency*, 2026. Available at GitHub PDF.

[2] B. Müller, K. Xue, and K. A. Anirudha, *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics*, 2026. Available at `https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/reality_as_consensus_protocol.pdf`.

[3] B. Müller, A. Osika, K. Xue, and B. Cassie, *Federated Echosahedral Screen Microphysics: Patch Hardware, Records, and Observer Synchronization in OPH*, 2026. Available at `https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/screen_microphysics_and_observer_synchronization.pdf`.

[4] 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`.

[5] M. H. A. Newman, "On theories with a combinatorial definition of 'equivalence'," *Ann. of Math. (2)* **43** (1942), 223–243.

[6] E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," *Commun. Math. Phys.* **28** (1972), 251–257.

[7] J. J. Bisognano and E. H. Wichmann, "On the duality condition for a Hermitian scalar field," *J. Math. Phys.* **16** (1975), 985–1007.

[8] J. J. Bisognano and E. H. Wichmann, "On the duality condition for quantum fields," *J. Math. Phys.* **17** (1976), 303–321.

[9] 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.

[10] H.-W. Wiesbrock, "Half-Sided modular inclusions of von-Neumann-Algebras," *Commun. Math. Phys.* **157** (1993), 83–92.

[11] W. G. Unruh, "Notes on black-hole evaporation," *Phys. Rev. D* **14** (1976), 870–892.

[12] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Phys. Rev. Lett.* **75** (1995), 1260–1263, arXiv:gr-qc/9504004.

[13] T. Jacobson, "Entanglement equilibrium and the Einstein equation," *Phys. Rev. Lett.* **116** (2016), 201101, arXiv:1505.04753.

[14] G. W. Gibbons and S. W. Hawking, "Cosmological event horizons, thermodynamics, and particle creation," *Phys. Rev. D* **15** (1977), 2738–2751.

[15] Planck Collaboration, "Planck 2018 results. VI. Cosmological parameters," *Astron. Astrophys.* **641** (2020), A6, arXiv:1807.06209.

[16] 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.

[17] E. H. Lieb and M. B. Ruskai, "Proof of the strong subadditivity of quantum-mechanical entropy," *J. Math. Phys.* **14** (1973), 1938–1941.

[18] D. Petz, "Sufficient subalgebras and the relative entropy of states of a von Neumann algebra," *Commun. Math. Phys.* **105** (1986), 123–131.

[19] D. Petz, "Sufficiency of channels over von Neumann algebras," *Quart. J. Math.* **39** (1988), 97–108.

[20] O. Fawzi and R. Renner, "Quantum conditional mutual information and approximate Markov chains," *Commun. Math. Phys.* **340** (2015), 575–611, arXiv:1410.0664.

[21] 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.

[22] S. Doplicher and J. E. Roberts, "A new duality theory for compact groups," *Invent. Math.* **98** (1989), 157–218.

[23] 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.

[24] T. Tannaka, "Über den Dualitätssatz der nichtkommutativen topologischen Gruppen," *Tohoku Math. J.* **45** (1938), 1–12.

[25] M. G. Krein, "A principle of duality for a bicompact group and a square block algebra," *Dokl. Akad. Nauk SSSR* **69** (1949), 725–728.

[26] H. Georgi and S. L. Glashow, "Unity of all elementary-particle forces," *Phys. Rev. Lett.* **32** (1974), 438–441.

[27] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak interactions with lepton-hadron symmetry," *Phys. Rev. D* **2** (1970), 1285–1292.

[28] E. Witten, "An SU(2) anomaly," *Phys. Lett. B* **117** (1982), 324–328.

[29] S. Dimopoulos, S. Raby, and F. Wilczek, "Supersymmetry and the scale of unification," *Phys. Rev. D* **24** (1981), 1681–1683.

[30] 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.

[31] D. J. Gross and W. Taylor, "Two-dimensional QCD is a string theory," *Nucl. Phys. B* **400** (1993), 181–208, arXiv:hep-th/9301068.

[32] F. Peter and H. Weyl, "Die Vollständigkeit der primitiven Darstellungen einer geschlossenen kontinuierlichen Gruppe," *Math. Ann.* **97** (1927), 737–755.

[33] 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.

[34] 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.

[35] W. Donnelly and A. C. Wall, "Entanglement entropy of electromagnetic edge modes," *Phys. Rev. Lett.* **114** (2015), 111603, arXiv:1412.1895.

[36] 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.

[37] 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.

[38] 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.

[39] Clay Mathematics Institute, "Yang–Mills & the Mass Gap." Available at `https://www.claymath.org/millennium/yang-mills-the-maths-gap/`.

[40] 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`.

[41] K. Osterwalder and R. Schrader, "Axioms for Euclidean Green's functions," *Communications in Mathematical Physics* **31** (1973), 83–112.

[42] K. Osterwalder and R. Schrader, "Axioms for Euclidean Green's functions II," *Communications in Mathematical Physics* **42** (1975), 281–305.

[43] Particle Data Group, "Review of Particle Physics," *Phys. Rev. D* **110** (2024), 030001. Available at https://pdg.lbl.gov/2024/download/db2024.pdf.

[44] Particle Data Group, "Electroweak Model and Constraints on New Physics," Available at https://pdg.lbl.gov/2024/reviews/rpp2024-rev-standard-model.pdf.