

Recovering Observer Spacetime and Einstein Dynamics from Overlap Consistency

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

July 26, 2026

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

Abstract

Observer-Patch Holography starts with finite patches that keep records, read their shared boundaries, and repair disagreement. This paper gives the full spacetime and gravity chain. A receipt-selected incidence limit produces the round celestial two-sphere and its cap geometry. Modular covariance gives Lorentz kinematics and the three-dimensional space of observer frames. Record germs produce a four-dimensional Lorentzian event manifold only under the stated population, separation, response-chart, cone, and causal receipts. Positive null translations and modular charges reconstruct a local conserved stress tensor. A fixed-cap generalized-entropy identity then gives the Einstein equation under one common scaling family, universal coupling, a vacuum reference, and independent scale readouts. The finite incidence and algebraic identities are source-derived; construction of a physical tower carrying every continuum receipt is work in progress.

Derivation map

The complete implication chain is

finite observer records and repaired quotient normal forms
⇒ receipt-selected round S^2 and cap normals
⇒ geometric modular flow, Lorentz action, and H^3 observer frames
⇒ conditional four-dimensional event manifold
⇒ positive translations and local conserved $\langle T_{ab} \rangle$
⇒ $G_{ab} + \Lambda g_{ab} = 8\pi G \langle T_{ab} \rangle$.

Each arrow below states its own finite, analytic, scaling, or physical antecedents. The celestial S^2 and the frame hyperboloid do not by themselves populate a spacetime. The event-manifold and Einstein conclusions therefore remain conditional on their named receipts.

Contents

1	Recovered Core: Foundations and Structural Branches	3
1.1	Overview	3
1.2	Model and Axioms	3
1.3	Information-Theoretic Tools	9

2	Finite Screen Spectrum Theorem Package	21
2.1	Overlap Consistency and Gluing	47
2.2	Modular Flow and Lorentz Kinematics	52
2.3	Gravity from Fixed-Cap Generalized-Entropy Stationarity	60
2.4	Einstein-chain status and falsification boundary	91
A	Technical Details for the Specialist Derivations	91
A.1	Lorentz and Einstein Bridge	91
A.2	Cosmological-Constant / Screen-Capacity Closure	93

1 Recovered Core: Foundations and Structural Branches

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

1.1 Overview

Overlap consistency first turns local patch states into a constrained gluing problem. The fixed-cutoff branch gives collar decompositions, edge centers, and finite repair dynamics on the declared physical quotient. The continuum geometric branch reads quotient normal forms as cap, stress, diamond, tetrad, and scale data. Modular flow on caps then gives Lorentz kinematics and a three-dimensional observer-frame hyperboloid. A four-dimensional event manifold follows only from the stated record-germ population, separation, response-chart, cone, and causal receipts. Positive null translations and modular charges reconstruct the local conserved stress tensor. Fixed-cap generalized-entropy stationarity, a uniform small-diamond limit, universal coupling, a vacuum reference, and independent scale readouts then imply the Einstein equation. Bare finite consensus stops at quotient normal forms, and the cap-normal H^3 chart stops at observer frames. Neither statement populates a spacetime.

The local pixel ratio $P \equiv a_{\text{cell}}/\ell_\star^2$, the cosmic record capacity $N_\star = \log D_\star$, and the selected no- G scale certificate enter only on their declared quantitative branches. The event-manifold and Einstein theorems do not depend on the Standard Model matter or Minimal Admissible Realization branches.

1.2 Model and Axioms

1.2.1 Observers and access model

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

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

An operational observer additionally carries overlap interface algebras and restriction maps, allowed update and repair instruments, and checkpoint data used for continuation. Observers are internal self-reading patterns in the global state. A patch and its compatible marginal supply a

geometric or algebra-state chart; operational observer status begins when the interface, record, repair, and checkpoint tests pass.

1.2.2 Screen, patches, and algebra net

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

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

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

1.2.3 Five OPH axioms

This paper states the five OPH axioms because both specialist derivations share their observer-overlap foundation. Branch-local hypotheses are written at the theorem that consumes them. The support-visible BW scaling theorem handles the scaling step. The event-manifold and Einstein theorems use the additional common-domain, stress, scaling, universal-coupling, vacuum-reference, and scale-readout receipts stated in their own hypotheses. They do not use Minimal Admissible Realization.

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

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

1.2.4 Closure values and theorem-local technical data

The quantitative branches discussed here use one incomplete declared-map coordinate, one proposed capacity coordinate, and one selected no- G scale certificate:

$$P \equiv a_{\text{cell}}/\ell_\star^2, \quad N_{\text{CRC}} \equiv \log \dim \mathcal{H}_{\text{tot}}, \quad \gamma_\star \equiv \frac{\ell_\star \nu_{\text{Cs}}}{c}, \quad B_\star \equiv \frac{3\pi}{\ell_\star^2}.$$

Here $P = P_\star$ is the local UV-area ratio selected by the unique root of that declared outer/inner self-read map. No measured coupling enters that solve. Its physical Thomson-limit interpretation requires the same-scheme hadronic spectral transport. The global coordinate is defined by

the source-only correctable public-record packet and the stable whole-fiber saturation equation $\mathfrak{F}_{r,0}(D_\star) = \{D_\star\}$, with $N_\star = \log D_\star$. The independent weak/Higgs common-load bridge gives $N_{\text{bridge}} = \pi \exp[6\pi/(P\alpha_U(P))]$. That bridge is $N_{\text{bridge}} \simeq 3.532 \times 10^{122}$, while the Planck- Λ central capacity is $N_\Lambda \simeq 3.313 \times 10^{122}$, a mismatch of about 6.6 percent. Capacity-level neutrino side estimates tied to that branch value are comparison-only bookkeeping terms separate from the rejected weighted-cycle comparison lane. The selected scale certificate supplies $\gamma_\star$, with $B_\star$ as the SI curvature display of the same branch. Once supplied, $\ell_\star^2 = 3\pi/B_\star$ becomes the usual Planck-area display after $G_{\text{SI}} = c^3\ell_\star^2/\hbar$. These are quantitative-branch quantities, not extra axioms.

Observation enters as an external test of the source-defined branches. In the OPH reverse-engineering reading, the universe is a closed mathematical fixed structure, and each readback object is specified independently of its measured target. For N , this object is set-valued before whole-fiber scalarization and its uniqueness theorem is the exact finite-size slack law, not a Banach contraction. Measurements may define comparison coordinates; they construct neither closure. The local map has a certified unique root. Global uniqueness requires the physical checkpoint packet and the regulator-stable condition $s(D_\star) = 0 < s(D)$ elsewhere.

The spacetime and Einstein derivation uses the same convention as a compression estimate. A quantitative row enters that estimate only when the declared source map has no dependency path from the measured target or a calibrated proxy. If p_i upper-bounds the conditional accidental hit probability of row i after previous accepted rows, then

$$P_{\text{acc}} \leq \prod_i p_i.$$

The evidence ledger contains no registered discriminating prediction. The electroweak checks use known targets and therefore have comparison-audit status. Other positive rows have theorem, conditional, comparison-only, or non-discriminating status.

If both missing physical identifications are supplied, the two dimensionless coordinates determine dimensionless geometry, not an SI scale by themselves. On the de Sitter branch,

$$N_\star = \frac{3\pi}{\Lambda_\star \ell_\star^2}, \quad P_\star = \frac{a_{\text{cell}}}{\ell_\star^2},$$

so

$$\Lambda_\star \ell_\star^2 = \frac{3\pi}{N_\star}, \quad \Lambda_\star a_{\text{cell}} = \frac{3\pi P_\star}{N_\star}.$$

Equivalently, for the scale product $\mathcal{B}_\ell := \Lambda_\star N_\star$,

$$\mathcal{B}_\ell \ell_\star^2 = 3\pi, \quad \mathcal{B}_\ell a_{\text{cell}} = 3\pi P_\star.$$

These invariants survive the rescaling $\ell_\star^2 \mapsto \lambda^2 \ell_\star^2$, $a_{\text{cell}} \mapsto \lambda^2 a_{\text{cell}}$, and $\Lambda_\star \mapsto \lambda^{-2} \Lambda_\star$, while $\mathcal{B}_\ell \mapsto \lambda^{-2} \mathcal{B}_\ell$. Thus the SI value of $\mathcal{B}_\ell$, the SI area $\ell_\star^2$, and G_{SI} require the selected scale certificate; they are not derived from $P_\star$ and $N_\star$ alone. If the selected no- G scale certificate supplies $\gamma_\star$, equivalently $B_\star$, then

$$\gamma_\star = \frac{\ell_\star \nu_{\text{Cs}}}{c}, \quad B_\star = \mathcal{B}_\ell = \Lambda_\star N_\star, \quad \ell_\star^2 = \frac{3\pi}{B_\star}.$$

The branch certificate used for the SI gravity row is

$$B_\star = 3.60787391468 \dots \times 10^{70} \text{ m}^{-2},$$

hence

$$\ell_\star^2 = 2.61228030237 \dots \times 10^{-70} \text{ m}^2.$$

Rounded capacity displays such as $N_\Lambda \simeq 3.313 \times 10^{122}$ are useful cosmological labels, but they are not high-precision inputs for the Newton row; with the Planck-2018 curvature benchmark, using only the rounded capacity would shift the Newton row at the 8.4×10^{-4} level.

The capacity normalization uses the de Sitter static-patch entropy [26]. If

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

denotes the bare horizon area ratio, then

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

Using the Planck-2018 late-time de Sitter scale [28] gives $N_{\text{patch}} \simeq 1.05 \times 10^{122}$, $N_{\text{scr}} \simeq 3.313 \times 10^{122}$, and $\Lambda \ell_P^2 \simeq 2.845 \times 10^{-122}$.

The controlled scaling step for the Lorentz, null-modular, and Einstein statements is the support-visible BW scaling theorem. Fixed-cutoff regulator bookkeeping, quasi-local propagation, endpoint-Lipschitz control, and collar decomposition are cited from their preceding theorems.

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

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

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

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

Rotationally invariant constraint branch. Constraint sets are $\text{SO}(3)$ -invariant on S^2 .

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

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

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

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

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

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

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

Geometric-subnet BW lift. The continuum Lorentz statement is asked only on the support-visible extracted 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 the special type-I generator form is $K_C = 2\pi B_C + Z_C$ with Z_C central. The spacetime and Einstein paper proves the support-visible BW cap automorphism only when one cofinal branch carries two independent inputs: the finite cap-normal certificate with its support-order, held-out oriented cross-ratio, geometric-flow, and independently normalized 2π -KMS comparison controls, plus the complete MGNS-1 algebra-state comparison package on that same branch. It also proves the cap-normal H^3 chart: $q(\Omega) = (1, \Omega)$, $n_C = (\cot \alpha, \csc \alpha \mathbf{c})$, $n_{gC} = \Lambda_g n_C$, and $H^3 \simeq \text{SO}^+(3, 1)/\text{SO}(3)$, with cap half-spaces but no populated-bulk or preferred-observer-point consequence. The spacetime and Einstein paper's geometry-producer packet produces only the finite cap-normal certificate from repair normal forms on a computable-receipt branch (spherical incidence, disk/mesh, cross-ratio, support-flow, and 2π -normalization receipts); it does not produce MGNS-1. Bare confluence underdetermines topology, dimension, framing, state, and normalization. Its null-net-standardness and event-manifold packets supply derived positive translations and a conditional Lorentzian event manifold, and its closure packets construct the local conserved stress tensor, the corrected generalized-entropy first law, the uniform small-diamond limit, and the composed branch-entry theorem. The composition requires one source-derived common-domain tower with certified tails, universal coupling, a source-derived vacuum reference, and independent scale readouts. Construction and certification of the tower are work in progress (Theorems 4.3c–4.3h).

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

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

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

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

Approximate modular covariance on the fixed-cutoff realized branch. Under the fixed-cutoff realized presentation, Lemma 2.6 (local Gibbs), the derived quasi-local

propagation bound above, and the collar double-scaling plus mixing hypotheses of Section 2.3, the modular flow $\sigma_t^{\omega, C}$ maps $\mathcal{A}(R)$ into a slightly thickened region algebra:

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

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

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

Continuum-limit heuristic. Define the induced region flow

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

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

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

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

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

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

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

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

Fixed-cutoff realized presentation. At a UV scale ℓ_{UV} , local patch algebras are type-I with finite-dimensional Hilbert spaces, and gauge-as-gluing is realized as a boundary action on a lifted finite-dimensional presentation. Section 2.3 proves the resulting fixed-cutoff collar package directly

from overlap consistency on this realized presentation, while distinguishing the lifted fixed-point algebra from the sector-preserving collar algebra used in the edge-center (EC) theorem.

External mathematical inputs include SSA and recovery theorems (Petz 1986, 1988; Fawzi and Renner 2015). The Einstein route also uses the fixed-volume area-variation identity for small geodesic balls and the local quadratic-polarization argument in the tensor upgrade. The null modular bridge and small-ball kernel are carried internally on the compact surface rather than imported from a separate EFT small-ball law.

Recovered-core theorem boundary. The Lorentz, null-modular, and Einstein recovered core runs through the support-visible automorphism theorem. Fixed-cutoff statements remain regulator statements. Full citations appear in the References.

1.2.5 Notation

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

1.3 Information-Theoretic Tools

1.3.1 Strong subadditivity and Markov states

For any tripartite state ρ_{ABC} ,

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

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

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

1.3.2 Approximate recovery

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

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

1.3.3 Collar refinement and sufficient mechanisms

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

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

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 $I(A_\delta : D_\delta \mid B_\delta) \rightarrow 0$ in the refinement limit. At fixed regulator scale, the same finite-dimensional realized presentation supplies the patch-net and overlap-gluing data used below. We therefore isolate that realized presentation first, then separate the exact-Markov and quantitative routes.

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

$$Q_r = \Sigma_r / \Gamma_r, \quad \pi_r : \Sigma_r \rightarrow Q_r.$$

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

$$N_r = n_r(Q_r).$$

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

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

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

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

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

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

$$(\mathcal{A}_{r,b}^{\text{phys}}, \tau_{r,b}^0), \quad \mathcal{A}_{r,b}^{\text{phys}} = z_{r,b} B(\tilde{\mathcal{H}}_r)^{G_r} z_{r,b},$$

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

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

If

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

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

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

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

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

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

and

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

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

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

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

is idempotent:

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

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

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

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

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

subject to

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

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

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

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

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

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

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

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

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

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

If the one-step defects are

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

then

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

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

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

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

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

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

then

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

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

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

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

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

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

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

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

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

gives detailed balance

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

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

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

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

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

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

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

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

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

Finite evidence accuracy. For bounded coarse observables O , if

$$\|\hat{\mu}_s - \mu_s\|_{\text{TV}} \leq \epsilon_{\text{samp}}$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- $E0$: seed noise, proposal noise, repair jitter,
- $E1$: conventional reference ensemble,
- $E2$: OPH-native quotient ensemble,
- $E3$: OPH vacuum,
- $E4$: OPH primordial field,
- $E5$: observable cosmological prediction.

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

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

$$\tilde{\mathcal{A}}(R) = \mathcal{B}(\tilde{\mathcal{H}}_R), \quad \tilde{\mathcal{H}}_R = \bigotimes_{i \in R} \tilde{\mathcal{H}}_i.$$

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

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

This is the fixed-cutoff realized presentation used throughout the collar analysis. The EC theorem works on the invariant-state realization and uses the sector-preserving induced collar algebra there; the entire fixed-point algebra on the unreduced tensor product is outside that theorem surface. When the consensus paper speaks of quotient repair, the physical repair law is defined on this fixed-point / overlap-invariant data and any representative-level map is only a lift of that quotient

update. Descent to the gauge quotient is therefore built into the physical-algebra formulation as quotient-level repair: the local map is $\text{locRep}_\lambda : Q \rightarrow Q$, and the global map is the finite quotient normal-form operator

$$\text{Rep}_\lambda = \overline{\text{nf}}_\lambda : Q \rightarrow Q.$$

On the declared fixed-cutoff collar branch, the local repair step itself is read from exact Markov splice or a declared Petz/Fawzi–Renner recovery channel. The declared fixed-cutoff branch adds repair completeness, together with control on the declared Petz domain where that branch is used. The consensus paper proves those clauses on a nontrivial rooted-tree packet-net export: finite packet labels live on a rooted tree, hidden labels are quotient-only, weighted parent-copy repair strictly lowers Φ , normal forms are exactly consistent packet assignments, and the classical full-support Petz channel is CPTP and trace-norm contractive on its positive-support-gap domain. Gauge-invariant observables factor through the quotient normal form on that verified branch. The declared touched-overlap acceptance contract makes Φ the finite-patch Lyapunov functional for primitive repair proposals on the scalar branch, but a proposal becomes a physical step only through the transactional acceptance layer: snapshot-current read sets, unchanged-outside-write validation, preserved boundary/sector/holonomy data, and exact well-founded descent. The read set is validation-complete: it contains every support of every measure term that can change under the write. Overlapping primitive proposals are aggregated by connected conflict component; the restriction-compatible union-collar glue supplies the canonical aggregate payload, while primitive members of that component do not commit independently. On the same fixed-cutoff quantum lift, the terminal expectation functional on each declared physical observable algebra is therefore unique even when microscopic representative lifts differ by gauge labels globally or by sector labels inside one quotient-local glued state. The unqualified coding statement at this stage is only a finite constraint-code statement: the overlap-net codewords are the globally consistent states $C = \Phi^{-1}(0)$. A bare graph does not determine code distance: the same graph can carry constant readouts with distance 1 or repetition constraints with distance $|V|$. Topological-code distance/min-cut, Knill–Laflamme resilience, spectral-gap convergence, BFT wall-clock liveness, and hardware search-work reduction therefore require their own certificates and are not imported by the core overlap graph. Exact descent gives termination. Confluence follows from an independently checked conflict-component diamond on the physical quotient, protected-support and protected-conflict completeness, and repair completeness. The theorem therefore gives one schedule-independent terminal quotient normal form from a fixed initial quotient state. A stronger same-boundary conclusion requires a preserved boundary/sector map and at most one consistent quotient extension in that boundary fiber. The layered functional carrier in the consensus paper proves the finite multi-edge, multi-step witness for $H_B \wedge H_{\text{fib}}$; the functional selected-fiber theorem gives the rooted obstruction-check form by constructing the single candidate from the boundary/root data and then checking all non-tree, sector, and holonomy obstructions. On that branch, multiple same-boundary candidate interiors can be present, inconsistent candidates are eliminated, and all surviving candidates share one quotient normal form. Normal-form hashes are public equality receipts after quotienting, not selectors among physically distinct minimizing endpoints. If accepted schedules terminate at different observer-facing quotient normal forms without a declared holonomy or higher-gauge obstruction, that repair law is outside the OPH consensus theorem. The neutral result of Ref. [1] packages this quantifier split exactly: under observation preservation and completeness, agreement modulo a silent equivalence of normal endpoints reached from all same-boundary sources is equivalent to injectivity of the induced boundary map on the consistent quotient. This cross-source criterion, same-source confluence, weak normalization, and all-schedule liveness are separate obligations. The named layered and functional carriers discharge the cross-source premise only on their stated domains. For sepa-

rated cofinal refinement systems, the consensus paper adds an inverse-limit bridge: if finite-stage restriction maps commute with the quotient normal-form maps and holonomy maps, and if compatible families are visibly separated on cofinal finite stages, then finite normal forms and holonomy obstructions assemble into unique refinement-limit classes with finite-stage witnesses for nonzero holonomy. It also adds the RG-facing comparison: when a chosen coarse-graining channel shadows the finite-stage normal forms and obstruction maps with declared defects ε^n and ε^h , reconciling first and then coarse-graining gives the same macroscopic law readout as coarse-graining first and then reconciling, up to $\max\{\varepsilon^n, \varepsilon^h\}$; exact naturality is the zero-defect case. Exact normal-form naturality is discharged by a repair-morphism witness, and holonomy naturality by the corresponding cochain map. For distributed implementations, the same paper proves that a worker run presents one finite universe only when it starts from one global carrier and each event projects to a legal monolithic repair path, a physical stutter, or a certified rollback to an earlier committed projection. The required certificate names the global graph, initial state, partition, cut interfaces, observer registry, run/config/code hashes, linearized commits, restart roots, final monolithic normal form, and recomputed readout. Executor order is separate from observer history. Observer-clock naturality additionally requires a history-augmented quotient with semantic event keys independent of worker ID, repair iteration, queue position, timestamp, and retry metadata; a global observer registry descended from local registry groupoids with disjoint patch/cap/future namespaces and explicit lineage arrows; a state-preserving observer-algebra extraction functor from terminal augmented normal forms; a support-visible cap chart map natural with that extraction; and an operational clock instrument whose readout is calibrated up to the declared affine reparameterization and residual bound. The listed objects or their proof-producing finite certificates are prerequisites. In their absence, modular-order/readback naturality is a named gate rather than a consequence of implementation bookkeeping. The same quotient boundary applies to neutral-bulk readout. For each shard s , raw record rows Σ_s must first be quotient by the declared presentation groupoid Γ_s and then sent through the terminal normal-form map n_s , producing

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

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

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

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

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

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

with channel isometries; refinement claims additionally require a cofinally vanishing distance tail modulus. Finite Euclidean claims require the double-centered Gram certificate

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

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

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

$$N : Q \rightarrow Q_{\text{nf}}$$

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

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

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

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

$$Q_r = \Sigma_r / \Gamma_r$$

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

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

and

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

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

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

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

over the affine constraint surface

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

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

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

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

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

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

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

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

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

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

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

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

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

then total-variation contraction gives

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

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

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

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

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

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

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

$$\text{settle}(q) \rightarrow n_r(q), \quad \text{sample}(\mu_r) \rightarrow q.$$

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

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

gives detailed balance

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

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

$$\frac{g(q - R_{\text{repair}}(q'))}{g(q' - R_{\text{repair}}(q))},$$

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

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

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

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

$$\mathbb{E}[q] = 0, \quad \mathbb{E}[q^\top K_r q] = d_r A_r, \quad d_r = \dim V_r.$$

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

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

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

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

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

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

and sample independent coefficients

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

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

$$\Sigma_r = C_{sr} \Sigma_s C_{sr}^\top, \quad K_r = A_r \Sigma_r^{-1}.$$

Equivalently, when eliminating fine variables from a precision matrix

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

the coarse precision is the Schur complement

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

not the principal block K_{cc} .

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

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

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

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

is stochastic and detailed-balanced with

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

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

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

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

2 Finite Screen Spectrum Theorem Package

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

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

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

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

$$f^\top M_r g \rightarrow \int_{S^2} f g d\Omega$$

with an explicit band-limited error bound

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

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

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

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

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

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

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

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

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

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

$$U^\top M'_r U = M_r, \quad q'_{0,r} = U q_{0,r}, \quad B'_r = U B_r,$$

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

Proof. Substitution gives

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

Subtracting from U gives the high-mode identity. □

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

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

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

$$K_r^{\text{diss}} = H_r - T_r^\top H_r T_r, \quad T_r^\top H_r T_r \preceq H_r;$$

or the reversible generator form

$$K_r^{\text{Dir}} = -\frac{1}{2}(L_r + L_r^{\dagger H}).$$

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

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

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

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

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

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

the unique entropy maximizer is

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

and

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

and the orientation-reversal coarea identity

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

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

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

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

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

$$\kappa_{\text{rep}}^{\text{edge}} = \frac{P_\star}{48(P_\star - \varphi)}.$$

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

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

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

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

then every positive measurable representative satisfies

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

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

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

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

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

Then

$$C_\ell^q = 4\pi Z_\star^2 \int d \ln k \Delta_\zeta^2(k) j_\ell^2(k R_\star),$$

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

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

Thus

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

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

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

then

$$|C_{\ell, W}^q - C_{\ell, \text{shell}}^q| \leq 4\pi Z_\star^2 \int d \ln k \Delta_\zeta^2(k) \left[2\delta_\ell(k) + \delta_\ell(k)^2 \right].$$

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

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

$$(q_r, S_{\text{scr}, r}, \theta, A_{q, r}, C_{\ell, r}^q, A_{\zeta, r}, \Delta_{\zeta, r}^2)$$

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

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

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

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

2.0.1 Source-only primordial bridge theorem

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

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

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

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

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

where

$$\text{BranchType} \in \{\text{FLATEXACT}, \text{FLATASSUMED}, \text{OPENCURVED}, \text{CLOSEDCURVED}, \text{UNRESOLVED}\}.$$

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

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

$$S_{\text{src}} = \sum_I S_I[g, \Psi_I], \quad T_{ab}^I = -\frac{2}{\sqrt{-g}} \frac{\delta S_I}{\delta g^{ab}}, \quad T_{ab}^{\text{tot}} = \sum_I T_{ab}^I,$$

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

$$\text{StressRead}_r : \text{nf}(\mathcal{Q}_r) \longrightarrow \{T_{r,ab}^I\}_I.$$

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

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

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

and StressRead_r descends to the physical quotient and converges. Then

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

On the recovered Einstein branch,

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

The evidence bundle must also report the non-circular comparison

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

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

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

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

and decompose

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

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

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

On an expanding branch, define

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

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

Theorem (exact total-curvature evolution). Let

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

and

$$\Gamma_a^{\text{eff}} = \nabla_a p_{\text{eff}} - \frac{\dot{p}_{\text{eff}}}{\dot{\rho}} \nabla_a \rho.$$

Then, exactly and without a long-wavelength approximation,

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

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

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

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

Relative to a frozen background,

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

On the uniform-total-density cut,

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

After removing the monopole and dipole,

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

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

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

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

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

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

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

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

and the intrinsic nonadiabatic pressure and transfer covectors obey

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

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

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

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

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

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

then

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

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

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

is required.

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

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

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

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

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

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

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

$$\Psi_\ell^{\text{BranchType}}(k) = \begin{cases} \int dr W(r) j_\ell(kr), & \text{BranchType} = \text{FLATEXACT or FLATASSUMED}, \\ \int dr W(r) \Phi_{\ell k}^{(-)}(r), & \text{BranchType} = \text{OPENCURVED}, \\ \int dr W(r) \Phi_{\ell k}^{(+)}(r), & \text{BranchType} = \text{CLOSEDCURVED}, \end{cases}$$

and for a thin shell

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

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

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

so the actual finite evidence receipt is

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

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

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

physical spatial k receipt

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

screen to physical k association receipt.

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

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

Homogeneity and isotropy on the safe band require

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

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

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

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

common primordial anomaly mode basis receipt.

Mode freezeout and common initial surface are separate gates:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

then

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

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

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

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

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

If the source refinement packet proves

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

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

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

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

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

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

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

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

so

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

with the complete amplitude conversion

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

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

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

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

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

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

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

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

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

Then

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

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

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

$$p_\star = p_0 + Q^{-1} A^\top (A Q^{-1} A^\top)^+ (C - A p_0).$$

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

Theorem (unrestricted tomographic uniqueness). For

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\sum w_i, \quad \sum w_i x_i, \quad \sum w_i x_i x_i^\top$$

or, for weighted response regression,

$$X^\top W X, \quad X^\top W y, \quad y^\top W y.$$

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

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

$$\mathcal{L}_r : V_r^{\text{screen}} \rightarrow \mathcal{H}_r^{(3)}$$

must define

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

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

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

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

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

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

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

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

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

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

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

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

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

- $E0$: seed noise, proposal noise, repair jitter,
- $E1$: conventional reference ensemble,
- $E2$: OPH-native quotient ensemble,
- $E3$: OPH vacuum,
- $E4$: OPH primordial field,
- $E5$: observable cosmological prediction.

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

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

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

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

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

and the sector-preserving collar algebra

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

with

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

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:

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

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

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

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

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

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

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

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

and the identification of the HJPW factors with the EC factors is a genuine additional state condition, the **Markov-split alignment hypothesis (MSA)**. It is not implied by exact Markovity: with four qubits A, B_L, B_R, D and $\rho = \Phi_{AB_R} \otimes \Phi_{B_L D}$ (Bell pairs across the cut, one center block), one has $I(A : D \mid B) = 0$ but $I(A : B_R) = 2 \ln 2$, so the state is exactly Markov yet not of the displayed EC-aligned form, which would force $I(A : B_R) = 0$. All literal Markov-modular equalities below therefore assume exact Markovity *plus* MSA, or an explicitly stated idealized limit that

supplies both; The spacetime and Einstein paper states MSA formally and records this counterexample. MSA is checkable state by state: The spacetime and Einstein paper proves it equivalent to a splitting of $\log \rho$ into commuting one-sided terms supported on Ab_L^α and $b_R^\alpha D$, to the existence of a ρ -preserving conditional expectation onto either one-sided edge algebra (a Takesaki commuting square over the collar center [35]), and to blockwise vanishing one-sided mutual information $I(Ab_L^\alpha : b_R^\alpha D) = 0$. On the MaxEnt branch this reduces the derivation of MSA to a sharp condition on the effective Hamiltonian: every cross-cut coupling must act through the central sector label, as it does in lattice-gauge-type regulators where the interface energy is a function of the conserved flux. On that declared **central-interface branch** The spacetime and Einstein paper derives both MSA and exact collar Markovianity as theorems (descent of one-sided invariants and boundary charges through the edge-center quotient), and it adopts the central-interface normal form as an explicit axiom-level branch clause stated with its MaxEnt axiom, so the literal Markov-modular identities hold on the declared branch without a separate state hypothesis. The clause is independent of the repair/consensus axioms: The spacetime and Einstein paper exhibits a finite regulator package that satisfies overlap consistency, schedule-independent transactional repair, and an exactly closed refinement-channel family while its retained constraint family carries a boundary-invariant noncentral cross-cut coupling, so overlap-consistent repair does not force the clause and it stays a permanent declared branch input. Off the declared branch MSA is carried as an explicit hypothesis.

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

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

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

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

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

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

Strict concavity of von Neumann entropy ensures uniqueness. Here the constrained operators are the N_{con} global sums $O_a := \sum_x O_a(x)$ plus the retained sector terms, so the exponent has one shared multiplier per density label. MaxEnt gives the Gibbs form; the displayed locality restriction on H_{sec} is separate and load-bearing. Had the

constraints instead been imposed per cell, $\langle O_a(x) \rangle = c_a(x)$, the multipliers would be cell-dependent fields $\lambda_a(x)$ and the family would grow with the lattice; that enlarged family belongs to the non-equilibrium approximation regime of the clarification above and is not the branch used in the refinement argument. QED.

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

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

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

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

$$\|\sigma - \omega_L(R_{\ell \rightarrow L}(\lambda))\|_1 \leq \sqrt{2 \varepsilon_{\ell \rightarrow L}(\lambda)},$$

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

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

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

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

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

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

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

The strong mixing premise is a boundary-anchor expansion

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

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

Theorem 2.5 (finite-range conditional Gibbs mixing gives collar recovery).

Under Lemma 2.6 and the strong conditional mixing premise,

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

with

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

Proof. Embedding the marginal logarithms in the tripartite algebra gives

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

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

For one uniform family, the sharp sufficient limit condition is

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

If

$$|\partial C|_{\text{UV}} \leq a(\ell_0/\ell)^p, \quad \delta(\ell) \geq \zeta(p + \eta)\ell \log(\ell_0/\ell), \quad \eta > 0,$$

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

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

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

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

This bound is the quantitative hinge for constructive gluing.

2.0.2 Concrete UV realization: quantum link models

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

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

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

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

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:

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

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 [53]. Exact Markovity depends on the state and is not forced by the decomposition alone; nor does exact Markovity align the HJPW factors with these edge-mode labels, which is the separate Markov-split alignment hypothesis of Section 2.3.

Dynamics and MaxEnt. The natural Hamiltonian is a 2+1D lattice gauge Hamiltonian on the screen worldvolume: plaquette ("magnetic") terms, electric terms on links, vertex Gauss terms as constraints, plus local matter couplings. In quantum link form this is finite-dimensional per link while behaving like gauge theory in the continuum limit. Then the MaxEnt assumption

becomes concrete: the MaxEnt state is a Gibbs state $\rho \propto e^{-\sum_i \lambda_i O_i}$ with quasi-local O_i , precisely the local-Gibbs regime.

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

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

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

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

2.0.3 Conformal-modular fixed-point microphysics

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

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

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

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

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

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

Derived propagation control on the same branch. Because H_{MaxEnt} is a finite-range or quasi-local sum on a finite type-I regulator net, standard Lieb–Robinson estimates [15] apply to the

automorphism group it generates, or to any branch generator lying in the same bounded-support algebraic closure. Thus there is a finite propagation velocity v_{LR} and constants C, ξ such that for local observables A_X, B_Y ,

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

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,

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Fixed-cutoff null-strip bridge.** The null-strip package proves transferred cut-center data, the theorem-local inherited left/right strip-split condition needed for the spatial-collar-type tensor decomposition, exact-or-controlled four-term strip additivity on one inherited strip model, renormalized endpoint-Lipschitz control up to the weak tail generator, and the derived half-sided modular pair whose Borchers–Wiesbrock consequence is an explicit positive null-translation generator on its Stone domain with the affine half-line modular relation; the same half-line family then fixes the generator/charge identification internally.
- **Derived half-sided modular inclusion.** Nested null half-line algebras satisfy half-sided modular inclusion on the declared geometric scaling branch by Corollary 5.2e; Borchers–Wiesbrock then identifies the positive null-translation generator on its Stone domain together with the affine half-line modular relation [19].
- **Weak continuity and finite variation.** The bounded-interval and half-line endpoint-Lipschitz control follow from the local MaxEnt branch and define the weak tail generator at fixed cutoff. The support-visible scaling theorem together with the extracted geometric cap pair supplies the continuum null-generator setting in which that weak-tail data can be matched to the geometric null modular action of the relativistic phase.

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

1. **Edge/cap label constraints:** fix the distribution of collar-sector labels, equivalently $\langle L_C \rangle$ for each cap size, giving the area term.
2. **Gauge charges:** fix boundary flux or charge operators.

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

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

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

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

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

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

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

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

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

Each ρ_n is faithful at finite n , and tensoring it with any fixed finite exact-Markov collar factor gives a full-rank exact-Markov collar family. Nevertheless $\lambda_{\min}(\rho_n) = e^{-n} \rightarrow 0$, 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, $\hat{\rho}_n$ the exact-Markov comparison marginal on the same collar model, and $\Delta_n = \|\rho_n - \hat{\rho}_n\|_1$. For $a > 0$, set $K_a(\rho) = -\log(\rho + a\mathbf{1})$. For every bounded collar observable O in the support-visible algebra,

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

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

$$\Delta_n/a_n \rightarrow 0, \quad d_{m,\delta}a_n \rightarrow 0, \quad \Delta_n |\log a_n| \rightarrow 0,$$

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

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

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

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

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

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

and a faithful detail state τ_j . At regulator r ,

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

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

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

The canonical refinement and coarse-graining maps are

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

and

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

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

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

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

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

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

The finite reference modular groups are also exactly compatible:

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

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

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

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

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

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

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

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

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

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

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

For a regularizer $\eta_s > 0$, define

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

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

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

Hence $\eta_s \rightarrow 0$ and $\epsilon_{s,r}/\eta_s \rightarrow 0$ give uniform-on-compact-time convergence on each fixed support-visible local algebra. For a countable cofinal local test tower, pointwise $\epsilon_{s,r} \rightarrow 0$ admits a diagonal subsequence and one cutoff schedule. The phrase “finite-stage modular action” below means this

transported, regularized local action; it does not mean the unregularized full finite-cap modular group.

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

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

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

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

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

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

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

2.1 Overlap Consistency and Gluing

The constructive part of overlap consistency is the tree-gluing theorem below. The structural part is the origin of the gluing redundancy itself. On the ordinary or central-defect branch, gauge-as-gluing is the finite-regulator overlap redundancy of local chart presentations: once the overlap algebras are realized in finite-dimensional charts, overlap-consistent rechartings of a cut form a compact unitary transition system. The collar theorem of Section 2.3 should therefore be read with its boundary group G_Σ understood as shorthand for this derived compact boundary action, not as a model-specific add-on. On the genuinely noncentral branch, the same weak gluing data are encoded by a compact crossed-module change system, so the fixed-cutoff collar theorem upgrades

to a higher-gauge statement rather than failing. The fixed-cutoff topological package is therefore closed on all three branches. A second structural point is UV underdetermination. If each patch Hilbert space is tensored with an inert finite ancillary factor and the observable patch algebras are embedded as $\mathcal{A}(P) \otimes \mathbf{1}$, then the physical observables, collar conditional mutual information, carried Markov errors, and quotient normal forms are unchanged. So the physical UV branch is fixed only modulo gauge or implementation hiding together with such ancillary stabilization, not a unique microscopic presentation; the unique theorem-grade object is the induced family of terminal expectation functionals on the declared physical observable algebras, not a preferred microscopic representative.

2.1.1 Constructive gluing on tree covers

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

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

with

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

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

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

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

2.1.2 Gauge-as-gluing and loops

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

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

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

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

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

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

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

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

with

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

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

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

2.1.3 Loop obstruction class (central defect)

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

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

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

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

2.1.4 Non-central obstruction (2-group cocycle)

When defects are not central, the natural coefficient data is a crossed module $(H \rightarrow G)$ 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 $\partial : H \rightarrow G$.

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

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

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

$$g_{ij} : P_{ij} \rightarrow G, \quad h_{ijk} : P_{ijk} \rightarrow H,$$

obeying the 2-cocycle conditions

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

and on quadruple overlaps,

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

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

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

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

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

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

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

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

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

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

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

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

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

The physical higher-gauge collar is

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

with

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

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

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

Theorem 3.4b (TransportabilityFromOverlapGluing). Fix a connected cut Σ and a finite regulator charting nerve N_Σ . Transport of a collar charge is defined by the path composite in the overlap recharting groupoid: along $p = (i_0 \rightarrow \cdots \rightarrow i_m)$, the charge block is moved by $U_p = U_{i_{m-1}i_m} \cdots U_{i_0i_1}$. This construction uses the overlap unitary transition system itself, not an added DHR transportability assumption.

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

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

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

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

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

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

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

2.2 Modular Flow and Lorentz Kinematics

2.2.1 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

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

3. the fixed-collar exact-Markov replacement modulus

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

where

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

Lemma 4.1a (EC-aligned exact Markov implies exact additivity up to a central term). If $I(A : D \mid B)_\omega = 0$ and ω satisfies the Markov-split alignment hypothesis of Section 2.3, then the state takes the EC-aligned HJPW block form

$$\omega_{ABD} = \bigoplus_{\alpha} p_{\alpha} \omega_{Ab_L^{\alpha}}^{(\alpha)} \otimes \omega_{b_R^{\alpha}D}^{(\alpha)},$$

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. By alignment, the HJPW blocks are the EC blocks. On each 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 in $Z(\mathcal{A}(B))$. Without alignment, exact Markovity cancels the four logarithms on the state's own HJPW blocks, but the resulting block constants need not lie in $Z(\mathcal{A}(B))$, so the central form used here is unavailable. 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,

$$\delta_{A:B:D}^M(\varepsilon) \longrightarrow 0 \quad (\varepsilon \downarrow 0).$$

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 $I(A : D | B)$ directly to an exact Markov state. The limit points are exact Markov states in the CMI-zero sense; identifying their factorization with the preselected EC factors is the Markov-split alignment hypothesis of Section 2.3, which is carried separately and is not produced by small CMI (the Bell-pair counterexample of Section 2.3 has zero CMI at fixed distance from every EC-aligned 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

$$d_M(\rho) \leq C_{A:B:D,\lambda_*} I(A : D | B)_\rho^{\theta_{A:B:D,\lambda_*}}.$$

This is a fixed-collar Łojasiewicz-type rate for the analytic function $I(A : D | B)$ 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 $I(A : D | B)_{\omega_\varepsilon} \leq \varepsilon$, and choose $\sigma_\varepsilon \in \mathfrak{M}_{A:B:D}$ with

$$\|\omega_\varepsilon - \sigma_\varepsilon\|_1 \leq \delta_{A:B:D}^M(\varepsilon).$$

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 $A \cup B$,

$$|\text{Tr}[X(\omega_\varepsilon - \sigma_\varepsilon)]| \leq \|X\|_\infty \delta_{A:B:D}^M(\varepsilon).$$

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

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

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

$$\|\omega_\varepsilon - \omega_\varepsilon^{\text{rec}}\|_1 \leq r_{\text{FR}}(\varepsilon), \quad r_{\text{FR}}(\varepsilon) := 2\sqrt{1 - e^{-\varepsilon}} \leq 2\sqrt{\varepsilon}.$$

This $O(\varepsilon^{1/2})$ term is the finite-stage observable error. The modulus $\delta_{A:B:D}^{\text{M}}(\varepsilon)$ 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 $\delta_{A:B:D}^{\text{M}}(\varepsilon_\delta) \rightarrow 0$, 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

$$\mathcal{P}_O\left(\{r_{\text{FR}}(\varepsilon_j)\}, \{\delta_j^{\text{M}}(\varepsilon_j)\}, \{\eta_j^{\text{reg}}\}\right),$$

where $\mathcal{P}_O \rightarrow 0$ 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-parameter 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.

2.2.2 Theorem: BW_{S²} on the extracted 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 $(\mathcal{A}_\infty^{\text{geo},\text{sv}}(C), \omega_\infty^{\text{geo},C})$. The branch theorem is the conditional implication

$$\text{FiniteCapBWCertificate} + \text{MGNS-1} \implies \sigma_t^{\omega_\infty^{\text{geo},C}} = \alpha_{\lambda_{\widehat{C}}(2\pi t)}.$$

The finite cap-normal certificate consists of cap-normal density and nondegeneracy, BW framing, support-order faithfulness, geometric support-flow group and continuity control, held-out oriented cross-ratio convergence, independently normalized geometric 2π -KMS comparison, and wrong-normalization separation. The independent MGNS-1 package supplies the algebra-state reference tower, common comparison maps, cap-family-uniform regularized modular transport, mixed-GNS convergence, and modular support covariance. The support-visible theorem below uses both inputs; it does not require a type-I continuum algebra or a full-algebra unregularized common spectral floor. Bare finite consensus implies neither input.

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

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

and, on each fixed faithful collar model,

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

where $\lambda_* > 0$ is the lower spectral bound for comparing 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

$$\delta_{A_\delta:B_\delta:D_\delta}^M(\varepsilon_\delta) \rightarrow 0 \quad (\delta \downarrow 0),$$

with $r_{\text{FR}}(\varepsilon_\delta)$ and η_δ^M carried explicitly.

For each cap C , let $\lambda_C(s) \subset \text{Conf}^+(S^2)$ 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 (Finite cap-net support-visible BW_{S^2} scaling theorem). Let $\hat{C} = (C, n_C, p_C^-, p_C^+)$ be a nondegenerate BW-framed cap. For every OPH-realized observer-supporting refinement branch satisfying the five OPH axioms, the derived fixed-cutoff regulator/collar/consensus package, the finite cap-normal BW certificate above, and independently the complete multiresolution reference-tower/MGNS-1 package of Theorems 2.6c–2.6f on that same branch, the support-visible geometric cap net admits a weak-*/GNS scaling-limit cap pair $(\mathcal{A}_\infty^{\text{geo},\text{sv}}(C), \omega_\infty^{\text{geo},C})$. Then the scaling-limit modular automorphism group is

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

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 $r_{\text{FR}}(\varepsilon_\delta)$, 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 the certificate's framed-cap and oriented cross-ratio rigidity clauses are applied. No separate cap-isotropy/ $\text{SO}(2)$ -equivariance selector, finite-cell Lorentz-invariance premise, or full-algebra unregularized common floor is used. 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

$$K_C = 2\pi B_C + Z_C, \quad Z_C \in Z(\mathcal{A}(C))_{\text{sa}};$$

otherwise the automorphism identity is the full statement and the action is outer. Consequently, under the paired same-branch finite cap-normal and complete MGNS-1 inputs, the transported regularized modular action on every fixed support-visible geometric local algebra converges uniformly on compact modular-parameter intervals to the geometric cap-dilation action. Without both inputs, this theorem makes no BW or finite-to-continuum modular claim. Even with them, it does not assert compact-time convergence of finite-regulator modular groups beyond the declared support-visible certificate.

Proof. Markov locality localizes the modular defect to the shrinking collar, with carried finite-stage errors $r_{\text{FR}}(\varepsilon_\delta)$ and δ^M . Proposition 2.6a shows why a full-algebra unregularized floor is unavailable in general. Theorem 2.6b supplies the matrix-element replacement on fixed collar models,

and Theorem 2.6f supplies the stronger compact-time group estimate only for transported regularized local actions with the stated reference-tower certificate. 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. The finite cap-normal certificate supplies support-order faithfulness, BW framing, held-out oriented cross-ratio rigidity, and compact-time equicontinuity, so the support flow is the framed cap-preserving conformal flow up to normalization:

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

for some $\kappa_C > 0$, with the certificate-selected frame understood in λ_C . The certificate then passes the independently normalized geometric 2π -KMS condition to $s \mapsto \alpha_{\lambda_C(s)}$. KMS uniqueness, together with wrong- β separation and nontriviality, fixes $\kappa_C = 2\pi$. Thus

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

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

Definition 4.2a (BW-branch observer-relative modular ordering). On the branch satisfying the hypotheses of Theorem 4.2, the modular automorphism parameter t of the extracted cap pair $(\mathcal{A}_\infty^{\text{geo},\text{sv}}(C), \omega_\infty^{\text{geo},C})$ supplies a dimensionless ordering for that observer's accessible algebra-state pair. Physical time requires an observer-readable transition, event correspondence, and calibrated clock instrument. No claim about arbitrary operational clocks, global time, or the full problem of time follows from the geometric modular-flow theorem alone.

Scope note. The theorem fixes the 2π normalization and the carried refinement limit on the support-visible extracted geometric cap pair under the paired finite cap-normal and independently complete same-branch MGNS-1 inputs. 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, BW framing, oriented cross-ratio rigidity, and geometric KMS normalization. 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. Compact-time convergence of finite-regulator modular groups is available only through the transported-local certificate of Theorem 2.6f; a black-box conformal-net BW reconstruction would be an additional validation route, not a hidden premise.

2.2.3 Theorem: cap-normal H^3 reconstruction

Theorem 4.3 (Cap-normal H^3 reconstruction, imported from the spacetime and Einstein derivation). On the support-visible round-cap BW branch, write the celestial null section as

$$q(\Omega) = (1, \Omega), \quad \Omega \in S^2.$$

For a nondegenerate oriented round cap

$$C(\mathbf{c}, \alpha) = \{\Omega : \mathbf{c} \cdot \Omega \geq \cos \alpha\}, \quad 0 < \alpha < \pi,$$

define

$$n_C = (\cot \alpha, \csc \alpha \mathbf{c}).$$

Then, for the time-first Lorentz form $\eta = -dt^2 + d\mathbf{x}^2$,

$$\eta(n_C, n_C) = 1, \quad \eta(n_C, q(\Omega)) = \frac{\mathbf{c} \cdot \Omega - \cos \alpha}{\sin \alpha}.$$

Thus $\eta(n_C, q(\Omega)) = 0$ exactly on ∂C , with positive sign in the cap interior and negative sign outside. For $g \in \text{Conf}^+(S^2)$ with Lorentz representative Λ_g ,

$$\Lambda_g q(\Omega) = \omega_g(\Omega) q(g\Omega), \quad n_{gC} = \Lambda_g n_C.$$

Consequently

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

oriented round caps form the de Sitter normal space

$$\text{Cap}_{\text{round}}^{\text{or}}(S^2) \cong dS_3^{\text{cap}} \cong \text{SO}^+(3, 1)/\text{SO}^+(2, 1),$$

and future unit timelike observer frames form the distinct hyperbolic space

$$H^3 \simeq \text{SO}^+(3, 1)/\text{SO}(3), \quad \dim H^3 = 3.$$

Each cap normal determines the geodesic half-space

$$\mathcal{H}_C^+ = \{u \in H^3 : \eta(n_C, u) \geq 0\},$$

not a preferred observer point.

Proof. This is the cap-normal reconstruction theorem in The spacetime and Einstein paper, which supplies the projective-null-cone model, the cap-normal bijection, Lorentz equivariance, the H^3 observer-frame hyperboloid, the ideal-boundary identification, and the cap/half-space duality. The BW theorem supplies the cap-preserving conformal subgroup and its 2π modular normalization; the cap-normal theorem supplies the explicit projective Lorentz representation and does not rederive BW from bare finite consensus. QED.

Corollary 4.3a (Frame-space boundary). The theorem proves the exact canonical Lorentz/ H^3 observer-frame hyperboloid on the stated branch. Its points are future unit timelike frames, not event positions. It does not populate an event base, construct a chart-blind neutral bulk, derive a physical curvature radius R_H , produce stress-energy, or close the Einstein branch-entry umbrella. Record-conditioned cap responses can estimate frame data under the certificate imported next; event location, object families, neutral bulk, scale, stress, and Einstein branch entry are separate receipts. Theorems 4.3c–4.3h below import The spacetime and Einstein paper's geometry-producer, null-net-standardness, event-manifold, stress, entropy-bridge, and composed branch-entry packets. Their composition is conditional on one source-derived common-domain tower with certified cofinal tails, universal coupling, a vacuum reference, and independent physical scale readouts. Construction and certification of such a tower are work in progress.

Theorem 4.3b (Conditional record-conditioned H^3 frame estimate). Fix a quotient-visible record token i and clock slice t . On the frame space of Corollary 4.3, suppose its record-conditioned modular cap responses

$$R_i(C, t, O) = \omega_{i,O}(\sigma_t^{C,O}(M_{C,0,O})), \quad \omega_{i,O}(A) = \frac{\omega_O(P_{i,O} A P_{i,O})}{\omega_O(P_{i,O})},$$

admit the calibrated frame-local factorization

$$y_{i,j} = F_j(X_i(t)) + e_{i,j}, \quad F_j(X) = \psi_j\left(\frac{\eta(X, n_j)}{R_H}\right),$$

on a compact frame region $\Omega \subset H_{R_H}^3$. Assume the finite response map has quantitative observability

$$\alpha \bar{d}(X, Y) \leq |F(X) - F(Y)|_W \leq L \bar{d}(X, Y), \quad \alpha > 0,$$

and the total error obeys $|e_i|_W \leq \sigma_i$. If $\mathcal{N}_\varepsilon$ is an ε -net of Ω and $\hat{X}_i(t)$ is a residual minimizer up to tolerance τ_i , then the conditioned frame support is

$$S_i(t) = B_H\left(\hat{X}_i(t), R_H \left[\frac{L}{\alpha}\varepsilon + \frac{2}{\alpha}\sigma_i + \frac{1}{\alpha}\tau_i\right]\right).$$

Exact continuous data have a unique minimizer. Finite noisy data produce the displayed frame ball; a unique finite frame value is licensed only when the certified residual gap $\Delta_{\text{loc},i}(t) > 0$. If that gap fails, the output is AMBIGUOUS. This theorem does not locate an event, populate the event base, resolve unlabeled mixtures, identify particle species, produce stress-energy, or establish chart-blind neutral bulk.

Theorem 4.3c (Quotient-intrinsic geometry producer and dimension-selection boundary). On a refinement tower of repaired quotient normal forms, the support-visible incidence complex: vertices the quotient-visible patch classes, simplices the jointly supported patch sets: is computable from the normal form alone and is invariant under repair schedule, gauge, and refinement. If the tower carries the computable receipts (i) spherical incidence (connected closed orientable combinatorial surface with Euler characteristic 2), (ii) disk/mesh nondegeneracy, (iii) modular cross-ratio Cauchy convergence, and (iv) independently normalized geometric 2π -KMS comparison with wrong-normalization separation, then the refinement limit produces the round S^2 , its conformal structure, the oriented round-cap family with cap normals, and every clause of the finite cap-normal BW certificate, with quantitative error envelopes. Conversely, there are finite repair systems with isomorphic rewrite structure whose incidence complexes realize S^2 , T^2 , the 2-skeleton of $\partial\Delta^4$, a wedge of two spheres, and any modular temperature β : repair confluence selects no topology, dimension, orientation framing, or normalization. The screen S^2 is therefore *derived on the receipt branch and underdetermined off it*; the receipts are the irreducible geometric branch content of Einstein branch entry. (Spacetime and Einstein paper, geometry-producer packet; machine receipts ship with the released geometry code.)

Theorem 4.3d (Null-net standardness, derived translations, and the four-translation assembly). On the producer branch, the null half-line/interval blow-up net at a produced cap boundary point lives on one common GNS space, and: (i) the reference vector is cyclic and separating for every half-line and interval algebra used, given the finite-stage faithfulness of the MaxEnt states, the mixed-GNS Cauchy clause, and one half-line cyclicity receipt, with interval cyclicity then derived by the Reeh–Schlieder/Borchers argument rather than assumed; (ii) the half-sided modular inclusion is derived from the produced geometric cap flow, without assuming any translation or dilation unitaries, and Borchers–Wiesbrock yields the positive null-translation generator on its Stone domain; (iii) on the exact-Markov/central-interface branch the reduced interval modular Hamiltonians split into local terms plus uniformly bounded interface blocks (endpoint-Lipschitz), while an explicit four-qubit nearest-neighbour Gibbs witness shows finite-range locality alone does *not* imply this (machine-verified); (iv) the two derived translation groups and the modular dilations generate Möbius covariance and the projective bounded-interval action; and (v) on the modular-intersection receipt branch the per-direction positive generators commute, transform with the null-vector weight

under the produced Lorentz action, satisfy the dependent-family linearity relation, and assemble into one four-parameter translation representation with joint spectrum in the closed future cone: a positive-energy Poincaré representation. (Spacetime and Einstein paper, null-net standardness packet.)

Theorem 4.3e (Conditional Lorentzian event manifold). Events are equivalence classes of record germs whose common-refinement distance tends to zero; cofinal box intersection is insufficient because it need not be transitive. Unconditionally the event space is a second-countable nonseparated uniform space. The manifold branch requires population plus germ realization (E1), separation (E2), a locally bi-Lipschitz response chart with an interior-ball receipt (E3), an affine overlap cocycle (E4), $C^{1,1}$ tetrads and inverse tetrads (E4'), a held-out quadratic-cone fit with inertia (1, 3) and wrong-signature controls (E5), and local semantic causal reachability (E6). The conformal celestial S^2 checks the inferred cone boundary; it does not force a quadratic cone. On these receipts the event base is a Lorentzian four-manifold, while H^3 is the fiber of future unit timelike frames over each event. Stable causality and global hyperbolicity remain named assumptions. Explicit countermodels show that $\dim H^3 = 3$ by itself determines neither event dimension nor manifold structure. This packet supplies, conditionally, the locally Lorentzian $d = 4$ regime consumed by the Einstein section. (Spacetime and Einstein paper, event-manifold packet; machine receipts ship with the released geometry code.)

Theorem 4.3f (Local conserved stress tensor from modular charges). A symmetric bilinear form on $(\mathbb{R}^4, \eta)$ is determined by its null-null values up to a multiple of η , and a directional charge family arises from a tensor iff it satisfies the finite dependent-family linearity relations; nine generic null directions reconstruct the trace-free part with Lipschitz stability. On the event region of Theorem 4.3e with the assembly branch of Theorem 4.3d, the smeared endpoint-derivative charges of the bounded-interval modular kernels satisfy these relations, so they define a semiclassical expectation tensor $\langle T_{ab} \rangle$: symmetric, local, Poincaré-covariant with one common normalization, equal on null projections to the positive Borchers generators on a common core, and weakly conserved ($\nabla^a \langle T_{ab} \rangle = 0$ a.e.) by translation invariance: with ambiguity exactly ϕg_{ab} plus improvement terms. Countermodel: per-direction positive generators violating one linearity relation admit no rank-two source (irreducible tomography residual); scalar collar CMI can never be promoted to a rank-two source. Universal coupling of this tensor to the geometric branch is the named receipt UC. (Spacetime and Einstein paper, stress packet; receipts ship with the released geometry code.)

Theorem 4.3g (Repaired entropy bridge and absolute Einstein equation). On the central-interface branch, write the state as a direct sum of bulk states tensored with maximally mixed edge factors. Its entropy splits exactly into $S_{\text{bulk}} = H(p) + \sum_{\alpha} p_{\alpha} S(\rho_{\alpha}^{\text{bulk}})$ and $S_{\text{edge}} = \sum_{\alpha} p_{\alpha} \log d_{\alpha}$. The finite first law is $\delta S = 2\pi \delta \langle B_C \rangle + \delta \langle Z_C \rangle$, with $Z_C = \sum_{\alpha} (\log d_{\alpha}) P_{\alpha}$. Hence $\delta S_{\text{bulk}} = 2\pi \delta \langle B_C \rangle$ and $\delta S_{\text{edge}} = \delta \langle Z_C \rangle$ throughout the declared central-interface class, including variations of the sector weights. The changing stress channel is supplied by the exact MaxEnt envelope identity $dS/dt = \lambda = 2\pi$, giving the Clausius balance $\delta S_{\text{edge}} = \lambda \delta t - 2\pi \delta \langle B_C \rangle$ on the coupled variation class. One common scaling family must carry the conditional-mixing, boundary-count, fixed-collar, sharp-rate, and $o(\ell^4)$ remainder bounds. Coverage of all local timelike directions follows from population plus the produced Lorentz orbit. The vacuum-reference premise fixes the first-variation integration constant. Independent scale readouts with trivial positive-rescaling stabilizer identify the structural coupling. These premises yield $G_{ab} + \Lambda g_{ab} = 8\pi G \langle T_{ab} \rangle$ per connected component. Without the vacuum reference, the result fixes the residue tensor only up to one covariantly constant metric term. (Spacetime and Einstein paper, entropy-bridge and small-diamond packets.)

Theorem 4.3h (Typed Einstein branch composition and realized status). Bare finite consensus is not Einstein-complete: the same finite-consensus reduct admits extensions in which

the Einstein equation holds and extensions in which it fails. The valid gravity result is conditional. Let one source-derived repaired quotient tower carry deterministic geometry, modular, event, stress, entropy, and scale readouts on a common domain, with commuting refinement diagrams. Assume the finite geometry and null-net receipts, cap-interior modular data, the event-manifold receipts, one certified cofinal family on which every relevant remainder is $o(\ell^4)$, the universal-coupling and vacuum-reference receipts, and two independent scale readouts with trivial positive-rescaling stabilizer. The stress-tomography, entropy, small-ball, timelike-polarization, Ward, and Bianchi results then compose to

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

on each connected event component. The exact finite layer includes the nine-direction tomography frame with determinant $8192/27$ and the entropy and MaxEnt identities. Conditional on the continuum diamond-kernel formula and the smooth fixed-volume small-ball area identity, exact arithmetic gives the coefficients $8\pi^2 \ell^4/15$ and $-4\pi \ell^4/15$; the continuum formulas themselves remain explicit analytic premises. The finite receipt programs verify algebra, evaluator behavior, manifest and deletion logic, and synthetic controls. They do not prove that the typed branch is inhabited. Realized-branch nonemptiness requires one antecedent-only source tower, certified asymptotic tails, universal coupling, a source-derived vacuum reference, an identifiable scale, and isolated full-tower countermodels for semantic receipt minimality. This construction is work in progress. (Spacetime and Einstein paper, composed branch-entry theorem.)

2.3 Gravity from Fixed-Cap Generalized-Entropy Stationarity

2.3.1 Cap first law

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

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

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

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

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

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

This displayed variation is the special type-I/fixed-central-sector form of $K_C = 2\pi B_C + Z_C$, where the central term has no first-order contribution on the chosen sector. In the generic continuum/non-type-I case, Section 4.2 does not supply an inner operator identity; the theorem surface there is the geometric modular automorphism statement, so the cap first-law formula is restricted to this special realization.

2.3.2 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. Bounded-interval formulas are discharged by the E0.5 affine/projective kernel below, 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

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

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

$$\begin{aligned} \tilde{\mathcal{H}}_{I_-} &\cong \bigoplus_{\alpha_-} W_{\alpha_-} \otimes \mathcal{H}_{i_-^{\alpha_-}}, \\ \tilde{\mathcal{H}}_J &\cong \bigoplus_{\alpha_-, \alpha_+} W_{\alpha_-}^* \otimes \mathcal{H}_{j^{\alpha_-, \alpha_+}} \otimes W_{\alpha_+}, \\ \tilde{\mathcal{H}}_{I_+} &\cong \bigoplus_{\alpha_+} W_{\alpha_+}^* \otimes \mathcal{H}_{i_+^{\alpha_+}}, \end{aligned}$$

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

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

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

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

For any 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 and EC-aligned on the inherited split (the strip form of the Markov-split alignment hypothesis of Section 2.3), 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 $I(I_- : I_+ \mid J)_\omega \leq \varepsilon$, choose $\tilde{\omega}_J \in \mathfrak{M}_{I_-:J:I_+}$, assumed EC-aligned on the inherited split wherever the central identification below is used, with

$$\|\omega - \tilde{\omega}_J\|_1 \leq \delta_{I_-:J:I_+}^M(\varepsilon).$$

Define

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

Then

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

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

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

If the relevant strip marginals are uniformly faithful with lower spectral bound $\lambda_* > 0$, then

$$\|\mathfrak{D}_J(\omega, \tilde{\omega}_J)\|_\infty \leq 4\lambda_*^{-1} \delta_{I_-:J:I_+}^M(\varepsilon).$$

Independently, Fawzi–Renner gives a recovered comparison state with trace-norm error

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

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

Proof. On the inherited split of Proposition 5.2a, the EC-aligned exact-Markov case is the same HJPW block calculation as in the spatial collar theorem: alignment identifies the HJPW blocks with the inherited blocks, the four logarithms cancel blockwise on the canonical model, and any residual bookkeeping term is central. Without alignment, HJPW factorizes the state only over its own state-dependent split, and the centrality of ΔK_J in $Z(\mathcal{A}(J))$ is not available. 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

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

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

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

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

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

for bounded intervals in a compact endpoint window, and

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

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

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

By isotony of the null interval net,

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

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 $a > 0$,

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

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

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

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

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

such that

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

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

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

In addition,

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

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

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

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,

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

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 discharged by E0.5 below, 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.

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

The effective-theory input here is the standard stress-tensor representation that names 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,

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

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 E0.5 bounded-interval kernel transports 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.

2.3.4 Localized generalized entropy from Markov + MaxEnt

Using the collar decomposition, exact Markovity, and MSA, the declared central-interface branch has the following edge-aligned Markov normal form. On the nonexact branch, Theorem 2.5 supplies a recovered comparison state and a vanishing CMI envelope only under strong conditional mixing; convergence to an exact Markov state is a fixed-collar statement, and alignment with the displayed edge factors remains a separate MSA condition. Under those conditions MaxEnt selection within each edge sector gives

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

The entropy splits as

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

Convention: Throughout this paper, "log" denotes the natural logarithm (ln), so entropies are measured in **nats** (1 nat = $1/\ln 2 \approx 1.443$ 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

$$S_{\text{bulk}}(C) := H(p_\alpha) + \sum_{\alpha} p_{\alpha} S(\rho_{\text{bulk},C}^{\alpha}),$$

and the central area operator

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

Then

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

Newton coupling as an edge-entropy area-law readout.

In the collar double-scaling limit, the edge contribution becomes extensive along the entangling surface $\Sigma = \partial C$:

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

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 geometric area-law coupling:

$$G_{\text{geom}} = \frac{a_{\text{cell}}}{4 \bar{\ell}(t)}$$

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:

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

On the OPH gravity branch this formula does not back-solve the scale. The selected no- G scale certificate supplies $\gamma_{\star} = \ell_{\star} \nu_{\text{CS}}/c$, equivalently $B_{\star} = 3\pi/\ell_{\star}^2$. The local cell readings are

$$a_{\text{cell}} = P \ell_{\star}^2, \quad \bar{\ell}_{\text{shared}} = P/4,$$

so the Newton area-law readout gives

$$G_{\text{geom}} = \frac{P \ell_{\star}^2}{4(P/4)} = \ell_{\star}^2.$$

The SI display is $G_{\text{SI}} = c^3 \ell_{\star}^2 / \hbar$.

2.3.5 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,

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

2.3.6 No-smuggling dependency discharge for the Einstein bridge

Theorem E0 (Einstein bridge dependency discharge). Write OPH5 for the five-axiom basis: the S^2 screen net, overlap consistency, local MaxEnt/refinement, recoverable generalized entropy, and MAR. Let one source-derived cofinal OPH regulator tower carry the typed common-domain geometry, modular, event, stress, entropy, and scale readouts of Theorem 4.3h, together with its finite, asymptotic, and physical-identification premises. On that branch, OPH5 and the additional receipts supply the inputs used by the Einstein bridge:

- geometry readout through quotient normal forms via the screen fold $\chi_{S,r} \circ n_r \circ \pi_r$;
- Lorentz/H3 readout from round cap pairs, since $\text{Conf}^+(S^2) \cong \text{PSL}(2, \mathbb{C}) \cong \text{SO}^+(3, 1)$ and $H^3 \simeq \text{SO}^+(3, 1)/\text{SO}(3)$;
- 2π -normalized support-visible BW flow on the extracted geometric subnet;
- the null generator/local null-stress charge identity on the same branch;
- the bounded-interval kernel of Lemma E0.5;
- fixed-cap generalized-entropy stationarity from the MaxEnt/refinement theorem;
- the standard fixed-volume small-ball area identity after geometry readout;
- $o(\ell^4)$ remainder control by Lemma E0.6; and
- all local timelike directions by the cap-pair coverage of Lemma E0.7.

Thus the finite consensus reduct supplies quotient normal forms. The typed branch receipts supply the geometric, modular, stress, entropy, coverage, tail, coupling, vacuum, and scale inputs needed by the Einstein bridge. Construction and certification of one source-derived tower carrying all of them are work in progress. Bare finite consensus does not imply the Einstein equation.

Lemma E0.5 (Bounded-interval kernel from null projective covariance). On the null blow-up of the support-visible geometric branch, the half-line generator/charge identity, endpoint-Lipschitz renormalized interval control, and affine/projective covariance give

$$K_{(a,b)}^{\text{null}} = 2\pi \int_a^b \frac{(v-a)(b-v)}{b-a} T_{kk}(v) dv + R_{(a,b)}.$$

The remainder $R_{(a,b)}$ is the carried endpoint/collar term of the same refinement family. In the small-ball limit this is the ball modular kernel used below.

Lemma E0.6 (Uniform remainder control). Let (r, ℓ_r, δ_r) be one declared cofinal family on which Theorem 2.5 holds uniformly over the event region and cap family, and set

$$\varepsilon_r = c |\partial C_r|_{\text{UV}} e^{-\delta_r/\xi_r}.$$

Assume that the fixed-collar Markov replacements have one modulus $\delta_r^{\text{M}} \leq C_{\text{M}} \varepsilon_r^\theta$, with $C_{\text{M}} < \infty$ and $\theta > 0$, and that for some

$$s > \max\{8, 4/\theta\}, \quad \omega_r \longrightarrow +\infty,$$

the same family satisfies

$$\frac{\delta_r}{\xi_r} - \log(c|\partial C_r|_{UV}) \geq s \log(\ell_0/\ell_r) + \omega_r.$$

If the interval-kernel residual is also $o(\ell_r^4)$ on this family, then

$$2\sqrt{\varepsilon_r} = o(\ell_r^4), \quad \delta_r^M = o(\ell_r^4),$$

and the endpoint, collar, Markov, and long-wavelength derivative remainders do not alter the Einstein coefficient. Indeed, $\varepsilon_r \leq (\ell_r/\ell_0)^s e^{-\omega_r}$; the two conclusions follow from $s/2 > 4$ and $s\theta > 4$. Mere pointwise convergence of a collar error does not supply this common-family rate, and no per-radius diagonal choice is used.

Lemma E0.7 (Cap-pair timelike coverage). On the cap-normal H^3 branch, every future timelike vector v has $r = \sqrt{-\eta(v, v)}$ and $u = v/r \in H^3$. For any unit $s \in u^\perp$, the vectors

$$q_\pm = u \pm s$$

are future null and

$$v = \frac{r}{2}(q_+ + q_-).$$

Thus normalized future null directions in observer skies cover every local timelike direction once the observer-frame chart is supplied. This does not say that two unnormalized projective rays alone select a unique observer frame; the observer normalization or equivalent causal-diamond scale data is also required.

Proposition E2 (Recovered-core stress closure). If the recovered stress tensor is reconstructed from the full compatible family of null and timelike modular charges on the recovered core, then $\nabla^a \langle T_{ab} \rangle = 0$. The overlap-compatible local charge family cancels on opposite faces of infinitesimal recovered-core diamonds; after the cap-pair tensor upgrade this cancellation is the covariant divergence-free condition.

2.3.7 Einstein equation from fixed-cap generalized-entropy stationarity

This subsection is a branch theorem. The finite consensus reduct supplies quotient normal forms; Theorem E0 discharges the OPH5 recovered-core inputs needed by the Einstein bridge: geometry readout, support-visible BW modular covariance, the null-stress bridge, bounded-interval transport, fixed-cap MaxEnt stationarity, the fixed-volume small-ball area variation, carried-remainder control, and the timelike scalar-to-tensor upgrade.

In the d=4 scaling regime, the small-ball bridge is internal to the derived gravity chain once the E0.5 bounded-interval kernel 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 geometric subnet, a sufficiently small cap is the tangent causal diamond D_ℓ , whose preserving conformal Killing field is

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

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 E0.5 affine/projective transport of the D4 half-line stress bridge, so the D3 cap-modular theorem and the D4 stress bridge together yield

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

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

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

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

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

2.3.8 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 local-diamond rest-frame relation says

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

for every such u . In local inertial coordinates,

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

so

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

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

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

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.

2.3.9 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 $\delta S_{\text{gen}} = 0$ 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; stable correctable-record closure plus identification of the record carrier with de Sitter horizon entropy fixes $\Lambda_* \ell_*^2 = 3\pi/N_*$, while the SI curvature scale also uses the selected OPH scale certificate.

2.3.10 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

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

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

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

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

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

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

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

These two comparisons do different jobs:

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

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

and

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

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

$$\|\Delta K_\delta(\omega) - \Delta K_\delta(\sigma_\delta)\|_\infty \leq \eta_\delta^M.$$

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

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

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

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 $\eta_\delta \rightarrow 0$ 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,

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

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

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

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

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

Then

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

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:

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

On a separate dark/anomaly continuation branch, this carried scalar remainder may be promoted to a repair stress only after a quotient-invariant source-localization receipt identifies the same collar operator with the small-ball Einstein remainder, and after a positive causal packet lift supplies the full tensor. In SI units that continuation has rest-energy normalization

$$T_A^{ab} u_a u_b = \frac{15\hbar c}{8\pi^2\ell^4} R_{r,\ell} + O(\epsilon_{\text{loc}}),$$

where $R_{r,\ell}$ is measured in nats and ℓ is the proper geodesic small-ball radius. The scalar quantity $I(A : D \mid B)$ alone is not a stress tensor: packet factorization, causal directions, an allocation rule, a rest frame, source-localization map, and transport law are additional continuation data.

The framework carries the Markov error through three distinct controls:

1. $r_{\text{FR}}(\varepsilon_\delta)$ controls bounded observables at one stage;
2. $\delta_{A_\delta:B_\delta:D_\delta}^{\text{M}}(\varepsilon_\delta)$ 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”.

2.3.11 Focusing/QNEC status on the null-modular branch

The focusing content of the recovered core enters through the Recoverable Generalized Entropy axiom and stays an axiom. This subsection records the status of the two focusing statements on the null-modular branch and the exact reason no internal derivation is claimed.

What the branch supplies. The fixed-cutoff null bridge of §5.2 supplies transferred cut-center data, exact-or-controlled strip additivity on the inherited strip model, the derived half-sided modular pair with its Borchers–Wiesbrock translation generator, and the half-line generator/charge identification $[K, P] = -i2\pi P$. These are kinematic inputs a QNEC statement consumes; they are not a proof of it.

Why relative-entropy monotonicity does not yield QNEC. Monotonicity under restriction orders the relative entropies $S(\rho_{R(\lambda)} \parallel \omega_{R(\lambda)})$ along nested null cuts; it fixes the sign of the first λ -derivative and says nothing about the second. A monotone function of the cut need not

be convex, so the step from monotonicity to $d^2 S_{\text{rel}}/d\lambda^2 \geq 0$ is unavailable. The entanglement first law $\delta S = \delta \langle K \rangle$ is a first-order state variation at fixed algebra and reference state; the QNEC second null shape variation is a different object, and expanding the first law to second order in the cut does not produce it. The published QNEC proofs supply exactly the missing machinery: null-surface locality and analyticity hypotheses in the free and super-renormalizable settings [29], and modular-flow, causality, and defect-OPE inputs in the general QFT setting [31]. None of those hypotheses is verified on the reconstructed patch net.

Why QFC is not derived from QNEC. The Quantum Focusing Conjecture is strictly stronger than QNEC: the foundational QFC statement obtains QNEC as its locally parallel limit [30], and no implication runs the other way. QNEC together with the Einstein branch (Theorem 5.1) and the Raychaudhuri identity controls the diagonal, stationary-cut part of the generalized-expansion variation and leaves the off-diagonal, shear, counterterm, and general nonstationary-cut content of QFC untouched. Focusing for S_{gen} therefore stands where the Recoverable Generalized Entropy axiom declares it and carries no independent internal support from this branch.

Claim boundary. The null-modular branch narrows what a future internal QNEC proof must construct: the half-sided modular structure is in place, and the open items are the QFT-side hypotheses named above, verified on the reconstructed net, plus the off-diagonal QFC content. Until such a proof lands as a dated artifact, every downstream use of focusing on this branch cites the axiom.

Theorem 5.1 (OPH Einstein-branch closure in the scaling regime). Under OPH5 on a recovered-core, support-visible scaling branch satisfying the E0 dependency discharge above, 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, the timelike scalar-to-tensor upgrade promotes it to the semiclassical Einstein equation modulo the expected Λg_{ab} ambiguity. The D6 capacity branch closes that metric term globally under stable whole-fiber correctable-record closure, the unique finite-size slack zero, and the horizon-record identification.

Proof sketch. The E0 discharge turns the OPH5 recovered-core branch into the exact hypothesis package required by the Jacobson-type bridge: quotient-safe geometry readout, 2π -normalized BW cap flow, local null stress charge, bounded-interval ball kernel with $o(\ell^4)$ remainder, fixed-cap generalized-entropy stationarity, the fixed-volume small-ball area identity, and all timelike directions. The cap first law identifies δS_C with $\delta \langle K_C \rangle$, while the null bridge together with the E0.5 bounded-interval kernel 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. Lemma E0.7 supplies all local timelike directions, the scalar-to-tensor theorem upgrades the relation to the tensor equation modulo Λg_{ab} , and the D6 capacity branch closes the remaining metric term under its separately typed stable-capacity and horizon-record premises. QED.

2.3.12 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 $d \rightarrow d'$. Finite same-boundary records and finite reconstruction thresholds do not by themselves supply physical radiation entropy, exterior time, evaporation flux, a geometric island, a radiative quotient, or an asymptotic waveform readout. 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 $\text{Tr}(\rho L_C) = \langle A \rangle / 4G_{\text{geom}}$, the area eigenvalues are

$$A_{\alpha} = 4G_{\text{geom}} \log d_{\alpha} = 4\ell_{\star}^2 \ln d_{\alpha},$$

where $\ell_{\star}^2 = \hbar G_{\text{SI}} / c^3$ is the scale-readout area displayed as 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 $d \rightarrow d'$ changes the area by

$$\Delta A = 4\ell_{\star}^2 \ln(d'/d).$$

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

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

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:

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

Structural condition: comb vs. generic discreteness. The log-integer *comb* structure requires the additional dynamical assumption that integer-ratio emission steps ($d \rightarrow d/k$) 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:

$$\frac{\Gamma}{\Delta E_k} \approx \frac{64\pi^2 p_0}{a \ln k},$$

a function of (a, k) . At $k = 2$ the declared range $a \in [1, 10]$ spans 1.8% to 18%; at $k = 3$, 1.1% to 11%. A few-percent value therefore requires the particle content and greybody model that pin a inside a narrower interval (for instance $a \approx 4-6$ at $k = 2$); no such model is selected here. The fraction is mass-independent at every (a, k) .

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

$$\text{Re } \omega \rightarrow \frac{c^3}{8\pi G M} \ln 3.$$

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

$$\Delta A = 4\ell_\star^2 \ln 3 \approx 4.39 \ell_\star^2$$

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 $\Delta f_k = (c^3 / 16\pi^2 G M) \ln k$:

- **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 = 10¹² 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 $64\pi^2 p_0 / (a \ln k)$, spanning 1.8%–18% at $k = 2$ over the declared $a \in [1, 10]$ (mass-independent at every (a, k)). Published 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:

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

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. A physical QNM or ringdown claim needs a separate exterior-background bridge, a radiative quotient, a finite operator converging to the continuum perturbation generator, future-horizon and future-null-infinity boundary conditions, and an asymptotic readout map. A finite repair spectrum or fitted clock scale is not that bridge.

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

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

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. The line pattern belongs to a continuation template outside recovered core.

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:

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

where $\Gamma_{\ell m}^{\text{GR}}$ 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:

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

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 [56] 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 the offset-subtracted ratio

$$\frac{f_{k,m} - m\Omega_H/(2\pi)}{f_{2,m} - m\Omega_H/(2\pi)} = \frac{\ln k}{\ln 2}$$

exactly, equivalently $x_k/x_2 = \ln k/\ln 2$ in the universal coordinate defined above, independent of remnant parameters. The full frequencies $f_{k,m}$ do not satisfy $f_k/f_2 = \ln k/\ln 2$ at nonzero Ω_H : the additive rotation offset breaks the ratio, and only the thermal comb pieces carry it. 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.

2.3.13 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:

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

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 ,

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

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.

2.3.14 Precision gravity predictions and experimental bounds

The pure Einstein linearization makes classical dispersion and polarization statements that can be confronted with the tightest available experimental bounds. This section translates those branch-scoped statements into the observables that experiments constrain; it does not infer a quantum graviton pole from the classical equation.

Speed of gravitational waves. On a suitable background, the two transverse-traceless classical modes of the pure Einstein linearization propagate on the same invariant null cone as the Maxwell modes:

$$\frac{c_{\text{GW}} - c_\star}{c_\star} = 0 \text{ exactly on this action/background branch.}$$

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

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

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

Einstein quadratic mass parameter. On the additional pure two-derivative Einstein–Hilbert action branch, the transverse-traceless quadratic operator has no Fierz–Pauli hard-mass parameter:

$$\mu_{\text{FP,hard}}^2 = 0.$$

This is an action-content statement, not a graviton-particle mass prediction from diffeomorphism redundancy alone. A quantum mass or pole claim requires Definition 6.18, while diffeomorphism-invariant Stueckelberg/bimetric completions and extra-field theories lie outside the pure-Einstein field-content branch.

Published bound (GW dispersion analysis, PDG 2025):

$$m_{\text{grav,disp}} \leq 1.76 \times 10^{-23} \text{ eV}/c^2 \quad (90\% \text{ credibility; model-dependent dispersion fit}).$$

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.

Published bound (GW170817 inspiral phasing, PDG 2025):

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

Tensor polarizations on the selected field-content branch. The pure Einstein linearization contains two tensor polarizations. It does not prove that every broader diffeomorphism-invariant EFT lacks additional scalar, vector, or massive modes. 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

2.3.15 State-observable error propagation and the dynamical bridge

On the additional pure-Einstein action/background branch, the framework states the classical equalities recorded in Section 2.3.14. Finite Markov/recovery control supplies a separate and more limited quantitative statement. This subsection therefore keeps two error ledgers separate. The *state-observable ledger* tracks expectation values of specified bounded local observables in a recovered state; the Markov/recovery machinery controls this ledger. The *dynamical-parameter ledger* tracks spectral and dispersion properties of an evolution generator: propagation speeds, particle-pole masses, and correlator poles. The trace-distance chain alone does *not* control that ledger.

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, writing ρ for the target state and σ for its recovered comparison state, the recovery map gives the trace-norm bound

$$\|\rho - \sigma\|_1 \leq \delta_k := 2\sqrt{1 - e^{-\varepsilon_k}} \leq 2\sqrt{\varepsilon_k}.$$

State–observable ledger. Trace distance gives immediate bounds on expectation-value errors for the declared operator class: bounded operators supported in the recovered region, compared at the time of recovery. 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. This inequality is exactly as strong as it looks and no stronger: the norm $\|O\|_\infty$ is the Lipschitz constant only for bounded O , and the comparison is an equal-time statement about states.

Exponential decay from the mixing hypothesis. Writing w for collar width, so it is not confused with the recovery error δ_k , the mixing assumption (Section 2.3) provides

$$I_\omega(A_w : D_w \mid B_w) \leq C_{\text{mix}} |\partial C|_{\text{UV}} e^{-w/\xi}.$$

Combining these gives an explicit precision dial for bounded local observables:

$$\delta_{\text{step}} \lesssim 2\sqrt{C_{\text{mix}} |\partial C|_{\text{UV}}} \cdot e^{-w/(2\xi)}.$$

For an operator normalized by $\|O\|_\infty \leq 1$, an independently chosen absolute expectation-error target τ_O is guaranteed at one step if $\delta_{\text{step}} \leq \tau_O$; the theorem therefore gives the sufficient CMI condition $\varepsilon \leq (\tau_O/2)^2$. The target τ_O belongs only to the state–observable ledger and cannot be set by importing a gravitational-wave speed tolerance.

What the dial does not certify: dynamical parameters. The gravitational-wave propagation speed and any graviton pole or dispersion-mass parameter are not bounded local observables. They are spectral/dispersion properties of the dynamical generator (poles of retarded correlators, or long-time wave-packet phase and group velocities), so the trace-distance inequality supplies neither their norm nor their Lipschitz constant. Two states can be close in trace distance at one time while being vacua of generators with different dispersion relations, and a small equal-time expectation error does not control phase accumulation over the $\sim 10^8$ -year propagation baseline of GW170817. Consequently, $\delta \lesssim 10^{-15}$ is *not* a theory-side certificate for the multimessenger bound $|c_{\text{GW}}/c - 1| \lesssim 10^{-15}$ of Section 2.3.14, and the recovery error must not be presented as the precision dial for that published bound. The ideal classical relation $c_{\text{GW}} = c_\star$ and the vanishing pure-Einstein hard parameter $\mu_{\text{FP,hard}}^2 = 0$ belong to the additional action-level branch recorded there; neither stability away from that branch nor a quantum graviton pole follows from the trace-distance chain.

Required bridge (dynamical-parameter ledger; open). A sufficient acceptance gate can be stated precisely, but its hypotheses have not been derived from the Markov remainder in this paper. Let $\mathcal{B} = [k_{\text{min}}, k_{\text{max}}]$, with $0 < k_{\text{min}} < k_{\text{max}}$, be the declared experimental wave-number band. Suppose a future construction supplies, for each tensor helicity, a C^1 self-adjoint finite-regulator generator $H_\nu^{\text{TT}}(k)$ and the GR reference generator $H_0^{\text{TT}}(k)$, expressed in angular-frequency units ($\hbar = 1$) on one specified common finite-dimensional helicity fiber or after a declared C^1 unitary identification of the two fibers. Assume that the reference TT tensor-mode eigenvalue $\omega_0(k) = c_\star k$ is simple in the helicity block and is separated from the rest of the spectrum by a uniform gap $g_{\mathcal{B}} > 0$. Define

$$M_{\mathcal{B}} := \frac{1}{c_\star} \sup_{k \in \mathcal{B}} \left\| \partial_k H_0^{\text{TT}}(k) \right\|_\infty$$

and the dimensionless dynamical error

$$\eta_{\text{dyn},\nu} := \max \left\{ \frac{1}{g_{\mathcal{B}}} \sup_{k \in \mathcal{B}} \|H_{\nu}^{\text{TT}}(k) - H_0^{\text{TT}}(k)\|_{\infty}, \frac{1}{c_{\star}} \sup_{k \in \mathcal{B}} \|\partial_k H_{\nu}^{\text{TT}}(k) - \partial_k H_0^{\text{TT}}(k)\|_{\infty} \right\}.$$

For $\eta_{\text{dyn},\nu} < 1/4$, let $\omega_{\nu}(k)$ be the continued isolated eigenvalue and define its group velocity by $c_{\text{GW},\nu}(k) := \partial_k \omega_{\nu}(k)$. The standard isolated-eigenvalue projector estimate and the Hellmann–Feynman formula then give, uniformly over $\mathcal{B}$,

$$\sup_{k \in \mathcal{B}} \left| \frac{c_{\text{GW},\nu}(k)}{c_{\star}} - 1 \right| \leq L_{\mathcal{B}} \eta_{\text{dyn},\nu}, \quad L_{\mathcal{B}} := 1 + 4M_{\mathcal{B}}.$$

Indeed, if $P_{\nu}(k)$ and $P_0(k)$ are the rank-one spectral projectors, then $\|P_{\nu} - P_0\|_{\infty} \leq 2\eta_{\text{dyn},\nu}$, hence $\|P_{\nu} - P_0\|_1 \leq 4\eta_{\text{dyn},\nu}$. Differentiating the isolated eigenvalue therefore bounds the group-velocity change by the derivative perturbation plus $4M_{\mathcal{B}}c\eta_{\text{dyn},\nu}$. This displays a dimensionless constant derived from a specified generator and uniform over the declared band. If $q := L_{\mathcal{B}}\eta_{\text{dyn},\nu} < 1$, then stationary propagation at fixed group velocity over path length D obeys the group-delay bound

$$|\Delta t| \leq \frac{D}{c_{\star}} \frac{q}{1 - q}.$$

The same bound applies to varying propagation only if it holds pointwise along the entire path; on a cosmological path, the redshifted wave number must remain inside $\mathcal{B}$. A wave-packet phase bound likewise requires uniform control of the inverse dispersion $k_{\nu}(\omega)$ on the corresponding frequency band. A model-dependent particle dispersion-mass bound is a separate spectral estimate at $k = 0$; it is not a corollary of the band-limited speed estimate.

The OPH-specific bridge requires three additional constructions:

1. construct H_{ν}^{TT} , or an equivalent retarded two-point kernel, from the finite and scaling OPH dynamics;
2. prove a norm-resolvent, quadratic-form, or correlator estimate of the form $\eta_{\text{dyn},\nu} \leq C_{\mathcal{B}}r_N$, where $r_N := \min(2, \sum_{k=2}^N \delta_k)$ is the accumulated trace-norm recovery bound of Theorem 3.1 for the associated N -step construction and $C_{\mathcal{B}}$ is derived rather than assumed; and
3. apply the resulting band-uniform dispersion estimate to the detector band and propagation baseline, including phase accumulation where used by the likelihood.

Only after the second item is proved may one combine the estimates to obtain

$$\sup_{k \in \mathcal{B}} \left| \frac{c_{\text{GW},\nu}(k)}{c} - 1 \right| \leq L_{\mathcal{B}} C_{\mathcal{B}} r_N.$$

No estimate $\eta_{\text{dyn},\nu} \leq C_{\mathcal{B}}r_N$ is proved here. Substituting δ_k , δ_{step} , or any equal-time trace-distance error for $\eta_{\text{dyn},\nu}$ without that bridge is invalid.

Precision summary. The framework provides:

1. The ideal pure-Einstein linearization statements $\mu_{\text{FP,hard}}^2 = 0$ and $c_{\text{GW}} = c_{\star}$, without promotion to a quantum graviton pole.
2. Translation of those ideal equalities into the specific observables experiments constrain (Section 2.3.14).
3. Explicit bounds on how far expectation values of bounded local observables in the recovered state can drift, using the conditional mutual information $\rightarrow$ trace distance $\rightarrow$ observable error chain.

It therefore provides quantitative error control on recovered states and bounded observables, but no present theory-side certificate for the published GW-speed or model-dependent particle dispersion-mass bounds. Dynamical-parameter certification remains in the separate open ledger above.

2.3.16 Dark-sector response from the modular anomaly

The D12 dark-sector continuation promotes scalar repair occupation to a canonical pair (n, θ) . Its lattice action is a proposed dynamical completion rather than a recovered-core theorem.

The candidate anomaly-tensor slot T_{ab}^{anom} is available only after the finite covariant source packet of Section 5.9 supplies source localization, CMI-to-modular-source matching when that route is used, rank-two reconstruction, conservation, normalization, and universal coupling. On that packet it supplies one structural ingredient for a possible dark-sector continuation without introducing additional particle species. Raw scalar collar CMI supplies only a recovery diagnostic.

Conditional source identification. Once the source packet has identified an anomalous local modular energy, its rest-frame normalization is

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

The quantity can be called dark only on a branch whose packet also establishes gravitational coupling and absence of an electromagnetic coupling. These properties motivate the dark-sector interpretation; neither they nor the observational identification follow from the collar-CMI theorem.

Conditional repair-charge action. The phase equation supplies a conserved repair current with explicit boundary flux. Its dilute homogeneous phase has $\rho_R \propto a^{-3}$, while the cubic condensed link energy gives $a_R = \sqrt{a_b a_0}$ and $v^4 = GM_b a_0$ on the spherical deep branch. The same action couples coherent matter through $q_* \chi_\nu^{\text{can}} S_{\text{coh}}^{\text{can}}$ and carries the field stress required for momentum closure.

Scope. The canonical pair, compact-phase completion, baryonic source map, dimensional couplings, full constitutive law, relativistic refinement limit, abundance, and physical likelihoods are work in progress. A compact neutral coherent source has no monopole and cannot support a closed device.

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

$$\bar{\rho}_A(a), \quad \bar{\rho}_{A,\text{eq}}(a), \quad w_A(a), \quad c_{s,A}^2(k, a), \quad \sigma_A(k, a), \quad Q_A^\mu, \quad B_A(k, a), \quad \Gamma_{\text{rec}}(k, a).$$

The first, source-side no-go is that $\bar{\rho}_A(a)$, $\bar{\rho}_{A,\text{eq}}(a)$, and $B_A(k, a)$ alone do not define a physical Boltzmann source. A physical source must be emitted by a finite covariant collar-packet parent

$$\mathcal{P}_r[X, g] = (C_r, Z_r, A_r, R_r, G_r, \pi_r[X], L_r[X], Q_r, D_r),$$

whose finite packet states split the anomaly and any repair recipient, whose repair generator and reaction channels close the total stress tensor, and whose causal response data set the pressure, sound speed, anisotropic stress, exchange current, and repair rate. The parent must also declare exactly one local source route: fixed-reference anomalous modular energy, or nonlinear CMI stress with a CMI-to-modular-source matching theorem/hypothesis and residual ledger. Raw collar CMI is otherwise a recoverability diagnostic, not the $15/(8\pi^2)$ BW source. B_A is gauge-invariant only when built from the anomaly-frame baryon density $n_b^{(A)} = -u_{A\mu} J_b^\mu$. If the exchange branch is nonzero,

explicit recipient stress and an equal-and-opposite exchange current are mandatory; otherwise the repair term is not a closed stress-energy source. A transition spectral number is only $\gamma_{\text{repair step}}$ without active-fiber, physical-clock, and common-parent response-pole receipts for a physical Γ_{rec} . A physical likelihood comparison also requires declared source, solver, dataset, covariance, and nuisance provenance.

Canonical finite covariant collar-packet parent. For a regulator r and background (X, g) , the source object is

$$\mathcal{P}_r[X, g] = (\mathcal{C}_r, g_r, e_r, U_r, Z_r, \omega_r, p_r, f_r, \Phi_r, \mathcal{R}_r, \mathcal{G}_r, \pi_r, \text{Read}_r).$$

Here $\mathcal{C}_r$ is a finite causal cell complex with oriented faces, cell four-volumes, causal adjacency, and finite parallel transports $U_{F \rightarrow c}$; e_r is the tetrad/local-frame data; $Z_r = \bigsqcup_s Z_{s,r}$ is the packet state space for anomaly, recipient, radiation, standard matter, and any interaction sector; ω_r, p_r, f_r are invariant weights, local momenta, and occupations; Φ_r is the face flux; $\mathcal{R}_r$ is the reaction channel list; $\mathcal{G}_r$ is the finite gauge/quotient data; π_r is the restriction/refinement map; and Read_r is the local observable readout. With signature $(-+++)$, every packet used in a kinetic stress branch must satisfy $g_{ab}p_z^a p_z^b = -m_z^2$, $p_z^0 > 0$, and $\omega_z f_{c,z} \geq 0$, with the invariant measure convention declared explicitly.

Scalar-to-stress nonuniqueness. A scalar collar row, even a finite row for $\rho_A(a)$, $\rho_{A,\text{eq}}(a)$, or $B_A(k, a)$, does not determine a stress tensor. Many inequivalent packet parents can have the same scalar density while differing in pressure, sound speed, anisotropic stress, exchange current, gauge-invariant perturbation variables, and causal response. Therefore no evidence bundle may promote scalar rows to physical CMB, lensing, growth, or S_8 inputs unless the parent supplies the missing stress and response variables.

Finite MaxEnt lift. When only finitely many one-cell source readings $e_x(u_i) = T_{ab}^A(x)u_i^a u_i^b$ are available, a stress lift is not unique until the branch supplies a finite selection rule. The admissible MaxEnt lift minimizes no hidden observational residual: it maximizes the declared packet entropy subject to the source readings, mass shells, conserved charges, local frame covariance, and refinement restrictions. The lift is physical only when held-out timelike probes pass a quadraticity/tomography test and the variational stress agrees with the packet-moment stress in the predeclared norm.

Packet stress and exchange closure. For a kinetic parent,

$$T_{I,r}^{ab}(c) = \frac{1}{V_c} \sum_{z \in Z_{I,r}(c)} \omega_z f_{c,z} p_z^a p_z^b, \quad J_{I,r}^a(c) = \frac{1}{V_c} \sum_{z \in Z_{I,r}(c)} \omega_z f_{c,z} p_z^a.$$

More general packet parents may add an internal stress Σ_z^{ab} , but then Σ_z^{ab} is part of the primitive parent data and must obey the same local-frame and refinement checks. The finite transport/reaction equation must imply

$$\text{Div}_r T_{I,r}^{ab} = Q_{I,r}^b, \quad \sum_I Q_{I,r}^b = 0.$$

If the repair exchange is nonzero, a row labelled “recipient” is not enough: the recipient stress T_R^{ab} , exchange current Q_R^b , and closure residual $Q_A^b + Q_R^b$ must be emitted and certified. If exchange is off, the parent instead emits a repair-exchange-off receipt.

Response, rate, and gauge receipts. The finite causal response must provide a finite domain of dependence, subluminal characteristic bound, retarded support residual, and response stability residual. The transition-matrix diagnostic may be called $\gamma_{\text{repair step}}$ until the active fiber, conserved-sector decomposition, physical clock $\Delta\tau_{\text{physical}}$, and common parent response pole are certified. Only then may one write

$$\Gamma_{\text{rec}} = -\text{Re } \lambda_{\text{active}} / \Delta\tau_{\text{physical}}.$$

Gauge checks must be stated in gauge-invariant or rest-frame variables; raw Newtonian/synchronous agreement is a diagnostic, not a gauge-independence receipt.

Certificate hierarchy. The finite parent certificate is the conjunction of primitive evidence, not a producer-declared Boolean. The source side must certify the route and entropy unit, modular nonadditivity, localization and stress tomography, finite-packet kinematics and mass shell, covariant transport and four-momentum, stress readout and variational agreement, local-frame and quotient invariance, total stress and exchange-current closure, cosmological gauge invariance, finite domain of dependence, subluminal and retarded response, stability under refinement, and the cold-dark-matter limit. Frozen source, solver, likelihood, and data hashes are later promotion receipts. They do not belong to the parent theorem and cannot rescue a failed parent.

The same firewall applies to the promoted CMB-source inputs

$$\eta_R, \quad \gamma_{\text{repair step}} \text{ or certified } \Gamma_{\text{rec}}, \quad A_\zeta, \quad q_{\text{IR}}, \quad \ell_{\text{IR}}, \quad B_A(k, a), \quad \rho_A(a), \quad N_{\text{CRC}}.$$

Each one must have a source-provenance DAG node recording the source report, parents, no-CMB-data flag, and measurement-use flags. Contradictory provenance, such as declaring no CMB data use while also fitting to Planck, fails closed. Reducers must pool additive sufficient statistics globally before any nonlinear estimate of tilt, amplitude, inverse precision, rank, condition number, isocurvature leakage, B_A , ρ_A , or Γ_{rec} ; raw transition diagnostics are $\gamma_{\text{repair step}}$. Shard-local nonlinear averages and shard-local `any()` likelihood rollups are diagnostic only. N_{CRC} is the D6 consensus invariant unless a separate additive capacity schema proves disjoint coverage. 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 $\Gamma_{\text{rec}}(k, a)$ after the physical-clock and response receipts pass, not fitted as a free environmental kernel to CMB, weak-lensing, SPARC, or cluster data. The compressed H_0 , Ω_m , σ_8 , and S_8 rows are plumbing diagnostics for a low- H_0 , Planck-like branch, not theorem-grade cosmology.

Finite cosmological source and prediction eligibility. The continuation-level physical cosmology gate is the conjunction

$$\text{CMB ready} = \text{SOURCE} \wedge \text{SCALE} \wedge \text{PARENT} \wedge \text{KERNEL} \wedge \text{INIT} \wedge \text{TRANSFER} \wedge \text{FREEZE} \wedge \text{LIKE}.$$

SOURCE is the source-only DAG with transitive measurement-taint rejection and globally pooled sufficient statistics; **SCALE** is the physical mode/clock/lift bridge; **PARENT** is the finite covariant dark/anomaly parent with packet kinematics, mass shell, packet stress readout, reaction-channel four-momentum, recipient stress and exchange-current closure for nonzero exchange, total stress closure, local-frame/quotient invariance, cosmological gauge invariance, finite domain of dependence, subluminal retarded response, response stability, refinement, and CDM-limit recovery receipts; **KERNEL** emits physical ρ_A , $\rho_{A,\text{eq}}$, B_A , and certified Γ_{rec} ; the source-only ρ_A row additionally requires the anomaly abundance selector

$$\mathcal{E}_{A,r} = \mathbb{E}_{\mu_{A,r}^{\text{rel}}} [\mathbb{L}_{A,r}]$$

and its no-data-use and refinement receipts; INIT supplies regular initial modes and Einstein constraints; TRANSFER supplies Boltzmann well-posedness and numerical convergence; FREEZE fixes the model generation; and LIKE is the official likelihood execution. While this conjunction is false, physical CMB, BAO, lensing, growth, RSD, and S_8 outputs are diagnostic or conditional, not recovered-core predictions.

This branch is a D12 completion contract, not a physical dark-matter theory.

Falsifiability. The repair-charge condensate requires a source-derived action, relativistic limit, and joint galaxy, Solar-System, lensing, cluster, and cosmological likelihood. These constructions are work in progress, so the falsification program contains no cosmological target or verdict.

2.3.17 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 2.2.2, the realized scaling-limit cap modular action on the extracted 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} .

Conditional theorem boundary: correctable public-record closure. Freeze a capacity carrier $\mathcal{H}_{\text{cap},r,D}$ with $D = \dim \mathcal{H}_{\text{cap},r,D}$ and $N = \log D$. For every terminal world q , compatible record atoms, endogenous reachability, a frozen publicness policy, and the complete source-supplied global checkpoint family define a compound confusability graph G_q . The exact public capacity and

whole-fiber readback are

$$M_0(q) = \alpha(G_q), \quad \mathfrak{F}_{r,0}(D) = \{M_0(q) : q \in \tilde{\Omega}_{r,D}\}.$$

When the entire nonempty terminal fiber has one common readback, define $M_0(\mathfrak{U}_N) := \hat{F}_{r,0}(e^N)$. The universe-level closure is then

$$\boxed{N = \log M_0(\mathfrak{U}_N)}.$$

The notation is typed: M_0 is always a multiplicative record count, while N is its logarithm. Stable direct closure and its exact finite-size selector are

$$\mathfrak{F}_{r,0}(D_\star) = \{D_\star\}, \quad s(D) := \log D - \log M_0(D), \quad s(D_\star) = 0 < s(D) \quad (D \neq D_\star).$$

On the exact reversible branch every checkpoint generator is injective, so

$$M_0(q) = |X_{\text{reach}}(q)|,$$

reducing the finite computation to exact CSP or model counting. This is the preferred first implementation of the source-only finite public checkpoint packet $\mathcal{C}_{r,D}^{\text{pub}}$.

The local null-data route is blind to metric-term shifts, so the Einstein branch is fixed locally only modulo $\Lambda_{g_{ab}}$. Conditional on stable direct closure and the independent horizon-record identification,

$$N_\star = \log D_\star = \frac{A_{\text{dS}}}{4\ell_\star^2}, \quad \Lambda_\star \ell_\star^2 = \frac{3\pi}{N_\star}.$$

With the selected scale certificate the same branch has the display

$$G_{ab} + \frac{3\pi}{GN_\star} g_{ab} = 8\pi G \langle T_{ab} \rangle.$$

The independent common screen/electroweak load-carrier identification tests the same closed N against the weak/Higgs load through

$$R_{\text{EW}}(P, N) = \alpha_U(P) \log\left(\frac{N}{\pi}\right) - \frac{6\pi}{P} = 0, \quad N_{\text{bridge}} = \pi \exp\left[\frac{6\pi}{P\alpha_U(P)}\right].$$

Neither physical bridge constructs the direct capacity map.

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

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

The observed late-time scale gives $N_{\text{patch}} \simeq 1.05 \times 10^{122}$, with a Planck- Λ central entropy capacity $N_\Lambda \simeq 3.313 \times 10^{122}$. The conditional electroweak bridge value $N_{\text{EW}} \simeq 3.532 \times 10^{122}$ is about 6.6 percent higher.

Fixed-capacity shock sign. The pure de Sitter shock parameter obeys

$$\mu^2 = (d-2)\kappa r_c = d-2 = \lambda_{\ell=1}(S^{d-2}),$$

so the cosmological constant cancels and the smooth isometry modes lie at zero shock eigenvalue. For finite sector dimensions, maximizing

$$S_{\text{gen}}(p, d) = -\sum_i p_i \log p_i + \sum_i p_i \log d_i$$

over p gives $p_i = d_i/M$ and $S_{\text{gen}}^{\text{max}} = \log M$. The associated logarithmic area observable $\mathcal{A} = \sum_i (d_i/M) \log d_i$ has symmetric-point Hessian

$$\text{Hess } \mathcal{A}|_{d_i=d} = \frac{1}{nd^2} \left(I - \frac{2}{n} J \right)$$

on the positive-real relaxation of the integer dimensions. Its curvature is positive on fixed- M tangent directions and negative in the homogeneous direction; the gradient at the symmetric point is nonzero. The physical sign comes from a one-sided transfer. Assigning a fraction f of a fixed total horizon–observer capacity to the observer uniformly depletes the horizon and gives

$$\Delta S_{\text{gen}}^{\text{max}} = \Delta \mathcal{A} = \log(1 - f) < 0.$$

The finite identity applies to admissible integer depletions, and its positive-real interpolation is strictly decreasing. Under the independent horizon-area and fixed-budget dictionaries, this is the de Sitter time-advance sign studied in Ref. [27]. It is a boundary maximum along the transfer direction, rather than an interior maximum at fixed horizon capacity. Identifying the finite capacity-ledger transfer with observer mass is an additional physical dictionary.

On the icosahedral carrier, the spectrum

$$\{-2, 0, 1 + 3/\sqrt{5}, 1 + \sqrt{5}\}$$

is conditional on two further statements: the A_5 rotation triplet is an exact gauge sector, and the shock kinetic term is the scaled nearest-neighbour port or edge-sector Laplacian. The regular line-graph identity makes the port and edge-sector low spectra equal. It adds the eigenvalue 10 with multiplicity 18 on the edge carrier. The shock coefficient, physical attachment, and a generator for sector-dimension shocks are open. The finite heat-bath repair generator acts at fixed sector dimensions and therefore supplies none of these objects. The sign mechanism does not rescue the static-patch trace conjecture [11]. That obstruction tests the cited positive cyclic trace interpretation. Nontrace horizon-screen descriptions, separately constructed dimension-changing generators, and source-derived kinetic operators beyond the scaled nearest-neighbour graph remain outside these classes. None of these boundaries disproves OPH; the pure-de-Sitter normalization, finite entropy and transfer identities, and regular line-graph theorem survive.

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 boundary. 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. On the conditional D6 branch, the dimensionless de Sitter capacity relation is the shared global constraint across overlap-consistent descriptions; its SI Λ display uses the selected scale certificate.

Summary. The model moves the holographic screen from “infinity” to an observer’s horizon. Direct N closure is a correctable public-record theorem candidate, not a checkpoint-fixed or marginal-information count. Horizon saturation and the weak/Higgs common-load equation are independent downstream commuting squares. The weak/Higgs bridge and the late-time de Sitter coordinate differ by about 6.6 percent; that comparison becomes physical only after the direct packet and both bridges share a certified carrier.

2.4 Einstein-chain status and falsification boundary

The finite support and algebraic packets establish the following statements without assuming a spacetime: repair-invariant incidence, the spherical branch receipts, cap-normal identities, the Lorentz action on the celestial screen, exact finite entropy splitting on the central-interface branch, and the finite tomography identities used to reconstruct a trace-free symmetric tensor from null projections.

The four-dimensional event manifold is conditional on the event receipts stated in Theorem 4.3e. The Einstein equation is conditional on a single source-derived common-domain tower carrying the modular, scaling, stress, universal-coupling, vacuum-reference, and scale-identification premises. Finite certificate programs verify their stated algebra and negative controls. They do not establish that one physical refinement tower carries every premise. That construction is work in progress.

The chain fails if any required receipt does not persist on one cofinal branch, if the held-out cone fit does not have inertia $(1, 3)$, if the directional charges violate tensor-linearity, if the small-diamond remainder is not $o(\ell^4)$, or if the scale readouts retain a nontrivial common rescaling. A three-dimensional observer-frame space is not evidence by itself for a four-dimensional event manifold.

A Technical Details for the Specialist Derivations

This appendix carries the gravity details that support the spacetime and Einstein derivation. It sharpens the Lorentz, Einstein, and cosmological-constant surfaces without widening the recovered core beyond its stated branch conditions.

A.1 Lorentz and Einstein Bridge

Theorem A.1 (Conformal group isomorphism). *The orientation-preserving conformal group of S^2 is isomorphic to the connected Lorentz group:*

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

Theorem A.2 (Support-visible geometric modular flow on caps on the extracted geometric subnet). *Assume the OPH axioms, the derived fixed-cutoff collar package, and the support-visible BW scaling theorem on the extracted geometric subnet. That theorem requires the finite cap-normal support/flow certificate and an independently complete MGNS-1 algebra-state package on the same tower. Then the scaling-limit modular automorphism group of the cap pair is geometric conformal dilation on that subnet, with the standard 2π normalization supplied by the finite cap-normal certificate. If the emitted scaling-limit cap algebra is type I, this may be written as $K_C = 2\pi B_C + Z_C$ with Z_C central; in the generic continuum case the theorem is the automorphism statement on a non-type-I cap algebra with outer geometric modular action. The collar replacements used in this branch are exact only at exact Markovity or in controlled fixed-collar families with $\delta^M \rightarrow 0$; small CMI supplies a Fawzi–Renner recovered comparison state, not a dimension-free one-shot exact Markov normal form.*

Definition A.3 (BW-branch observer-relative modular ordering). *On the branch satisfying the hypotheses of the preceding geometric modular-flow theorem, the modular automorphism parameter t of the extracted cap pair supplies a dimensionless ordering for that observer’s accessible algebra-state pair. Physical time requires an observer-readable transition, event correspondence, and calibrated clock instrument. No claim about arbitrary operational clocks, global time, or the full problem of time follows from the geometric modular-flow theorem alone.*

BW-side UV scaffold. The UV data on this branch are not a separate cap-isotropy or Euclidean-regularity selector. The fixed-cutoff cap algebras are type-I regulators, but the scaling-limit observer algebra may leave that class, and MaxEnt alone does not select the BW / canonical cap phase. The UV package is the realized transported geometric cap-local system together with the carried-collar schedule derived from the transported fixed-local-collar Markov/faithfulness datum on each fixed local collar model, followed by support-readable modular covariance and ordered cut-pair rigidity on the emitted scaling-limit geometric cap pair.

Null-strip completion boundary. On the D4 side, the fixed-cutoff strip package is more specific than the generic inherited-strip summary. The null cuts first transfer the same cut-center data as the spatial collar branch, which fixes the central sector-pair decomposition of the strip algebra. The stronger left/right tensor decomposition used by the null modular bridge then requires the extra inherited strip-split condition on the multiplicity spaces, together with the exact-or-controlled Markov hypotheses on one fixed inherited strip model. On that same fixed strip model, the renormalized half-line family is endpoint-Lipschitz, hence defines the weak tail generator, and on the scaling-limit geometric-cap branch the half-line blow-up net carries the derived half-sided modular pair. Borchers–Wiesbrock then supplies the positive null-translation generator on its Stone domain together with the affine half-line modular relation $K_a(\Omega) = K_0(\Omega) - 2\pi a P_\Omega$, and the same half-line family fixes the generator/charge identification internally. Bounded-interval formulas are downstream of the E0.5 affine/projective kernel in the Einstein bridge.

Lemma A.4 (Null data ambiguity). *If a symmetric tensor X_{ab} satisfies*

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

for every null vector k , then $X_{ab} = \phi g_{ab}$ for some scalar ϕ .

Corollary A.5 (Null modular data determine Einstein only up to the metric term). *Null modular data determine T_{ab} only up to ϕg_{ab} . Consequently the Einstein equation is fixed locally only up to Λg_{ab} .*

Theorem A.6 (Jacobson-type rest-frame relation). *In the local Lorentzian scaling regime, once the realized cap-label-preserving MaxEnt family satisfies the derived fixed-cap generalized-entropy stationarity theorem for admissible fixed-cap variations, the half-line generator/charge identification of the null bridge, the bounded-interval transport input used in the local Lorentzian regime, and the internal small-ball bridge yield the rest-frame first-variation relation that drives the spacetime and Einstein paper’s local Einstein branch; on the same scaling branch, if that rest-frame relation holds for all local observer four-velocities and all reference states, the spacetime and Einstein paper upgrades it to the full tensor equation by the explicit local quadratic-polarization argument in the main derivation.*

Theorem A.7 (Newton coupling from the scale certificate and edge entropy density). *The geometric coupling read by the Newton area law is*

$$G_{\text{geom}} = \frac{a_{\text{cell}}}{4\bar{\ell}(t)}.$$

On the OPH gravity row the scale is supplied first by $\gamma_\star = \ell_\star \nu_{\text{Cs}}/c$, equivalently $B_\star = 3\pi/\ell_\star^2$, as an independent scale certificate rather than a consequence of $P_\star$ and $N_\star$ alone. With $a_{\text{cell}} = P\ell_\star^2$ and $\bar{\ell}_{\text{shared}} = P/4$, this gives

$$G_{\text{geom}} = \ell_\star^2, \quad G_{\text{SI}} = \frac{c^3 \ell_\star^2}{\hbar}.$$

Here $\bar{\ell}_{\text{shared}}$ is a shared-cut density. It is not the logarithm of an independent cell Hilbert-space dimension and does not select a fixed primitive observer capacity.

A.2 Cosmological-Constant / Screen-Capacity Closure

Lemma A.8 (Vacuum energy blindness). *For any null vector k ,*

$$T_{kk}^{\text{vac}} = T_{ab}^{\text{vac}} k^a k^b = -\rho_{\text{vac}} g_{ab} k^a k^b = 0.$$

Vacuum-energy contributions therefore lie in the kernel of the null map $T_{ab} \mapsto T_{kk}$.

Proposition A.9 (Structural separation). *Within the structural split used here:*

1. *local modular/null data fix T_{ab} only up to ϕg_{ab} ; and*
2. *the dimensionless Λ -capacity relation is fixed by global screen capacity, not by local null data.*

So the large vacuum-energy bookkeeping of EFT is not itself the local quantity determining curvature on this branch.

Theorem A.10 (No local Λ prediction). *Within the null-modular reconstruction used here:*

1. *all T_{kk} data are unchanged under $T_{ab} \rightarrow T_{ab} + \phi g_{ab}$;*
2. *therefore Λ cannot be fixed by local overlap consistency alone; and*
3. *the dimensionless Λ -capacity relation requires the global screen-capacity closure beyond the local null data. The stable closure target is $\mathfrak{F}_{r,0}(D_\star) = \{D_\star\}$, with $N_{\text{CRC}} = \log D_\star$. The correctable public-record definition and conditional operational closure theorem close its operational public-section, correctable-code, capacity-bound, approximate-stability, and greatest-fixed-point implications under named premises. The full specification and receipt schema are carried*

by Observers Are All You Need. *The record-atom restrictions, endogenous reachability, frozen publicness policy, global checkpoint coupling, capacity-carrier representation, whole-fiber scalarization, confusability-reflecting extension and refinement packets, finite-size slack law with one physical zero, and horizon-record identification are work in progress. The electroweak G2 relation is an independent comparison after F and does not construct it. SI curvature values additionally use the selected scale certificate.*

Corollary A.11 (Benchmark cosmological readout on the D6 branch). *On the conditional D6 branch defined above, inserting the observed value $\Lambda \approx 1.09 \times 10^{-52} \text{ m}^{-2}$ gives*

$$\begin{aligned} r_{\text{dS}} &= \sqrt{\frac{3}{\Lambda}} \approx 1.66 \times 10^{26} \text{ m}, & t_{\Lambda} &= \frac{r_{\text{dS}}}{c} \approx 17.5 \text{ Gyr}, \\ N_{\text{patch}} &= \left(\frac{r_{\text{dS}}}{\ell_P} \right)^2 \approx 1.05 \times 10^{122}, & N_{\text{scr}} &= S_{\text{dS}} = \pi N_{\text{patch}} \approx 3.31 \times 10^{122}, \\ \Lambda \ell_P^2 &\approx 2.85 \times 10^{-122}. \end{aligned}$$

Proof. Immediate from the displayed static-patch and capacity relations. $\square$

Observed-age benchmark boundary. The branch quantity t_{Λ} is the de Sitter static-patch timescale. The usual cosmic age t_0 is a compare-only FLRW benchmark, with no additional D6 theorem output. On the flat Λ CDM benchmark

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_{\Lambda}}} \sinh^{-1} \left(\sqrt{\frac{\Omega_{\Lambda}}{\Omega_m}} \right)$$

one gets the standard comparison value $t_0 \approx 13.8 \text{ Gyr}$ at $H_0 \approx 67.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_{\Lambda} \approx 0.685$.

Lemma A.12 (FLRW curvature as visible scalar holonomy). *On a homogeneous-isotropic spatial slice with constant sectional curvature K , small spatial loop holonomy obeys*

$$\text{Hol}_{\square_{uv}} = \exp \left(K A_{\square} J_{uv} + O(A_{\square}^{3/2}) \right).$$

On a visibly separated OPH refinement system, the refinement-limit scalar spatial holonomy vanishes if and only if $K = 0$. Thus a flat FLRW branch is the zero-visible-spatial-holonomy branch.

Flatness boundary. This holonomy statement names a conditional cosmology bridge only. It does not add a D6 theorem output and does not solve the inflationary flatness or horizon problems. Selecting $K = 0$ requires an additional continuation theorem or premise: a direct flatness theorem, a conditional cosmological-minimal-holonomy selector on a clocked FLRW boundary, or an explicit flat-branch assumption. Minimal Admissible Realization (MAR) acts on low-energy gauge/matter packages and is not a cosmological flatness selector.

Theorem A.13 (D5–D6 cosmological-capacity closure stack). *Assume the spacetime and Einstein paper’s local Einstein branch, stable direct public-record closure and its capacity coordinate*

$$\mathfrak{F}_{r,0}(D_{\star}) = \{D_{\star}\}, \quad N_{\text{CRC}} = \log D_{\star},$$

and the independent horizon-record identification readout

$$N_{\text{CRC}} = S_{\text{dS}},$$

the standard de Sitter entropy relation

$$S_{\text{dS}} = \frac{A_{\text{dS}}}{4G} = \frac{3\pi}{G\Lambda},$$

and the standard static-patch formulas

$$r_{\text{dS}} = \sqrt{\frac{3}{\Lambda}}, \quad t_{\Lambda} = \frac{r_{\text{dS}}}{c}.$$

Then the cosmological-constant package is one local/global theorem stack: local null data fix the Einstein branch only modulo Λg_{ab} , and with the selected scale certificate the same branch has the global display

$$G_{ab} + \frac{3\pi}{GN_{\text{CRC}}} g_{ab} = 8\pi G \langle T_{ab} \rangle,$$

while the D6 closure itself fixes the entropy and dimensionless capacity relations

$$S_{\text{dS}} = N_{\text{CRC}}, \quad A_{\text{dS}} = 4GN_{\text{CRC}}, \quad r_{\text{dS}} = \sqrt{\frac{3}{\Lambda}}, \quad t_{\Lambda} = \frac{r_{\text{dS}}}{c},$$

and the observed cosmic age is a downstream FLRW benchmark, with no additional theorem output.

Scope boundary. The D6 hypotheses are exactly the D5 local Einstein branch, stable whole-fiber public-record closure $\mathfrak{F}_{r,0}(D_{\star}) = \{D_{\star}\}$, horizon-record identification $N_{\text{CRC}} = S_{\text{dS}}$, the standard de Sitter entropy relation, and the standard static-patch formulas. The local null-data route does not by itself determine the global capacity; that value is fixed only on a discharged stable-public-record branch with horizon-record identification. Conditional on that branch, the capacity closure fixes $\Lambda_{\text{CRC}} \ell_{\star}^2 = 3\pi/N_{\text{CRC}}$; the SI static-patch scale additionally requires the selected scale certificate. CMB kernels, inflation-replacement claims, H_0/S_8 branches and dark/anomaly growth kernels require separate continuation theorems or likelihood contracts. The baryogenesis anomaly/current theorem and gauge/deck no-go are separate from the global-capacity branch; the anomalous record attachment and physical CP-odd source generator are work in progress.

Theorem A.14 (Conditional OPH screen-spectrum theorem). *Let (S_r) be a cofinal finite spherical OPH screen system with a schedule-independent quotient-normal-form scalar q_r , removal of the background and dipole sector, and a target-free positive quadratic repair operator K_r . If $K_r \rightarrow K$, the source-selected finite scalar release energy satisfies $2E_{q,r}^{\text{src}}/d_r \rightarrow A_q$, and local MaxEnt is imposed at fixed expected quadratic release energy, then q_r converges in finite harmonic distributions to the centered Gaussian screen field with covariance $A_q K^{-1}$. Exact per-sample release energy would instead give a microcanonical ellipsoid. If the certified repair-scale measure is $t^{\theta/2} dt/\Gamma(1 + \theta/2)$, then*

$$K_{\text{asy}} = (-\Delta_{S^2})^{1+\theta/2}, \quad C_{\ell,\text{asy}}^q = A_q [\ell(\ell+1)]^{-1-\theta/2},$$

as the large- ℓ asymptotic model. The theorem-grade finite- ℓ conformal-shell family uses the normalized gamma-ratio precision

$$\kappa_{\ell}(\theta) = \frac{\Gamma(\ell+2+\theta/2)}{\Gamma(\ell-\theta/2)},$$

with

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

If q is additionally certified as the thin-shell pullback of a homogeneous curvature field with $\Delta_\zeta^2(k) = A_\zeta(k/k_\star)^{-\theta}$, the corresponding inverse amplitude relation is

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

Suppose the same source construction emits a strongly continuous full-collar survival cocycle with generator density $P_\star/24$, together with the orientation-reversal half-collar identity. Then

$$\theta = \frac{P_\star}{48}, \quad n_s = 1 - \frac{P_\star}{48}, \quad \kappa_{\text{rep}}^{\text{edge}} = \frac{P_\star}{48(P_\star - \varphi)}.$$

A finite one-step survival value determines θ through $-\log u_q(\log b)/\log b$; it does not supply the infinitesimal generator receipt. If a scale-natural source embedding transports this cocycle to the physical covariance through $D_s^{-1} C_\zeta D_s = e^{-\theta s} C_\zeta$, the source family is $\Delta_\zeta^2(k) = A_\zeta(k/k_\star)^{-\theta}$. A single shell has an infinite-dimensional radial kernel, including positive ambiguities. Complete radial cross-covariances provide the independent tomography route. A Fawzi–Renner/Markov-collar observable estimate gives an upper bound on release observables, not an amplitude equality. No physical TT/TE/EE statement follows without independent Boltzmann transfer and likelihood closure.

Screen-spectrum boundary. The conditional theorem fixes the exact angular family. Its radial theorem gives physical source-dilation and cross-covariance tomography as separate uniqueness routes. A finite OPH source DAG that emits the primitive collar ensemble, full-collar generator density, half-collar identity, conformal precision, and physical dilation-intertwiner or tomography receipt is work in progress. The thin-shell lift uses the gamma-ratio operator above; finite-width windows require a Bessel-kernel error certificate. Physical TT/TE/EE transfer and a frozen likelihood contract are separate continuation tasks.

Separate pixel-side benchmarks. The companion pixel numbers

$$a_{\text{cell}} \approx 1.63 \ell_\star^2, \quad \ell_{\text{UV}} = \sqrt{a_{\text{cell}}} \approx 1.28 \ell_\star, \quad \bar{\ell} \approx 0.408$$

belong to the separate pixel-closure/local-readout package, outside the D6 cosmological-parameter corollary itself. After the selected scale certificate is supplied, $\ell_\star$ is displayed as the usual Planck length.

References

- [1] B. Müller, D. Matscheko, and J. Hill, *Observation-Determined Normal Forms: Stability, Obstructions, and Refinement in Constraint and Rewrite Systems*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/observable_normal_forms.pdf.
- [2] B. Müller, A. Osika, M. Ponder, K. Xue, P. Nguyen, M. A. Visser, and D. Matscheko, *Recovering Observer Spacetime and Einstein Dynamics from Overlap Consistency*, 2026. Available at [GitHub PDF](#).
- [3] B. Müller, A. Osika, M. Ponder, K. Xue, P. Nguyen, M. A. Visser, and D. Matscheko, *Deriving Standard Model Gauge Structure from Observer Overlap Consistency*, 2026. Available at [GitHub PDF](#).

- [4] B. Müller, K. Xue, K. A. Anirudha, D. Matscheko, and J. Hill, *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/reality_as_consensus_protocol.pdf.
- [5] B. Müller, A. Osika, K. Xue, B. Cassie, M. A. Visser, and D. Matscheko, *Federated Echoshedral Screen Microphysics: Patch Hardware, Records, and Observer Synchronization in OPH*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/screen_microphysics_and_observer_synchronization.pdf.
- [6] B. Müller, A. Osika, M. Ponder, K. Xue, M. A. Visser, and D. Matscheko, *Deriving the Particle Zoo from Observer Consistency*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/deriving_the_particle_zoo_from_observer_consistency.pdf.
- [7] B. Müller, A. Osika, and D. Matscheko, *Theoretical Bounds on χ_ν in Observer-Patch Holography*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/chi_nu_susceptibility_bounds.pdf.
- [8] B. Müller and D. Matscheko, *Observer-Patch Holography and the Dark Matter Phenomenon*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/cosmology/oph_dark_matter_paper.pdf.
- [9] B. Müller, *Explaining the Yang–Mills Mass Gap with Observer-Patch Repair Dynamics*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/yang_mills_gap_clay_problem.pdf.
- [10] B. Müller, *Observer-Patch Holography as a String-Vacuum Selector: Observers, Clocks, Edge Strings, and the Bouchard-Donagi Witness*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/observer_patch_holography_as_string_vacuum_selector.pdf.
- [11] B. Müller, *The de Sitter Time-Advance Sign from a Finite Screen with Fixed Capacity*, 2026. Available at https://github.com/FloatingPragma/observer-patch-holography/blob/main/extra/de_sitter_time_advance_sign_from_fixed_screen_capacity.pdf.
- [12] B. Müller, *OPH-FPE: finite simulator and receipt engine for Observer-Patch Holography physics experiments*, 2026. Available at <https://github.com/muellerberndt/oph-physics-sim>.
- [13] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro, and T. Schwetz, *NuFit-6.0: Updated global analysis of three-flavor neutrino oscillations*, JHEP **12** (2024) 216, arXiv:2410.05380, with the NuFIT 6.1 (2025) profile-table release at <https://www.nu-fit.org/?q=node/309>.
- [14] M. H. A. Newman, “On theories with a combinatorial definition of ‘equivalence’,” *Ann. of Math. (2)* **43** (1942), 223–243.
- [15] E. H. Lieb and D. W. Robinson, “The finite group velocity of quantum spin systems,” *Commun. Math. Phys.* **28** (1972), 251–257.
- [16] J. J. Bisognano and E. H. Wichmann, “On the duality condition for a Hermitian scalar field,” *J. Math. Phys.* **16** (1975), 985–1007.

- [17] J. J. Bisognano and E. H. Wichmann, “On the duality condition for quantum fields,” *J. Math. Phys.* **17** (1976), 303–321.
- [18] R. Brunetti, D. Guido, and R. Longo, “Modular Structure and Duality in Conformal Quantum Field Theory,” *Commun. Math. Phys.* **156** (1993), 201–219, arXiv:funct-an/9302008.
- [19] H.-W. Wiesbrock, “Half-Sided modular inclusions of von-Neumann-Algebras,” *Commun. Math. Phys.* **157** (1993), 83–92.
- [20] H. Araki and L. Zsidó, “Extension of the structure theorem of Borchers and its application to half-sided modular inclusions,” *Rev. Math. Phys.* **17** (2005), 491–543, arXiv:math/0412061.
- [21] D. Guido, R. Longo, and H.-W. Wiesbrock, “Extensions of Conformal Nets and Superselection Structures,” *Commun. Math. Phys.* **192** (1998), 217–244, arXiv:hep-th/9703129.
- [22] H.-W. Wiesbrock, “Modular Intersections of von-Neumann-Algebras in Quantum Field Theory,” *Commun. Math. Phys.* **193** (1998), 269–285.
- [23] W. G. Unruh, “Notes on black-hole evaporation,” *Phys. Rev. D* **14** (1976), 870–892.
- [24] T. Jacobson, “Thermodynamics of spacetime: The Einstein equation of state,” *Phys. Rev. Lett.* **75** (1995), 1260–1263, doi:10.1103/PhysRevLett.75.1260, arXiv:gr-qc/9504004.
- [25] T. Jacobson, “Entanglement equilibrium and the Einstein equation,” *Phys. Rev. Lett.* **116** (2016), 201101, doi:10.1103/PhysRevLett.116.201101, arXiv:1505.04753.
- [26] G. W. Gibbons and S. W. Hawking, “Cosmological event horizons, thermodynamics, and particle creation,” *Phys. Rev. D* **15** (1977), 2738–2751, doi:10.1103/PhysRevD.15.2738.
- [27] Y. Chen, D. Stanford, H. Tang, and Z. Yang, “Negative shocks versus static patch holography,” arXiv:2607.14042 [hep-th] (2026). <https://arxiv.org/abs/2607.14042>
- [28] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* **641** (2020), A6, doi:10.1051/0004-6361/201833910, arXiv:1807.06209.
- [29] R. Bousso, Z. Fisher, J. Koeller, S. Leichenauer, and A. C. Wall, “Proof of the quantum null energy condition,” *Phys. Rev. D* **93** (2016), 024017, arXiv:1509.02542.
- [30] R. Bousso, Z. Fisher, S. Leichenauer, and A. C. Wall, “Quantum focusing conjecture,” *Phys. Rev. D* **93** (2016), 064044, arXiv:1506.02669.
- [31] S. Balakrishnan, T. Faulkner, Z. U. Khandker, and H. Wang, “A general proof of the quantum null energy condition,” *J. High Energy Phys.* **2019** (2019), 20, arXiv:1706.09432.
- [32] E. H. Lieb and M. B. Ruskai, “Proof of the strong subadditivity of quantum-mechanical entropy,” *J. Math. Phys.* **14** (1973), 1938–1941.
- [33] E. H. Wichmann, “Density matrices arising from incomplete measurements,” *J. Math. Phys.* **4** (1963), 884–896.
- [34] F. Hiai, M. Ohya, and M. Tsukada, “Sufficiency, KMS condition and relative entropy in von Neumann algebras,” *Pacific J. Math.* **96** (1981), 99–109.
- [35] M. Takesaki, “Conditional expectations in von Neumann algebras,” *J. Funct. Anal.* **9** (1972), 306–321.

- [36] D. Petz, “Sufficient subalgebras and the relative entropy of states of a von Neumann algebra,” *Commun. Math. Phys.* **105** (1986), 123–131.
- [37] D. Petz, “Sufficiency of channels over von Neumann algebras,” *Quart. J. Math.* **39** (1988), 97–108.
- [38] O. Fawzi and R. Renner, “Quantum conditional mutual information and approximate Markov chains,” *Commun. Math. Phys.* **340** (2015), 575–611, arXiv:1410.0664.
- [39] P. Hayden, R. Jozsa, D. Petz, and A. Winter, “Structure of states which satisfy strong subadditivity of quantum entropy with equality,” *Commun. Math. Phys.* **246** (2004), 359–374.
- [40] S. Doplicher and J. E. Roberts, “A new duality theory for compact groups,” *Invent. Math.* **98** (1989), 157–218.
- [41] S. Doplicher and J. E. Roberts, “Why there is a field algebra with a compact gauge group describing the superselection structure in particle physics,” *Commun. Math. Phys.* **131** (1990), 51–107.
- [42] T. Tannaka, “Über den Dualitätssatz der nichtkommutativen topologischen Gruppen,” *Tohoku Math. J.* **45** (1938), 1–12.
- [43] M. G. Krein, “A principle of duality for a bicomact group and a square block algebra,” *Dokl. Akad. Nauk SSSR* **69** (1949), 725–728.
- [44] H. Georgi and S. L. Glashow, “Unity of all elementary-particle forces,” *Phys. Rev. Lett.* **32** (1974), 438–441.
- [45] S. L. Glashow, J. Iliopoulos, and L. Maiani, “Weak interactions with lepton-hadron symmetry,” *Phys. Rev. D* **2** (1970), 1285–1292.
- [46] E. Witten, “An $SU(2)$ anomaly,” *Phys. Lett. B* **117** (1982), 324–328.
- [47] S. Dimopoulos, S. Raby, and F. Wilczek, “Supersymmetry and the scale of unification,” *Phys. Rev. D* **24** (1981), 1681–1683.
- [48] U. Amaldi, W. de Boer, and H. Fürstenau, “Comparison of grand unified theories with electroweak and strong coupling constants measured at LEP,” *Phys. Lett. B* **260** (1991), 447–455.
- [49] D. J. Gross and W. Taylor, “Two-dimensional QCD is a string theory,” *Nucl. Phys. B* **400** (1993), 181–208, arXiv:hep-th/9301068.
- [50] F. Peter and H. Weyl, “Die Vollständigkeit der primitiven Darstellungen einer geschlossenen kontinuierlichen Gruppe,” *Math. Ann.* **97** (1927), 737–755.
- [51] A. Bullivant, M. Calçada, Z. Kádár, P. Martin, and J. Faria Martins, “Topological phases from higher gauge symmetry in 3+1D,” *Phys. Rev. B* **95** (2017), 155118, arXiv:1606.06639.
- [52] S. Chandrasekharan and U.-J. Wiese, “Quantum link models: A discrete approach to gauge theories,” *Nucl. Phys. B* **492** (1997), 455–471, arXiv:hep-lat/9609042.
- [53] W. Donnelly and A. C. Wall, “Entanglement entropy of electromagnetic edge modes,” *Phys. Rev. Lett.* **114** (2015), 111603, arXiv:1412.1895.

- [54] F. Pastawski, B. Yoshida, D. Harlow, and J. Preskill, “Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence,” *JHEP* **06** (2015), 149, arXiv:1503.06237.
- [55] M. A. Levin and X.-G. Wen, “String-net condensation: A physical mechanism for topological phases,” *Phys. Rev. B* **71** (2005), 045110, arXiv:cond-mat/0404617.
- [56] D. Laghi, G. Carullo, J. Veitch, and W. Del Pozzo, “Quantum black hole spectroscopy: probing the quantum nature of the black hole area using LIGO-Virgo ringdown detections,” *Class. Quantum Grav.* **38** (2021), 095005, arXiv:2011.03816.
- [57] Clay Mathematics Institute, “Yang–Mills & the Mass Gap.” Available at <https://www.claymath.org/millennium/yang-mills-the-maths-gap/>.
- [58] A. Jaffe and E. Witten, “Quantum Yang–Mills Theory,” official Clay Mathematics Institute problem description. Available at <https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf>.
- [59] L. Wu, “Poincaré and transportation inequalities for Gibbs measures under the Dobrushin uniqueness condition,” *Annals of Probability* **34** (2006), 1960–1989, arXiv:math/0611635.
- [60] K. Osterwalder and R. Schrader, “Axioms for Euclidean Green’s functions,” *Communications in Mathematical Physics* **31** (1973), 83–112.
- [61] K. Osterwalder and R. Schrader, “Axioms for Euclidean Green’s functions II,” *Communications in Mathematical Physics* **42** (1975), 281–305.
- [62] Particle Data Group, “Review of Particle Physics,” *Phys. Rev. D* **110** (2024), 030001. Available at <https://pdg.lbl.gov/2024/download/db2024.pdf>.
- [63] Particle Data Group, “Electroweak Model and Constraints on New Physics,” Available at <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-standard-model.pdf>.