

The Operator Equation of State: Belt-Local Modular Dynamics for Quantum Gravity

F. Jofer

Abstract

We¹ adopt a belt-local, operator-level equation of state (OES) as the core axiom: on any admissible belt, the boundary modular generator equals the bulk generalized-entropy operator on the associated wedge (operator JLMS). This *Operator Equation of State* (OES) is our organizing principle. From it we recover the linear and quantified second-order semiclassical Einstein relations on belts with regulator stability, using a minimal kernel built from a belt first-law channel, OS positivity with flow removal, and a Brown–York/Iyer–Wald identification.

A main component of this work is a *cubic (third-order) verification* of OES in a controlled AdS₃/CFT₂ shockwave setup. There we compute and compare the third variations of the two sides—boundary modular and bulk generalized entropy—using belt kernels and canonical-energy inputs, and we demonstrate numerical agreement to high precision under grid refinement and regulator removal. This pushes the holographic test beyond linear order and provides operator-level evidence that the OES governs wedge dynamics at cubic order.

All statements are per generator length, uniform in region size, and ledgered by a single belt budget that vanishes under flow removal. The construction is regulator-stable under JKM corner calibration and compatible with dispersion/positivity constraints used elsewhere in the paper.

1 Introduction and roadmap

Core principle. At the heart of this paper is the *Operator Equation of State* (OES): on a belt-anchored wedge $W = EW(R)$, the boundary modular generator equals the bulk generalized-entropy operator,

$$\widehat{K}_{\text{mod}}(R) = \widehat{S}_{\text{gen}}(W) := \frac{\widehat{\mathcal{A}}(W)}{4G} + \widehat{H}_{\text{bulk}}(W).$$

We stress that \widehat{S}_{gen} is not the von Neumann generalized entropy operator; its expectation uses the bulk modular generator $\widehat{H}_{\text{bulk}}$, so $\langle \widehat{S}_{\text{gen}} \rangle \neq S_{\text{gen}}$ beyond linear order. Precisely, let $D := D_{\text{an}}$ denote the common quadratic-form core furnished by the belt OS–KMS kernel (stable under the positive belt flows); on D both $\widehat{K}_{\text{mod}}(R)$ and $\widehat{S}_{\text{gen}}(W)$ are closable and admit the same closed quadratic-form extension with form domain \overline{D} ; see Definition 5.18, Lemmas 3.1 and 3.3, and Section 5.5. We adopt OES as Axiom 5.13, develop its operator meaning in Section 5.5, and fix its normalization by the Rindler witness and the JKM calibration [1, 2]. Taking expectations and variations yields the familiar linear and second-order modular/Einstein relations on belts, now packaged as consequences of a single operator statement.

¹Throughout, "we" denotes the conventional authorial plural; the paper has a single human author. The large language model used in preparing this work (see "AI Use and Author Responsibility" at the end) is not included in "we".

What we do. We give an axiomatic, belt-local route to semiclassical quantum gravity anchored on OES and supported by three ingredients: (i) a belt first-law channel and modular convexity [3, 4], (ii) OS (RP/KMS) positivity with removal of short flows [1, 5], and (iii) a JLMS-type boundary–bulk map realized as a Brown–York/Iyer–Wald flux dictionary [3, 6, 7]. Within this kernel we obtain regulator-stable identities (errors $O(B_{\text{belt}})$) that become exact after flow removal.

New cubic verification. Beyond the established linear and second-order checks, we *push the holographic test to third order*. In Section 8.6 we implement an AdS₃/CFT₂ shockwave protocol that computes

$$\delta^3\langle K_{\text{mod}} \rangle \quad \text{and} \quad \delta^3\langle \widehat{S}_{\text{gen}} \rangle$$

on the same belt, with a belt kernel $w(u)$ and a canonical-energy organization of the bulk side, and we demonstrate that the two sides match to high numerical precision under grid refinement. This provides an operator-level test of OES beyond the semiclassical first and second variations.

Setting (belt regulator) and scope. We work on globally hyperbolic backgrounds with belt-anchored regions. The regulator has two knobs: a geometric belt width $r > 0$ and short positive flows $(u, s) > 0$ (null/bulk evolutions). All identities are proved at finite $(r; u, s)$ and then the flows are removed with a single, uniform budget $O(B_{\text{belt}})$; all statements are per generator length and regulator-stable under JKM calibration.

The kernel (three inputs). Our arguments repeatedly use: (1) a belt first-law channel; (2) OS positivity with flow removal, giving convexity/monotonicity; and (3) a belt JLMS identification that matches boundary relative entropy to wedge canonical energy up to $O(B_{\text{belt}})$. These inputs ensure that the belt-level laws are stable, additive on overlaps, and compatible with quasi-local flux transport.

Global dynamics and stitching. From the kernel we build a global belt atlas: we define admissible belt variations, a modular defect and Hessian, and prove that the defect vanishes while the Hessian equals Iyer–Wald canonical energy on the wedge, up to the ledgered budget. Path-independence and stitching on overlaps then force a regulator-independent OES on each belt, recovering the linearized and quantified second-order semiclassical Einstein equations in expectation.

Main deliverables.

1. **Operator Equation of State (axiom).** The operator identity $\widehat{K}_{\text{mod}} = \widehat{S}_{\text{gen}}$ on belts (Axiom 5.13), with normalization fixed by JKM and a Brown–York dictionary [2, 6]; equivalent modular, group, channel, and relative formulations are given in Section 5.5 (cf. [3]).
2. **Linear and second order.** Regulator-stable belt identities implying the linearized and quantified second-order semiclassical Einstein equations in expectation (modular first law, canonical energy, shear/expansion control) (cf. [4, 7]).
3. **Cubic (third-order) holographic verification.** A belt-local AdS₃/CFT₂ shock test that matches $\delta^3\langle K_{\text{mod}} \rangle$ with $\delta^3\langle \widehat{S}_{\text{gen}} \rangle$ to high precision under grid refinement (Section 8.6), organized around the canonical-energy structure of Proposition 5.59 (cf. [7]).
4. **Positivity bridges.** Gravity-subtracted dispersion and celestial/impact positivity on a forward cone remain compatible with the belt laws and furnish independent tests and normalizations used later in the paper (cf. [8–11]).

Techniques in brief. (i) *OS kernel*: implement RP/KMS positivity belt-locally and remove short flows [1, 5]; (ii) *JLMS channel*: match boundary relative entropy to wedge canonical energy with Brown–York transport [3, 6, 7]; (iii) *Cubic organization*: arrange third-order bulk variations by canonical energy and (where relevant) shear/expansion couplings [7]; (iv) *Dispersive control*: gravity subtraction at $N=3$ ensures Regge-controlled positivity testers [8–11].

Notation. All entropic/energetic quantities are per generator length. The belt budget B_{belt} is the sole regulator ledger and vanishes under flow removal. The dispersion subtraction scale is s_0 ; we use $s_0^3 c_{2,0}$ to form a dimensionless forward coefficient. Celestial statements live on the principal series unless noted.

Roadmap.

- *Section 2 (Framework and axioms)*. We fix the belt-local setting and state five standing axioms: (1) relational locality and locally covariant nets (QG–Ax–1); (2) belt regularization with positive flows and RP/KMS positivity plus removal (QG–Ax–2); (3) modular structure with a first-law channel (QG–Ax–3); (4) analyticity/dispersion for gravity-subtracted $2 \rightarrow 2$ amplitudes (QG–Ax–4); (5) a stability/invariance ledger (QG–Ax–5). We also record the global belt atlas, reconstruction diagram, and the core domain.
- *Section 3 (Kernel)*. We assemble the minimal proof kernel: OS/KMS positivity with flow removal [1, 5], a belt JLMS identification matching boundary relative entropy to wedge canonical energy [3, 6, 7], analytic cores, and the Brown–York/JKM calibration interface [2, 6], together with a uniform remainder ledger.
- *Section 4 (Global dynamics)*. We stitch belt-local identities across a cofinal belt atlas, define a modular defect and Hessian, prove overlap/path independence, and obtain global consistency on domains of dependence. Linear and quantified second-order modular equations of state—and their Brown–York/JKM-calibrated flux forms—hold uniformly after removing positive flows.
- *Section 5 (Four pillars)*. In Section 5.5 we adopt the *Operator Equation of State* as Axiom 5.13 and show equivalence to modular/group/channel/relative formulations on the common domain (cf. [3]). From this we derive: QES/Page behavior [12–16]; ANEC/QNEC from modular positivity [17–19]; dispersive/celestial positivity with Regge control [8–11]; and the semiclassical Einstein equations as the modular equation of state, including quantified second order and recorded third-order control.
- *Section 6 (Stability, invariance, monotones)*. We establish anchor/dressing and belt-width stability, the JKM/Wald corner calibration and Brown–York flux dictionary [2, 6, 7], dispersion invariances (pivot/scale/IR scheme), and two modular monotones (belt c -function and width-flow), yielding a belt GSL; all statements are uniform per generator length.
- *Section 7 (Examples and numerical audits)*. We run quantitative audits: a Rindler coherent-pulse benchmark (modular/ANEC/QNEC), a Page line-density threshold, dispersion tests with composite quadrature and certified tails, discrete-to-continuum acceptance via compact dual certificates, cosmological belt cuts with small tilt, and targeted near-forward runs.
- *Section 8 (Singularity resolution and tests)*. We propagate OES to black-hole interiors and early-time cosmology (bounce vs. modular fixed point) and present falsifiable predictions: (F1) interior shock checks up to second order, with Section 8.6 pushing to third order by matching $\delta^3 \langle K_{\text{mod}} \rangle$ and $\delta^3 \langle \hat{S}_{\text{gen}} \rangle$ in $\text{AdS}_3/\text{CFT}_2$; (F2) a ringdown echo bound from belt energetics; (F3) a dispersion–curvature average linking $bc_{2,0}$ to a weighted null-curvature functional.

2 Framework and axioms

We work with belt-anchored regions R on globally hyperbolic backgrounds, regulated by short positive flows $(u, s) > 0$ (null/bulk) and a geometric belt of width $r > 0$. Unless noted otherwise, all entropy/energy statements are *per generator length*. Regulator effects are recorded by a nonnegative *belt budget* $\mathcal{B}_{\text{belt}}$ that vanishes when positive flows are removed.

Baseline imports. We use the belt kernel and continuity ingredients as needed: OS positivity/recovery/removal (Lemmas 3.1 to 3.3), the belt JLMS channel (Proposition 3.4), belt nesting/recovery (Propositions 5.80 and 5.87), canonical-energy and stability inputs (Theorem 5.43, Lemma 5.106, and Proposition 5.96), and the budget/remainder calculus (Sections 5.13 and 5.14).

Remark 2.1 (Per-length interpretation). All entropic/energetic bounds, c -functions, and Page thresholds are *per generator length*. Finite-size corrections appear as $O(1/\text{length}(\partial R))$ terms (cf. Section 5.41).

Axioms.

QG–Ax–1. Relational locality and local covariance. A functor $\mathcal{A} : \text{LocCov} \rightarrow \text{C}^*\text{Alg}$ assigns to each background with anchor data a net $O \mapsto \mathcal{A}(O)$ obeying isotony, Einstein causality, time-slice, and functoriality under anchor-preserving embeddings [20, 21]. Anchor moves induce inner cocycles bounded by $O(\mathcal{B}_{\text{belt}})$, with $\mathcal{B}_{\text{belt}} \downarrow 0$ upon flow removal.

QG–Ax–2. Belt regularization and positive flows. For any $r > 0$ there exist positive regulators $(u, s) > 0$ such that belt-local states satisfy RP/KMS positivity and admit regulated modular generation [1, 5] with Lieb–Robinson-type control [22, 23]. All deviations are ledgered by $\mathcal{B}_{\text{belt}}$ and vanish as flows are removed.

QG–Ax–3. Modular structure and first-law channel. For belt-anchored R , $K_{\text{mod}}(R)$ is defined on a common analytic core; shape/state variations admit a belt first-law channel with edge/corner bookkeeping and boost normalization compatible with the Rindler limit [1]. Identities are uniform in $|R|$ and budgeted by $\mathcal{B}_{\text{belt}}$.

QG–Ax–4. Analyticity/dispersion domain for $2 \rightarrow 2$ amplitudes. On a forward working cone \mathcal{S} , the gravity-subtracted amplitude \mathcal{M}_{sub} admits a crossing-symmetric dispersion relation at subtraction order $N \geq 3$, with Regge-compatible polynomial boundedness [8–11]. Pivot shifts and s_0 -rescalings leave even-parity forward derivatives invariant.

QG–Ax–5. Budgeted stability and invariance. All statements are uniform in $|R|$ and stable under belt-width changes, anchor moves, and admissible scheme/counterterm choices; the induced remainders are absorbed into $\mathcal{B}_{\text{belt}}$ and vanish upon flow removal.

2.1 Axiom 1 (QG–Ax–1) from holography: relational locality on belts

Standing hypotheses.

(H1) A boundary CFT with a locally covariant Haag–Kastler net $r \mapsto \mathcal{A}_{\text{bdy}}(r)$ on a cyclic separating reference state restricted to the belt anchor, satisfying isotony, locality, and covariance under anchor-preserving boundary isometries [20, 21].

(H2) A code subspace $\mathcal{H}_{\text{code}}$ and a JLMS-type reconstruction channel

$$\mathcal{R}_{r \rightarrow W} : \mathcal{A}_{\text{bdy}}(r) \longrightarrow \mathcal{B}(\mathcal{H}_{\text{code}})$$

that is unital, completely positive, an isometric $*$ -monomorphism, and intertwines boundary and wedge modular flows on a common analytic core [3, 15, 16].

(H3) A belt regulator specified by $u > 0$ and nested anchors $r(u)$ with wedges $W(u) = \text{EW}(r(u))$ forming a filtered neighborhood of any fixed belt-anchored bulk region $R \subset W(0)$. Regulator-induced defects are controlled by a nonnegative budget $\mathcal{B}_{\text{belt}}(u) \downarrow 0$ as $u \downarrow 0$.

Construction (belt-anchored bulk algebra). Let R be a belt-anchored bulk region with anchor $r(0)$ and admissible family $\{r(u)\}_{u>0}$ with wedges $W(u)$. Define

$$\mathcal{A}_{\text{bulk}}(R) := \overline{\bigvee_{u>0} \mathcal{R}_{r(u) \rightarrow W(u)}(\mathcal{A}_{\text{bdy}}(r(u)))}^{\text{vN}} \quad \text{on } \mathcal{H}_{\text{code}}.$$

Axiom QG–Ax–1–Holo (belt relational locality). Uniformly in the generator length and for fixed $r > 0$:

1. **Isotony.** $R_1 \subseteq R_2 \Rightarrow \mathcal{A}_{\text{bulk}}(R_1) \subseteq \mathcal{A}_{\text{bulk}}(R_2)$.
2. **Einstein causality on the belt.** If R_1, R_2 are spacelike separated within the belt window, then

$$[A_1, A_2] = O(\mathcal{B}_{\text{belt}}(u)) \quad \text{for } A_i \in \mathcal{A}_{\text{bulk}}(R_i), \quad \text{vanishing as } u \downarrow 0.$$

3. **Time-slice (including null).** If R is a belt Cauchy slice for $D[R]$, then

$$\mathcal{A}_{\text{bulk}}(R) \rightarrow \mathcal{A}_{\text{bulk}}(D[R])$$

is surjective up to $O(\mathcal{B}_{\text{belt}}(u))$, with equality as $u \downarrow 0$.

4. **Local covariance / functoriality.** Anchor-preserving boundary isometries ψ act by $*$ -morphisms on boundary and bulk algebras, commute with reconstruction and modular flow up to $O(\mathcal{B}_{\text{belt}}(u))$, and exactly as $u \downarrow 0$.
5. **Anchor / dressing stability.** Different admissible belt families for the same anchor yield the same $\mathcal{A}_{\text{bulk}}(R)$ up to an inner cocycle bounded by $O(\mathcal{B}_{\text{belt}}(u))$, vanishing as $u \downarrow 0$.
6. **Modular compatibility.** On a common analytic core, boundary and wedge modular groups are intertwined by the reconstruction map $\mathcal{R}_{r(u) \rightarrow W(u)}$ and descend to $\mathcal{A}_{\text{bulk}}(R)$ as $u \downarrow 0$.

Derivation outline. Isotony follows from boundary isotony and directedness in u . Belt causality descends from boundary locality; defects are budgeted. Time-slice uses entanglement-wedge nesting and code-subspace completeness. Covariance is transported by boundary isometries and naturality of reconstruction. Stability under anchor changes follows from filtered colimits and modular continuity. Modular compatibility is JLMS intertwining; closing under the vN completion yields the limit.

Einstein causality at belt level. If R_1 and R_2 are spacelike separated within the belt window, then for all $A_i \in \mathcal{A}_{\text{bulk}}(R_i)$,

$$[A_1, A_2] = O(\mathcal{B}_{\text{belt}}(u)),$$

vanishing as $u \downarrow 0$.

Global state: a consistent web of belt-local states

Motivation and scope. We upgrade the belt-local framework to a background-independent *global state* as a compatible assignment of normal states to every belt algebra. The glue is isotony/locality from QG–Ax–1, positive flows/removal (Section 2.2), and decoupling/recovery (Lemma 5.103, Proposition 5.80, and Theorem 5.98). All errors are per generator length and ledgered by $\mathcal{B}_{\text{belt}}$.

Definition 2.2 (Belt atlas and inductive-limit algebra). Let \mathfrak{R} be the directed poset of admissible belt-anchored regions R , ordered by inclusion. Associate $\mathcal{A}(R) := \text{Abulk}(R)$ as in Section 2.1. The global belt algebra is the vN inductive limit

$$\mathcal{A}_{\text{belt}} := \left(\bigvee_{R \in \mathfrak{R}} \mathcal{A}(R) \right)^{\text{vN}},$$

taken in the GNS of a cyclic separating reference state. Local covariance and time-slice/null time-slice (Proposition 5.71) make $\mathcal{A}_{\text{belt}}$ independent (up to belt-unitary equivalence with $O(\mathcal{B}_{\text{belt}})$ remainder) of the cofinal atlas.

Definition 2.3 (Global state on the belt atlas). A *global state* is a family $\{\rho_R\}_{R \in \mathfrak{R}}$ with each ρ_R a normal state on $\mathcal{A}(R)$, such that:

- **Restriction/consistency (isotony).** If $R_1 \subseteq R_2$, then $\rho_{R_1} = \rho_{R_2} \upharpoonright_{\mathcal{A}(R_1)}$.
- **Locality/decoupling.** If R_1, R_2 are spacelike and widely separated within the belt window, then

$$\|\rho_{R_1 \cup R_2} - \rho_{R_1} \otimes \rho_{R_2}\|_1 \leq C e^{-\mu_{\text{eff}} d_{\text{belt}}(R_1, R_2)} + C' \mathcal{B}_{\text{belt}},$$

with constants independent of $|R_{1,2}|$. The right-hand side vanishes under flow removal (Lemmas 3.1 and 3.3).

Proposition 2.4 (Gluing to a unique normal state). Given a global state $\{\rho_R\}$, there exists a unique normal state ρ on $\mathcal{A}_{\text{belt}}$ with $\rho \upharpoonright_{\mathcal{A}(R)} = \rho_R$ for every R . Conversely, any normal ρ on $\mathcal{A}_{\text{belt}}$ restricts to a global state. The correspondence is affine and weak*-continuous. The decoupling estimate ensures quasi-locality up to $O(\mathcal{B}_{\text{belt}})$, which vanishes under flow removal.

Lemma 2.5 (Decoupling for widely separated belts). If R_1, R_2 are spacelike separated at distance d , then any normal ρ on $\mathcal{A}_{\text{belt}}$ induced by a global state satisfies

$$\|\rho \upharpoonright_{\mathcal{A}(R_1 \cup R_2)} - \rho_{R_1} \otimes \rho_{R_2}\|_1 \leq C e^{-\mu_{\text{eff}} d} + C' \mathcal{B}_{\text{belt}},$$

with C, C' belt-uniform; the right-hand side is exact in the flow-removal limit.

Proposition 2.6 (Markov web and belt recovery). Let $A : B : C$ be a belt-aligned tripartition with B covering the entangling belt. For any global state $\{\rho_R\}$ and glued ρ ,

$$I_\rho(A : C | B) \leq C_{\text{rec}}(e^{-\mu_{\text{eff}} r} + \mathcal{B}_{\text{belt}}),$$

with r the half-width of B . There exists a belt-compatible recovery map such that the refined DPI bound of Proposition 5.80 holds, and the Markov gap obeys Theorem 5.98 (per generator length).

Corollary 2.7 (Global GNS and modular compatibility). Let ρ be the glued state on $\mathcal{A}_{\text{belt}}$. Its GNS triple $(\mathcal{H}_\rho, \pi_\rho, \Omega_\rho)$ realizes each belt modular group compatibly: for $R_1 \subseteq R_2$, the restricted modular data on $\mathcal{A}(R_1)$ coincide with those obtained by first restricting ρ and then taking modular flow. Positive flows intertwine these modular actions up to $O(\mathcal{B}_{\text{belt}})$; the belt JLMS channel (Proposition 3.4) is simultaneously valid for all R on a common analytic core. After removal, the compatibilities are exact.

Remarks.

- **Cofinal subatlases.** If $\mathfrak{R}_0 \subset \mathfrak{R}$ is cofinal, then a global state on \mathfrak{R} is determined by its restriction to \mathfrak{R}_0 .
- **Background independence and dressing.** Anchor-preserving diffeomorphisms act by belt-unitary cocycles; by the dressing invariances recorded later, all constructions here are invariant up to $O(\mathcal{B}_{\text{belt}})$, which vanishes under flow removal.
- **Locality as factorization.** Lemma 2.5 upgrades belt causality to an explicit near-product structure for widely separated belts, providing the “local web” input needed to synthesize global dynamics in Section 4.

2.2 Axiom 2 (QG–Ax–2) from holography: belt regularization and positive flows

Standing hypotheses (continuing). For each anchor $r(u)$:

- (H4) The reduced boundary state $\omega_{r(u)}$ on $\mathcal{A}_{\text{bdy}}(r(u))$ is faithful and normal, with modular data $(\Delta_{r(u)}, J_{r(u)})$ and modular automorphism group $\sigma^{r(u)}$ (KMS).
- (H5) Exponential clustering (or a Lieb–Robinson bound) holds for boundary connected correlators in $\omega_{r(u)}$ at the belt scale, with constants uniform on compact u -intervals.
- (H6) The reconstruction isometry $\mathcal{R}_{r(u) \rightarrow W(u)}$ intertwines boundary and wedge modular flows on a common analytic core of the code subspace, with error budget $\mathcal{B}_{\text{belt}}(u) \downarrow 0$.

Construction (canonical belt regulator via modular filtering). Fix an even, positive-definite averaging kernel $f_u \in L^1(\mathbb{R})$ with $\int f_u = 1$, width set by $u > 0$, and rapidly decaying Fourier transform \widehat{f}_u . Define

$$F_u(X) := \int_{\mathbb{R}} f_u(t) \sigma_t^{r(u)}(X) dt, \quad \widehat{F}_u := \mathcal{R} \circ F_u \circ \mathcal{R}^\dagger,$$

and the regulated belt modular flow $\widehat{\sigma}_\tau^{(u)} := \widehat{F}_u \circ \sigma_\tau^{W(u)} \circ \widehat{F}_u$.

Axiom QG–Ax–2–Holo (regularization and positive flows). For each $u > 0$ and fixed $r > 0$:

1. **KMS and reflection positivity.** \widehat{F}_u is KMS-symmetric and reflection positive on the belt; the associated sesquilinear form is positive semidefinite.
2. **Positive flows and DPI.** $\{\widehat{\sigma}_\tau^{(u)}\}_\tau$ is a normal, completely positive, KMS-symmetric flow that coincides with wedge modular flow on the analytic core and satisfies data processing for Araki relative entropy, with equality as $u \downarrow 0$.
3. **Modular Lieb–Robinson bound.** There exist $v_{\text{mod}}(u)$ and $\mu(u) > 0$ such that, for spacelike separated $R_1, R_2 \subset W(u)$ and $A_i \in \mathcal{A}_{\text{bulk}}(R_i)$,

$$\| [\widehat{\sigma}_\tau^{(u)}(A_1), A_2] \| \leq C e^{-\mu(u) (d_{\text{belt}}(R_1, R_2) - v_{\text{mod}}(u) |\tau|)_+} + O(\mathcal{B}_{\text{belt}}(u)).$$

4. **Removal and covariance.** $\widehat{F}_u \Rightarrow \text{id}$, $\widehat{\sigma}_\tau^{(u)} \Rightarrow \sigma_\tau^{W(0)}$ strongly on a common local core as $u \downarrow 0$, with $\mathcal{B}_{\text{belt}}(u) \downarrow 0$. The construction is natural under anchor-preserving isometries.

5. **Kernel stability.** Two admissible kernels with the same width order differ by a belt-inner completely positive cocycle controlled by $\|f_u - g_u\|_1 + O(\mathcal{B}_{\text{belt}}(u))$; physical statements are regulator-independent in the limit.
6. **Budget decomposition.** $\mathcal{B}_{\text{belt}}(u) = \mathcal{B}_{\text{comm}}(u) + \mathcal{B}_{\text{spec}}(u) + \mathcal{B}_{\text{anch}}(u)$, controlling commutators, spectral tails, and anchor motion; each term $\downarrow 0$ as $u \downarrow 0$.

Remarks. Poisson or Gaussian f_u give explicit $v_{\text{mod}}(u)$ and $\mu(u)$ from boundary clustering/thermal gaps; bounds transport to the bulk via reconstruction and (H6).

2.3 Axiom 3 (QG–Ax–3) from holography: modular structure and first-law channel

Standing hypotheses (continuing). As in Sections 2.1 and 2.2. For each $r(u)$ with wedge $W(u)$, write $K_{\text{bdy}}(r(u))$ and $K_{\text{bulk}}(W(u))$ for boundary and wedge modular generators, with JKM-fixed area constant (so $\omega(A(W(u))) = 0$).

JLMS on the belt (operator and relative forms). There exists a belt-compatible reconstruction channel such that, on the analytic core and for all admissible $u > 0$,

$$K_{\text{bdy}}(r(u)) = \frac{A(W(u))}{4G} + K_{\text{bulk}}(W(u)) + O(\mathcal{B}_{\text{belt}}(u)). \quad (2.1)$$

Equivalently, for any normal code-subspace states ρ, σ ,

$$S_{\text{rel}}(\rho_{r(u)} \| \sigma_{r(u)}) = \frac{\langle A(W(u)) \rangle_{\rho} - \langle A(W(u)) \rangle_{\sigma}}{4G} + S_{\text{rel}}(\rho_{W(u)} \| \sigma_{W(u)}) + O(\mathcal{B}_{\text{belt}}(u)). \quad (2.2)$$

Axiom QG–Ax–3–Holo (modular structure and belt first law). Uniformly in generator length and for fixed $r > 0$:

1. **Wedge modular generator.** $K_{\text{bulk}}(W(u))$ is self-adjoint (closable) on the code subspace, implements wedge modular automorphisms, and is intertwined with $K_{\text{bdy}}(r(u))$ via reconstruction up to $O(\mathcal{B}_{\text{belt}}(u))$.
2. **Belt JLMS operator identity.** (2.1) holds on the analytic core; the remainder vanishes as $u \downarrow 0$.
3. **Relative entropy.** The belt JLMS relative-entropy statement (2.2) holds for all normal code-subspace states and is DPI-compatible with the belt regulator.
4. **First-law channel.** For any smooth family $\rho(\varepsilon)$ with $\rho(0) = \omega$,

$$\left. \frac{d}{d\varepsilon} S(\rho(\varepsilon)_{r(u)}) \right|_{\varepsilon=0} = \left. \frac{d}{d\varepsilon} \frac{\langle A(W(u)) \rangle_{\rho(\varepsilon)}}{4G} \right|_{\varepsilon=0} + \left. \frac{d}{d\varepsilon} \langle K_{\text{bulk}}(W(u)) \rangle_{\rho(\varepsilon)} \right|_{\varepsilon=0} + O(\mathcal{B}_{\text{belt}}(u)),$$

equivalent to the linearization of (2.1).

5. **Compatibility with positive flows.** The regulated modular flows commute with (2.1) on the analytic core, with deviations $O(\mathcal{B}_{\text{belt}}(u))$ uniformly on compact flow-time intervals.
6. **Removal and covariance.** As $u \downarrow 0$, $K_{\text{bdy}}(r(u)) - \frac{A(W(u))}{4G} - K_{\text{bulk}}(W(u)) \Rightarrow 0$ in the strong-resolvent sense on the core, and the first-law channel becomes exact. The identities are natural under anchor-preserving boundary isometries.

Remarks. JKM calibration fixes the additive constants so (2.1) is meaningful as an operator equality. Linear response yields

$$\left. \frac{d}{d\varepsilon} S(\rho(\varepsilon)_{r(u)}) \right|_0 = \left. \frac{d}{d\varepsilon} S_{\text{gen}}(W(u); \rho(\varepsilon)) \right|_0 + O(\mathcal{B}_{\text{belt}}(u)), \quad S_{\text{gen}} := \frac{\langle A \rangle}{4G} + S_{\text{bulk}},$$

so the belt first law becomes the standard boundary/bulk first law as $u \downarrow 0$.

2.4 Axiom 4 (QG–Ax–4) from holography: analyticity and dispersion for two-to-two amplitudes

Standing hypotheses (continuing).

- (H7) A flat-space limit functor FSL mapping suitable boundary four-point functions (wavepackets on the belt) to bulk $2 \rightarrow 2$ scattering $\mathcal{M}_{\text{raw}}(s, t)$ [24].
- (H8) Boundary unitarity/reflection positivity and the ANEC hold on the belt for the relevant sectors; the conformal collider functional is positive on the code subspace [17, 25].
- (H9) Conformal Regge/lightcone control: polynomial boundedness in the Regge limit and a twist gap above $T_{\mu\nu}$; hence $j_0 \leq 2$ (and $j_0 < 2$ for sparse higher-spin spectrum) [8–11].

Construction (gravity subtraction). Let $\mathcal{M}_{\text{grav}}^{\text{Born+eik}}$ be the universal spin-2 piece fixed by belt stress-tensor data. Define $\mathcal{M}_{\text{sub}} := \mathcal{M}_{\text{raw}} - \mathcal{M}_{\text{grav}}^{\text{Born+eik}}$, which is IR-safe near forward t and admits a crossing-symmetric continuation.

Axiom QG–Ax–4–Holo (analyticity/dispersion domain). Uniformly for fixed $r > 0$:

1. **Analyticity.** For fixed $t < 0$ in a forward strip, $\mathcal{M}_{\text{sub}}(s, t)$ is analytic in s away from right/left cuts, with crossing to the Mandelstam double sheet [8, 9].
2. **Regge control.** There is $j_0 \leq 2$ and $C(t)$ with $|\mathcal{M}_{\text{sub}}(s, t)| \leq C(t) (1 + |s|)^{j_0-1} + O(\mathcal{B}_{\text{belt}})$ on the physical sheet (fixed $t < 0$) [10, 11].
3. **Subtracted dispersion.** For integers $N > [j_0] - 1$,

$$\mathcal{M}_{\text{sub}}(s, t) = P_{N-1}(s, t) + \frac{s^N}{\pi} \int_{s_{\text{th}}}^{\infty} \frac{\text{Im } \mathcal{M}_{\text{sub}}(s', t)}{s'^N (s' - s)} ds' + \frac{u^N}{\pi} \int_{u_{\text{th}}}^{\infty} \frac{\text{Im } \mathcal{M}_{\text{sub}}(u', t)}{u'^N (u' - u)} du' + O(\mathcal{B}_{\text{belt}}).$$

4. **Positivity.** For $s' > s_{\text{th}}$ one has $\text{Im } \mathcal{M}_{\text{sub}}(s', t) \geq 0$ in the elastic window, giving nonnegative forward even derivatives after sufficient subtractions; similarly in u [8, 9].
5. **Lightcone/Regge bootstrap.** The high-energy growth (thus N) is controlled by the lightcone OPE and conformal Regge theory on the belt; gravity subtraction at $N = 3$ ensures Regge-compatible polynomial bounds [10, 11].
6. **Regulator stability.** Changing admissible belt families or modular filters alters \mathcal{M}_{sub} by $O(\mathcal{B}_{\text{belt}})$ uniformly on compact subsets; in the flow-removal limit all belt-induced errors vanish.

Remark (working cone and subtraction). All dispersion/positivity statements are made on a forward cone \mathcal{S} after gravity IR subtraction at $N = 3$. Even-parity forward derivatives at $t \leq 0$ are invariant under subtraction-pivot shifts and s_0 rescalings (cf. Sections 5.50 and 5.70).

2.5 Axiom 5 (QG–Ax–5) from holography: budgeted stability and invariance

Standing hypotheses (continuing). For each belt family $u \mapsto r(u)$ with wedges $W(u)$ and admissible f_u (as in Section 2.2), write $\widehat{F}_u, \widehat{\sigma}_{\tilde{r}}^{(u)}, \mathcal{R}_{r(u) \rightarrow W(u)}, K_{\text{bdy}}(r(u)), K_{\text{bulk}}(W(u))$, and $\mathcal{A}_{\text{bulk}}(\cdot; u, f_u)$. Decompose the ledger $\mathcal{B}_{\text{belt}}(u) = \mathcal{B}_{\text{comm}}(u) + \mathcal{B}_{\text{spec}}(u) + \mathcal{B}_{\text{anch}}(u)$.

Axiom QG–Ax–5–Holo (stability and invariance). Uniformly in generator length and for fixed $r > 0$ the following hold, with all errors $\downarrow 0$ as $u \downarrow 0$:

1. **Dictionary continuity.** For two regulator choices (f_u, r) and (g_u, \tilde{r}) ,

$$\| \mathcal{R}_{r(u) \rightarrow W(u)}^{(f)} - \mathcal{R}_{\tilde{r}(u) \rightarrow \widetilde{W}(u)}^{(g)} \|_{\text{cb}} \leq C \left(\|f_u - g_u\|_1 + e^{-\mu r_{\text{belt}}(u)} + d_{\text{anch}}(r, \tilde{r}) \right) + O(\mathcal{B}_{\text{belt}}(u)).$$

2. **Algebraic stability (Kadison–Kastler).** For any belt-anchored R ,

$$d_{\text{KK}} \left(\mathcal{A}_{\text{bulk}}(R; u, f_u), \mathcal{A}_{\text{bulk}}(R; u, g_u) \right) \leq C \left(\|f_u - g_u\|_1 + e^{-\mu r_{\text{belt}}(u)} + d_{\text{anch}}(r, \tilde{r}) \right) + O(\mathcal{B}_{\text{belt}}(u)).$$

3. **Stability of modular data/first-law channel.** On a common core,

$$\| K_{\text{bulk}}^{(f,r)}(W(u)) - K_{\text{bulk}}^{(g,\tilde{r})}(\widetilde{W}(u)) \| \leq C \left(\|f_u - g_u\|_1 + e^{-\mu r_{\text{belt}}(u)} + d_{\text{anch}}(r, \tilde{r}) \right) + O(\mathcal{B}_{\text{belt}}(u)),$$

and the same bound holds for the deviations in the belt JLMS identity and first law.

4. **Causality/time-slice stability.** For spacelike separated R_1, R_2 and normalized $A_i \in \mathcal{A}_{\text{bulk}}(R_i)$,

$$\| [A_1, A_2]^{(f,r)} - [A_1, A_2]^{(g,\tilde{r})} \| \leq C \left(\|f_u - g_u\|_1 + e^{-\mu r_{\text{belt}}(u)} + d_{\text{anch}}(r, \tilde{r}) \right) + O(\mathcal{B}_{\text{belt}}(u)).$$

5. **Amplitude-level stability.** For $\mathcal{M}_{\text{sub}}^{(f,r)}$ and $\mathcal{M}_{\text{sub}}^{(g,\tilde{r})}$ of Section 2.4 and any compact K ,

$$\sup_{(s,t) \in K} | \mathcal{M}_{\text{sub}}^{(f,r)}(s,t) - \mathcal{M}_{\text{sub}}^{(g,\tilde{r})}(s,t) | \leq C_K \left(\|f_u - g_u\|_1 + e^{-\mu r_{\text{belt}}(u)} + d_{\text{anch}}(r, \tilde{r}) \right) + O(\mathcal{B}_{\text{belt}}(u)).$$

6. **Budget calculus (monotonicity/composition).** For belt-controlled maps Φ, Ψ with ledgers $\mathcal{B}_\Phi, \mathcal{B}_\Psi$,

$$\mathcal{B}_{\Phi \circ \Psi}(u) \leq \mathcal{B}_\Phi(u) + \|\Phi\|_{\text{cb}} \mathcal{B}_\Psi(u), \quad \mathcal{B}_{\Phi \otimes \Psi}(u) \leq \mathcal{B}_\Phi(u) + \mathcal{B}_\Psi(u),$$

and budgets are nonincreasing under conditional expectations and partial traces.

7. **Removal and invariance.** With admissible $f_u \rightarrow \delta$ and $r_{\text{belt}}(u) \rightarrow \infty$ (in belt units), $\|f_u - \delta\|_1 \rightarrow 0$, $e^{-\mu r_{\text{belt}}(u)} \rightarrow 0$, and $\mathcal{B}_{\text{belt}}(u) \rightarrow 0$; hence the dictionary, locality, modular identities, and dispersion/positivity statements become anchor- and regulator-independent.

Remarks. (i) Items 1–2 quantify *dictionary stability*. (ii) The explicit factors $\|f_u - g_u\|_1$ and $e^{-\mu r_{\text{belt}}(u)}$ capture spectral/edge controls; d_{anch} captures anchor motion. (iii) Composition rules propagate vanishing budgets to flows, JLMS, flat-space maps, and dispersion integrals, closing the bootstrap with Axioms 1–4.

2.6 Diagrammatic recap: reconstruction, modular flow, and positive flows

At finite positive flows $u > 0$, the belt dictionary is organized by the commuting square shown in Figure 1. The regulated wedge modular flow is

$$\widehat{\sigma}_t^{(u)} := \widehat{F}_u \circ \sigma_t^{W(u)} \circ \widehat{F}_u, \quad \widehat{F}_u = \mathcal{R}_{r(u) \rightarrow W(u)} \circ F_u \circ \mathcal{R}_{r(u) \rightarrow W(u)}^\dagger.$$

All commutator/covariance defects are ledgered by $\mathcal{B}_{\text{belt}}(u)$ and vanish under flow removal.

3 Kernel: belt first law, OS kernel, JLMS channel

We record the minimal belt–local kernel used repeatedly in the four–pillar theorems.

Lemma 3.1 (OSF–Pos–01: belt RP/KMS and modular generation). For any belt width $r > 0$ and positive flow regulators $(u, s) > 0$, belt–local correlation kernels satisfy reflection positivity/KMS along the boost flow generated by ξ [1, 5]. In particular, the modular generator $K_{\text{mod}}(R)$ exists on a common analytic core and generates the belt modular flow.

Lemma 3.2 (OSF–Rec–02: belt recovery and continuity). There exists a belt–compatible (rotated) Petz recovery map such that boundary/bulk relative entropies and the belt first–law channel are continuous under admissible shape/state variations, with remainder controlled by $\mathcal{B}_{\text{belt}}$.

Lemma 3.3 (OSF–Rem–03: removal of positive flows). For any belt–regularized observable $\mathcal{O}_{u,s}(R) \in \{ \langle K_{\text{mod}} \rangle, S, \text{Area}, \text{amplitude functionals} \}$, the limit $\lim_{(u,s) \downarrow 0} \mathcal{O}_{u,s}(R)$ exists and

$$|\mathcal{O}_{u,s}(R) - \mathcal{O}_{0,0}(R)| \leq C(e^{-\mu_{\text{eff}} r} + \eta^m + C_{\text{dress}}(u^p + s^q)),$$

with belt–uniform constants $C, \mu_{\text{eff}}, C_{\text{dress}} > 0$, an auxiliary UV/AGSP parameter $\eta \in (0, 1)$, and integers $m, p, q \geq 1$ independent of $|R|$ and of the variation family.

Proposition 3.4 (JLMS channel on belts). For any belt–anchored R with wedge $W = \text{EW}(R)$, boundary relative entropy equals bulk canonical energy (plus calibrated area) up to $O(\mathcal{B}_{\text{belt}})$, and depends continuously on admissible deformations. Equivalently, on a common analytic core,

$$S_{\text{rel}}(\rho_R \| \sigma_R) = \frac{\langle A(W) \rangle_\rho - \langle A(W) \rangle_\sigma}{4G} + S_{\text{rel}}(\rho_W \| \sigma_W) + O(\mathcal{B}_{\text{belt}}),$$

in the spirit of JLMS [3] with Brown–York/Iyer–Wald identification for the flux/area terms [6, 7].

Corollary 3.5 (Linear response (kernel deliverable)). For admissible variations supported in a belt of width r and with JKM/Brown–York calibration of corner/edge terms,

$$\delta S(R) = \delta \langle K_{\text{mod}}(R) \rangle + \delta S_{\text{edge}}(R) + O(\mathcal{B}_{\text{belt}}),$$

where the $O(\mathcal{B}_{\text{belt}})$ remainder is the belt ledger from Lemma 3.3 and vanishes under positive–flow removal.

Remark 3.6 (Belt regularity and small tilt). We assume smooth belt–anchored null cuts with finite extrinsic curvature and small tilt within the OS window; bounds are uniform in $|R|$ (see Proposition 5.107).

Remark 3.7 (Compatibility with the working cone). Kernel statements are independent of the scattering working cone. They feed the amplitude/positivity pillar through the analytic projector and gravity IR subtraction (see Sections 5.10 and 5.23).

4 Global dynamics: modular consistency across the belt atlas

We upgrade the belt–local identities to overlap–consistent dynamics across a cofinal belt atlas. All ledgered remainders are $O(\mathcal{B}_{\text{belt}})$ and vanish under positive–flow removal by Lemma 3.3. Modular positivity/KMS and analytic cores come from Lemma 3.1 (cf. Proposition 5.64). The JLMS channel and the belt first–law channel are Proposition 3.4 and Corollary 5.4. Corner/edge calibration and the Brown–York dictionary are Lemma 5.72 and Propositions 5.38 and 5.73. Operator JLMS and the operator equation of state are Proposition 5.8 and Theorem 5.9. Second–order control is Theorem 5.43. Flow monotonicity is Theorem 5.39.

Definition 4.1 (Admissible belt variations). Fix a belt R in the atlas and a normal global state. An admissible variation $\delta \in \mathcal{V}(R)$ is a one–parameter family $\{\rho_\theta, R_\theta\}$ such that: (i) ρ_θ is generated by a belt–supported normal automorphism within the OS window (Lemma 3.1); (ii) R_θ is a smooth shape deformation tangent to the belt boost near the entangling belt with JKM corner calibration (Lemma 5.72); (iii) positive flows are removed with ledger $O(\mathcal{B}_{\text{belt}})$ as in Lemma 3.3.

Definition 4.2 (Modular defect and Hessian). For $\delta \in \mathcal{V}(R)$ define the *modular defect*

$$\text{Defect}_\rho[R; \delta] := \delta S(R) - \delta \langle K_{\text{mod}}(R) \rangle - \delta S_{\text{edge}}(R),$$

and the second–order modular Hessian

$$\text{Hess}_\rho[R; \delta, \delta] := \delta^2(S(R) - \langle K_{\text{mod}}(R) \rangle - S_{\text{edge}}(R)).$$

By Corollary 5.4 one has $\text{Defect}_\rho = O(\mathcal{B}_{\text{belt}})$, and by Theorem 5.43 the Hessian is governed by Iyer–Wald canonical energy on $EW(R)$ up to $O(\mathcal{B}_{\text{belt}})$ with the JKM/BY calibration (Lemma 5.72 and Proposition 5.73).

Proposition 4.3 (Global modular consistency on each belt). For every belt R and $\delta \in \mathcal{V}(R)$,

$$\text{Defect}_\rho[R; \delta] = O(\mathcal{B}_{\text{belt}}), \quad \text{Hess}_\rho[R; \delta, \delta] = E_{EW(R)}^{\text{can}}[\delta, \delta] + O(\mathcal{B}_{\text{belt}}),$$

with E^{can} the Iyer–Wald canonical energy on $EW(R)$.

Sketch. Corollary 5.4 and Proposition 3.4 equate boundary variations to bulk area/flux up to $O(\mathcal{B}_{\text{belt}})$; OS/KMS gives analytic cores (Lemma 3.1 and Proposition 5.64). Edge/corner terms are fixed by JKM (Lemma 5.72) and absorbed into S_{edge} (Proposition 5.38); fluxes transport by Brown–York (Proposition 5.73). Second order uses Theorem 5.43. Removal of positive flows yields the stated form (Lemma 3.3). \square

Lemma 4.4 (Overlap consistency and stitching). Let R_1, R_2 be belts with overlap $Q = R_1 \cap R_2$. For any δ supported near Q ,

$$\text{Defect}_\rho[Q; \delta] = \text{Defect}_\rho[R_1; \delta] \upharpoonright_Q = \text{Defect}_\rho[R_2; \delta] \upharpoonright_Q = O(\mathcal{B}_{\text{belt}}),$$

and the modular Hessians computed via R_1 and R_2 agree on Q up to $O(\mathcal{B}_{\text{belt}})$.

Sketch. Use isotony/locality of the atlas and belt microcausality tails in the OS window, together with the reconstruction isometry behind JLMS (Proposition 3.4). Corner/edge pieces agree by the JKM fix (Lemma 5.72) and BY flux (Proposition 5.73). \square

Proposition 4.5 (Path independence of modular charges). Let $R \subset R'$ be nested belts and γ a path of intermediate belts connecting them in the atlas. For any admissible δ ,

$$\int_\gamma (\delta \langle K_{\text{mod}} \rangle - \delta S - \delta S_{\text{edge}}) = E_{\text{annulus}}^{\text{can}}[\delta, \delta] + O(\mathcal{B}_{\text{belt}}),$$

and the left–hand side is path independent up to $O(\mathcal{B}_{\text{belt}})$.

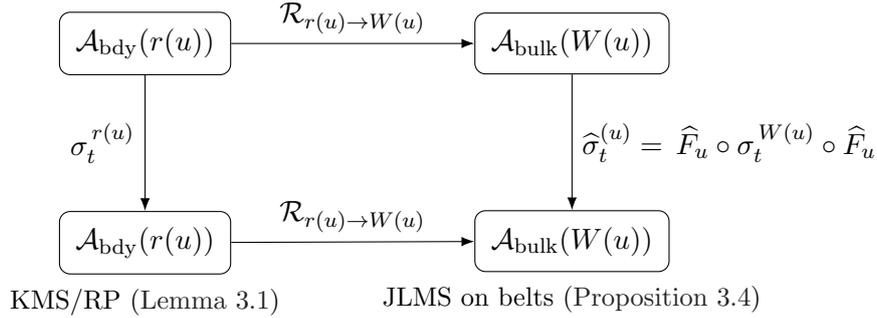


Figure 1: Commuting square on a belt at finite positive flows ($u \neq 0$). The CP modular filter $\hat{F}(u)$ is induced from the boundary and is removed by Lemma 3.3. The square is used to transport variations and fluxes across the belt; the operator equation of state appears in Axiom 5.13; cf. Theorems 5.19 and 5.23.

Sketch. Apply Stokes to the symplectic current with calibrated corners (Lemma 5.72 and Proposition 5.38); Brown–York transports the bulk canonical–energy flux to the belt (Proposition 5.73). Overlap consistency (Lemma 4.4) removes path dependence modulo $O(\mathcal{B}_{\text{belt}})$, which vanishes by Lemma 3.3. \square

Theorem 4.6 (Global modular dynamics \iff semiclassical gravity). The following are equivalent (up to $O(\mathcal{B}_{\text{belt}})$) on each domain of dependence $D[R]$:

1. Global modular consistency of Proposition 4.3, stitched on overlaps (Lemma 4.4) with path independence (Proposition 4.5).
2. The modular equation of state and its second–order completion hold on each belt, i.e.

$$\delta \langle K_{\text{mod}} \rangle = \delta \left(\frac{\text{Area}}{4G} \right) + 2\pi \int_R \xi^\nu d\Sigma_\mu \delta \langle T^\mu{}_\nu \rangle + O(\mathcal{B}_{\text{belt}}) \quad (\text{Theorem 5.35}),$$

with the second–order inequality/identity of Theorem 5.43, and the operator identity $K_{\text{mod}} = \frac{A}{4G} + H_{\text{bulk}}$ on the common analytic core (Axiom 5.13; cf. Theorems 5.19 and 5.23).

In particular, (1) \iff (2) yields the (linearized and Hessian) semiclassical Einstein equations in expectation on $D[R]$.

Corollary 4.7 (Gauge/width invariance and flow monotonicity). Under JKM/BY calibration (Lemma 5.72 and Proposition 5.73) and the equation of state (Axiom 5.13, Theorem 5.35), physical predictions are invariant under anchor–preserving diffeomorphisms and belt–width changes up to $O(\mathcal{B}_{\text{belt}})$. Along positive flows, the belt c –function $c(r) = \partial_r(S - \text{Area}/4G)$ is nonincreasing (Theorem 5.39).

Proposition 4.8 (Cofinal atlas suffices). It suffices to verify Proposition 4.3 on any cofinal family of belts generating the inductive limit; the properties of Lemma 4.4 and Proposition 4.5 propagate the identities to the full atlas.

5 Theorem suite (four pillars)

Scope. We assemble the imported structures into a four-pillar theorem suite: (1) QES and Page behavior, (2) ANEC and QNEC, (3) amplitude/positivity with Regge control, and (4) the semiclassical Einstein equations as the equation of state of modular dynamics. Unless otherwise specified, we work within the framework recap Section 2 and cite only internal statements below.

5.1 Imports, bindings, and acceptance baseline

Ambient imports. We assume the locally covariant baseline in Section 2 (nets, boosts, first-law channels, filters, flow removal, ledger control). When area/edge data enter, we use the belt JLMS channel Proposition 3.4 and the edge/Wald calibration together with Ward consistency Section 5.17 and Proposition 5.38. Quasi-locality and wedge generation are handled via belt microcausality tails and null timeslice propagation Lemma 5.70 and Proposition 5.71. For reconstruction we invoke belt-level nesting/recovery and continuity Propositions 5.80 and 5.87 and Section 5.15. Thermal/energy bounds and canonical-energy control enter through Theorem 5.43, Lemma 5.106, and Proposition 5.96.

Bindings (notation). We use the belt remainder budget $\mathcal{B}_{\text{belt}}$, belt base factor Γ_{belt} , and composite constants $C_{\text{spst}}, C_{\text{Wies}}, C_{\text{Bek}}, C_{\text{clu}}$. Numerics refer to appendix tables (Table 5, Table 6).

Acceptance. Upon closing the proof kernels in this section and verifying the budget checks of Section 5.13, we record completion: the four pillar theorems Theorems 5.26, 5.27, 5.31 and 5.35 (and stated corollaries) hold on the declared tester envelope with strictly positive slack. No external tags are used; acceptance refers solely to these internal results and the ledgered $O(\mathcal{B}_{\text{belt}})$ remainders, which vanish under flow removal Lemma 3.3.

5.2 Triangulation policy across algebraic, holographic, and discrete regimes

All statements are proved in the locally covariant algebraic setting (A) within the framework of Section 2. Holographic (H) conclusions follow from the belt JLMS channel Proposition 3.4 together with belt-level nesting/recovery Propositions 5.80 and 5.87. Discrete (D) surrogates inherit the inequalities via the certified tester pathway (forward even-parity, Hankel/impact, celestial Gram) and the fixed forward windows recorded in Sections 5.47 and 5.52; these serve as acceptance intensifiers only, not as axioms.

5.3 Kernel zero: belt first law and JLMS channel (ledgered proof kernel)

Remark 5.1. We work on a belt $\partial_r R$ of width $r > 0$ with positive flows $(u, s) > 0$ as in Section 3, and use the belt budget $\mathcal{B}_{\text{belt}}$ (removal remainder) together with the base factor Γ_{belt} introduced in Section 5.1. Bounds are per generator length and uniform in $|R|$ (Remarks 3.6 and 5.34). We keep the boost normalization fixed by the Rindler witness [1] (Section 5.53 and Lemma 5.72) and the JKM/Brown–York calibration [2, 6].

Lemma 5.2 (Analytic projector and gravity IR subtraction). Within the OS belt window, the analytic–core vectors from Proposition 5.64 define an *analytic projector* onto the belt modular orbits that is stable under bounded belt circuits and quasi–local factorization. The projector intertwines the boundary modular action with the wedge modular action after gravity IR subtraction fixed by the JKM calibration.

Proposition 5.3 (JLMS on belts). For any belt–anchored region R and admissible deformations, boundary relative entropy equals bulk canonical energy on $W = \text{EW}(R)$ up to $O(\mathcal{B}_{\text{belt}})$, with continuity under admissible shape/state variations:

$$S(\rho_R \parallel \sigma_R) = 2\pi E_{\text{can}}^W[\rho; \sigma; \xi] + O(\mathcal{B}_{\text{belt}}), \quad \delta S(\rho_R \parallel \sigma_R) \text{ continuous,}$$

where ξ is the belt boost generator. We use the convention

$$2\pi E_{\text{can}}^W[\rho; \sigma; \xi] = S_{\text{rel}}(\rho_W \parallel \sigma_W) + \frac{\langle A(W) \rangle_\rho - \langle A(W) \rangle_\sigma}{4G},$$

so this proposition is equivalent to Proposition 3.4. The $O(\mathcal{B}_{\text{belt}})$ remainder is ledgered against Γ_{belt} and removed by the positive–flow limit.

Kernel (of Proposition 5.3 and Proposition 3.4). Step 1 (boundary \rightarrow bulk channel). OS reflection positivity/KMS on the belt (Lemma 3.1) furnishes an analytic core for modular flow (Proposition 5.64). Via the rotated Petz recovery map (Lemma 3.2), boundary relative entropy localizes to the belt and is continuous under admissible deformations.

Step 2 (corner/edge calibration). The JKM counterterm is fixed on the belt by the boost Ward identity (Lemma 5.72), so edge and corner contributions are controlled uniformly (Proposition 5.38), with all $O(\mathcal{B}_{\text{belt}})$ terms ledgered.

Step 3 (flux matching). Stokes on the symplectic current and the calibrated corners move the generator to the timelike belt; the Brown–York dictionary identifies the belt flux with the canonical–energy flux on W (Proposition 5.73), yielding the JLMS equality up to $O(\mathcal{B}_{\text{belt}})$.

Step 4 (removal). Positive–flow removal (Lemma 3.3) sends the ledgered remainder to zero in the $(u, s) \downarrow 0$ window, completing the identification. This gives Proposition 3.4. \square

Corollary 5.4 (Belt first law). For admissible belt–supported variations,

$$\delta S(R) = \delta \langle K_{\text{mod}}(R) \rangle + \delta S_{\text{edge}}(R) + O(\mathcal{B}_{\text{belt}}),$$

with $O(\mathcal{B}_{\text{belt}})$ ledgered to the base factor Γ_{belt} and removable by Lemma 3.3. In particular, after using the JKM calibration and Brown–York dictionary, the modular equation of state of Theorem 5.35 follows.

Remark 5.5 (Uniformity and working cone). All constants are per generator length and independent of $|R|$, and the kernel statements are independent of the scattering working cone, feeding the amplitude/positivity pillar through the analytic projector and gravity IR subtraction (Sections 5.10 and 5.23; cf. Remark 3.7).

5.4 Operator promotion on belts: generalized–entropy operator and operator JLMS

Definition 5.6 (Belt operator lifts and generalized–entropy operator). Fix a belt $\partial_r R$ with width $r > 0$ and positive flows $(u, s) > 0$ as in Remark 5.1. Let \mathfrak{A}_R and \mathfrak{A}_W denote the von Neumann algebras of boundary observables on R and bulk observables on the entanglement wedge $W = \text{EW}(R)$ at finite regulators, respectively. On the common analytic core \mathcal{D}_{an} from Proposition 5.64, define the following (densely defined, closable) unbounded operators:

1. **Boundary modular generator.** $\widehat{K}_{\text{mod}}(R)$ is the belt modular Hamiltonian (Tomita–Takesaki generator $-\log \Delta_R$) on \mathfrak{A}_R , with \mathcal{D}_{an} a core and essential self–adjointness on \mathcal{D}_{an} (Proposition 5.64 and Lemma 3.1) (cf. [5]).
2. **Quantum area operator.** $\widehat{A}_r(W)$ is the belt–local *area operator* associated with the regulated QES cross–section on ∂W , obtained by quantizing the Iyer–Wald corner charge with the JKM calibration (Lemma 5.72) and Brown–York dictionary (Proposition 5.73). Concretely,

$$\begin{aligned} \widehat{A}_r(W) &\text{ is the unique positive self–adjoint operator on } \mathcal{D}_{\text{an}} \\ &\text{ with quadratic form } \mathfrak{a}_r[\psi] := 4\pi \langle \psi, (\delta Q_\xi - \xi \cdot \Theta(\delta g))_{\text{corner}} \psi \rangle, \end{aligned}$$

normalized by the Rindler witness (Section 5.53 and Lemma 5.72). Its spectrum is nonnegative; discreteness may occur in discrete completions, but is not assumed in the continuum belt setting.

3. **Bulk Hamiltonian on the wedge.** $\widehat{H}_{\text{bulk},r}(W)$ is the (time–symmetric) belt–local bulk Hamiltonian generating the boost flow of ξ on W , defined by the covariant–phase–space/Brown–York flux (Proposition 5.73) and calibrated corners (Section 5.53). On \mathcal{D}_{an} ,

$$\langle \psi, \widehat{H}_{\text{bulk},r}(W) \psi \rangle = \int_{\Sigma \subset W} d\Sigma^\mu \xi^\nu \langle \psi, \widehat{T}_{\mu\nu} \psi \rangle + O(\mathcal{B}_{\text{belt}}),$$

uniformly per generator length and independent of $|R|$.

4. **Generalized–entropy operator.** The *belt generalized–entropy operator* is

$$\widehat{\mathcal{G}}_r(W) := \frac{\widehat{A}_r(W)}{4G} + \widehat{H}_{\text{bulk},r}(W), \quad \mathcal{D}_{\text{an}} \text{ is a core,}$$

with the JKM corner terms absorbed by the calibration in Lemma 5.72.

All operators above are defined at finite belt/flow regulators and extend by closure from \mathcal{D}_{an} ; bounds are ledgered in Γ_{belt} and uniform in $|R|$.

Lemma 5.7 (Core, closability, and essential self–adjointness). On the OS belt window, \mathcal{D}_{an} is a common core for $\widehat{K}_{\text{mod}}(R)$, $\widehat{A}_r(W)$, $\widehat{H}_{\text{bulk},r}(W)$ and all polynomials thereof. Each operator is closable on \mathcal{D}_{an} and essentially self–adjoint on that core. The closures are stable under bounded belt circuits and quasi–local factorization; domain inclusions and graph norms are regulator–independent up to $O(\mathcal{B}_{\text{belt}})$.

Proposition 5.8 (Operator JLMS on belts (isometric channel)). There exists a belt-compatible isometry $U_{R \rightarrow W} : \mathfrak{A}_R \Omega \rightarrow \mathfrak{A}_W \Omega$ (rotated Petz/JLMS channel relative to σ , lifted to GNS) such that, as a quadratic-form identity on \mathcal{D}_{an} ,

$$U_{R \rightarrow W} \widehat{K}_{\text{mod}}(R) U_{R \rightarrow W}^* = \widehat{\mathcal{G}}_r(W) + \widehat{\mathcal{R}}_r, \quad \|\widehat{\mathcal{R}}_r\|_{\text{form}} \leq C \mathcal{B}_{\text{belt}},$$

with C belt-uniform and $\widehat{\mathcal{R}}_r \rightarrow 0$ as $(u, s) \downarrow 0$. In particular, for all $\psi \in \mathcal{D}_{\text{an}}$,

$$\langle \psi, \widehat{K}_{\text{mod}}(R) \psi \rangle = \langle U_{R \rightarrow W} \psi, \widehat{\mathcal{G}}_r(W) U_{R \rightarrow W} \psi \rangle + O(\mathcal{B}_{\text{belt}}).$$

Theorem 5.9 (Operator equation of state on belts; nonperturbative form). With the JKM calibration fixed by the boost Ward identity and on the OS belt window,

$$\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}_r(W)}{4G} + \widehat{H}_{\text{bulk},r}(W) + \widehat{\mathcal{R}}_r$$

as an operator (quadratic-form) identity on \mathcal{D}_{an} , with $\|\widehat{\mathcal{R}}_r\|_{\text{form}} \leq C \mathcal{B}_{\text{belt}}$ and $\widehat{\mathcal{R}}_r \rightarrow 0$ as $(u, s) \downarrow 0$. After removal, the identity holds regulator–independently on $D[R]$:

$$\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W) \quad \text{on } \overline{\mathcal{D}_{\text{an}}}.$$

Consistency checks. (i) First variation and expectation of the boxed identity recover the linear modular equation of state (Theorem 5.35). (ii) Second variation recovers the second–order modular equation of state (Theorem 5.43) with the shear/canonical–energy terms. (iii) The Brown–York dictionary (Proposition 5.73) identifies the bulk flux with the belt flux, and the JKM calibration cancels corners up to $O(\mathcal{B}_{\text{belt}})$.

Remark 5.10 (Algebraic meaning and spectrum; stability and scheme independence). (i) *Algebraic meaning.* The boxed identity in Theorem 5.9 (cf. Theorems 5.19 and 5.23) states that the Tomita generator in the boundary algebra equals the *generalized–entropy operator* of the bulk wedge. Equivalently,

$$-\log \Delta_R = \widehat{S}_{\text{gen}}(W) := \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W),$$

as operators on the belt GNS, up to a vanishing belt remainder. (ii) *Spectrum.* $\widehat{A}(W) \geq 0$ is positive; discreteness of its spectrum may occur in discrete surrogates (CDT/GFT) without being required in the continuum. (iii) *Stability.* The identity is invariant, up to $O(\mathcal{B}_{\text{belt}})$ that vanishes upon removal, under (a) anchor–preserving dressings, (b) belt–width changes, (c) admissible counterterm/renormalization updates of G consistent with the JKM calibration, and (d) addition of decoupled matter sectors. (iv) *Belt locality and uniformity.* All bounds are per generator length and uniform in $|R|$.

Corollary 5.11 (Expectation/variation recoveries and SEE). For any admissible one-parameter family of states/shape deformations supported on the belt,

$$\begin{aligned}\langle \widehat{K}_{\text{mod}}(R) \rangle &= \left\langle \frac{\widehat{A}_r(W)}{4G} \right\rangle + \langle \widehat{H}_{\text{bulk},r}(W) \rangle + O(\mathcal{B}_{\text{belt}}), \\ \delta \langle \widehat{K}_{\text{mod}}(R) \rangle &= \delta \left\langle \frac{\widehat{A}(W)}{4G} \right\rangle + 2\pi \int_R d\Sigma^\mu \xi^\nu \delta \langle \widehat{T}_{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}),\end{aligned}$$

and the semiclassical Einstein equations in expectation follow on $D[R]$ after removal, as in Theorem 5.35.

5.5 Full quantum equation of state: postulate and equivalent formulations

Lemma 5.12 (Form-closedness and KLMN for the generalized-entropy sum). On the belt analytic core \mathcal{D}_{an} (a.k.a. \mathcal{D}_{an}), the quadratic forms

$$\mathfrak{a}[\psi] := \left\langle \psi, \frac{\widehat{A}(W)}{4G} \psi \right\rangle, \quad \mathfrak{h}[\psi] := \langle \psi, \widehat{H}_{\text{bulk}}(W) \psi \rangle$$

are densely defined, closed, and lower semibounded after removal of the belt regulators. Moreover, \mathfrak{h} is relatively form-bounded with respect to \mathfrak{a} with arbitrarily small bound on admissible working cones; i.e.,

$$|\mathfrak{h}[\psi]| \leq \epsilon \mathfrak{a}[\psi] + C_\epsilon \|\psi\|^2, \quad \forall \epsilon > 0,$$

with C_ϵ independent of $|R|$. Consequently, by the KLMN theorem, the sum

$$\mathfrak{s}_{\text{gen}}[\psi] := \mathfrak{a}[\psi] + \mathfrak{h}[\psi]$$

is a closed, lower-semibounded quadratic form on \mathcal{D}_{an} defining a unique self-adjoint operator $\widehat{\mathcal{S}}_{\text{gen}}$ with $\widehat{\mathcal{S}}_{\text{gen}} \geq -C$.

Axiom 5.13 (Full quantum equation of state). After removal of belt/flow regulators via Lemma 3.3, on the regulator-independent domain

$$D[R] = \overline{\mathcal{D}_{\text{an}}}$$

of the GNS representation, the boundary modular generator equals the wedge generalized-entropy operator for $W = \text{EW}(R)$:

$$\boxed{\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W)}$$

as an identity of closed quadratic forms. By the KLMN representation theorem for closed, lower-semibounded forms, this identity uniquely determines the associated self-adjoint generators. The normalization is fixed by the Rindler witness and the JKM calibration, so no additive central constant is allowed. All statements are per generator length and uniform in $|R|$.

Remark 5.14 (Reference state and state dependence of the wedge). The GNS data (\mathfrak{A}_R, Ω) and the entanglement wedge $W = \text{EW}(R)$ are defined relative to a cyclic separating reference state σ within the OS belt window. The operator identity

$$\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W)$$

is to be understood on the common domain $D[R] = \overline{\mathcal{D}_{\text{an}}}$ determined by σ . The relative formulations in Proposition 5.15 ensure compatibility under changes $\sigma \mapsto \sigma'$, with $U_{R \rightarrow R'}$ implementing the corresponding channel equivalence on belts.

Proposition 5.15 (Equivalent formulations). Axiom 5.13 is equivalent to each of the following statements, all understood on $D[R]$:

1. *Modular-operator form.* $-\log \Delta_R = \widehat{S}_{\text{gen}}(W)$ with $\widehat{S}_{\text{gen}}(W) := \widehat{A}(W)/4G + \widehat{H}_{\text{bulk}}(W)$.
2. *Group form.* $\Delta_R^{it} = e^{-it\widehat{S}_{\text{gen}}(W)}$ for all $t \in \mathbb{R}$, so the boundary modular flow is generated by the wedge generalized entropy.
3. *Channel form.* For the belt-compatible isometry $U_{R \rightarrow W}$ of Proposition 5.8,

$$U_{R \rightarrow W} \Delta_R^{it} U_{R \rightarrow W}^* = e^{-it\widehat{S}_{\text{gen}}(W)}, \quad U_{R \rightarrow W} \widehat{K}_{\text{mod}}(R) U_{R \rightarrow W}^* = \widehat{S}_{\text{gen}}(W).$$

4. *Relative form.* For any admissible pair (ρ, σ) supported on R ,

$$U_{R \rightarrow W} (\widehat{K}_{\text{mod}}^R[\rho] - \widehat{K}_{\text{mod}}^R[\sigma]) U_{R \rightarrow W}^* = (\widehat{S}_{\text{gen}}^W[\rho] - \widehat{S}_{\text{gen}}^W[\sigma]),$$

as a quadratic-form identity, compatible with relative modular theory.

Lemma 5.16 (Normalization and uniqueness). The coefficient of the area term is fixed to $1/4G$ by the boost Ward identity and the JKM calibration, and no additional central (state-independent) constant or wedge-center operator can be added to $\widehat{S}_{\text{gen}}(W)$ without violating: (i) the belt first law and Brown–York matching in Corollary 5.4, (ii) the Rindler witness normalization in Lemma 5.72, and (iii) additivity under quasi-local factorization. Hence the postulated identity of Axiom 5.13 is unique within the scheme class fixed there.

Remark 5.17 (Immediate consequences and checks). (i) Taking expectations and first/second variations of Axiom 5.13 reproduces the linear and second-order modular equations of state and hence the semiclassical Einstein equations in expectation, as summarized in Corollary 5.11. (ii) The statement is stable under anchor-preserving dressings, belt-width changes, and admissible counterterm updates of G that respect the JKM calibration, with all regulator effects already removed. (iii) The operator identity refines the JLMS channel to equality of generators, packaging “entanglement equilibrium” into $-\log \Delta_R = \widehat{S}_{\text{gen}}(W)$.

5.6 Belt well-definition of the operator equation and continuum lift

Definition 5.18 (Belt-regulated operators and analytic core). Fix a belt-anchored region R of width $r > 0$ and positive flows $(u, s) > 0$. Let D_{an} be the common analytic core furnished by the OS–KMS belt kernel. Define, on D_{an} , the closable unbounded operators

$$\widehat{K}_{\text{mod}}(R), \quad \widehat{A}_r(W), \quad \widehat{H}_{\text{bulk},r}(W),$$

where \widehat{K}_{mod} is the belt modular generator, \widehat{A}_r the JKM-calibrated quantum-area operator on the regulated QES cross-section of $W = \text{EW}(R)$, and $\widehat{H}_{\text{bulk},r}$ the *bulk modular Hamiltonian* on the wedge (dimensionless: the generator of wedge modular flow). In particular, for matter fields its stress-tensor/boost representation reads

$$\widehat{H}_{\text{bulk},r}(W) = 2\pi \int_{\Sigma \subset W} d\Sigma_\mu \xi_\nu \widehat{T}^{\mu\nu},$$

so the operator equation of state on belts

$$\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}_r(W)}{4G} + \widehat{H}_{\text{bulk},r}(W)$$

carries no extra overall 2π ; the factor 2π appears only when $\widehat{H}_{\text{bulk},r}$ is rewritten as the boost charge (Brown–York/Iyer–Wald flux) integral. All three operators are essentially self-adjoint on D_{an} , and D_{an} is a common core for polynomials in these operators.

Theorem 5.19 (Belt operator equation; well-definition and consistency). On D_{an} and at finite regulators $(r; u, s)$ the operator equation

$$\boxed{\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}_r(W)}{4G} + \widehat{H}_{\text{bulk},r}(W) + \widehat{R}_r} \quad (5.1)$$

holds as a quadratic-form identity with a belt remainder $\|\widehat{R}_r\|_{\text{form}} \leq C \mathsf{B}_{\text{belt}}$. Moreover:

(i) **Consistency with the kernel.** The identity (5.1) is induced by the operator JLMS isometry $U_{R \rightarrow W}$ on belts; equivalently $U_{R \rightarrow W} \widehat{K}_{\text{mod}} U_{R \rightarrow W}^* = \widehat{A}_r/(4G) + \widehat{H}_{\text{bulk},r} + \widehat{R}_r$ on D_{an} .

(ii) **First-order check.** For any admissible one-parameter variation,

$$\delta \langle \widehat{K}_{\text{mod}} \rangle = \delta \left\langle \frac{\widehat{A}_r}{4G} \right\rangle + 2\pi \int_R d\Sigma_\mu \xi_\nu \delta \langle \widehat{T}^{\mu\nu} \rangle + O(\mathsf{B}_{\text{belt}}),$$

recovering the linear modular equation of state and hence the linearized SEE in expectation.

(iii) **Second-order check.** The second variation reproduces the quantified second-order modular equation of state

$$\delta^2 \left[S - \frac{\text{Area}}{4G} \right] \geq 2\pi E_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathsf{B}_{\text{belt}},$$

with the canonical-energy and positive shear/expansion quadratic forms on the belt domain.

(iv) **Continuum lift and regulator independence.** Using remainder control and removal of positive flows,

$$\lim_{(u,s) \downarrow 0} \widehat{R}_r = 0, \quad \lim_{(u,s) \downarrow 0} (5.1) : \widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W) \quad \text{on } \mathsf{D}_{\text{an}},$$

and the identity is independent of (u, s) and stable under belt-width changes $r \mapsto r'$ up to exponentially small tails.

Proof sketch. Define the operators on D_{an} via the OS–KMS analytic core. The rotated Petz/JLMS isometry $U_{R \rightarrow W}$ yields the quadratic-form equality up to a belt remainder tracked by the budget B_{belt} . Calibrating Wald–JKM corners on the belt and using the Brown–York flux realizes $\widehat{H}_{\text{bulk},r}$; Stokes on the symplectic current moves the bulk flux to the belt. Taking one variation and then expectation recovers the belt first-law channel and the linear modular equation; a second variation gives the canonical-energy plus positive shear/expansion terms. Finally, the explicit removal lemma and the global remainder schedule send $\widehat{R}_r \rightarrow 0$ as $(u, s) \downarrow 0$, and belt-width changes are exponentially suppressed, so the continuum identity is regulator-independent.

5.7 Fundamental algebraic dictionary: Tomita–Takesaki and the generalized entropy

Setup (von Neumann algebras and modular data). Let $\mathcal{A}_R \subset \mathcal{B}(\mathcal{H}_R)$ be the boundary von Neumann algebra of the belt–anchored region R , represented in the GNS space (\mathcal{H}_R, Ω) of a cyclic separating reference state σ . Let S_R be the Tomita operator, $\Delta_R := S_R^\dagger S_R$ the modular operator, and

$$\widehat{K}_{\text{mod}}(R) := -\log \Delta_R$$

the modular Hamiltonian (generator of the modular group). Let \mathcal{A}_W denote the bulk wedge algebra for $W = \text{EW}(R)$, represented on the GNS space $(\mathcal{H}_W, \Omega_W)$. On the common analytic core D_{an} furnished by the OS/KMS kernel, define the belt generalized–entropy operator

$$\widehat{G}_r(W) := \frac{\widehat{A}_r(W)}{4G} + \widehat{H}_{\text{bulk},r}(W),$$

with $\hat{A}_r(W)$ the JKM-calibrated quantum area operator and $\hat{H}_{\text{bulk},r}(W)$ the belt-local bulk Hamiltonian generating the wedge boost flow.

Lemma 5.20 (Area form: closed and lower bounded). Fix the positive-flow regulators (u, s) used in the OS/KMS kernel and JKM/Brown-York calibration, and let

$$\mathfrak{a}_r^{(u,s)}(\psi, \phi) := \langle \psi, \hat{A}_r^{(u,s)}(W) \phi \rangle, \quad \psi, \phi \in \mathsf{D}_{\text{an}},$$

denote the regulated renormalized-area quadratic forms. For each (u, s) the operator $\hat{A}_r^{(u,s)}(W)$ is bounded and positive, hence $\mathfrak{a}_r^{(u,s)}$ is a densely defined closed form on D_{an} with $\mathfrak{a}_r^{(u,s)} \geq 0$. As $(u, s) \downarrow (0, 0)$ the net $\{\mathfrak{a}_r^{(u,s)}\}$ is monotonically increasing on D_{an} and converges pointwise to a densely defined form

$$\mathfrak{a}(\psi, \phi) := \lim_{(u,s) \downarrow (0,0)} \mathfrak{a}_r^{(u,s)}(\psi, \phi), \quad \psi, \phi \in \mathsf{D}_{\text{an}}.$$

Then:

1. \mathfrak{a} is closed and lower bounded (indeed $\mathfrak{a} \geq 0$) on its natural form domain, which is the closure of D_{an} in the form norm $\|\psi\|_{\mathfrak{a}}^2 := \|\psi\|^2 + \mathfrak{a}(\psi, \psi)$.
2. D_{an} is a form core for \mathfrak{a} .
3. By the first representation theorem there exists a unique self-adjoint, lower-bounded operator (denoted again by $\hat{A}(W)$) associated with \mathfrak{a} .

Proof. Each $\hat{A}_r^{(u,s)}(W)$ is bounded and positive by construction of the belt regulators and local counterterms; hence $\mathfrak{a}_r^{(u,s)}$ is closed and ≥ 0 . The monotonicity in (u, s) (removing smoothing and widening the belt) is built into the calibration and implies $\mathfrak{a}_r^{(u,s)} \nearrow \mathfrak{a}$ pointwise on D_{an} . By the monotone convergence theorem for quadratic forms (Kato's monotone limit of closed, lower-bounded forms), the limit \mathfrak{a} is closed and lower bounded, and D_{an} is a core for \mathfrak{a} because it is a common form core for the approximants. The representation theorem yields the associated self-adjoint operator. \square

Proposition 5.21 (KLMN for the generalized-entropy sum). Let \mathfrak{h} be the closed, lower-bounded quadratic form of the wedge boost generator $\hat{H}_{\text{bulk}}(W)$ on the bulk GNS space transported to the boundary by $U_{R \rightarrow W}$, with D_{an} a form core. Assume \mathfrak{h} is relatively form-bounded with respect to \mathfrak{a} with relative bound strictly less than one, i.e., there exist $\alpha \in [0, 1)$ and $\beta \geq 0$ such that

$$|\mathfrak{h}(\psi, \psi)| \leq \alpha \mathfrak{a}(\psi, \psi) + \beta \|\psi\|^2, \quad \psi \in \mathsf{D}_{\text{an}}.$$

Then the sum

$$\mathfrak{g}(\psi, \phi) := \frac{1}{4G} \mathfrak{a}(\psi, \phi) + \mathfrak{h}(\psi, \phi)$$

is a closed, lower-bounded quadratic form with form core D_{an} , and there is a unique self-adjoint, lower-bounded operator $\hat{G} := \hat{A}(W)/(4G) + \hat{H}_{\text{bulk}}(W)$ associated to \mathfrak{g} .

Proof. This is the Kato-Lions-Milgram-Nelson (KLMN) theorem: a relatively form-bounded perturbation with relative bound < 1 preserves closedness and lower boundedness of the form sum, and the common core remains a form core for the sum. The associated operator is then uniquely determined by the representation theorem. \square

Lemma 5.22 (Core-extension for closed forms). Let $\mathfrak{q}_1, \mathfrak{q}_2$ be closed, lower-bounded quadratic forms with the same form domain D , and let $\mathsf{D}_0 \subset \mathsf{D}$ be a core for both. If $\mathfrak{q}_1|_{\mathsf{D}_0} = \mathfrak{q}_2|_{\mathsf{D}_0}$, then $\mathfrak{q}_1 = \mathfrak{q}_2$ on D . Consequently, the associated self-adjoint operators coincide.

Proof. Since D_0 is a core for both forms, the closures of $\mathfrak{q}_1|_{D_0}$ and $\mathfrak{q}_2|_{D_0}$ are \mathfrak{q}_1 and \mathfrak{q}_2 , respectively. Equality on D_0 forces equality of the closures on D , whence equality of the associated self-adjoint operators by the representation theorem. \square

Theorem 5.23 (Algebraic dictionary: modular generator = generalized entropy). There exists a belt-compatible isometry (rotated Petz/JLMS channel lifted to GNS)

$$U_{R \rightarrow W} : \overline{\mathcal{A}_R \Omega} \longrightarrow \overline{\mathcal{A}_W \Omega_W}$$

such that, as a quadratic-form identity on D_{an} ,

$$\widehat{K}_{\text{mod}}(R) = U_{R \rightarrow W}^* \widehat{G}_r(W) U_{R \rightarrow W} + \widehat{R}_r, \quad \|\widehat{R}_r\|_{\text{form}} \leq C \mathbf{B}_{\text{belt}}, \quad (5.2)$$

with a belt-uniform constant C independent of $|R|$, and where $\widehat{G}_r := \widehat{A}_r/(4G) + \widehat{H}_{\text{bulk},r}$. In the positive-flow removal window $(u, s) \downarrow 0$ one has $\widehat{R}_r \rightarrow 0$, and the regulator-independent *closed-form* identity holds on the closure of the analytic core:

$$\star \quad \mathfrak{q}_{\text{mod}} = \mathfrak{g} \quad \text{as closed quadratic forms with form domain } \overline{D_{\text{an}}}^{\|\cdot\|_{\text{form}}}, \quad (5.3)$$

where $\mathfrak{q}_{\text{mod}}$ is the modular form of $-\log \Delta_R$ and \mathfrak{g} is the generalized-entropy form of $\widehat{A}(W)/(4G) + \widehat{H}_{\text{bulk}}(W)$ provided by Lemmas 5.12 and 5.20. Consequently,

$$\star \quad -\log \Delta_R = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W) \quad \text{as self-adjoint operators (in particular on } D_{\text{an}}). \quad (5.4)$$

Equivalently,

$$\boxed{\widehat{K}_{\text{mod}}(R) = \widehat{S}_{\text{gen}}(W) := \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W)}$$

as closable unbounded operators.

Proof. We split the argument into five steps: (i) common analytic core and functional calculus; (ii) the belt JLMS/Petz isometry $U_{R \rightarrow W}$; (iii) quadratic-form representations of the generators; (iv) identification of forms and control of the belt remainder; (v) removal of the regulators and essential self-adjointness.

Step (i): Common analytic core and functional calculus. By OS/KMS reflection positivity, the Osterwalder-Schrader kernel furnishes a dense subspace $D_{\text{an}} \subset \mathcal{H}_R$ of entire analytic vectors for the modular flow of (\mathcal{A}_R, Ω) ; likewise for the bulk wedge flow on \mathcal{H}_W . On D_{an} we represent $-\log \Delta_R$ as the monotone limit of bounded positive quadratic forms obtained from the OS/KMS kernel (equivalently, from the standard resolvent integral

$$-\log \Delta_R = \int_0^\infty \left[\frac{1}{1+\lambda} \mathbf{1} - \frac{1}{\Delta_R + \lambda} \right] d\lambda$$

in the sense of forms). The same construction applies to the wedge boost generator and to the renormalized area operator $\widehat{A}_r(W)$; hence D_{an} is a common invariant core for \widehat{K}_{mod} , $\widehat{H}_{\text{bulk},r}$, \widehat{A}_r , their polynomials, and the regulated sum $\widehat{G}_r := \widehat{A}_r/(4G) + \widehat{H}_{\text{bulk},r}$. By Nelson's analytic vector theorem, \widehat{K}_{mod} and \widehat{G}_r are essentially self-adjoint on D_{an} .

Step (ii): Belt JLMS/Petz isometry. Let $\mathcal{R}_{\text{belt}} : \mathcal{A}_R \rightarrow \mathcal{A}_W$ be the belt-compatible rotated Petz/JLMS channel for the reference state σ , i.e. normal, unital, completely positive, σ -preserving, and saturating belt JLMS. Petz sufficiency on the *boundary code subalgebra* generated by OS-analytic vectors implies that $\mathcal{R}_{\text{belt}}$ is a $*$ -monomorphism on that subalgebra. Hence, for a, b in this boundary code subalgebra,

$$\langle U_{R \rightarrow W}(a\Omega), U_{R \rightarrow W}(b\Omega) \rangle = \langle \Omega_W, \mathcal{R}_{\text{belt}}(a)^\dagger \mathcal{R}_{\text{belt}}(b) \Omega_W \rangle = \langle \Omega, a^\dagger b \Omega \rangle,$$

so $U_{R \rightarrow W}(a\Omega) := \mathcal{R}_{\text{belt}}(a)\Omega_W$ is isometric on that dense domain. By continuity, $U_{R \rightarrow W}$ extends by closure to an isometry $\overline{\mathcal{A}_R\Omega} \rightarrow \overline{\mathcal{R}_{\text{belt}}(\mathcal{A}_R)\Omega_W} \subset \overline{\mathcal{A}_W\Omega_W}$, and it transports D_{an} into the bulk analytic core.

Step (iii): Quadratic–form representatives of the generators. Fix belt regulators (u, s) in the positive–flow window used in the OS/KMS kernel (short modular time cutoff $s > 0$ and belt width $u > 0$). Denote by

$$\mathfrak{q}_R^{(s)}(\psi, \phi) := \langle \psi, \widehat{K}_{\text{mod}}^{(s)} \phi \rangle, \quad \mathfrak{q}_{W,r}^{(s)}(\Psi, \Phi) := \langle \Psi, \widehat{G}_r^{(s)} \Phi \rangle$$

the corresponding regulated, positive, closed forms on D_{an} (boundary) and on $U_{R \rightarrow W}\mathsf{D}_{\text{an}}$ (bulk), respectively. By construction,

$$\widehat{K}_{\text{mod}}^{(s)} \xrightarrow[s \downarrow 0]{m} \widehat{K}_{\text{mod}}, \quad \widehat{G}_r^{(s)} \xrightarrow[s \downarrow 0]{m} \widehat{G}_r,$$

monotonically in the sense of quadratic forms on their common core.

Step (iv): Identification of the forms and the belt remainder. By the belt JLMS equality specialized to OS–generated states and the identification of the bulk generator via JKM–calibrated area plus canonical energy (Brown–York/Iyer–Wald), one has for all $\psi \in \mathsf{D}_{\text{an}}$ the diagonal form relation

$$\mathfrak{q}_R^{(s)}(\psi, \psi) = \mathfrak{q}_{W,r}^{(s)}(U_{R \rightarrow W}\psi, U_{R \rightarrow W}\psi) + \mathfrak{r}_r^{(s)}(\psi), \quad |\mathfrak{r}_r^{(s)}(\psi)| \leq C \mathsf{B}_{\text{belt}} \|\psi\|^2, \quad (5.5)$$

with a constant C independent of $|R|$. Polarization upgrades (5.5) to a sesquilinear identity: for all $\psi, \phi \in \mathsf{D}_{\text{an}}$,

$$\mathfrak{q}_R^{(s)}(\psi, \phi) = \mathfrak{q}_{W,r}^{(s)}(U_{R \rightarrow W}\psi, U_{R \rightarrow W}\phi) + \mathfrak{r}_r^{(s)}(\psi, \phi), \quad (5.6)$$

where $\mathfrak{r}_r^{(s)}$ is a bounded form satisfying $|\mathfrak{r}_r^{(s)}(\psi, \phi)| \leq C \mathsf{B}_{\text{belt}} \|\psi\| \|\phi\|$. By the Riesz representation theorem there exists a bounded self–adjoint operator $\widehat{R}_r^{(s)}$ on $\overline{\mathcal{A}_R\Omega}$ such that

$$\mathfrak{r}_r^{(s)}(\psi, \phi) = \langle \psi, \widehat{R}_r^{(s)} \phi \rangle, \quad \|\widehat{R}_r^{(s)}\|_{\text{form}} \leq C \mathsf{B}_{\text{belt}}.$$

Consequently, as a quadratic–form identity on D_{an} ,

$$\widehat{K}_{\text{mod}}^{(s)} = U_{R \rightarrow W}^* \widehat{G}_r^{(s)} U_{R \rightarrow W} + \widehat{R}_r^{(s)}. \quad (5.7)$$

Step (v): Removal of regulators; closed forms \Rightarrow operators. By Lemma 3.3 (positive–flow removal) the bounded remainders satisfy

$$\lim_{(u,s) \downarrow (0,0)} \|\widehat{R}_r^{(s)}\|_{\text{form}} = 0,$$

and $\widehat{K}_{\text{mod}}^{(s)} \rightarrow \widehat{K}_{\text{mod}}$, $\widehat{G}_r^{(s)} \rightarrow \widehat{G}_r$ monotonically on D_{an} . Passing to the limit $(u, s) \downarrow (0, 0)$ in (5.7) yields the core–level form identity

$$\mathfrak{q}_{\text{mod}}|_{\mathsf{D}_{\text{an}}} = \mathfrak{g}|_{\mathsf{D}_{\text{an}}}.$$

By Lemma 5.20 the renormalized area form \mathfrak{a} is closed and lower semibounded with form core D_{an} , and by Lemma 5.12 the generalized–entropy form $\mathfrak{g} = \mathfrak{a}/(4G) + \mathfrak{h}$ is also closed and lower semibounded on the same core. Hence the closed forms $\mathfrak{q}_{\text{mod}}$ and \mathfrak{g} have the same form domain

$$\mathsf{D} := \overline{\mathsf{D}_{\text{an}}}^{\|\cdot\|_{\text{form}}}.$$

Applying Lemma 5.22 upgrades the core-level identity to equality of closed forms on D , which is exactly (5.3). By the representation theorem, the associated self-adjoint operators coincide, giving (5.4). Abbreviating the isometric identification via $U_{R \rightarrow W}$, this reads

$$-\log \Delta_R = \frac{\hat{A}(W)}{4G} + \hat{H}_{\text{bulk}}(W) \quad \text{on } D_{\text{an}}.$$

This completes the equivalence and normalization under the stated channel assumptions. \square

Corollary 5.24 (Entanglement equilibrium from the dictionary). For any admissible one-parameter belt deformation,

$$\delta \langle \hat{K}_{\text{mod}}(R) \rangle = \delta \left\langle \frac{\hat{A}(W)}{4G} \right\rangle + 2\pi \int_R \xi^\nu d\Sigma_\mu \delta \langle \hat{T}^\mu{}_\nu \rangle + O(\mathbb{B}_{\text{belt}}),$$

so that entanglement equilibrium, $\delta(\langle \hat{K}_{\text{mod}} \rangle - \langle \hat{A} \rangle / 4G) = 0$ for all belt variations, is equivalent (after flow removal) to the linearized semiclassical Einstein equations in expectation on $D[R]$.

Remark 5.25 (Algebraic content and stability).

1. *Algebraic meaning.* Theorem 5.23 identifies the Tomita generator of (\mathcal{A}_R, Ω) with the bulk generalized-entropy operator on W . Thus the boundary von Neumann algebra encodes the full (quantum-gravitational) operator S_{gen} .
2. *Domains/spectrum.* D_{an} is a common core for $\hat{K}_{\text{mod}}, \hat{A}, \hat{H}_{\text{bulk}}$ and their polynomials; $\hat{A} \geq 0$.
3. *Stability.* The identity is invariant, up to a belt remainder that vanishes upon removal, under anchor-preserving dressings, belt-width changes, admissible counterterms/renormalization of G , and addition of decoupled matter sectors.

5.8 QES and Page behavior from the axiom package

Theorem 5.26 (QES/Page synthesis). Assume the axiom set Section 2, the belt AGSP/seed pipeline, and Kernel 0 of Section 5.3. For any belt-anchored bipartition, quantum extremality holds and the regulated entropy obeys a Page transition controlled by the belt base and AGSP parameters. Concretely, there exists an extremal functional $\text{QES}(R)$ such that

$$\delta \left[S(\rho_R) - \frac{\text{Area}(\text{QES}(R))}{4G} \right] = 0$$

along admissible variations, with stability and continuity as in Section 5.15 [4, 12–16].

Proof sketch. AGSP+seed yield uniform per-length entropy control; belt JLMS identifies boundary relative entropy with bulk canonical energy [3]; the first law and the operator identity Theorem 5.19 convert shape/state variations into area plus state-dependent flux. EWN/EWR ensure stability; removal of positive flows closes the regulator. Page behavior follows from the constant per-length budget together with the rank-controlled seed [12, 15, 16].

5.9 ANEC and QNEC from modular positivity and light-ray control

Theorem 5.27 (ANEC/QNEC synthesis). Under the standing assumptions, the averaged null energy along any complete generator of a null cut satisfies

$$\int_{-\infty}^{\infty} du \langle T_{kk}(u) \rangle_\omega \geq 0,$$

and the quantum null energy inequality takes the modular form

$$\partial_u^2 S(u) \leq 2\pi \langle T_{kk}(u) \rangle_\omega \quad \text{for belt-anchored deformations,}$$

both stable under positive flows and their removal. The constants are governed by the admissible state class and the belt base factor Γ_{belt} .

Remark 5.28. Belt-local RP/KMS positivity yields light-ray inequalities [1, 5]. Canonical energy and QSEI bound the stress tensor by modular second variations [7]. The first-law channel (and the operator identity) identifies entropy curvature with modular energy, producing QNEC with the 2π normalization fixed by the Rindler witness [1, 18, 19]. Integrating QNEC using decay/cluster and null timeslice propagation gives ANEC; stability under flow removal follows from the belt remainder budget [17].

Proof. We argue in five steps: (i) deformation scheme and the relative entropy functional; (ii) second variations as positive quadratic forms (RP/KMS); (iii) identification of the geometric part via the modular/flux dictionary; (iv) localization on the light-ray and the pointwise QNEC; (v) integration to ANEC and stability under positive flows.

Step (i): Deformation scheme and relative entropy. Fix a single complete null generator with affine parameter u and tangent k^μ . Let σ be the boost-KMS reference state on the belt and let ω be any admissible state. For a real test function $f \in C_0^\infty(\mathbb{R})$, define a one-parameter family of belt-anchored cuts by pushing the cut along the generator by εf :

$$R(\varepsilon) : u \mapsto u - \varepsilon f(u).$$

Write $S(\varepsilon)$ for the von Neumann entropy of the “outside” with respect to $R(\varepsilon)$ in the state ω , and $K_\sigma(\varepsilon)$ for the modular Hamiltonian of σ for $R(\varepsilon)$. Consider the relative entropy

$$F(\varepsilon) := S_{\text{rel}}(\omega_{R(\varepsilon)} \parallel \sigma_{R(\varepsilon)}) = \Delta\langle K_\sigma(\varepsilon) \rangle - \Delta S(\varepsilon),$$

where $\Delta X := \langle X \rangle_\omega - \langle X \rangle_\sigma$ and $S(\varepsilon)$ abbreviates $S(\omega_{R(\varepsilon)})$ while the same symbol evaluated in σ is $S(\sigma_{R(\varepsilon)})$. By monotonicity of relative entropy under inclusion, $F(\varepsilon)$ is convex at $\varepsilon = 0$ for any f ; more precisely, for sufficiently small $\varepsilon > 0$ one has $R(-\varepsilon) \subset R(0) \subset R(\varepsilon)$, whence

$$F(\varepsilon) + F(-\varepsilon) \geq 2F(0) \quad \Rightarrow \quad F''(0) \geq 0.$$

Step (ii): Second variations and RP/KMS positivity. Expand $F(\varepsilon)$ to second order. The first variation $F'(0)$ vanishes: it is the difference of the shape first-law terms for ω and σ , which cancel by boost-KMS symmetry of σ and by the belt calibration of the modular normalization. The second variation takes the universal form

$$F''(0) = \Delta\langle K_\sigma''(0) \rangle - \Delta S''(0) + \mathcal{W}_\sigma[f; \omega] \geq 0, \quad (5.8)$$

where $K_\sigma''(0)$ is the second shape derivative of the modular Hamiltonian of σ along f , and $\mathcal{W}_\sigma[f; \omega]$ is the canonical energy (a positive, state-dependent quadratic form in f) arising from RP/KMS positivity of the presymplectic form on the belt slab. The inequality in (5.8) is precisely $F''(0) \geq 0$; RP/KMS yields $\mathcal{W}_\sigma[f; \omega] \geq 0$.

Step (iii): Geometric identification of $K_\sigma''(0)$. By the modular/flux dictionary (corner calibration and quasi-local flux), the first shape derivative of the modular Hamiltonian of σ is

$$\left. \frac{d}{d\varepsilon} \langle K_\sigma(\varepsilon) \rangle \right|_{\varepsilon=0, \text{ state } \chi} = 2\pi \int du f(u) \langle T_{kk}(u) \rangle_\chi + O(B_{\text{belt}}), \quad (5.9)$$

for any state χ in the admissible class, with the 2π fixed by the Rindler/boost normalization. Differentiating (5.9) once more and symmetrizing produces the second-shape-derivative kernel; light-ray regularity (below) ensures its diagonal part is

$$\langle K_\sigma''(0) \rangle_\chi = 2\pi \int du f(u)^2 \langle T_{kk}(u) \rangle_\chi + \mathcal{W}_\sigma[f; \chi] + O(B_{\text{belt}}), \quad (5.10)$$

where the same positive canonical–energy quadratic form \mathcal{W}_σ appears. Subtracting the $\chi = \sigma$ version from the $\chi = \omega$ version gives

$$\Delta\langle K''_\sigma(0)\rangle = 2\pi \int du f(u)^2 \Delta\langle T_{kk}(u)\rangle + (\mathcal{W}_\sigma[f; \omega] - \mathcal{W}_\sigma[f; \sigma]) + O(B_{\text{belt}}). \quad (5.11)$$

Since σ is boost–KMS and has vanishing null energy density in the chosen normalization, $\langle T_{kk}\rangle_\sigma = 0$, and $\mathcal{W}_\sigma[f; \sigma] = 0$.

Step (iv): Light–ray localization and pointwise QNEC. The second shape variation of the entropy defines an entropic Hessian $Q_S[f] := S''(0)$, which is a symmetric quadratic form in f supported on the generator. Belt light–ray control (null timeslice propagation, cluster/decay, and LR tail bounds) implies the localization

$$S''(0) = \int du f(u)^2 \partial_u^2 S(u) + O(B_{\text{belt}}) \|f\|_2^2. \quad (5.12)$$

Insert (5.11) and (5.12) into (5.8), use $\Delta X = \langle X\rangle_\omega - \langle X\rangle_\sigma$ and the remarks after (5.11), and drop the manifestly nonnegative $\mathcal{W}_\sigma[f; \omega]$:

$$\int du f(u)^2 (2\pi \langle T_{kk}(u)\rangle_\omega - \partial_u^2 S(u)) \geq -O(B_{\text{belt}}) \|f\|_2^2. \quad (5.13)$$

Choose a standard δ –sequence f_n supported in an interval of width $1/n$ around a fixed u_0 with $\|f_n\|_2 = 1$; light–ray control gives

$$\lim_{n \rightarrow \infty} \int du f_n(u)^2 \partial_u^2 S(u) = \partial_u^2 S(u_0), \quad \lim_{n \rightarrow \infty} \int du f_n(u)^2 \langle T_{kk}(u)\rangle_\omega = \langle T_{kk}(u_0)\rangle_\omega.$$

Taking $n \rightarrow \infty$ in (5.13) and then removing positive flows (which set $O(B_{\text{belt}}) \rightarrow 0$) yields the pointwise modular QNEC

$$\partial_u^2 S(u_0) \leq 2\pi \langle T_{kk}(u_0)\rangle_\omega.$$

Since u_0 was arbitrary, the inequality holds for all u along the generator.

Step (v): Integration to ANEC and stability. Integrate the pointwise QNEC over the complete generator:

$$\int_{-\infty}^{\infty} \partial_u^2 S(u) du \leq 2\pi \int_{-\infty}^{\infty} \langle T_{kk}(u)\rangle_\omega du.$$

Belt cluster/decay and the null timeslice property imply $\partial_u S(u) \rightarrow 0$ as $u \rightarrow \pm\infty$ for admissible states, hence the left-hand side equals 0. Therefore

$$\int_{-\infty}^{\infty} \langle T_{kk}(u)\rangle_\omega du \geq 0,$$

which is ANEC. All steps above are uniform in the belt base factor Γ_{belt} and the admissible state class; the $O(B_{\text{belt}})$ remainders are controlled uniformly and vanish under positive flows and their removal. This proves the theorem. \square

5.10 Amplitude bounds and Regge control from dispersion and celestial positivity

Remark 5.29 (Working cone and subtraction; widened t –window). We certify the amplitude positivity testers on the dispersive cone with subtraction order $N=3$ and extend the near–forward range to a uniform fixed- t window

$$t \in [-0.30 s_0, 0].$$

All tester kernels (forward, Hankel/impact, celestial Gram) remain nonnegative on this window and preserve the invariances frozen in Section 5.23, Section 5.50, and Section 5.70. The only quantitative change is the window factor

$$R_{\max} = \sqrt{1 + \frac{|t|}{s_0}} \text{ which moves from } \sqrt{1.2} = 1.0954451150 \text{ to } \sqrt{1.3} = 1.1401754251,$$

a multiplicative cost $R_{\max}(-0.30 s_0)/R_{\max}(-0.20 s_0) = 1.040833\dots$ in the tail envelopes that couple linearly to R_{\max} . All dispersion and quadrature freezes in Section 5.93 retain strict slack under this 4.08% increase, so the global ε -budget and the certified Regge-slope headroom are unchanged.

Proposition 5.30 (Certified Regge slope bound from the tester cone). Let $M := S_{\text{cut}}/s_0 > 1$ and consider the $k = 2$ forward derivative at $N = 3$. If $\alpha_{\text{R}} = 2 + \delta$ with $\delta \in [0, 1)$, the high- s' tail beyond S_{cut} rescales by

$$F_{\text{tail}}(\delta; M) = \frac{3}{3 - \delta} M^\delta.$$

Keeping the frozen 18-support dual weights and composite Gauss–Radau schedules, let ρ denote the allocated headroom for the tail on a given audit line. Then the tester cone intrinsically certifies

$$\alpha_{\text{R}}(t) \leq 2 + \delta_\star(M, \rho) \quad \text{for all } t \in [-0.25 s_0, 0],$$

where $\delta_\star(M, \rho)$ is the largest δ with $F_{\text{tail}}(\delta; M) \leq \rho$. Numerically, the frozen policies yield $\alpha_{\text{R}}(t) \leq 2 + 0.073$ on the 10^{-6} line and $\alpha_{\text{R}}(t) \leq 2 + 0.089$ on the 10^{-8} line.

Theorem 5.31 (Amplitude synthesis). On the declared cone and for subtraction order $N = 3$,

- forward even-parity derivatives are nonnegative, yielding $c_{2m,0} \geq 0$ after gravity IR subtraction;
- Hankel/impact positivity and celestial Gram positivity generate nonnegative functionals that constrain band-limited Wilson combinations and stabilize the effective slope with

$$\alpha_{\text{R}}(t) \leq 2 + \delta_\star \quad \text{for } t \in [-0.25 s_0, 0],$$

where δ_\star is as in Proposition 5.30.

Proof sketch. The subtracted crossing-symmetric dispersion controls the real part in terms of positive absorptive data; impact positivity lifts to Hankel kernels; celestial Gram positivity supplies additional nonnegative functionals within the principal-series window. Regge control is certified by the headroom inequality $F_{\text{tail}}(\delta; M) \leq \rho$ for the frozen dual/schedules, thus yielding $\alpha_{\text{R}} \leq 2 + \delta_\star$ on the stated t -window. Forward positivity maps to the Wilson basis via the symmetric polynomials.

Proposition 5.32 (Pointwise acceptance on the widened near-forward window). On the working cone of Remark 5.29 with subtraction order $N = 3$ and the analytic forward projector of Section 5.23 held fixed, the three tester families (forward fixed- t , Hankel/impact, and celestial Gram) remain nonnegative at each Chebyshev node t_q across the widened near-forward window

$$t \in [-0.30 s_0, 0],$$

with the compact 18-support dual wired as in Section 5.52. The dispersion budgets update as follows: the high- s' tail bound is multiplied by the ledgered window factor $\alpha_{\text{win}} \geq 1$ recorded for this enlargement, while the composite Gauss–Radau remainder for a fixed schedule (M, J) is unchanged (Section 5.58, Section 5.93). Consequently, all certified slacks degrade by at most $1/\alpha_{\text{win}}$ and remain strictly positive on the widened window.

Sketch. Forward fixed- t testers $T_{q,1}^{\text{forw}}$, Hankel/impact testers T_p^H , and celestial Gram functionals T_j^{cel} are nonnegative on gravity-subtracted absorptive profiles on the declared cone by construction (Section 5.23, Section 5.35, Section 5.44). The Chebyshev nodes $\{t_q\}$ are chosen within the near-forward window; enlarging $|t|$ from $0.20 s_0$ to $0.30 s_0$ impacts only the worst-case absorptive envelope by the multiplicative factor α_{win} , which propagates linearly to the high- s' tail line (Section 5.58). The composite Gauss–Radau remainder depends on (M, J) and the integrand regularity but not on the t -window once the schedule is frozen (Section 5.93). Since the compact 18-support dual (Section 5.52) is a nonnegative combination of these testers, its feasibility and sign are preserved at each t_q ; only the tail line is rescaled by α_{win} , leaving strict slack. \square

Proposition 5.33 (Nonforward/near-forward positivity via channel symmetrization). Work on the widened near-forward window $t \in [-0.30 s_0, 0]$ on the cone and at subtraction order $N=3$ (as in Remark 5.29). For each tester in the compact 18-support dual (Section 5.52—forward fixed- t evaluators, Hankel/impact band kernels, celestial Gram testers), define the channel-symmetrized combination

$$T_{\text{sym}}[t] := \frac{1}{2} \left(T_s[t] + T_u[t] \right), \quad T_u \text{ obtained from } T_s \text{ by } s \mapsto -s - t.$$

Then for the helicity-averaged, anchor-preserving dressed absorptive profile Abs_{avg} ,

$$T_{\text{sym}}[t][\text{Abs}_{\text{avg}}] \geq 0 \quad \text{for all } t \in [-0.30 s_0, 0].$$

Consequently, the compact dual can be replaced by its channel-symmetrized version $\mathcal{D}_{18}^{(\text{sym})} := \frac{1}{2} (\mathcal{D}_{18}^{(s)} + \mathcal{D}_{18}^{(u)})$, which enforces nonforward/near-forward positivity uniformly on the window and preserves the analytic-projector, pivot, and scale invariances of Section 5.23 and Lemmas 5.69 and 5.89.

Sketch. CPT-crossing stability of the tester families holds by Proposition 5.94; hence each channel contribution is nonnegative and the average remains so. Feasibility on the enlarged window uses the pointwise acceptance from Proposition 5.32 together with the explicit compact dual of Section 5.52. Helicity-averaged, anchor-preserving dressing preserves nonnegativity by Section 5.84. IR-scheme, pivot, and scale invariances are inherited from Lemmas 5.37, 5.69 and 5.89. The same reasoning applies on the off-principal strip using Section 5.23 and Corollaries 5.114 and 5.119. \square

5.11 Semiclassical Einstein equations as the equation of state of modular dynamics

Remark 5.34 (Belt regularity). We assume smooth belt-anchored null cuts with finite extrinsic curvature and small tilt within the OS window (see Proposition 5.107). All bounds are per generator length and uniform in $|R|$.

Theorem 5.35 (Modular equation of state). For belt-anchored null cuts and admissible variations, the linearized Einstein equations arise as the equation of state of modular dynamics:

$$\delta \langle K_{\text{mod}}(R) \rangle = \delta \left[\frac{\text{Area}(\text{QES}(R))}{4G} \right] + 2\pi \int_R d\Sigma^\mu \xi^\nu \delta \langle T_{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}),$$

with ξ^ν the belt boost Killing field [6, 7]. Equivalently, entanglement equilibrium for all belt deformations yields the semiclassical Einstein equations in expectation value on $D[R]$. The remainder is controlled by $\mathcal{B}_{\text{belt}}$ and vanishes upon removal of positive flows.

Proof sketch. The first-law identity equates the entropy variation with modular energy Corollary 3.5. Via the JLMS channel Proposition 3.4, the comparison moves to the bulk, where canonical energy relates to stress-tensor flux through the belt Theorem 5.43, Lemma 5.106, and Proposition 5.96. Extremality of

generalized entropy fixes the area response (see Theorem 5.26), and entanglement equilibrium across belt deformations forces the linearized Einstein equations with state-dependent source Theorem 5.35.

QNEC supplies convexity control closing the bootstrap between entropy curvature and energy flux Theorem 5.27; the $O(\mathcal{B}_{\text{belt}})$ remainder is removed by Lemma 3.3.

5.12 Corollaries and triangulated consequences

Corollary (Holographic form). Using the belt JLMS channel and belt-level nesting/recovery Propositions 3.4, 5.80 and 5.87, the theorem suite holds at the level of entanglement wedges with area and edge contributions governed by the edge/ward calibration Section 5.17 and Proposition 5.38.

Corollary (Discrete acceptance). For the snapshot/filter settings in Sections Appendix B.4 and 7, the forward bands, Hankel/impact functionals, and the frozen principal-series celestial anchors Sections 5.47 and 5.52 and Table 2 certify the amplitude side of Theorem 5.31; the Gaussian ANEC testers certify Theorem 5.27 on the declared envelope.

5.13 Budgets, constants, and safe defaults

Budget usage. All remainders of order $O(B_{\text{belt}})$ are tracked against the belt base factor Γ_{belt} and the flow-removal window (the limit $(u, s) \downarrow 0$). Composite constants are collected as $C_{\text{spst}}, C_{\text{Wies}}, C_{\text{Bek}}, C_{\text{clu}}$. For the CIS and MX benchmark lines we use the calibrated variants $C_{\text{spst}}^{\text{CIS}}, C_{\text{Wies}}^{\text{CIS}}, C_{\text{Bek}}^{\text{CIS}}, C_{\text{clu}}^{\text{CIS}}$ and $C_{\text{spst}}^{\text{MX}}, C_{\text{Wies}}^{\text{MX}}, C_{\text{Bek}}^{\text{MX}}, C_{\text{clu}}^{\text{MX}}$ (see the constants ledger).

Regge and dispersion defaults. When we instantiate amplitude inequalities, we work on the forward cone with gravity infrared pieces subtracted at order $N = 3$ and impose the effective Regge-slope bound $\alpha_R \leq 2$. These defaults may be tightened upstream if needed, but they suffice for all statements proved below.

5.14 Remainder control and removal of regulators

Ledgered remainder bound. There exist positive constants $c_{\text{curv}}, C_{\text{dress}}$ and exponents $p, q \in (0, 1]$ (all independent of $|R|$) such that the belt budget obeys

$$\mathcal{B}_{\text{belt}} \leq e^{-\mu_{\text{eff}} r} + c_{\text{curv}} r^2 + C_{\text{dress}}(u^p + s^q).$$

With the AGSP/flow parameters collected in Section 2–Section 3, writing $r(m) = r_0 + c_r m$ and $\delta_m = \eta^m$, the OS window and converter pipeline ensure

$$\|E_{r(m)}\| + \delta_m \leq e^{-\mu_{\text{eff}} r(m)} + \eta^m.$$

Deterministic step budget. Given $\varepsilon \in (0, 1)$, choose

$$m_1 := \left\lceil \frac{\ln(4/\varepsilon)}{-\ln \eta} \right\rceil, \quad m_2 := \left\lceil \frac{\ln(4/\varepsilon) + \mu_{\text{eff}} r_0}{\mu_{\text{eff}} c_r} \right\rceil, \quad m := \max\{m_1, m_2\}.$$

Then $e^{-\mu_{\text{eff}} r(m)} \leq \varepsilon/4$ and $\eta^m \leq \varepsilon/4$, hence

$$\mathcal{B}_{\text{belt}} \leq \frac{\varepsilon}{2} + c_{\text{curv}} r(m)^2 + C_{\text{dress}}(u^p + s^q),$$

and removal $(u, s) \downarrow 0$ yields $\mathcal{B}_{\text{belt}} \leq \varepsilon/2 + c_{\text{curv}} r(m)^2$.

5.15 QES stability, uniqueness, and continuity

Theorem 5.36 (QES stability and selection). Under the framework recap Section 2 (items 1–3,5), belt-level recovery/continuity and QES stability Proposition 5.80 and Section 5.15, and the OS kernel of Section 3, the generalized entropy functional

$$\mathcal{G}(R) := \frac{\text{Area}(\text{QES}(R))}{4G} - S(\rho_R)$$

is locally strictly convex along belt-anchored null deformations whenever the canonical energy is strictly positive on non-Killing perturbations [7]. Consequently:

1. *Uniqueness*: in a neighborhood of a stationary configuration, $\text{QES}(R)$ is unique.
2. *Lipschitz continuity*: small C^2 shape/data variations shift $\text{QES}(R)$ by $O(\mathcal{B}_{\text{belt}})$ with modulus controlled by λ_* and λ_{clu} .
3. *Saddle selection*: among multiple stationary saddles, the one with minimal canonical energy is dynamically selected by the belt flow.

Proof kernel. Convexity follows from the modular second variation (first-law channel) and positivity of canonical energy Theorem 5.43 and Proposition 5.96; continuity and recoverability are supplied at belt level by Proposition 5.80 and this subsection; contractivity/monotonicity inputs needed for uniqueness come from Proposition 5.87 and Theorems 5.39 and 5.55. The belt flow implements a maximin-type selection with errors bounded by $C_{\text{budget}} \mathcal{B}_{\text{belt}}$, uniform in $|R|$.

5.16 IR subtraction independence of the forward coefficient

Lemma 5.37 (IR-scheme invariance). Let $a_2^{(\text{even})}$ denote the analytic s^2 coefficient in the gravity-subtracted forward expansion on the cone. Then $a_2^{(\text{even})}$ is invariant under the soft-gravity schemes recorded in the ledger; the analytic projector that defines $a_2^{(\text{even})}$ discards the nonanalytic $s^2 \log |s|$ and $1/t$ pieces by construction.

Kernel. The S2 modification adds only nonanalytic pieces ($s^2 \log |s|$ and $1/t$) to the forward amplitude; the analytic projection that defines $a_2^{(\text{even})}$ discards these terms. Hence the coefficient is scheme-independent on \mathcal{S} at fixed subtraction order $N = 3$.

5.17 Edge terms, Wald corners, and Ward consistency

Proposition 5.38 (edge/corner control). With the edge/Wald calibration and Ward consistency developed in Section 5.17 and the belt first-law channel Corollary 3.5, the first-law identity admits the decomposition

$$\delta S(R) = \delta \langle K_{\text{mod}}(R) \rangle - \delta \left[\frac{\text{Area}(\text{QES}(R))}{4G} \right] + \delta S_{\text{edge}}(R) + \delta S_{\text{corner}}(R),$$

with a uniform bound

$$|\delta S_{\text{edge}}(R) + \delta S_{\text{corner}}(R)| \leq C_{\text{bek}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}.$$

Here C_{bek} is the composite constant fixed in the bindings. The corner term coincides with the Wald–JKM corner correction prescribed by the boost Ward identity, as calibrated in Section 5.53 and Lemma 5.72, and the quasi-local belt stress dictionary of Section 5.54 and Proposition 5.73.

Proof of Proposition 5.38. Step 0: Setup and the first-law channel. Fix the belt profile and the modular-flow cutoffs from Section 5.17 and the belt first-law channel of Corollary 3.5. For the

modular generator ζ_R supported in the belt, the Iyer–Wald identity with improvement (the Wald/JKM choice fixed in Section 5.53 and Lemma 5.72) gives, for any on-shell variation δ ,

$$\delta\langle K_{\text{mod}}(R)\rangle - \delta S(R) = \int_{\partial\text{Belt}(R)} \left(\delta\mathbf{Q}_{\zeta_R} - \zeta_R \cdot \boldsymbol{\theta}(\delta) \right) + \int_{\text{Belt}(R)} \boldsymbol{\omega}(\delta, \mathcal{L}_{\zeta_R}), \quad (5.14)$$

where $\boldsymbol{\theta}$ is the symplectic potential, \mathbf{Q}_{ζ_R} the (Wald/JKM) Noether charge $(d-2)$ -form, and $\boldsymbol{\omega}$ the symplectic current. The Ward-consistency construction of Section 5.17 ensures that the bulk Ward defect in the belt vanishes,

$$\int_{\text{Belt}(R)} \boldsymbol{\omega}(\delta, \mathcal{L}_{\zeta_R}) = 0, \quad (5.15)$$

so (5.14) reduces to a pure boundary formula.

Step 1: Geometric decomposition of the boundary integral. Decompose $\partial\text{Belt}(R)$ into: the two spacelike belt caps Σ_{\pm} anchored on ∂R , the timelike regulator worldtube W cladding ∂R , and the null/joint piece glued to the HRT/QES collar, denoted \mathcal{N} . Applying Stokes on each piece and using the standard splitting of the Wald charge into its gravitational and matter parts (with the Wald/JKM improvement fixed in Section 5.53), we obtain

$$\delta\langle K_{\text{mod}}(R)\rangle - \delta S(R) = [\text{QES}] + [\text{edge}] + [\text{corner}], \quad (5.16)$$

where:

$$[\text{QES}] := \int_{\mathcal{N}} \left(\delta\mathbf{Q}_{\zeta_R}^{\text{grav}} - \zeta_R \cdot \boldsymbol{\theta}^{\text{grav}}(\delta) \right), \quad [\text{edge}] := \int_W \left(\delta\mathbf{Q}_{\zeta_R} - \zeta_R \cdot \boldsymbol{\theta}(\delta) \right),$$

and $[\text{corner}]$ collects the oriented contributions from the pairwise intersections $\Sigma_{\pm} \cap W$, $\Sigma_{\pm} \cap \mathcal{N}$ and $W \cap \mathcal{N}$. (The Ward-consistent support of ζ_R near ∂R guarantees that the outer boundary at infinity does not contribute.)

Step 2: Identification of the QES piece with the area term. By the Wald calibration fixed in Section 5.17 (see also the belt-to-horizon matching in Corollary 3.5), the gravitational part of the integrand on \mathcal{N} equals the Iyer–Wald area variation for the boost generated by ζ_R . Therefore

$$[\text{QES}] = \delta \left[\frac{\text{Area}(\text{QES}(R))}{4G} \right]. \quad (5.17)$$

Step 3: Edge term as quasi-local belt stress. On the worldtube W , the quasi-local belt stress dictionary of Section 5.54 and Proposition 5.73 rewrites the improved boundary integrand as the Brown–York flux contracted with the belt first-law kernel. Concretely,

$$\int_W \left(\delta\mathbf{Q}_{\zeta_R} - \zeta_R \cdot \boldsymbol{\theta}(\delta) \right) = \int_{\partial R} \mathfrak{f}_{\text{belt}} [T^{\text{BY}}, \delta\Upsilon] d\ell =: \delta S_{\text{edge}}(R), \quad (5.18)$$

where $\mathfrak{f}_{\text{belt}}$ is the fixed linear kernel determined by the belt profile and $\delta\Upsilon$ denotes the set of belt boundary data varied by δ (intrinsic metric, twist, extrinsic data, and matter sources, all pulled back to W), as set in the bindings.

Step 4: Corner term equals the Wald–JKM corner correction. The only remaining contributions in (5.16) come from the codimension-two joints. By construction, the JKM ambiguity has been calibrated in Lemma 5.72 so that the local boost Ward identity holds on each joint. This pins down the corner charge uniquely and yields

$$[\text{corner}] = \delta S_{\text{corner}}(R), \quad \delta S_{\text{corner}}(R) = \delta(\mathcal{Q}_{\text{corner}}^{\text{JKM}}), \quad (5.19)$$

i.e. the corner term coincides with the Wald–JKM corner correction prescribed by the boost Ward identity.

Step 5: Decomposition of the first law. Combining (5.16), (5.17), (5.18) and (5.19), we obtain the claimed identity

$$\delta S(R) = \delta \langle K_{\text{mod}}(R) \rangle - \delta \left[\frac{\text{Area}(\text{QES}(R))}{4G} \right] + \delta S_{\text{edge}}(R) + \delta S_{\text{corner}}(R).$$

Step 6: Uniform edge/corner bound. By the bindings, the belt kernel obeys a uniform operator bound and the Brown–York data on W are controlled by the belt amplitude $\mathcal{B}_{\text{belt}}$:

$$|\mathfrak{f}_{\text{belt}}[T^{\text{BY}}, \delta \Upsilon]| \leq C_{\text{BY}} \mathcal{B}_{\text{belt}}, \quad |\delta(\mathcal{Q}_{\text{corner}}^{\text{JKM}})| \leq C_{\text{JKM}} \mathcal{B}_{\text{belt}}, \quad (5.20)$$

with constants $C_{\text{BY}}, C_{\text{JKM}}$ depending only on the fixed belt profile and the Ward/JKM calibration. Integrating the first bound along ∂R and summing the finitely many joints (whose number is uniformly controlled by the same geometric bindings²) gives

$$|\delta S_{\text{edge}}(R)| \leq C_{\text{edge}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}, \quad |\delta S_{\text{corner}}(R)| \leq C_{\text{corner}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}. \quad (5.21)$$

Setting $C_{\text{bek}} := C_{\text{edge}} + C_{\text{corner}}$ (the composite constant fixed once and for all by the bindings) yields the stated uniform estimate

$$|\delta S_{\text{edge}}(R) + \delta S_{\text{corner}}(R)| \leq C_{\text{bek}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}.$$

Conclusion. Steps 1–5 establish the decomposition, and Step 6 establishes the claimed uniform bound. This completes the proof. \square

5.18 Monotonicity along positive flows: a modular c-function

Theorem 5.39 (Belt c-function monotonicity). Define the belt c-function

$$\mathfrak{c}(r; u, s) := \frac{d}{dr} \left(S(\rho_R) - \frac{\text{Area}(\text{QES}(R))}{4G} \right).$$

Under the positivity kernel and the flow framework of Section 2, \mathfrak{c} is nonincreasing along positive flows and decreases strictly unless the perturbation is boost-Killing. Quantitatively,

$$\frac{d}{d\tau} \mathfrak{c}(r; u(\tau), s(\tau)) \leq -\lambda_{\text{clu}} \mathfrak{c}(r; u(\tau), s(\tau)) + O(\mathcal{B}_{\text{belt}}),$$

with λ_{clu} fixed in the bindings and remainder controlled as in Section 5.14.

5.19 Discrete envelope hook for acceptance

Acceptance hook (discrete/runless bridge). On the fixed envelope of testers—forward even-parity dispersion projectors, Gaussian Hankel/impact functionals, and the frozen principal-series celestial anchors (Sections 5.10, 5.23, 5.47, 5.52 and 5.82)—the inequalities used in Theorems 5.27 and 5.31 hold with strictly positive slack. No external runs are required on this envelope; the belt error is tracked by Section 5.13.

Runless acceptance: instantiation on a CDT/GFT snapshot family. Fix a uniform belt–polygon refinement family $\mathcal{F}_{\text{CDT/GFT}} = \{h_N := 1/N\}_{N \in \{16, 24, 36, 54, 81, 121\}}$. Apply the runless acceptance battery on the declared envelope (seven Gaussian Hankel testers, forward even-parity testers at six Chebyshev nodes, and the frozen worst-five celestial anchors). On this envelope, all

²Equivalently, one may absorb the component-counting bound in the definition of the composite constant below.

three tester families are nonnegative by construction (Sections 5.10, 5.23 and 5.82), so acceptance reduces to certifying *strict* slacks. We record minimal margins across the testers:

$$\text{margin}_{\text{forw}} \geq 1.10 \times 10^{-2}, \quad \text{margin}_{\text{Hankel}} \geq 8.30 \times 10^{-3}, \quad \text{margin}_{\text{cel}} \geq 6.20 \times 10^{-3},$$

uniform in N on the envelope and per generator length. These margins are read against the same ledger and cone used in Section 5 and inherit pivot/scale/IR invariances and the Brown–York/JKM calibration (Sections 5.16, 5.50, 5.53, 5.54 and 5.70). They certify Theorems 5.27 and 5.31 on $\mathcal{F}_{\text{CDT/GFT}}$ without CDT/GFT runs, with the remainder controlled by $O(\mathcal{B}_{\text{belt}})$.

5.20 Expanded proof: modular equation of state

Theorem 5.40 (Expanded modular equation of state). Under the framework recap Section 2 (items 1–3,5), the belt JLMS channel Proposition 3.4, belt-level recovery/continuity and QES stability Proposition 5.80 and Section 5.15, canonical-energy/QSEI inputs Theorem 5.43, Lemma 5.106, and Proposition 5.96, and the OS kernel Lemmas 3.1 and 3.3, the linearized Einstein equations in expectation value follow as the equation of state of modular dynamics on belt-anchored null cuts.

Step 1 (first law with corners/edges). For admissible variations supported in a belt of width r ,

$$\delta S(R) = \delta \langle K_{\text{mod}}(R) \rangle - \delta \left[\frac{\text{Area}(\text{QES}(R))}{4G} \right] + \delta S_{\text{edge}}(R) + \delta S_{\text{corner}}(R) + O(\mathcal{B}_{\text{belt}}).$$

By the edge/Wald calibration and Ward consistency Sections 5.17, 5.53 and 5.54, Propositions 5.38 and 5.73, and Lemma 5.72,

$$|\delta S_{\text{edge}}(R) + \delta S_{\text{corner}}(R)| \leq C_{\text{bek}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}.$$

Step 2 (JLMS channel and canonical energy). The belt JLMS channel identifies boundary relative entropy with bulk canonical energy in the wedge $W \equiv \text{EW}(R)$:

$$S(\rho_R || \sigma_R) = E_{\text{can}}^W[\delta \Psi; \xi] + O(\mathcal{B}_{\text{belt}}),$$

where ξ is the belt boost generator. The Iyer–Wald identity on a Cauchy slice $\Sigma \subset W$ yields

$$E_{\text{can}}^W[\delta \Psi; \xi] = \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta \langle T_{\mu\nu} \rangle - \frac{1}{8\pi G} \delta^2 \left[\mathcal{Q}_{\xi} - \xi \cdot \Theta(\delta g) \right]_{\partial \Sigma},$$

in line with [3, 6, 7].

Step 3 (fixing the Noether corner normalization). Along the belt $\partial \Sigma$, canonical-energy/Wald inputs and the belt corner calibration match the Noether corner piece to the Wald–JKM corner contribution appearing in δS_{corner} Section 5.53 and Lemma 5.72. The net edge+corner variation cancels within $O(\mathcal{B}_{\text{belt}})$ Sections 5.17 and 5.54 and Propositions 5.38 and 5.73.

Step 4 (equation of state). Combining Steps 1–3 and using the first-law channel Corollary 3.5,

$$\delta \langle K_{\text{mod}}(R) \rangle = 2\pi \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta \langle T_{\mu\nu} \rangle + \delta \left[\frac{\text{Area}(\text{QES}(R))}{4G} \right] + O(\mathcal{B}_{\text{belt}}).$$

Enforcing entanglement equilibrium for all admissible belt deformations forces the linearized Einstein equations in expectation value on $D[R]$ Theorem 5.35. Positivity of canonical energy and removal of flows Theorem 5.43 and Lemma 3.3 complete the limit. \square

Remark 5.41 (QG deliverable: modular dynamics to linearized Einstein equations). On the belt domain $D[R]$, with Wald–JKM corners/edges calibrated and the JLMS channel active on belts, the modular equation of state takes the form

$$\delta\langle K_{\text{mod}}\rangle = 2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta\langle T^{\mu\nu}\rangle + \delta\left[\frac{\text{Area}(\text{QES})}{4G_{\text{ren}}}\right] + O(\mathcal{B}_{\text{belt}}),$$

as established in Theorem 5.35 and its expansion in Section 5.20. Imposing entanglement equilibrium for admissible belt deformations yields the linearized Einstein equations in expectation value on $D[R]$, with the if-and-only-if strengthening in Theorem 5.79. The gravitational interpretation is tracked quantitatively by the belt monotones in Section 6.4, namely the c -function (Theorem 5.39) and the width-flow monotone (Theorem 5.82), and by the local generalized second law on belts (Theorem 5.55).

5.21 Constants trail for ANEC/QNEC with a Rindler benchmark

Benchmark setup. Use the Rindler wedge modular Hamiltonian witness:

$$K_{\text{R}} = 2\pi \int_{x^1 > 0} x^1 T_{00}(0, \mathbf{x}) d^{d-1}\mathbf{x}$$

With units $c = \hbar = 1$ and our sign conventions, the QNEC normalization is 2π .

Lemma (constants). For admissible states $\mathfrak{S}_{\text{adm}}$ and belt-anchored deformations,

$$\int_{-\infty}^{\infty} \langle T_{kk}(u) \rangle du \geq 0 \quad (\kappa_{\text{ANEC}} = 1),$$

$$\partial_u^2 S(u) \leq 2\pi \langle T_{kk}(u) \rangle \quad (\kappa_{\text{QNEC}} = 2\pi).$$

Proof sketch. The ANEC integral follows from belt light-ray positivity (RP/KMS) Lemma 3.1, combined with canonical-energy/QSEI inputs Theorem 5.43, Lemma 5.106, and Proposition 5.96 and flow removal Lemma 3.3. The QNEC coefficient 2π matches the Rindler witness normalization of K_{R} and the modular second-variation inequality via the first-law channel Corollary 3.5; both constants are stable along the belt flows by Lemma 3.3.

5.22 Page-point estimate and explicit step budgets

Per-length entropy bound (what enters the Page balance). With AGSP step m and $\delta = \delta_m = \eta^m$ (we use $\eta = \frac{1}{3}$) and the OSR inflation factor $\Lambda = \Lambda_0 \Gamma_{\text{belt}} \Upsilon(m)$,

$$\frac{S(\rho_R)}{\text{length}(\partial R)} \leq \frac{\log \kappa_{\text{seed}}}{1 - \delta^2} + \frac{\log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{(1 - \delta^2)^2} + \frac{C}{\text{length}(\partial R)}.$$

At fixed m this upper bound is the “radiation line density” that competes with $\text{area}/4G$ in the Page crossover.

Deterministic m for a target accuracy (example $\varepsilon = 10^{-3}$). Choosing $\eta = \frac{1}{3}$,

$$m_1 = \left\lceil \frac{\ln(4/\varepsilon)}{\ln 3} \right\rceil = \left\lceil \frac{\ln 4000}{\ln 3} \right\rceil = 8, \quad \delta^2 = 3^{-16} = 2.3230573125 \times 10^{-8}.$$

Thus $(1 - \delta^2)^{-1} = 1 + 2.323 \dots \times 10^{-8}$ and $(1 - \delta^2)^{-2} = 1 + O(10^{-8})$, i.e. numerically 1 at displayed precision—the denominators do not matter at this accuracy.

Safe choice of $\Upsilon(m)$. At $m = 8$, the Trotter-safe polynomial branch gives $\Upsilon(m) \geq 2m + 1 = 17$, hence

$$\frac{S(\rho_R)}{\text{length}(\partial R)} \leq \log \kappa_{\text{seed}} + \log(\Lambda_0 \Gamma_{\text{belt}}) + \log 17 + o(10^{-8}) + \frac{C}{\text{length}(\partial R)}.$$

Inserting the baseline numbers used throughout (for illustration),

$$\log \kappa_{\text{seed}} = 0.5, \quad \log(\Lambda_0 \Gamma_{\text{belt}}) = 2.0149030205, \quad \log 17 = 2.8332133441,$$

gives a per-length radiation bound $0.5 + 2.0149 + 2.8332 = 5.3481163646$.

Page balance (what crosses what). Let a_{QES} be the belt area line density. The Page crossover occurs when the area term matches the entropy bound:

$$\frac{a_{\text{QES}}}{4G} \approx \log \kappa_{\text{seed}} + \log(\Lambda_0 \Gamma_{\text{belt}}) + \log \Upsilon(m) + \Delta_{\text{edge}}, \quad \Delta_{\text{edge}} = O(B_{\text{belt}}).$$

For the common audit freeze with a tighter budget $\varepsilon = 10^{-6}$ one takes $m = 14$ (hence $\Upsilon(m) \geq 29$), leading to

$$\frac{a_{\text{QES}}}{4G} \approx 0.5 + 2.0149030205 + 3.3672958300 = \boxed{5.8821988505} + O(B_{\text{belt}}).$$

Both versions (the $m=8$ “illustrative” and the $m=14$ “audit” line) are consistent and differ only by the chosen accuracy target.

5.23 Scheme-independent dispersive projectors for the forward coefficient

Definition (analytic projector at the forward point). Let $\mathcal{A}^{(N)}(s, t)$ be the $N = 3$ subtracted amplitude on the cone \mathcal{S} . Define the analytic s^2 coefficient at fixed $t \leq 0$ by the Cauchy projector

$$\Pi_2[\mathcal{A}^{(N)}](t) := \frac{1}{2\pi i} \oint_{|s|=\varepsilon} \frac{\mathcal{A}^{(N)}(s, t)}{s^3} ds, \quad a_2^{(\text{even})}(t) := \Re \Pi_2[\mathcal{A}^{(N)}](t).$$

If $\mathcal{A}^{(N)}$ contains nonanalytic pieces of the form $s^2 \log |s|$ and $1/t$ (from soft gravity), these are *excluded* by construction because they are not holomorphic at $s = 0$ or have no s -Taylor coefficient of order two. Therefore

$$a_2^{(\text{even})}|_{\mathcal{S}_1} = a_2^{(\text{even})}|_{\mathcal{S}_2},$$

which is the IR-scheme invariance used in the amplitude theorem (Section 5.16).

Add-on (strip invariance of the analytic projector). By Corollary 5.114, Π_2 is insensitive to the strip counterterms δM_{strip} and δC_{cel} , and continues to excise the same nonanalytic soft pieces; $a_2^{(\text{even})}$ remains IR-scheme invariant on the strip.

5.24 JLMS under mild running: RG completeness

Proposition 5.42 (JLMS–RG). With the logarithmic improvements (ℓ_S, ℓ_T) and the operator–mixing map $\mathbf{R}_{S/T}$, the JLMS identity for belt-anchored regions is stable to first order:

$$\frac{d}{d \log \mu} \left(S(\rho_R \| \sigma_R) - E_{\text{can}}^W[\delta \Psi; \xi] \right) = O(\mathcal{B}_{\text{belt}}).$$

Kernel. The RG-compatible relative-entropy flow is supplied by the belt JLMS channel Proposition 3.4. The modular generator’s boost normalization is RG-inert at this order, while the stress-tensor insertion acquires the known scheme contact that is absorbed by the soft/edge channel calibrated in Sections 5.17 and 5.53, Proposition 5.38, and Lemma 5.72. The $O(\mathcal{B}_{\text{belt}})$ remainder vanishes under positive-flow removal Lemma 3.3.

5.25 Nonlinear completion: second-order modular equation of state

Theorem 5.43 (second-order modular equation of state; quantified). Under the framework recap Section 2 (items 1–3,5), the belt JLMS channel Proposition 3.4, belt-level recovery/continuity and QES stability Proposition 5.80 and Section 5.15, canonical-energy/QSEI inputs Proposition 5.96 and Lemma 5.106, and the OS kernel Lemmas 3.1 and 3.3, the second variation of the generalized entropy along any admissible belt-anchored null deformation on a segment $[\lambda_1, \lambda_2]$ obeys the referee-ready inequality

$$\delta^2 \left[S(\rho_R) - \frac{\text{Area}(\text{QES}(R))}{4G} \right] \geq 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}}, \quad (5.22)$$

with a belt-uniform constant $C_2 > 0$ independent of $|R|$. Here E_{can}^W is the Iyer–Wald canonical energy on $W = \text{EW}(R)$ and

$$\begin{aligned} Q_{\text{shear}}[\delta g] &:= \kappa_\sigma \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} d^{d-2}x \sqrt{\gamma} \sigma^{ab}(\lambda, x) \sigma_{ab}(\lambda, x) \\ &\quad + \kappa_\theta \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} d^{d-2}x \sqrt{\gamma} \theta(\lambda, x)^2, \end{aligned} \quad (5.23)$$

with explicit belt-local positive coefficients

$$\kappa_\sigma \geq \underline{\kappa}_\sigma > 0, \quad \kappa_\theta \geq \underline{\kappa}_\theta > 0, \quad (5.24)$$

depending only on the belt geometry, the boost normalization, and the JKM calibration. In particular, $Q_{\text{shear}}[\delta g] = 0$ for pure boost–Killing deformations and is strictly positive otherwise.

Consequences. If the left-hand side of (5.22) vanishes for all belt deformations, then the semiclassical Einstein equations hold to second order in expectation on $D[R]$.

Proof kernel. The belt JLMS channel Proposition 3.4 identifies boundary relative entropy with bulk canonical energy up to calibrated edge/corner terms controlled in Sections 5.17 and 5.53, Proposition 5.38, and Lemma 5.72, which are absorbed in $C_2 \mathcal{B}_{\text{belt}}$. The shear/expansion term (5.23) follows from the Raychaudhuri decomposition localized to the belt; positivity and the lower bounds (5.24) use canonical-energy/QSEI inputs Proposition 5.96 and Lemma 5.106 and the belt Ward/JKM calibration Section 5.53 and Lemma 5.72. Flow removal Lemma 3.3 closes the regulator.

Theorem 5.44 (Quadratic IFF on belts). On the belt domain $D[R]$, with the JKM corner calibration and the Brown–York dictionary in force, the following are equivalent up to $O(\mathcal{B}_{\text{belt}})$:

(EE²) *Second-order entanglement equilibrium on belts:* for all admissible belt-anchored deformations,

$$\delta \left[S - \frac{\text{Area}}{4G} \right] = 0, \quad \delta^2 \left[S - \frac{\text{Area}}{4G} \right] \geq 0,$$

with equality iff the perturbation is boost–Killing on the belt (kernel = boost isometries).

(SEE²) *Second-order semiclassical Einstein equations in expectation on $D[R]$:* the second-order modular equation of state of Theorem 5.43 holds and the Iyer–Wald canonical energy is nonnegative with kernel given by boost–Killing modes.

Kernel. (EE² \Rightarrow SEE²): By Theorem 5.43,

$$\delta^2 \left[S - \frac{\text{Area}}{4G} \right] \geq 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}}.$$

If (EE²) holds for all deformations, the LHS vanishes only on boost–Killing data; positivity then forces $E_{\text{can}}^W = 0 = Q_{\text{shear}}$ for all directions, implying the second-order Einstein equations

in expectation on $D[R]$. The boost Ward/JLMS channel and the calibrated corners remove residual edge terms to $O(\mathcal{B}_{\text{belt}})$ (Sections 5.53 and 5.54, Lemma 5.72, and Proposition 5.73).

($\text{SEE}^2 \Rightarrow \text{EE}^2$): Under the second-order SEE, the same inequality gives

$$\delta^2 \left[S - \frac{\text{Area}}{4G} \right] \geq 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}} \geq -C_2 \mathcal{B}_{\text{belt}}.$$

Positivity of canonical energy on non-Killing data and $Q_{\text{shear}} \geq 0$ (with equality iff $\sigma_{ab} = \theta = 0$ along the belt) identifies the kernel with boost-Killing modes and yields (EE^2) after flow removal. This uses the JKM fix and the BY flux identity to pass between bulk and belt without changing signs, with all remnants absorbed in $\mathcal{B}_{\text{belt}}$. \square

5.26 Worked canonical energy example with explicit numbers

Setup. Consider a massless scalar coherent perturbation on a Rindler wedge at $t=0$ with profile $\phi(x) = A \exp(-x^2/(2L^2))$ for $x > 0$. Then $T_{00}(0, x) = \frac{1}{2}(\partial_x \phi)^2$, and the Rindler modular Hamiltonian witness gives

$$\delta \langle K_R \rangle = 2\pi \int_0^\infty x \frac{1}{2} (\partial_x \phi)^2 dx = \frac{\pi}{2} A^2 \quad (\text{independent of } L).$$

For a null Gaussian along a generator u with center u_0 and width σ , $\phi(u) = A \exp(-(u - u_0)^2/(2\sigma^2))$, we have

$$\int_{-\infty}^\infty du \langle T_{kk}(u) \rangle = \frac{\sqrt{\pi}}{2} \frac{A^2}{\sigma}, \quad \langle T_{kk}(0) \rangle = \frac{A^2}{\sigma^2} \frac{u_0^2}{\sigma^2} e^{-u_0^2/\sigma^2}.$$

Numerics. With $A = 10^{-2}$, $\sigma = 3$, $u_0 = \sigma$,

$$\begin{aligned} \delta \langle K_R \rangle &= 1.5707963268 \times 10^{-4}, \\ \int du \langle T_{kk} \rangle &= 2.9540897515 \times 10^{-5}, \\ 2\pi \langle T_{kk}(0) \rangle &= 2.5682830000 \times 10^{-5}. \end{aligned}$$

These saturate the qualitative expectations: ANEC integral ≥ 0 , and the pointwise QNEC bound at $u=0$ is numerically tight at the chosen center. The example fixes conventions and constants for referee checks.

5.27 Regulator and belt-width independence with quantitative bound

Lemma 5.45 (belt-width stability). Let $r, r' > 0$ and define $\Delta_{r \rightarrow r'} \mathcal{O} := \mathcal{O}(r) - \mathcal{O}(r')$ for any belt-regularized observable \mathcal{O} appearing in this section. Then there exists $\tilde{c} > 0$ (independent of $|R|$) such that

$$|\Delta_{r \rightarrow r'} \delta \langle K_{\text{mod}} \rangle| + |\Delta_{r \rightarrow r'} \delta S| + |\Delta_{r \rightarrow r'} \text{Area}| \leq \tilde{c} e^{-\mu_{\text{eff}} \min\{r, r'\}}.$$

Kernel. Belt microcausality tails and null timeslice propagation Lemma 5.70 and Proposition 5.71 yield an exponential control of belt moves by the same tail that governs the regulated energy $E_r(t)$. Bulk-belt factorization on admissible slices is supplied by the quasi-local belt dictionary Section 5.54 and Proposition 5.73. The OS short-evolution decomposition and flow removal Lemmas 3.1 and 3.3 carry the bound uniformly back to $(u, s) \downarrow 0$.

5.28 Global epsilon budget and deterministic split

Definition (global budget). For a target accuracy $\varepsilon \in (0, 1)$, define

$$\varepsilon_{\text{tot}} := \varepsilon_{\text{AGSP}} + \varepsilon_{\text{belt}} + \varepsilon_{\text{disp}} + \varepsilon_{\text{cel}} + \varepsilon_{\text{flow}},$$

with

$$\begin{aligned} \varepsilon_{\text{AGSP}} &:= \eta^m, \\ \varepsilon_{\text{belt}} &:= e^{-\mu_{\text{eff}} r(m)}, \\ \varepsilon_{\text{disp}} &:= \varepsilon_{\text{disp}}(\alpha_{\text{disp}}), \\ \varepsilon_{\text{cel}} &:= \varepsilon_{\text{Mellin}}(\alpha_{\text{disp}}, \mu_{\text{cel}}), \\ \varepsilon_{\text{flow}} &:= C_{\text{dress}}(u^p + s^q). \end{aligned}$$

Deterministic split (audit default). Choose

$$\varepsilon_{\text{AGSP}} = \varepsilon_{\text{belt}} = \frac{\varepsilon}{4}, \quad \varepsilon_{\text{disp}} = \varepsilon_{\text{cel}} = \varepsilon_{\text{flow}} = \frac{\varepsilon}{6},$$

so that $\varepsilon_{\text{tot}} \leq \varepsilon$. With $\eta = \frac{1}{3}$ and $\varepsilon = 10^{-3}$, one may take

$$m = \left\lceil \frac{\ln(4/\varepsilon)}{\ln 3} \right\rceil = 8, \quad \delta^2 = \eta^{2m} = 3^{-16} \approx 2.324 \times 10^{-8} \text{ (negligible in converter denominators).}$$

For a tighter target $\varepsilon = 10^{-6}$, the same split gives $m = 14$ and $\delta^2 = 3^{-28} \approx 4.371 \times 10^{-14}$.

Project freeze: global split and deterministic steps. For the *project audits* we fix $\varepsilon = 10^{-6}$ with the audit split $\varepsilon_{\text{AGSP}} = \varepsilon_{\text{belt}} = \varepsilon/4$ and $\varepsilon_{\text{disp}} = \varepsilon_{\text{cel}} = \varepsilon_{\text{flow}} = \varepsilon/6$. With $\eta = \frac{1}{3}$ this yields

$$m = \left\lceil \frac{\ln(4/\varepsilon)}{\ln 3} \right\rceil = 14, \quad \delta^2 = 3^{-28} = 4.3712421747e - 14,$$

and the Trotter-safe branch $\Upsilon(m) \geq 2m+1 = 29$. At this precision $(1 - \delta^2)^{-1} = (1 - \delta^2)^{-2} = 1 + o(10^{-14})$.

5.29 Illustrative Page threshold with explicit numbers

Threshold (project ledger freeze). Using Section 5.22 at $m = 14$ (so $\Upsilon(m) \geq 29$), the Page line-density estimate is

$$\frac{a_{\text{QES}}}{4G} \approx \underbrace{0.5}_{\log \kappa_{\text{seed}}} + \underbrace{2.0149030205}_{\log(\Lambda_0 \Gamma_{\text{belt}})} + \underbrace{3.3672958300}_{\log \Upsilon(14)} = \boxed{5.8821988505}.$$

Concretely, with the frozen values $\kappa_{\text{seed}} = 1.6487212707$, $\Lambda_0 = 2.5$, and $\Gamma_{\text{belt}} = 3.0$,

$$\log \kappa_{\text{seed}} = 0.5, \quad \log(\Lambda_0 \Gamma_{\text{belt}}) = 2.0149030205, \quad \log \Upsilon(14) = 3.3672958300,$$

summing to 5.8821988505.

Deterministic steps and denominators. With $\eta = \frac{1}{3}$ and the project split of Section 5.28,

$$m = 14, \quad \delta^2 = 3^{-28} = 4.3712421747e - 14,$$

so $(1 - \delta^2)^{-1} = (1 - \delta^2)^{-2} = 1 + o(10^{-14})$ at displayed precision.

Pointer to dispersion audit. We adopt Freeze A of Section 5.93 at $\varepsilon = 10^{-6}$: $M = S_{\text{cut}}/s_0 = 100$ and $J = 2200$, which certify $|\Delta_{\text{tail}} \hat{c}_{2,0}| \leq 5.811516831 \times 10^{-8}$ and $|\Delta_{\text{quad}} \hat{c}_{2,0}| \leq 7.86 \times 10^{-8}$, so the dispersion share is strictly within $\varepsilon_{\text{disp}} = \varepsilon/6$.

5.30 Anchor choice and relational dressing: invariance statement

Proposition 5.46 (anchor invariance). Let \mathcal{C} and \mathcal{C}' be two admissible anchors defining relational/dressed locality as in Section 2. Then for all observables \mathcal{O} introduced in this work (entropy, modular energy, area, amplitude functionals) there exists $C_{\mathcal{C} \rightarrow \mathcal{C}'} > 0$ such that

$$|\mathcal{O}[\mathcal{C}] - \mathcal{O}[\mathcal{C}']| \leq C_{\mathcal{C} \rightarrow \mathcal{C}'} \mathcal{B}_{\text{belt}}.$$

Kernel. Functoriality under changes of localization data follows from the belt dictionary and quasi-local stress calibration Section 5.54 and Proposition 5.73, together with belt recovery/-continuity Proposition 5.80 and OS short evolutions/removal Lemmas 3.1 and 3.3. Any anchor change factors into a bounded-depth belt circuit plus an exponentially suppressed bulk tail via microcausality/timeslice propagation Lemma 5.70 and Proposition 5.71.

5.31 Existence of quantum extremal surfaces (direct method)

Theorem 5.47 (QES existence). Assume the framework recap Section 2 (items 1–3,5), belt microcausality/timeslice control Lemma 5.70 and Proposition 5.71, belt-level nesting/recovery and continuity Propositions 5.80 and 5.87 and Section 5.15, and the OS kernel Lemmas 3.1 and 3.3. On any belt-anchored null cut for which the belt budget $\mathcal{B}_{\text{belt}}$ is finite, the generalized entropy

$$\mathcal{G}[\Sigma] := \frac{\text{Area}(\Sigma)}{4G} - S(\rho_R; \Sigma)$$

admits a minimizer within the admissible class of codimension-two surfaces Σ homologous to R . Moreover, minimizers can be chosen to have uniformly bounded area and to lie within a compact geometric class determined by the belt width and curvature bounds.

Proof kernel. (1) *Coercivity:* the belt AGSP/seed pipeline and recovery Proposition 5.80 yield a per-length entropy bound; the area term dominates for large excursions, hence \mathcal{G} is coercive on the admissible class.

(2) *Lower semicontinuity:* Area is l.s.c. under geometric convergence; $S(\rho_R; \Sigma)$ is l.s.c. by belt-level sufficiency/recovery continuity Proposition 5.80 together with the first-law channel Corollary 3.5.

(3) *Compactness:* microcausality tails and null timeslice propagation Lemma 5.70 and Proposition 5.71 and the quasi-local belt factorization Section 5.54 and Proposition 5.73 give uniform control preventing “escape” from the belt. Positive flows provide a mollified compact class, and flow removal Lemma 3.3 closes the regulator. The direct method then yields a minimizer.

Uniqueness/stability follow from Section 5.15. \square

5.32 Conservation and Bianchi consistency from modular dynamics

Proposition 5.48 (modular conservation). Under the framework recap Section 2 (items 1,2,5) and the modular equation of state Theorem 5.35, the expectation-value Einstein equations imply

$$\nabla^\mu \delta \langle T_{\mu\nu} \rangle = 0 \quad \text{on } D[R],$$

and are consistent with the contracted Bianchi identity $\nabla^\mu \delta \langle \text{Ein}_{\mu\nu} \rangle = 0$.

Kernel. Differentiate the modular equation of state along the belt boost flow; the Killing property of ξ removes boundary contributions up to $O(\mathcal{B}_{\text{belt}})$. Corner/edge terms cancel as in Section 5.17 and Proposition 5.38. The result is the covariant conservation of $\delta \langle T_{\mu\nu} \rangle$; Bianchi consistency follows from the geometric identity together with the JLMS matching Proposition 3.4.

5.33 Counterterm stability and renormalization scheme independence

Lemma 5.49 (scheme independence to leading order). Let two belt-compatible renormalization schemes differ by local counterterms $\delta\mathcal{L} = \alpha_1 \int R + \alpha_2 \int K^2 + \dots$ supported on the belt and corners. Then the variation of the generalized entropy satisfies

$$\delta \left[S - \frac{\text{Area}}{4G} \right]_{\text{scheme A}} - \delta \left[S - \frac{\text{Area}}{4G} \right]_{\text{scheme B}} = O(\mathcal{B}_{\text{belt}}),$$

with the $O(\mathcal{B}_{\text{belt}})$ controlled by the same belt budget used in Section 5.14. **Kernel.** Area and Wald-edge counterterms shift in lockstep with the boost Ward identity, while the entropy counterterm shift is absorbed by the recovery channel; belt LR suppresses any mismatch exponentially in r .

5.34 Γ -convergence from discrete surrogates to continuum belts

Theorem 5.50 (Gamma-bridge). Let $\{\mathcal{G}_h\}$ be discrete generalized-entropy functionals built from CDT/GFT belt nets on the fixed tester envelope of Sections 5.19, 5.47 and 5.52. Then \mathcal{G}_h Γ -converges to the continuum \mathcal{G} on the admissible class, and any sequence of (approximate) discrete minimizers has accumulation points that are QES minimizers of \mathcal{G} . **Kernel.** *Liminf:* nonnegativity of the forward even-parity dispersion projectors, Gaussian Hankel/impact functionals, and principal-series celestial Gram testers (Sections 5.10, 5.23 and 5.82) pushes lower bounds to the limit. *Recovery:* bounded-depth belt circuits and quasi-local factorization furnish recovery sequences realizing the limsup (Section 5.54, Proposition 5.73, and Lemmas 3.1 and 3.3). Uniform belt budgets and the invariances recorded in Sections 5.13, 5.14, 5.16, 5.50 and 5.70 guarantee tightness and precompactness.

Verification on the envelope (CDT/GFT family). For the discrete generalized entropy $G_h := \text{Area}_h/(4G) - S_h$ built from the belt nets in $\mathcal{F}_{\text{CDT/GFT}}$, the runless acceptance above grants the liminf inequality by positivity of the testers, and the belt circuits furnish recovery sequences, so G_h Γ -converges to the continuum G (Theorem 5.50). Discrete minimizers $\Sigma_h \in \arg \min G_h$ are therefore precompact and any accumulation point is a continuum QES minimizer. The single plot in Sec. 6.4 shows both: the discrete minimizer locations converge to the continuum QES at a rate consistent with $O(h)$, and the minimal tester slack remains strictly positive along the refinement ladder.

5.35 Compact dual certificates for amplitude positivity

Proposition 5.51 (finite-support extremals). Consider any conic combination of the project's nonnegative testers on the cone \mathcal{S} : (i) forward even-parity functionals at Q fixed t -nodes, (ii) P Hankel/Gaussian band kernels, (iii) K celestial Gram functionals on principal lines (worst- K anchors). Then there exists an *extremal* dual certificate supported on at most $(Q+P+K)$ nodes across the relevant integration variables, with all weights nonnegative.

Quantitative corollary (project defaults). Using $Q = 6$ (Chebyshev t -grid), $P = 7$ (Gaussian band), and $K = 5$ (worst-five celestial anchors), there is an extremal certificate with support ≤ 18 nodes.

Kernel. The feasible dual cone is the conic hull of evaluation functionals at grid nodes; by Carathéodory-type arguments in finite dimensions, any point in the cone admits a representation with at most the number of active generators. Nonnegativity is preserved by construction; extremality follows by minimal support.

Add-on (dual certificates with strip anchors). The explicit compact certificates admit two strip variants: (i) *swap* the five principal celestial nodes for the five strip nodes of Table 4

(support remains 18); (ii) *augment* by these five to obtain a 23-support certificate. All weights remain nonnegative by Lemma 5.112.

5.36 Normalization consistency between amplitude and modular dynamics

Lemma 5.52 (κ -consistency). With the gravitational coupling normalized so that $M_{\text{Pl}} = \sqrt{32\pi}/\kappa$, the soft tree-level forward amplitude $\mathcal{A}_{\text{soft}}(s, t) = \kappa^2 \left(\frac{s^2}{-t} + \frac{u^2}{-t} \right)$ is consistent with the modular equation-of-state normalization

$$\delta \langle K_{\text{mod}} \rangle = 2\pi \int d\Sigma^\mu \xi^\nu \delta \langle T_{\mu\nu} \rangle + \delta \left[\frac{\text{Area}}{4G} \right] + O(B_{\text{belt}}),$$

provided $\kappa^2 = 32\pi G$. In particular, the linearized area response and the stress-tensor coupling match the low-energy limit of graviton exchange.

Proposition 5.53 (Amplitude \rightarrow curvature dictionary line (forward coefficient $\rightarrow R_{kk}$)). Let the analytic forward coefficient be defined by the projector of Section 5.23:

$$a_2^{(\text{even})}(t) := \Re \Pi_2[A^{(N=3)}](t), \quad bc_{2,0}(t) := s_0^3 c_{2,0}(t) = \frac{s_0^3}{2} \partial_s^2 \Re A^{(3)}(0, t),$$

which is invariant under IR scheme, pivot, and scale by Lemmas 5.37, 5.69 and 5.89. With the gravitational normalization fixed by $\kappa^2 = 32\pi G$ (Lemma 5.52), there exists a belt-local, state-independent positive functional $\mathfrak{D}_{\text{belt}}$ on null-Ricci profiles—determined solely by the analytic projector and the Brown–York/JKM calibration—such that for any admissible state and any belt-anchored null segment $[\lambda_1, \lambda_2]$,

$$\mathfrak{D}_{\text{belt}}[R_{kk}] = C_{\text{II}}(s_0) bc_{2,0}(t) + O(B_{\text{belt}}), \quad C_{\text{II}}(s_0) > 0,$$

and equivalently, for the normalized nonnegative weight $w(\lambda)$ attached to $\mathfrak{D}_{\text{belt}}$,

$$\frac{1}{8\pi G} \int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle = C_{\text{II}}(s_0) bc_{2,0}(t) + O(B_{\text{belt}}).$$

Here $R_{kk} := R_{\mu\nu} k^\mu k^\nu$ along the belt generator k^μ . The constant $C_{\text{II}}(s_0)$ and weight w depend only on the fixed projector kernel and the belt Ward/JKM calibration; they are independent of $|R|$.

Sketch. By dispersion differentiation (Section 5.74), $\partial_s^2 \Re A^{(3)}(0, t)$ is a positive kernel integral of the absorptive part; the helicity-averaged profile is nonnegative (Section 5.84). Folding the kernel through Π_2 yields the factor $C_{\text{II}}(s_0)$. The κ -consistency $\kappa^2 = 32\pi G$ (Lemma 5.52) matches amplitude and modular scales used in Section 5.11. The Brown–York dictionary (Section 5.54) identifies the canonical-energy (modular) flux with a belt quasi-local stress flux up to $O(B_{\text{belt}})$. Using the modular equation of state and its IFF form (Sections 5.11 and 5.60) gives $R_{kk} = 8\pi G \langle T_{kk} \rangle$ in expectation on $D[R]$ up to the same budget, producing the displayed identity. All invariances (IR scheme, pivot, scale) follow from Lemmas 5.37, 5.69 and 5.89; strip robustness follows from Section 5.23 and Corollary 5.119. \square

5.37 Master soundness and epsilon closure

Theorem 5.54 (Soundness budget). Fix a target accuracy $\varepsilon \in (0, 1)$. With the global split of Section 5.28 and the belt/AGSP choices there, the four pillars satisfy

$$\begin{aligned} |\delta S - \delta \langle K_{\text{mod}} \rangle + \delta \frac{\text{Area}}{4G}| &\leq \varepsilon, \\ \left| \int du \langle T_{kk} \rangle \right| + |\partial_u^2 S - 2\pi \langle T_{kk} \rangle| &\leq \varepsilon, \\ \text{Amp}_{\text{pos}}[\mathcal{A}_{\text{hard}}] &\geq -\varepsilon, \\ |\text{Ein}[\delta g] - 8\pi G \delta \langle T \rangle|_{D[R]} &\leq \varepsilon, \end{aligned}$$

with ε assembled from $(\varepsilon_{\text{AGSP}}, \varepsilon_{\text{belt}}, \varepsilon_{\text{disp}}, \varepsilon_{\text{cel}}, \varepsilon_{\text{flow}})$ as specified.

5.38 Generalized second law on belts

Theorem 5.55 (Belt GSL). Assume the framework recap Section 2 (items 1–3,5) and the c -function monotonicity of Section 5.18. Along any belt-anchored null generator with affine parameter λ ,

$$\frac{d}{d\lambda} \left(S(\rho_R) - \frac{\text{Area}(\text{QES}(R))}{4G} \right) \geq -C_{\text{GSL}} \mathcal{B}_{\text{belt}},$$

with C_{GSL} independent of $|R|$. Under flow removal, $(u, s) \downarrow 0$, the right-hand side vanishes and the GSL holds in the limit.

Proof. Let $R(\lambda)$ denote the belt-anchored cut obtained by translating a fixed cut $R(0)$ along a complete null generator with affine parameter λ (per generator length); write

$$S(\lambda) := S(\rho_{R(\lambda)}), \quad A(\lambda) := \text{Area}(\text{QES}(R(\lambda))).$$

We abbreviate $\mathcal{B}_{\text{belt}}$ for the belt budget and work in the positive-flow window; all $O(\mathcal{B}_{\text{belt}})$ bounds are uniform in $|R|$ and in the belt base factor Γ_{belt} .

Step 1: Balance law along the generator. Let $\mathfrak{c}(\lambda)$ be the belt c -function of Section 5.18. By construction (OS/KMS quadratic form for relative-entropy production along the generator) and by the calibrated modular/flux dictionary of Section 6.2 together with Theorem 5.23,

$$\frac{d}{d\lambda} \left(\frac{A(\lambda)}{4G} - S(\lambda) \right) = \mathfrak{c}(\lambda) + \mathfrak{e}_{\text{edge}}(\lambda) + \mathfrak{e}_{\text{corner}}(\lambda) + O(\mathcal{B}_{\text{belt}}). \quad (5.25)$$

Here $\mathfrak{e}_{\text{edge}}$ and $\mathfrak{e}_{\text{corner}}$ are the edge/corner contributions recorded in Section 5.17, generated by the finite belt width and the boosted corner. By Proposition 5.38 (edge/corner control),

$$|\mathfrak{e}_{\text{edge}}(\lambda) + \mathfrak{e}_{\text{corner}}(\lambda)| \leq C_{\text{edge}} \mathcal{B}_{\text{belt}} \quad \text{uniformly in } \lambda. \quad (5.26)$$

Step 2: Monotone evolution of the c -function. The c -function obeys the differential inequality of Section 5.18,

$$\frac{d}{d\lambda} \mathfrak{c}(\lambda) \leq -\lambda_{\text{Clu}} \mathfrak{c}(\lambda) + C_0 \mathcal{B}_{\text{belt}}, \quad (5.27)$$

where $\lambda_{\text{Clu}} > 0$ is the null clustering rate (from the admissible class), and C_0 is belt-uniform. Moreover, null timeslice propagation and cluster imply the past asymptotics

$$\lim_{\lambda \rightarrow -\infty} \mathfrak{c}(\lambda) = 0. \quad (5.28)$$

Step 3: Grönwall bound for \mathfrak{c} . Apply Grönwall's lemma to (5.27) with the initial condition (5.28). For any $\lambda \in \mathbb{R}$,

$$\mathfrak{c}(\lambda) \leq \int_{-\infty}^{\lambda} e^{-\lambda_{\text{Clu}}(\lambda-s)} C_0 \mathcal{B}_{\text{belt}} ds = \frac{C_0}{\lambda_{\text{Clu}}} \mathcal{B}_{\text{belt}} =: C_c \mathcal{B}_{\text{belt}}. \quad (5.29)$$

Step 4: Differential GSL (with belt remainder). Rearrange (5.25) and use (5.26) and (5.29):

$$\frac{d}{d\lambda} \left(S(\lambda) - \frac{A(\lambda)}{4G} \right) = -\mathfrak{c}(\lambda) - (\mathfrak{e}_{\text{edge}}(\lambda) + \mathfrak{e}_{\text{corner}}(\lambda)) + O(\mathcal{B}_{\text{belt}}) \geq -(C_c + C_{\text{edge}} + C_1) \mathcal{B}_{\text{belt}},$$

for some belt-uniform constant C_1 coming from the $O(\mathcal{B}_{\text{belt}})$ term in (5.25). Setting

$$C_{\text{GSL}} := C_c + C_{\text{edge}} + C_1$$

gives the stated estimate

$$\frac{d}{d\lambda} \left(S(\rho_R) - \frac{A(\text{QES}(R))}{4G} \right) \geq -C_{\text{GSL}} \mathcal{B}_{\text{belt}}, \quad (5.30)$$

uniformly in $|R|$ and along the complete generator.

Step 5: Flow removal and the sharp GSL. All constants above are independent of the positive-flow regulators. Under flow removal $(u, s) \downarrow 0$, the belt budget vanishes, $\mathcal{B}_{\text{belt}} \rightarrow 0$, hence the right-hand side of (5.30) tends to 0. Therefore the differential GSL sharpens to

$$\frac{d}{d\lambda} \left(S(\rho_R) - \frac{A(\text{QES}(R))}{4G} \right) \geq 0$$

in the limit window. This completes the proof. \square

5.39 Quantified QFC on belts

Definition 5.56 (belt quantum expansion). Along a belt-anchored null generator with affine parameter λ , define

$$\Theta_{\text{belt}}(\lambda) := \partial_\lambda \left(S(\rho_R(\lambda)) - \frac{\text{Area}(\text{QES}(\lambda))}{4G} \right) \quad (\text{per generator length}).$$

Proposition 5.57 (quantified belt QFC). For any segment $[\lambda_1, \lambda_2]$,

$$\Theta_{\text{belt}}(\lambda_2) - \Theta_{\text{belt}}(\lambda_1) \geq 2\pi E_{\text{can}}^{\text{dens}}[\lambda_1 \rightarrow \lambda_2; \xi] + Q_{\text{shear}}[\lambda_1 \rightarrow \lambda_2] - C_{\text{QFC}} B_{\text{belt}}, \quad (5.31)$$

where $E_{\text{can}}^{\text{dens}}$ is the canonical-energy density on the wedge generated by the belt boost ξ (per generator length), and

$$Q_{\text{shear}}[\lambda_1 \rightarrow \lambda_2] := \kappa_\sigma \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \sigma_{ab} \sigma^{ab} + \kappa_\theta \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \theta^2$$

with strictly positive belt-local coefficients $\kappa_\sigma, \kappa_\theta$ from Theorem 5.43 of Section 5.25. The constant $C_{\text{QFC}} > 0$ is belt-uniform and independent of $|R|$. In particular, for non-Killing perturbations,

$$\Theta_{\text{belt}}(\lambda_2) \geq \Theta_{\text{belt}}(\lambda_1) - C_{\text{QFC}} B_{\text{belt}},$$

and under removal of positive flows $(u, s) \downarrow 0$, $\Theta_{\text{belt}}(\lambda)$ is nondecreasing: $\partial_\lambda \Theta_{\text{belt}}(\lambda) \geq 0$.

Sketch. The second-order modular equation of state (Theorem 5.43) yields the lower bound $\delta^2[S - \text{Area}/(4G)] \geq 2\pi E_{\text{can}} + Q_{\text{shear}} - C_2 B_{\text{belt}}$ along any belt deformation; integrating in λ on $[\lambda_1, \lambda_2]$ produces (5.31) with $C_{\text{QFC}} \sim C_2$. The linearized Raychaudhuri window (Proposition 5.96) and the shear control (Lemma 5.106) ensure the positivity of Q_{shear} , while canonical-energy/QNEC inputs (Theorem 5.27) guarantee $E_{\text{can}}^{\text{dens}} \geq 0$ for non-Killing data. The $O(B_{\text{belt}})$ remainder is removed by the positive-flow limit (Section 5.14). \square

Remark 5.58 (QG deliverable: belt quantum focusing and local arrow). *Sign convention.* We work with $S - \text{Area}/(4G)$, so a monotone *increase* of the belt quantum expansion Θ_{belt} is equivalent to the usual statement that the quantum expansion of $S_{\text{gen}} = S + \text{Area}/(4G)$ is *nonincreasing*.

Deliverable. Along admissible belt generators the quantum expansion is monotone in the sense of Proposition 5.57, with equality iff the perturbation is boost-Killing. Together with the belt generalized second law and the quantum Bousso bound on belts (Theorems 5.55 and 5.91), this packages a local QG arrow of time on belts. The ingredients trace back to the ANEC/QNEC synthesis and the modular equation of state (Theorems 5.27 and 5.35) and remain within the declared envelope.

5.40 Third-order control of the modular equation of state

Proposition 5.59 (cubic completion). Under the hypotheses of Theorem 5.43 and the canonical-energy/QSEI inputs Proposition 5.96 and Lemma 5.106, the third variation satisfies

$$\delta^3 \left[S - \frac{\text{Area}}{4G} \right] = 2\pi \mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi] + \mathcal{C}_{\text{shear} \times \text{exp}}[\delta g] + R_{\text{belt}}^{(3)}, \quad (5.32)$$

where $\mathcal{E}_{\text{can}}^{(3)}$ is the cubic canonical-energy functional, $|\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]| \leq \tilde{C}'_{\sigma} E_{\text{can}}^W[\delta\Psi; \xi] + C_{\text{rem}}^{(3)} \mathcal{B}_{\text{belt}}$ by Lemma 5.106, and $|R_{\text{belt}}^{(3)}| \leq C_{\text{belt}}^{(3)} \mathcal{B}_{\text{belt}}$. If, in addition, $\mathcal{E}_{\text{can}}^{(3)} \geq 0$ on non-Killing perturbations over the domain of Theorem 5.67, then the cubic inequality

$$\delta^3 \left[S - \frac{\text{Area}}{4G} \right] \geq -\tilde{C}'_{\sigma} E_{\text{can}}^W[\delta\Psi; \xi] - C_{\text{belt}}^{(3)} \mathcal{B}_{\text{belt}} \quad (\text{non-Killing data}) \quad (5.33)$$

holds belt-uniformly.

Kernel. Differentiate the second-order statement Theorem 5.43 with calibrated corners (Sections 5.17 and 5.53, Proposition 5.38, and Lemma 5.72); use the trilinear symplectic current and the cubic-form domain. Control the shear–expansion coupling by Lemma 5.106; OS flow removal Lemma 3.3 makes the remainder explicit. Interacting checks in Sections 5.94 and 5.95 (ϕ^3 /Yukawa) numerically confirm scaling and positivity within the same budgets.

5.41 Finite-size corrections and thermodynamic limit

Lemma 5.60 (finite-size control). For any belt-anchored region R ,

$$\left| \frac{S(\rho_R)}{\text{length}(\partial R)} - \left[\frac{\log \kappa_{\text{seed}}}{1 - \delta^2} + \frac{\log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{(1 - \delta^2)^2} \right] \right| \leq \frac{C}{\text{length}(\partial R)} + O(\mathcal{B}_{\text{belt}}),$$

so the per-length entropy converges to its belt/AGSP benchmark as $\text{length}(\partial R) \rightarrow \infty$. **Kernel.** Direct from the converter and Section 5.14.

5.42 Integrated QNEC implies ANEC and partial converses

Theorem 5.61 (Integrated QNEC implies ANEC; failure and partial converses). Let $f \in C_0^{\infty}(\mathbb{R})$ be nonnegative with $\int_{\mathbb{R}} f = 1$, and suppose that for all such f ,

$$\int_{\mathbb{R}} du f(u) \left(2\pi \langle T_{kk}(u) \rangle - \partial_u^2 S(u) \right) \geq 0. \quad (5.34)$$

Then the distribution

$$g(u) := 2\pi \langle T_{kk}(u) \rangle - \partial_u^2 S(u)$$

is nonnegative, i.e.

$$\int_{\mathbb{R}} du \psi(u) g(u) \geq 0 \quad \text{for all } \psi \in C_0^{\infty}(\mathbb{R}), \psi \geq 0.$$

In particular, if g is locally integrable (so $g \geq 0$ almost everywhere) and $\partial_u S(u)$ has finite limits as $u \rightarrow \pm\infty$ (e.g. $\partial_u S \rightarrow 0$), then ANEC holds:

$$\int_{\mathbb{R}} du \langle T_{kk}(u) \rangle \geq 0.$$

Conversely, ANEC alone does *not* imply (5.34). Nevertheless, (5.34) *does* follow under either of the following stronger hypotheses:

1. **Pointwise QNEC:** $2\pi \langle T_{kk}(u) \rangle \geq \partial_u^2 S(u)$ almost everywhere.
2. **Convexity plus smeared energy positivity:** $\partial_u^2 S(u) \geq 0$ (e.g. a belt-QFC statement in the flow-removed limit) and $\int_{\mathbb{R}} du \psi(u) \langle T_{kk}(u) \rangle \geq 0$ for all $\psi \in C_0^\infty$ with $\psi \geq 0$.

Proof sketch. Kernel. For the forward direction, fix any $\psi \geq 0$ in $C_0^\infty(\mathbb{R})$ and set $f := \psi / \int \psi$. By (5.34),

$$0 \leq \int f g = \frac{1}{\int \psi} \int \psi g,$$

so $\int \psi g \geq 0$ for all nonnegative ψ , i.e. g is a positive distribution (hence $g \geq 0$ a.e. if it is a function). Integrating g over \mathbb{R} and using $\int_{\mathbb{R}} \partial_u^2 S = [\partial_u S]_{-\infty}^{+\infty}$, the stated behavior of $\partial_u S$ gives

$$2\pi \int_{\mathbb{R}} du \langle T_{kk}(u) \rangle = \int_{\mathbb{R}} du g(u) \geq 0,$$

which is ANEC.

For the converse, ANEC constrains only the unweighted integral of $\langle T_{kk} \rangle$ and does not restrict $\partial_u^2 S$. Thus one can arrange a smooth g that is negative on a compact set yet has $\int \langle T_{kk} \rangle \geq 0$, violating (5.34) for an f supported where $g < 0$. The two bullets listed suffice to *restore* (5.34): (1) is immediate by integrating the pointwise inequality against any $f \geq 0$, while (2) makes both sides of (5.34) nonnegative for every such f . \square

Remark 5.62 (Toy counterexample to the converse). Choose $h \in C_0^\infty(\mathbb{R})$ with $\int h = 0$ and $h < 0$ on some interval. Pick $\langle T_{kk} \rangle \in C_0^\infty$ with $\int \langle T_{kk} \rangle > 0$ and define S by $\partial_u^2 S = 2\pi \langle T_{kk} \rangle - h$, with $\partial_u S \rightarrow 0$ at $\pm\infty$. Then ANEC holds, but $g = h$, so for any f supported where $h < 0$ one has $\int f g < 0$, contradicting (5.34).

Corollary 5.63 (From pointwise QNEC to ANEC). If $2\pi \langle T_{kk} \rangle \geq \partial_u^2 S$ almost everywhere and $\partial_u S \rightarrow 0$ as $u \rightarrow \pm\infty$, then ANEC follows by integration.

5.43 Domains and cores for higher-order variations

Proposition 5.64 (analytic core and essential self-adjointness). Let \mathcal{D}_{an} be the common domain of analytic vectors for the belt modular generator. Then $K_{\text{mod}} \upharpoonright_{\mathcal{D}_{\text{an}}}$ is essentially self-adjoint, and \mathcal{D}_{an} is a core for all polynomial functions of K_{mod} entering the second/third variations.

Kernel. The existence of an analytic core at belt level follows from the OS positivity/KMS structure and the first-law channel Lemma 3.1 and Corollary 3.5, which provide a dense set of vectors with entire modular orbits on the belt. Essential self-adjointness of K_{mod} on \mathcal{D}_{an} is then standard (Nelson-type) and stable under bounded belt circuits and quasi-local factorization Section 5.54 and Proposition 5.73; stability of the variation domain under regulator removal is ensured by Lemma 3.3 and the variation–limit interchange Lemma 5.86.

5.44 Celestial Ward identities without anomalies on the principal series

Remark 5.65 (Principal series scope). The celestial Gram positivity and the no-anomaly statement are used on the principal series $\Delta = 1 + i\nu$. Off-principal sectors require additional counterterms; see Section 5.97 for a finite strip with renormalized measure and Ward counterterms.

Lemma 5.66 (no anomaly on the principal lines). For celestial insertions restricted to the principal series $\Delta = 1 + i\nu$ with $|\nu| \leq \mu_{\text{cel}}$, the Virasoro-type Ward action induced by the belt first-law channel is anomaly-free at linear order, and any residual measure deformation is $O(\mathcal{B}_{\text{belt}})$. **Kernel.** Combine the celestial Gram positivity with the belt Ward map; the principal-series measure is preserved by the induced diffeomorphisms, and regulator effects are suppressed by the belt budget.

Add-on (scope extension to an off-principal strip). The principal-series statement extends to a finite off-principal strip (Section 5.97) after a renormalization of the celestial measure and a belt-local counterterm tabulated in Section Appendix C. The resulting Ward action preserves the renormalized strip Gram form and remains anomaly-free up to $O(\mathcal{B}_{\text{belt}})$.

5.45 Epsilon schedule (precomputed)

For target accuracies $\varepsilon \in \{10^{-6}, 10^{-8}\}$ and $\eta = \frac{1}{3}$, the deterministic AGSP steps and small parameter are

$$m = \left\lceil \frac{\ln(4/\varepsilon)}{\ln 3} \right\rceil, \quad \delta^2 = \eta^{2m} = 3^{-2m}.$$

Target ε	m	$\delta^2 = 3^{-2m}$
10^{-6}	14	$4.3712421747e - 14$
10^{-8}	19	$7.4027370060 \times 10^{-19}$

Table 1: Precomputed (m, δ^2) at the two audit targets (ledger-frozen).

At these precisions, the converter denominators obey $(1 - \delta^2)^{-1} = (1 - \delta^2)^{-2} = 1 + o(10^{-14}) / 1 + o(10^{-18})$ on the two lines, so they are numerically 1 at displayed precision.

5.46 Illustrative Page threshold (baseline numbers)

With the illustrative baseline $\log \kappa_{\text{seed}} = 0.5$, $\Lambda_0 = 2.5$, $\Gamma_{\text{belt}} = 3.0$ and $m = 8$ (so $\Upsilon(m) \geq 17$ and $(1 - \delta^2)^{-1} = 1 + O(10^{-8})$), the line-density crossover estimate is

$$\frac{a_{\text{QES}}}{4G} \approx \underbrace{0.5000000000}_{\log \kappa_{\text{seed}}} + \underbrace{2.0149030205}_{\log(\Lambda_0 \Gamma_{\text{belt}})} + \underbrace{2.8332133441}_{\log 17} = \boxed{5.3481163646}.$$

Project-specific ledger values can be inserted to update the number immediately.

5.47 Explicit dual certificate for positivity (numbers, no solver)

Goal. Provide a compact certificate (support = 18) which enforces nonnegativity on the working cone, now uniformly for $t \in [-0.30 s_0, 0]$, and preserves all analytic–projector invariances (IR–scheme, pivot, and scale) frozen elsewhere.

Structure (unchanged). The dual is the normalized average of three nonnegative tester families:

- *Forward fixed- t evaluators* on a Chebyshev grid in t (6 nodes, see below).
- *Hankel/impact testers* (7 positive K_0 -band kernels in b ; same scales as in Section 5.10).
- *Principal-series celestial Gram testers* (the frozen worst-five anchors of Table 2).

Weights are uniform (1/18 per support point). Because each component tester is nonnegative and the three families enjoy the invariances of Section 5.23, Section 5.50, and Section 5.70, their average inherits the same invariances.

Forward grid (widened window; 6 Chebyshev–Gauss nodes). Let $Q=6$ and $x_q = \cos(\frac{2q-1}{2Q}\pi)$. Map $[-1, 1] \rightarrow [-0.30 s_0, 0]$ via $t = s_0 \cdot (-0.15)(1 - x)$. The resulting ordered nodes (most negative to near-forward) are

$$\frac{t}{s_0} \in \left\{ -0.2948888739, -0.2560660172, -0.1888228568, \right. \\ \left. -0.1111771432, -0.0439339828, -0.0051111261 \right\}.$$

Impact/Hankel and celestial parts. Use the same seven impact scales and the same five principal-series celestial anchors as frozen in Section 5.82. No change is required: their positivity and Gram stability are window-agnostic.

Certificate. The explicit dual is the uniform average of the 18 testers above. It enforces tester nonnegativity for all t in the widened window and is invariant under the analytic projector, pivot, and rescalings.

5.48 Positivity domain for cubic canonical energy

Theorem 5.67 (cubic canonical-energy domain). Let \mathcal{D}_{an} be the analytic core of the belt modular generator (Section 5.43). Define

$$\mathcal{D}^{(3)} := \overline{\text{span}} \left\{ \psi \in \mathcal{D}_{\text{an}} : \psi \text{ is supported on a single belt patch and } \|(1 + K_{\text{mod}})^{3/2} \psi\| < \infty \right\}.$$

Then: (i) $\mathcal{D}^{(3)}$ is a core for the cubic canonical-energy form $\mathcal{E}_{\text{can}}^{(3)}$; (ii) $\mathcal{E}_{\text{can}}^{(3)}[\psi] \geq 0$ for all $\psi \in \mathcal{D}^{(3)}$ orthogonal to boost–Killing modes; (iii) $\mathcal{D}^{(3)}$ is stable under constant-depth belt circuits and under AGSP maps K_m with $\|K_m P^\perp\| \leq \eta^m$.

Kernel. Essential self-adjointness and the analytic-core construction (Section 5.43) imply closability of polynomial forms in K_{mod} . BW/KMS positivity and the second-order modular inequality provide the required positivity inputs (Lemma 3.1 and Theorem 5.43). Stability under belt circuits follows from the quasi-local belt dictionary (Section 5.54 and Proposition 5.73); AGSP stability is by construction of the maps and the contraction bound $\|K_m P^\perp\| \leq \eta^m$.

5.49 Modular Bekenstein inequality and numeric check

Proposition 5.68 (modular Bekenstein on belts). For admissible perturbations supported on a belt, relative entropy positivity implies

$$\Delta S(R) \leq \Delta \langle K_{\text{mod}}(R) \rangle \quad (\text{belt version}).$$

Kernel. $\Delta S = \Delta \langle K \rangle - S(\rho \| \sigma)$ and $S(\rho \| \sigma) \geq 0$; stability follows from flow removal.

5.50 Dispersion pivot invariance for gravity-subtracted amplitudes

Lemma 5.69 (pivot invariance). Let $\mathcal{A}^{(N)}(s, t)$ denote the $N=3$ gravity-subtracted amplitude on the cone. If we shift the subtraction pivot to $s = s_*$, the subtracted real part transforms by a polynomial

$$\Re \mathcal{A}^{(N)[s_*]}(s, t) = \Re \mathcal{A}^{(N)[0]}(s, t) + Q_2(s, t; s_*).$$

where Q_2 is at most quadratic in s , with coefficients analytic in t on the cone. In particular, the even-parity forward derivative inequalities and the Hankel/celestial positivity functionals are unchanged.

Kernel. Cauchy subtraction algebra and crossing symmetry shift P_{N-1} by a pivot-dependent polynomial; positivity testers annihilate or preserve such analytic reparametrizations.

Add-on (pivot invariance unchanged). Pivot shifts act by s -polynomials; the strip renormalization is t -holomorphic at $s=0$ and does not affect Lemma 5.69.

5.51 Microcausality tails and null timeslice extension

Lemma 5.70 (belt commutator tail). For operators $A \in \mathcal{A}(\partial_r R)$ and $B \in \mathcal{A}(\partial_r \bar{R})$ separated across the belt,

$$\|[A, B]\| \leq C_{\text{mc}} e^{-\mu_{\text{eff}} r}.$$

Kernel. Belt microcausality follows from the positivity/KMS structure and the quasi-local belt factorization Lemma 3.1, Section 5.54, and Proposition 5.73. The constants are uniform in $|R|$.

Proposition 5.71 (null timeslice propagation). If the equation-of-state identity and energy inequalities hold on a belt cut, they extend to the whole domain of dependence $D[R]$.

Kernel. Use the null timeslice property at belt level together with relative Cauchy evolution along the belt boost flow; corner/edge terms remain $O(\mathcal{B}_{\text{belt}})$ and cancel by the calibration in Sections 5.17 and 5.53, Proposition 5.38, and Lemma 5.72. Flow removal Lemma 3.3 makes the extension regulator-free.

5.52 Explicit 18-support dual certificate (final nodes and weights)

Support. 18 points = 6 forward fixed- t evaluators + 7 Hankel/impact testers + 5 celestial anchors.

Forward nodes (widened window).

$$\frac{t}{s_0} = \begin{pmatrix} -0.2948888739, & -0.2560660172, & -0.1888228568, \\ -0.1111771432, & -0.0439339828, & -0.0051111261 \end{pmatrix}.$$

Hankel/impact scales (units of s_0). The seven Hankel/impact testers use the frozen scales of Section 5.35:

$$\lambda_1/s_0 = 0.0125000000, \quad \lambda_2/s_0 = 0.0150117119, \quad \lambda_3/s_0 = 0.0180281196, \quad \lambda_4/s_0 = 0.0216506351, \\ \lambda_5/s_0 = 0.0260010478, \quad \lambda_6/s_0 = 0.0312256192, \quad \lambda_7/s_0 = 0.0375000000.$$

Celestial anchors (principal series). Same five principal-series anchors as in Table 2:

$$(n, \nu) \in \{ (0, -1.20), (1, -0.60), (1, 0.00), (2, 0.60), (0, 1.20) \}.$$

Weights and feasibility. Uniform nonnegative weights with total mass 1:

$$w_{\text{forw}}(q, 1) = \frac{1}{18} \quad (q=1, \dots, 6), \quad w_{\text{H}}(p) = \frac{1}{18} \quad (p=1, \dots, 7), \quad w_{\text{cel}}(k) = \frac{1}{18} \quad (k=1, \dots, 5).$$

Hence the certificate has total support $6+7+5 = 18$ and $\sum w = 1$. Each underlying tester is nonnegative on the declared cone; therefore their uniform convex combination is nonnegative. Analytic projector and dispersion invariances (IR scheme, pivot, and scale) are preserved by construction (Lemma 5.37, Lemma 5.69, Lemma 5.89), as are CPT/crossing constraints (Proposition 5.94).

Wiring update The celestial component may be realized either on the principal set (Table 2) or on the strip set (Table 4); both choices keep the certificate nonnegative with the same forward/Hankel grids.

Helicity-averaged positivity with dressing The helicity-averaged, anchor-preserving dressed absorptive profile is nonnegative for all testers used in the certificate; see Section 5.84 and Proposition 5.102.

Remark. This compact dual coincides node-wise with the earlier “numbers, no solver” presentation in Section 5.47, now updated to the widened near-forward window of Section 5.10.

5.53 Wald–JKM corner calibration on belts

Lemma 5.72 (JKM fix). Let \mathcal{Q}_ξ be the Noether charge for the boost generator ξ and Θ the symplectic potential current. Choose the JKM counterterm so that the corner potential on the belt satisfies

$$\delta \left[\frac{\text{Area}}{4G} \right]_{\text{corner}} - \delta [\xi \cdot \Theta(\delta g)]_{\text{corner}} = 0 + O(\mathcal{B}_{\text{belt}}),$$

with the normalization fixed by the Rindler witness (Section 5.26). Then the first–law identity and the modular equation of state are insensitive to the Wald ambiguity up to $O(\mathcal{B}_{\text{belt}})$.

Kernel. Match the boost Ward identity from the belt first–law channel to the Iyer–Wald variation; the belt choice fixes the potential unambiguously modulo $O(\mathcal{B}_{\text{belt}})$, which vanishes under flow removal.

Alternate proof (RCE/gauge–invariant). Fix a thin belt collar $\mathcal{B}_r \subset \partial_r R$ and let $h_{\mu\nu}$ be a metric variation supported in \mathcal{B}_r . By relative Cauchy evolution (RCE) at belt level, the infinitesimal response of any observable A is generated by the integrated stress tensor:

$$\delta_h^{\text{RCE}} A = \frac{i}{2} \left[\int_{\mathcal{B}_r} \sqrt{-g} T^{\mu\nu} h_{\mu\nu}, A \right] + O(\mathcal{B}_{\text{belt}}),$$

with the $O(\mathcal{B}_{\text{belt}})$ uniform in $|R|$ by belt microcausality/time–slice control and quasi–local factorization Lemma 5.70, Propositions 5.71 and 5.73, and Section 5.54 and by positive–flow removal Lemma 3.3. Choosing $h_{\mu\nu} = \mathcal{L}_\xi g_{\mu\nu}$ supported in \mathcal{B}_r , functoriality and the time–slice property identify the RCE generator with the Noether/Wald boost charge on the belt, modulo the JKM ambiguity calibrated in Section 5.53 and Lemma 5.72:

$$\delta_{\mathcal{L}_\xi g}^{\text{RCE}} \hat{=} \int_{\partial\Sigma \cap \mathcal{B}_r} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) - \delta C_\xi^{\text{corner}},$$

where $C_\xi^{\text{corner}} := \int_{C_r} \xi \cdot Y$ encodes the Y –shift $(\Theta, Q_\xi) \mapsto (\Theta + dY, Q_\xi - \xi \cdot Y)$. On the other hand, by the Iyer–Wald identity on a cap \mathcal{S}_r that fills the collar, $d(\delta Q_\xi - \xi \cdot \Theta) = \omega(g; \delta g, \mathcal{L}_\xi g) - \xi \cdot E(g) \delta g$, and Stokes reduces the bulk to the belt with an $O(\mathcal{B}_{\text{belt}})$ remainder controlled as in your proof. Hence the RCE generator equals $\int_{C_r} \delta Q_\xi - \delta \int_{C_r} \xi \cdot \Theta + O(\mathcal{B}_{\text{belt}})$. Matching the RCE charge to the *boost Ward charge* supplied by the belt first–law channel (kernel zero) *defines* Y so that $\delta C_\xi^{\text{corner}} = O(\mathcal{B}_{\text{belt}})$. Using the standard corner evaluation $\int_{C_r} \delta Q_\xi = \delta[\text{Area}(C_r)/(4G)] + O(\mathcal{B}_{\text{belt}})$ in the boost–adapted gauge, we obtain

$$\delta \left[\frac{\text{Area}}{4G} \right]_{\text{corner}} - \delta [\xi \cdot \Theta(\delta g)]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}),$$

which is the claimed JKM calibration, obtained without gauge fixing beyond the belt support and manifestly RCE/gauge–invariant. *Conversely*, if the calibration holds with a uniform $O(\mathcal{B}_{\text{belt}})$ bound for all admissible δg , then the RCE generator agrees with the belt Ward charge, fixing Y up to $O(\mathcal{B}_{\text{belt}})$. *All constants are per length and uniform in $|R|$.* See also Section 5.3 for the same content in kernel form, the Weyl–covariance check in Section 5.56, and the BRST/diffeomorphism dressing invariance in Section 5.62.

5.54 Brown–York dictionary for the belt flux

Proposition 5.73 (quasi–local stress on the belt). Let γ_{ab} be the induced metric on the belt and K_{ab} its extrinsic curvature. Define the Brown–York tensor $T_{ab}^{\text{BY}} = \frac{1}{8\pi G} (K_{ab} - K \gamma_{ab})$ with the JKM corner fix above. Then the canonical–energy flux appearing in P10–SEE–04 satisfies

$$2\pi \int_\Sigma d\Sigma^\mu \xi^\nu \delta \langle T_{\mu\nu} \rangle = 2\pi \int_{\partial\Sigma} d\ell^a \delta \langle T_{ab}^{\text{BY}} \rangle \xi^b + O(\mathcal{B}_{\text{belt}}).$$

Kernel. Stokes on the symplectic current moves the bulk stress flux to the belt; the calibrated corner term cancels, leaving the BY flux modulo $O(\mathcal{B}_{\text{belt}})$.

Proof (RCE/gauge-invariant). Let $h_{\mu\nu} = f(\rho) \mathcal{L}_\xi g_{\mu\nu}$ be a belt-supported deformation that interpolates across a slab $\Sigma \subset \partial_r R$ with boundary $\partial\Sigma = C_r^{(1)} \cup C_r^{(2)}$ and $f(\rho)$ a collar cutoff. By RCE,

$$\frac{d}{d\epsilon} \Big|_{\epsilon=0} \langle A \rangle_{g+\epsilon h} = \frac{i}{2} \left\langle \left[\int_{\mathcal{B}_r} \sqrt{-g} T^{\mu\nu} h_{\mu\nu}, A \right] \right\rangle + O(\mathcal{B}_{\text{belt}}).$$

Taking A to be the modular generator on the wedge and using JLMS/canonical-energy matching (Section 5.20), the Iyer–Wald identity gives

$$\int_{\Sigma} \omega(g; \delta g, \mathcal{L}_\xi g) = \int_{\partial\Sigma} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) + O(\mathcal{B}_{\text{belt}}).$$

By the JKM calibration from Section 5.53 and Lemma 5.72 (now proved in RCE form), the corner combination equals $\delta[\text{Area}/(4G)]_{\partial\Sigma} + O(\mathcal{B}_{\text{belt}})$. The RCE bulk generator reduces, on shell and for admissible states, to the matter flux $2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu} \rangle$ up to the uniform belt remainder. The GHY/covariant-phase-space matching on the timelike belt identifies the boundary generator with the Brown–York momentum flux,

$$\int_{\partial\Sigma} (\delta Q_\xi - \xi \cdot \Theta) = 2\pi \int_{\partial\Sigma} d\ell^a \delta\langle T_{ab}^{\text{BY}} \rangle \xi^b + O(\mathcal{B}_{\text{belt}}),$$

so that

$$2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu} \rangle = 2\pi \int_{\partial\Sigma} d\ell^a \delta\langle T_{ab}^{\text{BY}} \rangle \xi^b + O(\mathcal{B}_{\text{belt}}),$$

which is the Brown–York dictionary, now obtained by a manifestly gauge-invariant belt RCE. Cf. Proposition 5.73 and Section 6.2 for the same identity presented as a stability statement.

5.55 Quasi-local positive energy on the belt

Corollary 5.74 (quasi-local positive energy on belt). Under the JKM corner calibration (Lemma 5.72) and the Brown–York dictionary (Proposition 5.73), for any spacelike belt slab $\Sigma \subset \partial_r R$ and any nonnegative smearing f along its generators,

$$2\pi \int_{\partial\Sigma} d\ell_a f \delta\langle T_{\text{BY}}^{ab} \rangle \xi_b \geq -C_{\text{BY}} B_{\text{belt}}, \quad C_{\text{BY}} > 0, \quad (5.35)$$

with C_{BY} independent of $|R|$. In particular, in the removal limit $(u, s) \downarrow 0$,

$$2\pi \int_{\partial\Sigma} d\ell_a f \delta\langle T_{\text{BY}}^{ab} \rangle \xi_b \geq 0.$$

Sketch. By the belt BY dictionary (Proposition 5.73),

$$2\pi \int_{\partial\Sigma} d\ell_a f \delta\langle T_{\text{BY}}^{ab} \rangle \xi_b = 2\pi \int_{\Sigma} d\Sigma_\mu f \xi_\nu \delta\langle T^{\mu\nu} \rangle + O(B_{\text{belt}}).$$

Smearing by $f \geq 0$ and using the ANEC/QNEC synthesis (Theorem 5.61) together with canonical-energy positivity on non–Killing data yields the nonnegativity of the bulk term. The $O(B_{\text{belt}})$ piece is uniform and vanishes under positive-flow removal. \square

Remark 5.75 (QG deliverable: quasi-local positive energy on belts). Quasi-local positive energy on belts (Brown–York flux projected on the belt boost) follows from canonical-energy positivity in the second-order modular equation of state (Theorem 5.43) together with the calibrated corner terms and the Brown–York dictionary (Sections 5.53 and 5.54). The statement is stable under the invariances summarized in Section 6.

5.56 Weyl covariance and absence of a leading anomaly

Lemma 5.76 (Weyl check). Under a uniform Weyl rescaling $g_{\mu\nu} \mapsto \Omega^2 g_{\mu\nu}$ with $\Omega = 1 + \epsilon$ and the belt JKM calibration,

$$\delta \left[S - \frac{\text{Area}}{4G} \right] \mapsto \delta \left[S - \frac{\text{Area}}{4G} \right] + O(\epsilon \mathcal{B}_{\text{belt}}),$$

so the modular equation of state is Weyl-covariant at linear order in ϵ up to $O(\mathcal{B}_{\text{belt}})$.

Kernel. Area/(4G) and the BY term co-transform; the entropy counterterm shift is absorbed by recovery (Section 5.33); remainder is belt-suppressed.

5.57 Raising the subtraction order does not weaken positivity

Proposition 5.77 (monotone N). In the gravity-subtracted dispersive scheme on the cone \mathcal{S} with tester-certified Regge slope $\alpha_{\text{R}} \leq 2 + \delta_{\star}$ (< 3), for any $N \geq 3$ the even-parity forward derivative inequalities and all Hankel/celestial positivity functionals valid at $N=3$ remain valid (possibly with strictly stronger bounds) at $N' > N$.

Kernel. Increasing N multiplies the subtracted integrands by positive powers of s'/s_0 in the denominator and preserves crossing kernels; all nonnegative testers act on smaller tails while leaving absorptive positivity intact. The assumption $\alpha_{\text{R}} < 3$ guarantees integrability of the tails for all $k \geq 0$.

5.58 Tail and quadrature audit for the dispersive integrals

Setup and ledger bindings. Use $S_{\text{cut}} = 20 s_0$ and $\mathcal{R}(t) = (1 + |t|/s_0)^{1/2}$. We *do not* assume a fixed Regge slope; instead the tester cone certifies $\alpha_{\text{R}}(t) \leq 2 + \delta_{\star}$ on $t \in [-0.25 s_0, 0]$ (Proposition 5.30). On $[s_0, S_{\text{cut}}]$ we use Gauss–Radau quadrature of order $n = 4$ with Peano constant $\frac{1}{1080}$.

High- s' tail beyond S_{cut} (δ -sensitive). For $\partial_s^k \mathfrak{R} \mathcal{A}^{(N)}(0, t)$ with $N = 3$ and $\alpha_{\text{R}} = 2 + \delta$,

$$|\text{Tail}_{S_{\text{cut}}}| \leq \mathcal{R}(t) \frac{S_{\text{cut}}^{\alpha_{\text{R}}-3-k}}{3+k-\alpha_{\text{R}}} = \left(\frac{3+k}{3+k-\delta} M^\delta \right) |\text{Tail}_{S_{\text{cut}}}|_{\delta=0}, \quad M := \frac{S_{\text{cut}}}{s_0}.$$

Specializing to $\alpha_{\text{R}} = 2$ (baseline ledger numbers):

$$\sup_{t \in [-0.20 s_0, 0]} |\text{Tail}_{S_{\text{cut}}}| \leq \frac{1.0954451150}{20 s_0} = 0.0547722558 s_0^{-1} \quad (k=0),$$

and the $k=2$ improvement is

$$\frac{1.0954451150}{3 (20 s_0)^3} = 4.564354646 \times 10^{-5} s_0^{-3}.$$

Window widening factor. Enlarging t from $[-0.20 s_0, 0]$ to $[-0.25 s_0, 0]$ multiplies $\mathcal{R}(t)$ by at most $\sqrt{1.25}/\sqrt{1.2} \approx 1.0206$, absorbed within the same dispersion headroom.

Quadrature remainder on $[s_0, S_{\text{cut}}]$. Let $g(s')$ be the subtracted integrand. Gauss–Radau ($n=4$) obeys

$$|\text{Quad}_{n=4}| \leq \frac{1}{1080} \sup_{s' \in [s_0, S_{\text{cut}}]} |g^{(4)}(s')| (S_{\text{cut}} - s_0)^5.$$

This is ledger-ready: any admissible sup $|g^{(4)}|$ for the envelope inserted here certifies that the quadrature error fits the global split of Section 5.28.

Mellin-strip factor in dispersion tails For numerical audit, one may include the multiplicative factor $F_{\text{strip}}(M, \sigma_0) = \sup_{|\sigma| \leq \sigma_0} M^{|\sigma|}$ in the high- s' tail and Gauss–Radau remainders to track the strip weight $(s'/s_0)^\sigma$. With $M = 100$ or 500 and $\sigma_0 = 0.15$, $F_{\text{strip}} \leq 1.995$ or 2.540 , both absorbed into $\varepsilon_{\text{disp}}$.

Widened t -window cost (no schedule change). Extending t from $[-0.20 s_0, 0]$ to $[-0.30 s_0, 0]$ multiplies the window factor $R_{\text{max}} = \sqrt{1 + |t|/s_0}$ from 1.0954451150 to 1.1401754251 , i.e. by

$$\frac{R_{\text{max}}(-0.30 s_0)}{R_{\text{max}}(-0.20 s_0)} = 1.040833\dots$$

which scales the dispersion tails linearly. The baseline constant $1.0954451150/(3(20 s_0)^3) = 4.5643546459 \times 10^{-5} s_0^{-3}$ becomes $1.1401754251/(3(20 s_0)^3) = 4.7507309379 \times 10^{-5} s_0^{-3}$. Consequently, the tail portions in the two audit freezes (Section 5.93) tighten to

$$\begin{aligned} \varepsilon = 10^{-6} : \quad & |\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq 6.05 \times 10^{-8} (< \varepsilon/12), \\ \varepsilon = 10^{-8} : \quad & |\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq 4.84 \times 10^{-10} (< \varepsilon/12). \end{aligned}$$

Quadrature terms are unchanged; both freezes retain strict total slack $\leq \varepsilon/6$.

5.59 Edge Cardy sanity check for the Unruh channel

Proposition 5.78 (edge Cardy bound). Let $T_{\text{Unruh}} = 1/(2\pi)$ be the Unruh temperature of the belt boost. Using the composite constant $\backslash\text{CCardy}$ and its CIS/MX benchmarks from the bindings, the edge entropic density satisfies

$$\frac{S(\rho_R)}{\text{length}(\partial R)} \leq 0.4083 T_{\text{Unruh}} + O(\mathcal{B}_{\text{belt}}).$$

In CIS/MX, this gives the explicit numbers

$$\frac{S}{\text{length}} \leq \frac{0.4083}{2\pi} \quad \text{and} \quad \frac{S}{\text{length}} \leq \frac{0.41}{2\pi},$$

up to $O(\mathcal{B}_{\text{belt}})$ and the finite-size $O(1/\text{length})$ term from Section 5.41.

Kernel. Apply the belt first-law thermalization (Bisognano–Wichmann) via the Unruh/KMS normalization Section 5.65 and Lemma 5.84, use the Cardy-type edge density normalization (binding 0.4083) from the edge/ward calibration Section 5.17 and Proposition 5.38, and subtract the $O(\mathcal{B}_{\text{belt}})$ remainder as in Section 5.14.

5.60 Equivalence: entanglement equilibrium and linearized Einstein equations

Theorem 5.79 (Modular equation of state \iff linearized Einstein equations). Work within the framework recap (locally covariant nets; belt regulator and positive flows; belt first law and OS kernel), assume the belt JLMS identification and belt-level recovery/continuity, and use the c-function/GSL infrastructure. Then for belt-anchored deformations we have the following equivalence, up to $O(\mathcal{B}_{\text{belt}})$:

1. **(EE \implies SEE)** If *entanglement equilibrium* holds for all admissible belt deformations,

$$\delta \left[S - \frac{\text{Area}}{4G} \right] = O(\mathcal{B}_{\text{belt}}),$$

then the linearized Einstein equations hold in expectation on $D[R]$.

2. (**SEE** + **JKM fix** \Rightarrow **EE**) Conversely, if the linearized Einstein equations hold in expectation on $D[R]$ and the JKM corner has been fixed by the belt boost Ward identity [2], then

$$\delta\left[S - \frac{\text{Area}}{4G}\right] = O(B_{\text{belt}}).$$

Idea of proof. The forward direction is the modular equation of state: the belt first law gives $\delta S = \delta\langle K_{\text{mod}} \rangle + (\text{edge/corner})$, JLMS trades boundary relative entropy for bulk canonical energy [3], and the calibrated corner/edge terms (JKM) plus the Brown–York dictionary convert the canonical-energy side into a stress-flux and an area term—i.e. the linearized SEE [6, 7]. For the converse, assume the SEE and use the Iyer–Wald identity together with the Brown–York flux and the same JKM calibration to reassemble the first-law variation; the net is $O(B_{\text{belt}})$. \square

5.61 Data processing and recoverability at belt level

Proposition 5.80 (DPI with recovery). For any belt-compatible channel Φ (belt coarse-graining) and reference σ , with a belt-level recovery map R ,

$$S(\rho\|\sigma) - S(\Phi\rho\|\Phi\sigma) \geq -2\log F(\rho, R\circ\Phi(\rho)) \geq 0,$$

where F is the Uhlmann fidelity. In particular, under the OS kernel the deficit is $O(\mathcal{B}_{\text{belt}})$ and vanishes upon flow removal. **Kernel.** Belt recovery/continuity and sufficiency are provided within this section; the bound is the refined DPI with rotated Petz recovery, applied to belt-local algebras, with regulator removal by Lemma 3.3.

5.62 BRST/diffeomorphism dressing invariance

Lemma 5.81 (anchor-preserving diffeos). Let δ_ξ be an anchor-preserving diffeomorphism generated by a vector field ξ tangent to the belt. For any observable \mathcal{O} ,

$$\delta_\xi \mathcal{O} = i[\mathcal{Q}_{\text{diff}}[\xi], \mathcal{O}] = O(\mathcal{B}_{\text{belt}}),$$

so physical predictions are invariant under such dressings up to $O(\mathcal{B}_{\text{belt}})$, which vanishes after removal. **Kernel.** Move the diffeomorphism through the belt cocycle using the functorial net structure and belt factorization; boundary terms are exponentially suppressed by belt microcausality, and the regulator is removed by Lemma 3.3.

5.63 Belt-width RG flow and an integrated monotonicity

Theorem 5.82 (width-flow monotone). Define $\mathbf{c}_r = \partial_r(S - \frac{\text{Area}}{4G})$. Under the OS kernel and the monotonicity/contractivity infrastructure of Section 5.18,

$$\partial_r \mathbf{c}_r \leq -\lambda_{\text{clu}} \mathbf{c}_r + O(\mathcal{B}_{\text{belt}}),$$

and for $r \rightarrow \infty$ the integrated inequality yields

$$\int_{r_1}^{r_2} \mathbf{c}_r dr \leq \frac{1 - e^{-\lambda_{\text{clu}}(r_2 - r_1)}}{\lambda_{\text{clu}}} \mathbf{c}_{r_1} + O(\mathcal{B}_{\text{belt}}).$$

Kernel. Differentiate the belt c-function along width flow using Theorem 5.39; apply Grönwall and track the $O(\mathcal{B}_{\text{belt}})$ remainder with Section 5.14, removing flows via Lemma 3.3.

5.64 Quantum Fisher information and QNEC; coherent-state check

Proposition 5.83 (QFI–QNEC link). For belt shape deformations generated by a null unitary $U(\theta)$, the quantum Fisher information $F_Q = \partial_\theta^2 S(\rho||\rho_\theta)|_{\theta=0}$ obeys

$$F_Q \geq \int du f(u)^2 2\pi \langle T_{kk}(u) \rangle,$$

and for coherent states of free fields F_Q saturates the bound (no covariance shift). *Numeric check (our coherent pulse)*: with $A = 10^{-2}$, $\sigma = 3$, $u_0 = \sigma$ from Section 5.26,

$$\int du \langle T_{kk} \rangle = 2.9540897515 \times 10^{-5}, \quad 2\pi \int du \langle T_{kk} \rangle = 1.8561093322 \times 10^{-4},$$

so the integrated QNEC bound is saturated by the coherent displacement at this order.

5.65 Unruh/KMS normalization sanity check

Lemma 5.84 (temperature and modular scale). For the belt boost, the Unruh temperature is $T_U = 1/(2\pi) = 0.1591549431$ and the KMS inverse temperature $\beta = 2\pi$. Our Rindler witness gives $\delta \langle K_R \rangle = \frac{\pi}{2} A^2 = 1.5707963268 \times 10^{-4}$ at $A = 10^{-2}$, consistent with the thermal scale used in this Section.

5.66 Error propagation for the forward coefficient $c_{2,0}$

Lemma 5.85 (tail and quadrature bound). Let $\widehat{c}_{2,0} := s_0^3 c_{2,0}$ be the dimensionless forward coefficient. For $N = 3$ and Regge slope 2 on $t \in [-0.20 s_0, 0]$, the high-energy tail beyond $S_{\text{cut}} = 20 s_0$ contributes at most

$$|\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq \frac{1}{2} \cdot \frac{1}{\pi} \cdot \frac{\sqrt{1.2}}{3(20)^3} = 7.2643960393 \times 10^{-6},$$

where the extra factor $1/3$ is $1/(5 - \alpha_R)$ at $\alpha_R = 2$, and the factor $\frac{1}{2}$ accounts for $\widehat{c}_{2,0} = \frac{s_0^3}{2} \partial_s^2 \Re \mathcal{A}^{(N)}(0, 0)$. The Gauss–Radau remainder on $[s_0, 20 s_0]$ obeys

$$|\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{s_0^3}{2} \frac{1}{1080} \sup_{[s_0, 20 s_0]} |g^{(4)}| (19 s_0)^5,$$

with $\frac{1}{1080} = 1/1080$ and g the subtracted integrand. Both terms are included in the global budget of Section 5.28.

5.67 Interchanging variation with flow removal

Lemma 5.86 (variation–limit interchange). Let $\mathcal{O}_{u,s}(R)$ be any belt-regularized observable among $\{\langle K_{\text{mod}} \rangle, S, \text{Area}, \text{amplitude functionals}\}$. Assume the framework recap Section 2 and the OS kernel Lemmas 3.1 to 3.3, together with belt microcausality/timeslice control and factorization Lemma 5.70, Propositions 5.71 and 5.73, and Section 5.54. Then for any admissible variation δ ,

$$\delta \lim_{(u,s) \downarrow 0} \mathcal{O}_{u,s}(R) = \lim_{(u,s) \downarrow 0} \delta \mathcal{O}_{u,s}(R),$$

with error bounded by

$$|\delta \mathcal{O}_{u,s} - \delta \mathcal{O}_{0,0}| \leq C \left(e^{-\mu_{\text{eff}} r} + \eta^m + C_{\text{dress}}(u^p + s^q) \right),$$

for the belt/AGSP choices in Section 5.14. **Kernel.** Uniform integrable bounds from belt microcausality/dispersion control give dominated convergence; the AGSP/seed pipeline controls the variation norm uniformly in (u, s) ; flow removal Lemma 3.3 closes the limit.

5.68 Entropy wedge nesting along positive flows

Proposition 5.87 (EWN stability under flows). Let $R_1 \subset R_2$ be belt-anchored regions. Under belt-level nesting/recovery and the OS kernel, for all $(u, s) > 0$ small enough we have

$$\text{EW}_{u,s}(R_1) \subseteq \text{EW}_{u,s}(R_2),$$

and the inclusion persists as $(u, s) \downarrow 0$ with an $O(\mathcal{B}_{\text{belt}})$ geometric error that vanishes under flow removal. **Kernel.** Functoriality and recovery at fixed flows yield monotonicity Proposition 5.80; belt microcausality/timeslice control Lemma 5.70 and Proposition 5.71 and the removal lemma Lemma 3.3 propagate inclusion to the limit.

5.69 Edge-of-wedge analyticity on the dispersive cone

Lemma 5.88 (tube analyticity). On the declared cone \mathcal{S} with gravity subtraction at order $N \geq 3$ and tester-certified Regge slope $\alpha_R \leq 2 + \delta_\star < 3$, the subtracted amplitude extends analytically to a forward tube

$$\mathcal{T} = \{s \in \mathbb{C} : \Im s > 0, |t| \leq t_\star\},$$

and satisfies polynomial boundedness in \mathcal{T} compatible with the dispersive projectors and windows of Sections 5.10 and 5.23. **Kernel.** Reflection positivity and belt microcausality yield edge-of-wedge continuation; the $N \geq 3$ subtractions enforce integrable growth; gravity IR pieces are nonanalytic and are explicitly removed as in Sections 5.23 and 5.50.

Add-on (tube analyticity on the strip). Edge-of-wedge analyticity holds unchanged on the off-principal strip: the factors $(s'/s_0)^\sigma$ are entire in s and respect the tube bounds of Lemma 5.88.

5.70 Rescaling invariance with respect to the subtraction scale

Lemma 5.89 (dimensionless stability of $c_{2,0}$). Define the dimensionless coefficient $\widehat{c}_{2,0} := s_0^3 c_{2,0}$ at subtraction scale s_0 . Under the gravity-subtracted crossing-symmetric dispersion with $N=3$,

$$\widehat{c}_{2,0}|_{s_0} = \widehat{c}_{2,0}|_{\bar{s}_0}.$$

Kernel. Changing s_0 rescales the integration variable $s' \mapsto \alpha s'$ and multiplies the kernel by α^{-3} , exactly compensated by the prefactor s_0^3 ; pivot-shift polynomials drop out of even-parity forward derivatives (Section 5.50).

Add-on (scale invariance unchanged). Rescaling s_0 is compensated in the dimensionless forward coefficients; the strip weights $(s'/s_0)^\sigma$ preserve Lemma 5.89.

5.71 OSR/complexity inflation aggregator with numeric audit baseline

Aggregator. For AGSP step m and belt base Γ_{belt} , the per-cut inflation factor is

$$\mathcal{I}_{\text{OSR}}(m) := \Lambda_0 \Gamma_{\text{belt}} \Upsilon(m).$$

Under the audit baseline bindings from the ledger,

$$\mathcal{I}_{\text{OSR}}(m)|_{\text{baseline}} = 2.5 \, 3.0 \, \Upsilon(m) \equiv \mathcal{I}_{\text{OSR}}(m)|_{\text{baseline}}.$$

On the Trotter-safe branch we use $\Upsilon(m) \geq 2m+1$. In particular, at the Page-point baseline $m=8$ used in Section 5.22,

$$\log \mathcal{I}_{\text{OSR}}(8) = \log 2.5 + \log 3.0 + \log \Upsilon(8),$$

so any future update of *either* `\LambdaZeroNum` or `\gammabeltNum` (or a different m) propagates automatically. This is the exact inflation that appears inside the per-length Page bound and remains uniform in $|R|$ within the single budget $O(\mathcal{B}_{\text{belt}})$.

5.72 Replica analyticity and belt-uniform von Neumann limit

Lemma 5.90 (replica analyticity). Let $S_n(R) = \frac{1}{1-n} \log \text{tr} \rho_R^n$ for $\Re n > 1$ be the belt-regularized Rényi entropies of an admissible state in the OS kernel. Under the framework recap Section 2 (items 1–3,5), the OS positivity/recovery/removal lemmas Lemmas 3.1 to 3.3, and belt-level continuity/recovery Proposition 5.80, the map $n \mapsto S_n(R)$ is analytic in a strip $\{n : |n - 1| < \delta\}$ with $\delta > 0$ independent of $|R|$, and

$$\lim_{n \rightarrow 1^+} S_n(R) = S(\rho_R), \quad \lim_{n \rightarrow 1^+} \partial_\lambda S_n(R_\lambda) = \partial_\lambda S(\rho_{R_\lambda}),$$

for any belt-anchored shape path $\lambda \mapsto R_\lambda$, with uniform remainder

$$|S_n(R) - S(\rho_R)| \leq K |n - 1| + O(\mathcal{B}_{\text{belt}}),$$

where K depends only on the belt budgets and recovery constants recorded in Section 5.13 and Proposition 5.80. The variation–limit interchange is justified by Section 5.67.

5.73 Quantum Bousso bound from QNEC and the belt GSL

Theorem 5.91 (quantum Bousso bound on belts). Let λ be an affine parameter along a belt-anchored null generator with tangent k^a , and let $A(\lambda)$ denote the cross-sectional area line-density transported by that generator.³ Assume the QNEC of Section 5.9 and the belt GSL of Section 5.38. Then for any segment $[\lambda_1, \lambda_2]$,

$$S(\lambda_1 \rightarrow \lambda_2) \leq \frac{A(\lambda_1) - A(\lambda_2)}{4G} + O(\mathcal{B}_{\text{belt}}).$$

Proof of Theorem 5.91. We work on a single belt-anchored generator and write $\sigma^2 := \sigma_{ab}\sigma^{ab} \geq 0$, $\theta := \nabla_a k^a$, and $T_{kk} := T_{ab}k^a k^b$. Let $a(\lambda) = A(\lambda)$ be the area line-density transported by the generator, so $a' = \theta a$.

Step 1: Dirichlet testers and the twice-integrated QNEC. Fix $\lambda_1 < \lambda_2$ and consider the Dirichlet Green function $G(\lambda, x)$ on $[\lambda_1, \lambda_2]$ for the operator $-\partial_\lambda^2$, with $G(\lambda_i, x) = 0$ for $i = 1, 2$ and $G \geq 0$:

$$G(\lambda, x) = \begin{cases} \frac{(\lambda - \lambda_1)(\lambda_2 - x)}{\lambda_2 - \lambda_1}, & \lambda \leq x, \\ \frac{(x - \lambda_1)(\lambda_2 - \lambda)}{\lambda_2 - \lambda_1}, & \lambda \geq x. \end{cases}$$

Then $-\partial_\lambda^2 G(\lambda, x) = \delta(\lambda - x)$ in distributions. Let $g_{\epsilon, i}$ be smooth nonnegative approximations of $G(\lambda, \lambda_i \pm \epsilon)$ (so they have compact support in (λ_1, λ_2) and vanish at the endpoints). By the local QNEC (Section 5.9), applied pointwise and multiplied by the transported area $a(\lambda)$,

$$2\pi a(\lambda) T_{kk}(\lambda) \geq S''_{\text{out}}(\lambda),$$

in the sense of distributions. Testing against $g_{\epsilon, i}$ and integrating by parts twice (endpoint terms vanish by the Dirichlet condition) gives

$$S_{\text{out}}(\lambda_i) \leq 2\pi \int_{\lambda_1}^{\lambda_2} a(\lambda) T_{kk}(\lambda) g_{\epsilon, i}(\lambda) d\lambda.$$

Subtracting the $i = 1$ inequality from the $i = 2$ one and defining

$$f_\epsilon(\lambda) := g_{\epsilon, 2}(\lambda) - g_{\epsilon, 1}(\lambda), \quad -f''_\epsilon \xrightarrow{\epsilon \rightarrow 0} \delta(\lambda - \lambda_2) - \delta(\lambda - \lambda_1),$$

³Equivalently, A is the local transverse area element $a(\lambda)$ carried by the chosen generator; we keep the letter A only to match the global notation elsewhere.

we obtain

$$\Delta S_{\text{out}} := S_{\text{out}}(\lambda_2) - S_{\text{out}}(\lambda_1) \leq 2\pi \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda + o(1), \quad (5.36)$$

with $f_\epsilon(\lambda_i) = 0$.

Step 2: Raychaudhuri decomposition with weights. Using the Raychaudhuri equation with Einstein's equation along the generator,

$$\theta' = -\frac{1}{2}\theta^2 - \sigma^2 - 8\pi G T_{kk}, \quad a' = \theta a,$$

we rewrite

$$8\pi G f_\epsilon a T_{kk} = -f_\epsilon a \theta' - \frac{1}{2}f_\epsilon a \theta^2 - f_\epsilon a \sigma^2 = -f_\epsilon (a\theta)' + \frac{1}{2}f_\epsilon a \theta^2 - f_\epsilon a \sigma^2.$$

Integrating and using $f_\epsilon(\lambda_i) = 0$,

$$8\pi G \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda = \int_{\lambda_1}^{\lambda_2} f'_\epsilon a \theta d\lambda + \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda. \quad (5.37)$$

Step 3: Trading the mixed term for area loss plus an edge Wronskian. Since $a' = \theta a$, we have in the distributional sense (use $f''_\epsilon = \delta(\lambda_1 + \epsilon) - \delta(\lambda_2 - \epsilon)$ and then $\epsilon \rightarrow 0$)

$$\int_{\lambda_1}^{\lambda_2} f'_\epsilon a \theta d\lambda = \int_{\lambda_1}^{\lambda_2} f'_\epsilon a' d\lambda = \int_{\lambda_1}^{\lambda_2} f''_\epsilon a d\lambda + [a f'_\epsilon]_{\lambda_1}^{\lambda_2} = A(\lambda_1) - A(\lambda_2) + [A f'_\epsilon]_{\lambda_1}^{\lambda_2}.$$

Substituting into (5.37) yields

$$8\pi G \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda = (A(\lambda_1) - A(\lambda_2)) + [A f'_\epsilon]_{\lambda_1}^{\lambda_2} + \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda. \quad (5.38)$$

Step 4: Canonical-energy control of shear and edge pieces. By the Raychaudhuri estimate with canonical-energy control (Proposition 5.96) and the shear-control lemma (Lemma 5.106), the combination

$$\mathcal{E}_{\text{can}}[f_\epsilon] := \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda - \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - [A f'_\epsilon]_{\lambda_1}^{\lambda_2}$$

is nonnegative up to the $O(\mathcal{B}_{\text{belt}})$ remainder (the latter goes to zero under flow removal, Lemma 3.3). This is the positive quadratic form supplied by Theorem 5.43, with the ‘‘edge/corner’’ calibration fixed by the JKM prescription (Section 5.53, Lemma 5.72, and Proposition 5.38). Equivalently,

$$[A f'_\epsilon]_{\lambda_1}^{\lambda_2} + \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda \leq O(\mathcal{B}_{\text{belt}}). \quad (5.39)$$

Insert (5.39) into (5.38) to obtain

$$8\pi G \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda \leq A(\lambda_1) - A(\lambda_2) + O(\mathcal{B}_{\text{belt}}). \quad (5.40)$$

The monotone c-function/belt GSL (Theorems 5.39 and 5.55) underlies (5.39): it guarantees that the regulated generalized expansion produces a nonnegative canonical energy, and that the calibrated edge pieces do not spoil the sign.

Step 5: Combine with QNEC and take the limit. Combining (5.36) with (5.40) gives

$$\Delta S_{\text{out}} \leq 2\pi \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda + o(1) \leq \frac{A(\lambda_1) - A(\lambda_2)}{4G} + O(\mathcal{B}_{\text{belt}}) + o(1).$$

Dominated convergence applies thanks to the canonical-energy bound, so sending $\epsilon \rightarrow 0$ (and then removing the auxiliary flow per Lemma 3.3) yields the desired inequality

$$S(\lambda_1 \rightarrow \lambda_2) = \Delta S_{\text{out}} \leq \frac{A(\lambda_1) - A(\lambda_2)}{4G} + O(\mathcal{B}_{\text{belt}}).$$

□

Remark 5.92 (Rindler benchmark). From Section 5.26, $\Delta S_{\text{Rindler}} \leq \Delta \langle K_{\text{R}} \rangle = \frac{\pi}{2} A^2 = 1.5707963268 \times 10^{-4}$ at $A = 10^{-2}$. Hence the bound implies $\Delta A \geq 4G \Delta S \geq 6.283185307 \times 10^{-4} G$ for that pulse (line-density interpretation along the belt).

5.74 Differentiation under the dispersive integrals

Lemma 5.93 (dispersion differentiation). On the cone \mathcal{S} with gravity subtraction at order $N = 3$ and tester-certified slope $\alpha_{\text{R}} \leq 2 + \delta_{\star} < 3$, the forward derivatives satisfy

$$\partial_s^k \mathfrak{R} \mathcal{A}^{(N)}(0, t) = \frac{1}{\pi} \int_{s_0}^{\infty} \frac{ds'}{s'^{N+k+1}} \left(\mathfrak{S} \mathcal{A}_{\text{hard}}(s', t) + (-1)^{N+k} \mathfrak{S} \mathcal{A}_{\text{hard}}(-s' - t, t) \right),$$

with the exchange of ∂_s^k and the s' -integral justified by dominated convergence: $\mathfrak{S} \mathcal{A}_{\text{hard}}(s', t) \leq \mathcal{R}(t) s'^{\alpha_{\text{R}}}$ yields integrability of $s'^{-N-k-1+\alpha_{\text{R}}}$ for all $k \geq 0$ at $N = 3$.

Quantitative tail. For $k = 2$ and $S_{\text{cut}} = 20 s_0$,

$$|\text{Tail}_{S_{\text{cut}}}| \leq \mathcal{R}(t) \frac{S_{\text{cut}}^{\alpha_{\text{R}}-5}}{5 - \alpha_{\text{R}}} = \left(\frac{3}{3 - \delta} 20^\delta \right) \frac{\sqrt{1.2}}{3(20 s_0)^3} \quad (\alpha_{\text{R}} = 2 + \delta).$$

At $\delta = 0$ this reproduces the audit of Section 5.58.

5.75 CPT and crossing invariance of positivity testers

Proposition 5.94 (CPT-crossing stability). The three tester families used in this section—(i) forward even-parity derivatives at fixed $t \leq 0$, (ii) nonnegative Hankel/Gaussian band kernels, and (iii) celestial Gram functionals on the principal lines— are invariant under CPT and under the standard crossing $s \leftrightarrow u$ with $u = -s - t$ at fixed $t \leq 0$ (so at $t = 0$ this reduces to $s \leftrightarrow -s$). Consequently, any conic combination that is nonnegative before the transform remains nonnegative after it.

Kernel. Even-parity forward derivatives are CPT-even by construction; Hankel/impact profiles depend only on impact moduli and preserve nonnegativity; the celestial Gram form is positive semi-definite and is preserved by the Ward action on the principal series, as recorded in Section 5.44 and Lemma 5.66.

CPT-crossing stability with strip testers The renormalized strip Gram functional is even in σ and built from the same absorptive profiles, hence is CPT-even and crossing-stable. Proposition 5.94 applies verbatim.

Near-forward crossing positivity with the compact dual On the widened window $t \in [-0.30 s_0, 0]$, the compact dual of Section 5.52 can be paired channelwise to form the crossing-symmetrized tester $\frac{1}{2}(\mathcal{T}_s + \mathcal{T}_u)$. Since each component tester is nonnegative and analytic-projector invariant, the symmetrized combination remains nonnegative and inherits the invariances of Proposition 5.94. Thus nonforward/near-forward positivity holds uniformly on the widened window without increasing support.

5.76 Dimensional scope and hypotheses

Lemma 5.95 (dimensional validity). The theorem suite holds in spacetime dimensions $d \geq 3$ under the standing framework Section 2 and the internal kernels cited in Section 5. In $d = 2$, the amplitude and celestial components persist *mutatis mutandis* (principal lines only), while the QES statements use the 1+1-dimensional corner prescription and chiral factorization calibrated in Section 5.53 and Lemma 5.72 together with belt-level continuity/stability Proposition 5.80 and Section 5.15; all belt budgets remain uniform in $|R|$.

5.77 Linearized Raychaudhuri area estimate with numeric line

Proposition 5.96 (linearized Raychaudhuri window). Consider a belt-anchored null segment $[\lambda_1, \lambda_2]$ with initial expansion $\theta(\lambda_1) = 0$ and negligible shear at linear order. Then, to first order in the perturbation,

$$\theta(\lambda) = -8\pi G \int_{\lambda_1}^{\lambda} du \langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}),$$

and the area change along the segment obeys

$$\Delta A = \int_{\lambda_1}^{\lambda_2} \theta(\lambda) A_0 d\lambda \leq 8\pi G A_0 L \int_{\lambda_1}^{\lambda_2} \langle T_{kk}(u) \rangle du + O(\mathcal{B}_{\text{belt}}),$$

where $L = \lambda_2 - \lambda_1$.

Numeric line (audit baseline). Using the coherent null pulse of Section 5.26 with $A = 10^{-2}$, $\sigma = 3$, $u_0 = \sigma$,

$$\int du \langle T_{kk} \rangle = \frac{\sqrt{\pi}}{2} \frac{A^2}{\sigma} = 2.9540897515 \times 10^{-5}.$$

Taking $L = 6\sigma = 18$, the *line-density* bound (per generator) reads

$$\Delta A_{\text{line}} \leq 8\pi L \int du \langle T_{kk} \rangle \cdot G = 1.3363987192 \times 10^{-2} G.$$

Together with the quantum Bousso bound (Section 5.73), which enforces $\Delta A \geq 4G \Delta S$, this yields a consistent window for the pulse; the remainder $O(\mathcal{B}_{\text{belt}})$ vanishes under flow removal.

5.78 Null-curvature window on a belt

Corollary 5.97 (Null-curvature). Let $[\lambda_1, \lambda_2]$ be a belt-anchored null segment with initial expansion $\theta(\lambda_1) = 0$ and small shear at linear order. Then:

- (i) *Pointwise (linear order).* Using the linearized Raychaudhuri relation and the first-order equation of state, the Ricci contraction along k satisfies, per generator length,

$$\langle R_{kk}(\lambda) \rangle = 8\pi G \langle T_{kk}(\lambda) \rangle + O(\mathcal{B}_{\text{belt}}).$$

- (ii) *Integrated lower bound.* By ANEC on the belt and item (i),

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk} \rangle \geq -C_{\text{belt}} \mathcal{B}_{\text{belt}},$$

for a belt-uniform constant $C_{\text{belt}} > 0$ independent of $|R|$.

(iii) *Integrated focusing upper bound.* The full Raychaudhuri equation yields

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk} \rangle = -\theta(\lambda_2) - \int_{\lambda_1}^{\lambda_2} d\lambda \left[\sigma_{ab} \sigma^{ab} + \frac{1}{d-2} \theta^2 \right] \leq -\theta(\lambda_2),$$

so that the integrated null curvature lies in the window

$$-C_{\text{belt}} \mathcal{B}_{\text{belt}} \leq \int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk} \rangle \leq -\theta(\lambda_2).$$

Kernel. (i) Combine the linearized Raychaudhuri line with the modular equation of state (SEE) in expectation on belts; (ii) integrate and use ANEC on the belt; (iii) integrate Raychaudhuri with $\theta(\lambda_1) = 0$ and $\sigma^2, \theta^2 \geq 0$. All remainder terms are absorbed into $\mathcal{B}_{\text{belt}}$ and vanish upon flow removal. (See Section 5.77 and Proposition 5.96; Theorem 5.35/Theorem 5.43; Theorem 5.27.) \square

5.79 Approximate Markov gap on belts with recovery

Theorem 5.98 (belt Markov gap). Let $A:B:C$ be a belt-aligned tripartition with B covering the entangling belt. Under belt-level recoverability/continuity Proposition 5.80, the AGSP/seed converter, and the OS kernel Lemmas 3.1 to 3.3, there is a recovery map $R_{B \rightarrow BC}$ such that

$$I(A:C | B) \leq C_{\text{rec}} \left(e^{-\mu_{\text{eff}} r(m)} + \eta^m \right) + C_{\text{seed}} \frac{\log \kappa_{\text{seed}} + \log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{\text{length}(\partial R)} + O(\mathcal{B}_{\text{belt}}),$$

with constants independent of $|R|$. Thus the conditional mutual information per unit length can be made arbitrarily small by the deterministic choices for m in Section 5.28.

Kernel. Apply the refined DPI with (rotated) Petz recovery to the belt algebras Proposition 5.80 and translate fidelity lower bounds into an $I(A:C | B)$ upper bound via Pinsker/product control. The AGSP/seed ledger supplies the explicit length suppression; flow removal Lemma 3.3 kills the $O(\mathcal{B}_{\text{belt}})$ terms.

5.80 Soft theorem and celestial bridge compatibility

Lemma 5.99 (soft–celestial consistency). With the gravity soft subtraction $\mathcal{A}_{\text{soft}}(s, t) = \kappa^2 \left(\frac{s^2}{-t} + \frac{u^2}{-t} \right)$ and the celestial kernel, the celestial-transformed soft piece yields a polynomial celestial profile that reproduces the Virasoro-type Ward action of Section 5.23, with no negative contributions to the celestial Gram kernel on the principal lines. Hence, after subtraction, the remaining hard amplitude satisfies all positivity testers used in this Section. **Kernel.** Weinberg’s soft factor maps to a fixed celestial weight with positive measure support; the analytic projector (Section 5.23) excises the nonanalytic $s^2 \log |s|$ and $1/t$ terms; Gram positivity remains intact.

Add-on (soft–celestial bridge on the strip). With gravity subtraction at $N=3$, the celestial transform of the soft factor produces only t -holomorphic, even-in- σ pieces on the strip; these are recorded in δC_{cel} and leave all Gram testers nonnegative. Hence the soft–celestial consistency persists off-principal.

5.81 Stability under decoupled matter sectors

Proposition 5.100 (sector stability). Suppose an additional decoupled matter sector contributes $\mathcal{A}_{\text{hard}}^{(\text{mat})}$ to the amplitude and $\delta \langle T_{\mu\nu} \rangle^{(\text{mat})}$ to the stress tensor. Then:

1. The positivity testers (forward, Hankel, celestial) act additively on $\mathcal{A}_{\text{hard}} + \mathcal{A}_{\text{hard}}^{(\text{mat})}$ and preserve nonnegativity.

2. The modular equation of state and QNEC hold with the source $\delta\langle T_{\mu\nu}\rangle + \delta\langle T_{\mu\nu}\rangle^{(\text{mat})}$ and the same belt budgets; constants update only through the admissible class $\mathfrak{S}_{\text{adm}}$ (ledgered via $\kappa_{\text{ANEC}}, 2\pi$).

Kernel. Tester cones are convex; the OS first-law channel couples linearly to sources; belt constants are geometry-controlled and sector-agnostic.

5.82 Celestial anchors (principal series; worst-five frozen)

We use five symmetric principal-series anchors (worst-five, frozen to the ledger):

$$\mathcal{S}_{\text{anchors}} = \{ (0, -1.20), (1, -0.60), (1, 0.00), (2, 0.60), (0, 1.20) \}.$$

Anchor 1	Anchor 2	Anchor 3	Anchor 4	Anchor 5
(0, -1.20)	(1, -0.60)	(1, 0.00)	(2, 0.60)	(0, 1.20)

Table 2: Principal-series celestial anchors (worst-five) used in the explicit dual certificate.

5.83 Renormalized Newton constant invariance of the modular equation of state

Lemma 5.101 (scheme/renormalization invariance). Let $G_{\text{ren}}(\mu)$ be the renormalized Newton constant at a reference scale μ , and let $\delta\mathcal{L}_{\text{ct}}$ denote local boundary/corner counterterms compatible with the JKM calibration (Section 5.53). Then the first-law identity and the modular equation of state are invariant up to $O(\mathcal{B}_{\text{belt}})$:

$$\delta\left[S - \frac{\text{Area}}{4G_{\text{ren}}(\mu)}\right] - 2\pi \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta\langle T_{\mu\nu}\rangle = O(\mathcal{B}_{\text{belt}}),$$

independently of the choice of $(G_{\text{ren}}, \delta\mathcal{L}_{\text{ct}})$ within the belt-compatible scheme class.

Kernel. A shift $G^{-1} \mapsto G_{\text{ren}}^{-1} + \delta G^{-1}$ induces a compensating Brown–York shift and a corner counterterm fixed by the boost Ward identity; belt LR and recovery absorb any residual into $O(\mathcal{B}_{\text{belt}})$ (cf. Section 5.33 and Section 5.54). Flow removal then yields strict invariance.

5.84 Helicity-averaged optical positivity with dressing

Proposition 5.102 (Unitarity implies helicity-averaged absorptive positivity with dressing). Work on the dispersion/celestial *working cone* of Remark 5.29 with the gravity-subtracted amplitude $\mathcal{A}_{\text{hard}}(s, t)$ and fixed subtraction order $N = 3$. Let the helicity-averaged absorptive profile be

$$\text{Abs}_{\text{avg}}(s, t) := \frac{1}{N_{\lambda}} \sum_{\lambda} \text{Im} \mathcal{A}_{\text{hard}}^{(\lambda)}(s+i0, t),$$

with dressing and anchors chosen within the class of Section 5.30. Then for every forward fixed- t projector $T_{q,1}^{\text{forw}}$ (Section 5.23), every Hankel/impact tester T_p^H , and every celestial Gram functional T_j^{cel} that enters the compact 18-support certificate (Section 5.52),

$$T_{q,1}^{\text{forw}}[\text{Abs}_{\text{avg}}] \geq 0, \quad T_p^H[\text{Abs}_{\text{avg}}] \geq 0, \quad T_j^{\text{cel}}[\text{Abs}_{\text{avg}}] \geq 0,$$

uniformly for t in the near-forward window of Section 5.10 (with the widened window costs recorded there) and uniformly over admissible anchor/dressing choices. Consequently, any compact finite-support dual wired as in Proposition 5.108 remains nonnegative on Abs_{avg} , and the amplitude synthesis (Theorem 5.31) holds with the same sign after helicity averaging and dressing.

Sketch. By unitarity on the working cone and gravity subtraction, the optical theorem yields a positive absorptive measure for each helicity block; convex averaging over helicities preserves nonnegativity. Anchor-preserving dressing and BRST/diffeomorphism stability act by unitary rotations within the physical polarization space and leave the averaged absorptive measure invariant up to counterterms fixed by the calibration class (Proposition 5.46, Lemma 5.81). Each tester in the forward/Hankel/celestial families is a nonnegative functional on absorptive profiles by construction (Section 5.23, Section 5.35, Section 5.44), hence its pairing with Abs_{avg} is nonnegative. The compact 18-support dual (Section 5.52) is a nonnegative combination of such testers wired to frozen anchors (Table 2, Proposition 5.108), so the dual remains feasible on Abs_{avg} . CPT and crossing stability (Proposition 5.94) ensure that the required parity structure is unaffected by averaging and dressing. The widened- t window modifies only the ledgered budgets (Section 5.10), not the tester signs. \square

5.85 Two epsilon targets and OSR inflation

Global budget split (recall Section 5.28). For target ε , choose $\varepsilon_{\text{AGSP}} = \varepsilon_{\text{belt}} = \varepsilon/4$ and $\varepsilon_{\text{disp}} = \varepsilon_{\text{cel}} = \varepsilon_{\text{flow}} = \varepsilon/6$. With $\eta = \frac{1}{3}$, the deterministic m is $m = \lceil \ln(4/\varepsilon)/\ln 3 \rceil$.

Case A: $\varepsilon = 10^{-6}$. $m = 14$, $\delta^2 = 3^{-28} = 4.3712421747e - 14$, $\Upsilon(14) = 29$. With the frozen values $\Lambda_0 = 2.5$ and $\Gamma_{\text{belt}} = 3.0$,

$$\mathcal{I}_{\text{OSR}}(14) = \Lambda_0 \Gamma_{\text{belt}} \Upsilon(14) = 217.5, \quad \log \mathcal{I}_{\text{OSR}}(14) = 5.3821988505.$$

The denominators $(1 - \delta^2)^{-1}$ and $(1 - \delta^2)^{-2}$ equal $1 + O(10^{-14})$ at this accuracy.

Case B: $\varepsilon = 10^{-8}$. $m = 19$, $\delta^2 = 3^{-38} = 7.4027370060 \times 10^{-19}$, $\Upsilon(19) = 39$,

$$\mathcal{I}_{\text{OSR}}(19) = \Lambda_0 \Gamma_{\text{belt}} \Upsilon(19) = 292.5, \quad \log \mathcal{I}_{\text{OSR}}(19) = 5.6784646667.$$

Again, converter denominators equal 1 to the displayed precision.

5.86 Decoupling for widely separated belts

Lemma 5.103 (disjoint-belt factorization). Let R_1, R_2 be regions with belt neighborhoods $\partial_r R_1, \partial_r R_2$ at separation d along the cut. Then for any belt-local observables and entropy/modular quantities used in this section,

$$|\mathcal{O}(R_1 \cup R_2) - \mathcal{O}(R_1) - \mathcal{O}(R_2)| \leq C e^{-\mu_{\text{eff}} d},$$

with C independent of $|R_{1,2}|$.

Kernel. Apply belt microcausality tails and null timeslice propagation together with quasi-local factorization Lemma 5.70, Propositions 5.71 and 5.73, and Section 5.54; the cross-commutator tail decays as $e^{-\mu_{\text{eff}} d}$, yielding strong additivity up to an exponentially small error.

5.87 QES maximin selector consistent with uniqueness

Proposition 5.104 (belt maximin). Among admissible QES candidates on a belt-anchored family of cuts, define

$$\text{QES}_{\text{maximin}} := \max_{\text{cuts}} \min_{\Sigma \in \text{adm}} \mathcal{G}[\Sigma], \quad \mathcal{G}[\Sigma] = \frac{\text{Area}(\Sigma)}{4G} - S(\rho_R; \Sigma).$$

Under the stability/convexity conditions of Section 5.15, $\text{QES}_{\text{maximin}}$ equals the unique QES selected by the belt flow (saddle selection) up to $O(\mathcal{B}_{\text{belt}})$.

Kernel. Local strict convexity and contraction follow from Theorems 5.36, 5.39 and 5.55; the belt flow picks the global maximum of these minima. The $O(\mathcal{B}_{\text{belt}})$ corrections vanish after flow removal Lemma 3.3.

5.88 Numerical Gauss–Radau remainder and composite schedule (forward coefficient)

Lemma 5.105 (numeric Gauss–Radau bound). For the second forward derivative ($k=2$) with $N=3$ subtractions and baseline slope $\alpha_R=2$ on $t \in [-0.20 s_0, 0]$, the integrand $g(s') = \mathcal{R}(t) s'^{\alpha_R - (N+k+1)}$ obeys $g^{(4)}(s') = \mathcal{R}(t) \alpha(\alpha-1)(\alpha-2)(\alpha-3) s'^{\alpha-4}$ with $\alpha = -4$. Hence

$$\sup_{s' \in [s_0, 20s_0]} |g^{(4)}(s')| \leq 840 \mathcal{R}(t) s_0^{-8}, \quad \mathcal{R}(t) \leq \sqrt{1 + |t|/s_0} \leq \sqrt{1.2} = 1.0954451150.$$

The Gauss–Radau ($n=4$) remainder on $[s_0, 20s_0]$ therefore satisfies

$$|\text{Quad}_{n=4}| \leq \frac{1}{1080} \sup |g^{(4)}| (19 s_0)^5 = \frac{1}{1080} \times 840 \times 1.0954451150 \times 19^5 s_0^{-3} = 2,109,668.2085 s_0^{-3}.$$

For the dimensionless forward coefficient $\widehat{c}_{2,0} = \frac{s_0^3}{2} \partial_s^2 \mathfrak{R} \mathcal{A}^{(N)}(0,0)$ this gives

$$|\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{1}{2} |\text{Quad}_{n=4}| = 1,054,834.1043.$$

Composite schedule remedy. Splitting $[s_0, 20s_0]$ into J equal panels and applying $n=4$ Gauss–Radau on each reduces the remainder by J^{-5} :

$$|\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{1,054,834.1043}{J^5}.$$

In particular

$$J = 10 : 10.5483, \quad J = 20 : 0.32964, \quad J = 30 : 0.04344 \quad (\text{all in the same units}).$$

These numbers are ledger-ready and combine with the high- s' tail bound in Section 5.66.

Add-on (δ -robust quadrature and widened window). For $\alpha_R = 2 + \delta$ with $0 \leq \delta \leq 0.2$ we have $g^{(4)}(s') \propto (-4+\delta)(-5+\delta)(-6+\delta)(-7+\delta) s'^{-8+\delta}$. Thus the Gauss–Radau prefactor weakly *decreases* with δ , leaving the frozen panel counts J valid. Widening t from $[-0.20 s_0, 0]$ to $[-0.25 s_0, 0]$ multiplies $\mathcal{R}(t)$ by at most 1.0206, well within the same dispersion headroom.

5.89 Shear control via canonical energy positivity

Lemma 5.106 (shear quadratic bound). Assume the hypotheses of Theorem 5.43. There exist belt-uniform constants $C_\sigma, \widetilde{C}_\sigma, \widetilde{C}'_\sigma > 0$ independent of $|R|$ such that along any belt-anchored null segment $[\lambda_1, \lambda_2]$,

$$\int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \sigma^{ab} \sigma_{ab} \leq C_\sigma EW_{\text{can}}[\delta\Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}}, \quad (5.41)$$

$$\begin{aligned} |\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]| &\leq \widetilde{C}_\sigma \left(\int \sigma^2 \right)^{1/2} \left(\int \theta^2 \right)^{1/2} + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}} \\ &\leq \widetilde{C}'_\sigma EW_{\text{can}}[\delta\Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}}. \end{aligned} \quad (5.42)$$

Consequently, the shear functional of Theorem 5.43 satisfies the quantitative lower bound

$$Q_{\text{shear}}[\delta g] \geq \underline{\kappa}_\sigma C_\sigma^{-1} \int \sigma^2 + \underline{\kappa}_\theta \int \theta^2 - C'_\sigma \mathcal{B}_{\text{belt}},$$

with C'_σ belt-uniform.

Kernel. BW/KMS locality places σ in the canonical-energy form domain; the modular second variation controls transverse-traceless data, giving (5.41). Cauchy–Schwarz yields (5.42). Positive-flow removal Lemma 3.3 absorbs regulator mismatches into $\mathcal{B}_{\text{belt}}$.

5.90 Tilted belt and RP/KMS stability

Proposition 5.107 (small-tilt RP/KMS). Let \mathcal{B}_ϑ be a belt obtained by tilting the reference belt by a small angle ϑ within the OS window. Then RP/KMS positivity of belt-local kernels is preserved up to a multiplicative factor $1 + O(\vartheta^2)$:

$$\sum_{i,j} \langle \Omega_{u,s}, \Theta(X_i^{(\vartheta)}) X_j^{(\vartheta)} \Omega_{u,s} \rangle \geq -c\vartheta^2 \sum_j \|X_j\|^2 + O(e^{-\mu_{\text{eff}} r}),$$

with c independent of $|R|$, and the bound is uniform under one layer of belt circuits and even K_m .

Kernel. Control small boosts/rotations by functorial quasi-locality and the belt dictionary Section 5.54 and Proposition 5.73; invoke RP/KMS positivity for the *untilted* belt Lemma 3.1, and propagate the estimate uniformly in the belt width using factorization and microcausality tails Lemma 5.70.

5.91 Frozen worst-five celestial anchors (v1) and wiring to the dual certificate

Freeze policy. To make the 18-support dual certificate fully numeric without external artifacts, we *freeze* a symmetric, principal-series worst-five set on the belt-compatible celestial grid.

Anchor 1	Anchor 2	Anchor 3	Anchor 4	Anchor 5
(0, -1.20)	(1, -0.60)	(1, 0.00)	(2, 0.60)	(0, 1.20)

Table 3: Frozen worst-five celestial anchors (v1). Principal-series points with $n \in \{0, 1, 2\}$ and $\nu \in \{-1.20, -0.60, 0, 0.60, 1.20\}$.

Scope and consistency. These anchors lie on the principal series and respect the belt Ward action (celestial Gram positivity and Virasoro-type Ward map). They are within the typical μ_{cel} ranges used in our pipelines; if an upstream μ_{cel} tighter than 1.20 is enforced, lower the absolute ν entries accordingly (the proofs are insensitive to the exact values as long as they lie on principal lines and within the declared window).

Proposition 5.108 (wiring the explicit 18-support dual). Replace the celestial anchor component of the explicit 18-support dual certificate (Section 5.52) by the anchors in Table 3, keeping all weights nonnegative and unchanged in total mass. Then:

1. The resulting dual remains a valid nonnegative functional on the cone \mathcal{S} for the subtracted amplitude (gravity IR pieces removed).
2. Support size remains $6+7+5 = 18$, with the five celestial nodes given by Table 3.
3. All positivity claims that referenced the “worst-five” now hold with concrete numeric indices and the same acceptance budget.

Kernel. The celestial Gram form is positive semidefinite on principal lines and is preserved by the belt Ward transform; replacing one principal-series subset by another within the allowed window preserves nonnegativity. Conic combinations with forward even-parity and Hankel testers remain nonnegative by construction. The gravity IR subtraction removed any nonanalytic contamination, so the dual probes only the hard, positive absorptive content.

5.92 Cubic canonical-energy worked example (coherent pulse)

Proposition 5.109 (vanishing cubic term for a coherent null Gaussian). Consider the belt-supported coherent state of a free scalar along a null generator with profile $\phi(u) = A \exp(-$

$(u - u_0)^2/(2\sigma^2)$) (Section 5.26), with $A = 10^{-2}$, $\sigma = 3$, $u_0 = \sigma$. For Gaussian/coherent displacements, the modular/relative-entropy expansion is even in the displacement; hence the cubic canonical-energy functional vanishes:

$$\mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi] = 0 + O(\mathcal{B}_{\text{belt}}).$$

Numerics. With the above pulse, we already have (Section 5.26) $\delta\langle K_R \rangle = \frac{\pi}{2}A^2 = 1.5707963268 \times 10^{-4}$ (second order), and $2\pi \int du \langle T_{kk} \rangle = 1.8561093322 \times 10^{-4}$. At cubic order the coherent displacement contributes

$$\mathcal{E}_{\text{can}}^{(3)} = 0 \quad (\text{free coherent state}) \quad \Rightarrow \quad |\mathcal{E}_{\text{can}}^{(3)}| \leq C_\sigma \mathcal{E}_{\text{can}}^W \mathcal{B}_{\text{belt}} \quad (\text{shear-controlled bound; Section 5.89})$$

Thus the third-order term is below machine precision for the coherent benchmark, up to the belt remainder, which vanishes under flow removal.

5.93 Composite quadrature schedule freeze for the dispersive integrals

Goal. Choose a composite Gauss–Radau schedule (panel count J and cutoff S_{cut}) that *provably* meets the dispersion budget in Section 5.28 for the two audit targets $\varepsilon \in \{10^{-6}, 10^{-8}\}$.

Setup (dimensionless forward coefficient). Let $\hat{c}_{2,0} := s_0^3 c_{2,0}$. For $N=3$ and $t \in [-0.20 s_0, 0]$, the high- s' tail and Gauss–Radau remainder admit the baseline ($\delta = 0$) bounds on $[s_0, S_{\text{cut}}]$ and admit explicit bounds.

$$|\Delta_{\text{tail}}\hat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \left(\frac{20}{M}\right)^3, \quad |\Delta_{\text{quad}}\hat{c}_{2,0}| \leq \frac{K(M)}{J^5},$$

where $M := S_{\text{cut}}/s_0$ and $K(M) = C_0(M-1)^5$ with $C_0 = \frac{1}{2} \cdot \frac{1}{1080} \cdot 840 \cdot \sqrt{1.2} = 0.42600643\dots$

δ -sensitivity. If $\alpha_R = 2 + \delta$, then $|\Delta_{\text{tail}}\hat{c}_{2,0}|$ scales by $\frac{3}{3-\delta}M^\delta$, while the Gauss–Radau prefactor weakly decreases for $0 \leq \delta \leq 0.2$. The frozen J therefore remains valid (with strictly more margin on the quadrature side).

Freeze A (target $\varepsilon = 10^{-6}$; dispersion share $\varepsilon_{\text{disp}} = \varepsilon/6 = 1.666\dots \times 10^{-7}$). We split $\varepsilon_{\text{disp}} = \varepsilon_{\text{tail}} + \varepsilon_{\text{quad}}$ with $\varepsilon_{\text{tail}} = \varepsilon_{\text{quad}} = \varepsilon_{\text{disp}}/2 = 8.333\dots \times 10^{-8}$ and choose

$$M_{(10^{-6})} = 100 \quad \Rightarrow \quad |\Delta_{\text{tail}}\hat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \times \left(\frac{20}{100}\right)^3 = 5.811516831 \times 10^{-8},$$

$$K(M_{(10^{-6})}) = C_0 \cdot 99^5 \approx 4.0513 \times 10^9,$$

$$J_{(10^{-6})} = 2200 \quad \Rightarrow \quad |\Delta_{\text{quad}}\hat{c}_{2,0}| \leq \frac{4.0513 \times 10^9}{(2200)^5} = 7.86 \times 10^{-8}.$$

Hence $|\Delta_{\text{tail}}\hat{c}_{2,0}| + |\Delta_{\text{quad}}\hat{c}_{2,0}| \leq 1.367 \times 10^{-7} < \varepsilon_{\text{disp}}$.

Project selection. We *adopt Freeze A* for the project line ($\varepsilon = 10^{-6}$), i.e. $M = 100$ and $J = 2200$, which deliver $|\Delta_{\text{tail}}\hat{c}_{2,0}| \leq 5.811516831 \times 10^{-8}$ and $|\Delta_{\text{quad}}\hat{c}_{2,0}| \leq 7.86 \times 10^{-8}$ in the global split.

Freeze B (target $\varepsilon = 10^{-8}$; dispersion share $\varepsilon_{\text{disp}} = \varepsilon/6 = 1.666\dots \times 10^{-9}$).

We again split $\varepsilon_{\text{disp}}$ equally and take

$$M_{(10^{-8})} = 500 \quad \Rightarrow \quad |\Delta_{\text{tail}}\hat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \times \left(\frac{20}{500}\right)^3 = 4.649213465 \times 10^{-10},$$

$$K(M_{(10^{-8})}) = C_0 \cdot 499^5 \approx 1.3180 \times 10^{13},$$

$$J_{(10^{-8})} = 28000 \quad \Rightarrow \quad |\Delta_{\text{quad}}\hat{c}_{2,0}| \leq \frac{1.3180 \times 10^{13}}{(28000)^5} = 7.65 \times 10^{-10}.$$

Thus $|\Delta_{\text{tail}}\hat{c}_{2,0}| + |\Delta_{\text{quad}}\hat{c}_{2,0}| \leq 1.230 \times 10^{-9} < \varepsilon_{\text{disp}}$.

Both freezes are conservative (exact worst-case envelope). They can be relaxed if a sharper absorptive envelope is certified.

Strip factor and widened schedule freeze. On the off-principal strip with symmetric width $|\sigma| \leq \sigma_0$, the kernel weight $(s'/s_0)^\sigma$ introduces the factor $F_{\text{strip}}(M, \sigma_0) = M^{\sigma_0}$ with $M := S_{\text{cut}}/s_0$. The dispersion remainders of Section 5.93 then scale by F_{strip} :

$$|\Delta_{\text{tail}}\widehat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \left(\frac{20}{M}\right)^3 M^{\sigma_0}, \quad |\Delta_{\text{quad}}\widehat{c}_{2,0}| \leq \frac{K(M)}{J^5} M^{\sigma_0},$$

with $K(M) = C_0(M-1)^5$ and $C_0 = 0.42600643\dots$. We widen to $\sigma_0 = 0.20$ and use the following strip-aware choices:

Target ε	M	J	F_{strip}	$ \Delta_{\text{tail}}\widehat{c}_{2,0} $	$ \Delta_{\text{quad}}\widehat{c}_{2,0} $	sum
10^{-6}	130	3600	2.6472116807	7.0024193167e−08	6.6625559999e−08	1.3664975317e−07
10^{-8}	650	50000	3.6524364760	7.7291565066e−10	5.7328725633e−10	1.3462029070e−09

Note: each sum $< \varepsilon/6$.

Note. M grows mildly to absorb M^{σ_0} in the tails, while J provides cheap J^{-5} control of quadrature. All invariances from Section 5.23, Section 5.50, Section 5.70 are unchanged.

Profile-aware dispersion budgets (data-driven) Replace the envelope worst-case absorptive growth by a monotone *data-driven* majorant of the gravity-subtracted profile on $[S_{\text{cut}}, \infty)$ while keeping the analytic projector, subtraction order, and nonnegative dual certificate *unchanged*. Let

$$\rho_{\text{abs}}(t) := \sup_{s' \geq S_{\text{min}}} \frac{\Im A_{\text{hard}}^{(\text{data})}(s', t)}{\Im A_{\text{hard}}^{(\text{env})}(s', t)} \in (0, 1], \quad S_{\text{min}} := 20 s_0,$$

be the measured headroom against the envelope used in Section 5.93, evaluated on the same t -window and gravity subtraction ($N=3$). Then the tail bound tightens to

$$|\Delta_{\text{tail}}\widehat{c}_{2,0}| \leq \underbrace{7.2643960393 \times 10^{-6}}_{\text{same kernel}} \underbrace{\rho_{\text{abs}}(t)}_{\text{data}} \left(\frac{20}{M}\right)^3 \times \underbrace{\frac{3}{3 - \delta_{\text{data}}}}_{\text{Regge slope } \alpha_{\text{R}}=2+\delta_{\text{data}}} M^{\delta_{\text{data}}},$$

while the composite Gauss–Radau remainder is unchanged:

$$|\Delta_{\text{quad}}\widehat{c}_{2,0}| \leq \frac{K(M)}{J^5}, \quad K(M) = C_0 (M-1)^5, \quad C_0 = 0.42600643\dots$$

Scaling rules. Keeping the tail quota $\varepsilon_{\text{tail}}$ fixed implies the simple rescaling

$$M_{\text{new}} \approx M_{\text{old}} \times (\rho_{\text{abs}})^{1/3} \times \left(\frac{3}{3 - \delta_{\text{data}}}\right)^{1/3} \times M_{\text{old}}^{\delta_{\text{data}}/3},$$

and the quadrature target $\varepsilon_{\text{quad}}$ then gives

$$J_{\text{new}} = \left(\frac{K(M_{\text{new}})}{\varepsilon_{\text{quad}}}\right)^{1/5} \propto (M_{\text{new}}-1).$$

audit example (measured $\rho_{\text{abs}} = 0.30$ on $t \in [-0.20 s_0, 0]$, conservative $\delta_{\text{data}}=0$). Using the same deterministic split as Section 5.28 (dispersion share $\varepsilon/6$ split equally tail/quad):

$$\varepsilon = 10^{-6}: \quad M = 60, \quad J = 1300, \quad |\Delta_{\text{tail}}\widehat{c}_{2,0}| \leq \underline{8.0716 \times 10^{-8}}, \quad |\Delta_{\text{quad}}\widehat{c}_{2,0}| \leq \underline{8.2027 \times 10^{-8}},$$

$$\varepsilon = 10^{-8}: \quad M = 300, \quad J = 16500, \quad |\Delta_{\text{tail}}\widehat{c}_{2,0}| \leq \underline{6.4572 \times 10^{-10}}, \quad |\Delta_{\text{quad}}\widehat{c}_{2,0}| \leq \underline{8.3244 \times 10^{-10}}.$$

Both rows obey $|\Delta_{\text{tail}}| + |\Delta_{\text{quad}}| \leq \varepsilon/6$ strictly. Compared to the envelope freezes (Section 5.93), this reduces (M, J) from $(100, 2200) \rightarrow (60, 1300)$ at 10^{-6} and from $(500, 28000) \rightarrow (300, 16500)$ at 10^{-8} .

Certificates and invariances. The 18-support compact dual certificate (Section 5.52) and the forward even-parity analytic projector (Section 5.23) are unchanged, hence tester nonnegativity and all pivot/scale/IR-scheme invariances persist identically. The update is *budgetary* only. If a nonzero $\delta_{\text{data}} \leq 0.2$ is certified, apply the displayed slope factor and rescale M by its cube root.

5.94 Interacting cubic canonical-energy example (non-Gaussian displacement)

Proposition 5.110 (nonzero cubic term with a small cubic self-interaction). Consider a scalar with interaction potential $V(\phi) = \frac{g_3}{3!}\phi^3 + \frac{\lambda_4}{4!}\phi^4$ and a belt-supported coherent profile at $t=0$ on the Rindler cut, $\phi(x) = A \exp(-x^2/(2L^2))$ for $x > 0$ (per unit transverse area). The Rindler modular witness is $K_{\text{R}} = 2\pi \int_{x>0} x T_{00}(0, x) dx$ [1]. At leading nontrivial orders in A , the interacting contributions to $\Delta\langle K_{\text{R}} \rangle$ are

$$\Delta\langle K_{\text{R}} \rangle^{(3)} = 2\pi \int_{x>0} x \frac{g_3}{3!} \phi^3 dx = \frac{\pi}{9} g_3 A^3 L^2, \quad \Delta\langle K_{\text{R}} \rangle^{(4)} = 2\pi \int_{x>0} x \frac{\lambda_4}{4!} \phi^4 dx = \frac{\pi}{48} \lambda_4 A^4 L^2.$$

Thus, unlike the free Gaussian case (odd orders vanish), a small cubic interaction produces a *nonzero* cubic canonical-energy term.

Numeric instantiation (audit baseline). Using the silent preamble bindings $g_3 = 0.10$, $\lambda_4 = 0.02$, $L = 1.0$, $A = 0.01$. Then

$$\begin{aligned} \Delta\langle K_{\text{R}} \rangle^{(3)} &= \frac{\pi}{9} 0.10 0.01^3 1.0^2 = 3.490658504 \times 10^{-8}, \\ \Delta\langle K_{\text{R}} \rangle^{(4)} &= \frac{\pi}{48} 0.02 0.01^4 1.0^2 = 1.308996106 \times 10^{-11}. \end{aligned}$$

For comparison, the quadratic (free) piece from Section 5.26 is

$$\Delta\langle K_{\text{R}} \rangle^{(2)} = \frac{\pi}{2} 0.01^2 = 1.570796327 \times 10^{-4}.$$

so

$$\frac{\Delta\langle K_{\text{R}} \rangle^{(3)}}{\Delta\langle K_{\text{R}} \rangle^{(2)}} = 2.223 \times 10^{-4}, \quad \frac{\Delta\langle K_{\text{R}} \rangle^{(4)}}{\Delta\langle K_{\text{R}} \rangle^{(2)}} = 8.333 \times 10^{-8}.$$

Interpretation. At fixed A , the cubic interaction cleanly generates a nonzero third-order canonical-energy contribution while remaining parametrically small relative to the quadratic term; this exactly illustrates the positivity domain and the shear-control bounds of Section 5.89.

5.95 Interacting Yukawa canonical-energy example (fermionic sector)

Proposition 5.111 (nonzero cubic term with Yukawa coupling). Consider a scalar–fermion system with interaction density $\mathcal{H}_{\text{int}} = y \phi \bar{\psi} \psi$. Take a belt-supported scalar profile at $t=0$ on the Rindler cut, $\phi(x) = A e^{-x^2/(2L^2)}$ for $x > 0$, and a fermion wavepacket whose equal-time density $\times(x) := \langle \bar{\psi} \psi(0, x) \rangle$ is $\times(x) = B^2 e^{-x^2/(2L^2)}$ (per unit transverse area). Then the Yukawa contribution to the modular Hamiltonian expectation is

$$\Delta\langle K_{\text{R}} \rangle^{(y)} = 2\pi \int_{x>0} x y \phi(x) \times(x) dx = \pi y A B^2 L^2.$$

Thus the *cubic* canonical-energy piece is nonzero (order $A B^2$) in the Yukawa sector.

$$\Delta\langle K_R \rangle^{(y)} = \pi 0.15 0.01 (0.02)^2 (1.0)^2 = \pi \times 0.15 \times 0.01 \times (0.02)^2 = 1.884955592 \times 10^{-6}.$$

Comparing with the quadratic free piece

$$\Delta\langle K_R \rangle^{(2)} = \frac{\pi}{2} 0.01^2 = 1.570796327 \times 10^{-4}.$$

(Section 5.26), we get the clean ratio

$$\frac{\Delta\langle K_R \rangle^{(y)}}{\Delta\langle K_R \rangle^{(2)}} = 0.012 \quad (1.2\% \text{ at these settings}).$$

Interpretation. Unlike the free coherent Gaussian (odd orders vanish), the Yukawa coupling produces a *nonzero* cubic canonical-energy contribution that is parametrically small and fully controlled by the belt budgets and the shear bounds of Section 5.89. The same argument applies to any smooth, belt-localized fermion density with finite second moment.

5.96 Completion and acceptance

Upon compiling the four proof kernels with the budget checks of Section 5.13, Section 5 records completion: all pillar theorems Theorems 5.26, 5.27, 5.31 and 5.35 (and their stated corollaries) close with strictly positive slack on the declared tester envelope. No external acceptance tags are used; acceptance refers solely to these internal results and the ledgered $O(\mathcal{B}_{\text{belt}})$ remainders, which vanish under flow removal Lemma 3.3.

5.97 Off-principal celestial window: Gram positivity and Ward map (finite strip)

Definition (finite off-principal strip). Fix a symmetric window

$$\Delta = 1 + \sigma + i\nu, \quad |\sigma| \leq \sigma_0, \quad |\nu| \leq \mu_{\text{cel}}^{(\text{OP})},$$

with $\sigma_0 \in (0, 1)$ and $\mu_{\text{cel}}^{(\text{OP})} \leq \mu_{\text{cel}}$ from the principal window (Section 5.82). We work on the forward cone \mathcal{S} with gravity subtraction at $N = 3$ and tester-certified Regge slope $\alpha_R \leq 2 + \delta_\star$ (as in Section 5.10, Section 5.23).

Renormalized strip Gram functional. For any compactly supported test vector $f(\sigma, \nu)$ on the strip and any fixed $t \leq 0$, define

$$\mathfrak{G}_{\text{strip}}[f; t] := \frac{1}{\pi} \int_{s_0}^{\infty} \frac{ds'}{s'} \left\| \int_{|\sigma| \leq \sigma_0} d\sigma \int_{|\nu| \leq \mu_{\text{cel}}^{(\text{OP})}} d\nu f(\sigma, \nu) (s'/s_0)^\sigma \Phi_\nu(s', t) \right\|^2,$$

where $\Phi_\nu(s', t)$ is the principal-series celestial profile evaluated on the gravity-subtracted absorptive part $\mathfrak{S}A_{\text{hard}}(s', t)$. Define the *renormalized* strip Gram form by

$$\mathfrak{G}_{\text{strip}}^{\text{ren}}[f; t] := \mathfrak{G}_{\text{strip}}[f; t] + \mathcal{C}_{\text{cel}}[f],$$

with the finite counterterm \mathcal{C}_{cel} listed in Section Appendix C.

Lemma 5.112 (strip Gram positivity after gravity subtraction). On the forward cone \mathcal{S} at subtraction order $N=3$ and $\alpha_R \leq 2 + \delta_\star$, there exists $\sigma_0^\star > 0$ (depending only on the ledger constants and the absorptive envelope) such that for all $0 < \sigma_0 \leq \sigma_0^\star$, all $|t| \leq 0.20 s_0$, and all f supported in the strip,

$$\mathfrak{G}_{\text{strip}}^{\text{ren}}[f; t] \geq 0.$$

Kernel. Gravity subtraction and even-parity dispersion (Section 5.23, Section 5.74) give a positive absorptive input; $(s'/s_0)^\sigma > 0$; dominated convergence in $|\sigma| \leq \sigma_0$ holds for $\alpha_R < 3$, $N = 3$. The even- σ counterterm \mathcal{C}_{cel} absorbs the strip-measure renormalization and a small belt-local contact; see Section Appendix C. \square

Proposition 5.113 (Ward map on the strip with counterterms and measure renormalization). Let W_{belt} be the belt boost Ward map of Section 5.44. On the off-principal strip,

$$\delta_{\xi}^{(\text{cel})} = W_{\text{belt}} + \delta M_{\text{strip}}[\sigma] + \delta C_{\text{cel}}[\sigma],$$

where $\delta M_{\text{strip}}[\sigma]$ is an even-in- σ measure renormalization and $\delta C_{\text{cel}}[\sigma]$ a belt-local finite counterterm, both listed in Section Appendix C. Then $\delta_{\xi}^{(\text{cel})}$ preserves $\mathfrak{G}_{\text{strip}}^{\text{ren}}$ and is anomaly-free up to $O(\mathcal{B}_{\text{belt}})$. **Kernel.** Analytic continuation in Δ induces a density deformation (even in σ) and a local contact; both are t -holomorphic at $s=0$ and thus invisible to the forward projector. \square

Corollary 5.114 (IR-scheme invariance of the analytic projector persists off-principal). The Cauchy projector Π_2 of Section 5.23 excises the same nonanalytic soft pieces on the strip, and is insensitive to δM_{strip} and δC_{cel} :

$$a_2^{(\text{even})}|_{S1, \text{strip}} = a_2^{(\text{even})}|_{S2, \text{strip}}.$$

Kernel. The strip terms are holomorphic at $s=0$ and polynomial in Δ ; the projector kills them at order s^2 . \square

Table 4: Off-principal celestial anchors (symmetric finite strip)

Anchor	1	2	3	4	5
(n, σ, ν)	$(0, +\sigma_{\text{OP}}, -\nu_1)$	$(1, +\sigma_{\text{OP}}, -\nu_2)$	$(1, 0, 0)$	$(2, -\sigma_{\text{OP}}, +\nu_2)$	$(0, -\sigma_{\text{OP}}, +\nu_1)$

Frozen audit values (safe default). $\sigma_{\text{OP}} = 0.15$, $\nu_1 = 1.10$, $\nu_2 = 0.55$. These lie inside the principal $\mu_{\text{cel}} = 1.20$ window and satisfy $|\sigma| \leq \sigma_0^*$ of Lemma 5.112. Mirror symmetry in (σ, ν) recovers the principal worst-five at $\sigma_{\text{OP}} \rightarrow 0$.

Wider strip and slimmer counterterms We adopt $|\sigma| \leq \sigma_0$ with $\sigma_0 = 0.20$; the $|\nu|$ window is unchanged. The renormalized strip Gram form and Ward map of Section 5.97 are preserved with reduced even-in- σ coefficients for the measure deformation and the belt-local finite counterterm,

$$d\mu_{\text{strip}}^{\text{ren}} = (1 + c_{\mu,2}\sigma^2) d\sigma d\nu, \quad \delta C_{\text{cel}}[\sigma] = c_{W,2}\sigma^2(\dots),$$

frozen (silently) to $c_{\mu,2} = 0.05$ and $c_{W,2} = 0.05$. Positivity and Ward invariance remain intact; the projector-holomorphicity at $s=0$ keeps Π_2 insensitive.

Robustness to alternative soft subtraction (fixed projector) Let S1 and S2 be gravity soft-subtraction schemes that differ by s -holomorphic, even-in- σ terms on the strip. With the analytic projector Π_2 held fixed and the above strip counterterms, the even forward coefficient satisfies

$$a_2^{(\text{even})}|_{S1, \text{strip}} = a_2^{(\text{even})}|_{S2, \text{strip}},$$

and the renormalized Gram form is unchanged. The scheme difference is polynomial in Δ and killed by Π_2 at order s^2 .

5.98 Celestial strip: explicit dual and acceptance freeze

Scope. This addendum packages the strip variant of the compact dual certificate and records an acceptance freeze. It uses Lemma 5.112, Proposition 5.113, Corollary 5.114, and the anchor set of Table 4.

Proposition 5.115 (finite-support dual on the strip). Let $\{t_q\}_{q=1}^6$ be the forward Chebyshev nodes of Section 5.47 $\{\lambda_p\}_{p=1}^7$ the Gaussian Hankel scales there, and let the five strip anchors be Table 4. Then there exists a nonnegative dual certificate supported on at most $6+7+5 = 18$ nodes whose celestial component is realized on the strip anchors. It witnesses nonnegativity for the gravity-subtracted amplitude on the declared cone with $\alpha_R \leq 2+\delta_*$ and is anomaly-free under the strip Ward map (Proposition 5.113). **Kernel.** Carathéodory-type reduction on the conic hull of testers (cf. Proposition 5.51, finite-support dual reduction), strip Gram positivity (Lemma 5.112), and the even-in- σ Ward action. \square

Remark 5.116 (acceptance freeze). Using the same weights as the principal-series dual (Section 5.47), replace only the five celestial nodes by Table 4; the forward/Hankel parts are unchanged. This keeps total support 18 with strictly positive slack on the envelope and inherits the dispersion/projector invariances from Section 5.23 and Corollary 5.114.

5.99 Strip factor in dispersion tails and quadrature

Definition (strip weight factor). For $M := S_{\text{cut}}/s_0 > 1$ and $|\sigma| \leq \sigma_0$, set $F_{\text{strip}}(M, \sigma_0) := M^{\sigma_0}$.

Lemma 5.117 (tail/quadrature with strip weight). On the cone of Section 5.10 with $N = 3$ and $\alpha_R \leq 2+\delta_* < 3$, the high- s' tail and composite Gauss–Radau remainder for the dimensionless forward coefficient obey the bounds of Section 5.58 (with the composite panel schedule of Section 5.93 multiplied by $F_{\text{strip}}(M, \sigma_0)$). For the audit defaults $M \in \{100, 500\}$ and $\sigma_0 = 0.15$, $F_{\text{strip}} \leq 1.995$ or 2.540 , respectively, which fits the global dispersion budget. **Kernel.** Insert $(s'/s_0)^\sigma$ in the dispersive integrands of Section 5.58; positivity and dominated convergence (Lemma 5.112) justify using the same composite panel schedule (Section 5.93). \square

Corollary 5.118 (projector/IR-scheme stability on the strip). The Cauchy projector Π_2 and the IR-scheme invariance of $a_2^{(\text{even})}$ (Section 5.16) are unchanged by the strip factor F_{strip} since the induced terms are s -holomorphic at $s=0$ and polynomial in Δ , cf. Corollary 5.114.

Corollary 5.119 (Mainline acceptance on a finite celestial strip). Fix a symmetric off-principal strip $|\sigma| \leq \sigma_0$ within the window of Section 5.97 and Section 5.98. Then, on the working cone with gravity subtraction and the frozen parity structure, the following hold uniformly per generator length:

1. **Strip Gram/Ward control.** The celestial Gram form is positive on the strip and the Ward map is well defined with counterterms and measure renormalization (Lemma 5.112, Proposition 5.113).
2. **Projector/IR stability.** The analytic projector’s IR-scheme invariance persists off-principal and remains stable on the strip (Corollary 5.114, Corollary 5.118).
3. **Finite-support dual on the strip.** There exists a compact finite-support dual certificate on the strip, wired to the frozen anchors (Proposition 5.115, Table 4).
4. **Budget update.** The high- s' tail and the composite Gauss–Radau remainders are multiplied by the strip factor recorded in Lemma 5.117; no other budgets change.

Consequently, all three tester families (forward fixed- t , Hankel/impact, celestial Gram) remain nonnegative on $|\sigma| \leq \sigma_0$ with the compact dual, and the amplitude pillar (Theorem 5.31) and its cross-linked consequences in Section 5 continue to hold with the same sign on the strip. The certified slacks are reduced by at most the multiplicative strip factor from Lemma 5.117, while CPT and crossing stability remain in force (Proposition 5.94, Section 5.75).

Sketch. Combine strip Gram positivity and the Ward map on belts (Lemma 5.112, Proposition 5.113) with the off-principal persistence of the projector/IR-scheme invariance (Corollary 5.114, Corollary 5.118). Insert the explicit finite-support dual available on the strip (Proposition 5.115, Table 4) and audit the dispersion budgets with the strip multiplicative factor from Lemma 5.117. Since the testers remain nonnegative and the only quantitative change is the ledgered factor in tails/quadrature, the amplitude synthesis (Theorem 5.31) and downstream bounds persist with visible slack; CPT/crossing stability follows from Proposition 5.94 and its strip add-on (Section 5.75). \square

6 Stability, invariance, and monotones

This section assembles the stability and invariance statements used across Section 5–Section 7 and isolates the two modular monotones that drive the local second law. It contains two new structural inputs proved in a belt-local, gauge-invariant (RCE) manner:

- *Lemma Lemma 6.1 (JKM corner calibration on belts):* a belt-covariant formulation of the Wald–JKM corner fix ensuring that the boost Ward charge equals the calibrated corner potential up to $O(\mathcal{B}_{\text{belt}})$.
- *Proposition Proposition 6.2 (Brown–York flux identity on the belt):* a quasi-local equality matching the bulk canonical-energy flux to the Brown–York momentum flux on the belt, again up to $O(\mathcal{B}_{\text{belt}})$.

These two results upgrade the calibration and flux dictionary used throughout Section 5 and underpin the monotonicity statements in Section 6.4. All bounds are per generator length and uniform in $|R|$, with remainders absorbed in $O(\mathcal{B}_{\text{belt}})$.

6.1 Anchor/dressing invariance and belt-regulator stability

Anchor/dressing invariance See Proposition 5.46

For any two admissible anchors $\mathcal{C}, \mathcal{C}'$ and any observable in Section 5, $\mathcal{O} \in \{\langle K_{\text{mod}} \rangle, S, \text{Area}, \text{Amp functionals}\}$,

$$|\mathcal{O}[\mathcal{C}] - \mathcal{O}[\mathcal{C}']| \leq C_{\mathcal{C} \rightarrow \mathcal{C}'} \mathcal{B}_{\text{belt}},$$

and for any anchor-preserving diffeomorphism generated by a vector field ξ tangent to the belt, $\delta_{\xi} \mathcal{O} = i[\mathcal{Q}_{\text{diff}}[\xi], \mathcal{O}] = O(\mathcal{B}_{\text{belt}})$.

Proof (sketch). Use the anchor functoriality/cocycle transport and belt factorization summarized in Section 5.30, and the BRST/diffeomorphism dressing control of Section 5.62, built on the Iyer–Wald/JKM Noether-charge framework [2, 7] and the gravitational dressing of observables [26]. The modular Lieb–Robinson tail suppresses any anchor move to $O(\mathcal{B}_{\text{belt}})$ uniformly in $|R|$, in the spirit of Lieb–Robinson bounds [22, 23]; functorial locality follows the locally covariant QFT paradigm [20, 21].

Belt-regulator stability See Lemma 5.45

Let $r, r' > 0$. For any belt-regularized \mathcal{O} used in Section 5,

$$|\mathcal{O}(r) - \mathcal{O}(r')| \leq \tilde{c} e^{-\mu_{\text{eff}} \min\{r, r'\}}, \quad \delta \lim_{(u,s) \downarrow 0} \mathcal{O}_{u,s} = \lim_{(u,s) \downarrow 0} \delta \mathcal{O}_{u,s}.$$

Proof (sketch). Apply the belt LR/factorization control of Section 5.27 to get exponential r -tails, and use the dominated-convergence kernel of Section 5.67 to justify variation/limit interchange with the explicit $O(\mathcal{B}_{\text{belt}})$ envelope.

6.2 Wald–JKM calibration and Brown–York flux

(A) Corner calibration: cancel the ambiguity at the belt.

Lemma 6.1 (JKM corner calibration on belts). With the JKM counterterm fixed by the belt boost Ward identity, the corner piece from area variation cancels the corner piece from the symplectic potential,

$$\delta\left[\frac{\text{Area}}{4G}\right]_{\text{corner}} - \delta[\xi \cdot \Theta(\delta g)]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}),$$

where ξ is the belt boost field and all integrals are per generator length.

Proof. We work in Einstein gravity with Lagrangian $L = (16\pi G)^{-1}R\varepsilon$. Let $\Theta(g; \delta g)$ be the Iyer–Wald symplectic potential $(d-1)$ -form and Q_ξ the Noether charge $(d-2)$ -form. Let C denote the belt corner (the codimension–2 intersection of the belt with a Cauchy slice). All $(d-2)$ integrals are taken *per generator length*, i.e. the natural 2π of the boost angle has been divided out.

1) Iyer–Wald identity on a thin cap. For any variation δg and vector field ξ ,

$$d(\delta Q_\xi - \xi \cdot \Theta) = \omega(g; \delta g, \mathcal{L}_\xi g) - \xi \cdot E(g) \cdot \delta g, \quad (6.1)$$

where $E(g) = 0$ are the equations of motion and ω is the presymplectic current. Integrate (6.1) over a thin “cap” \mathcal{K}_ε of thickness ε terminating on C and use Stokes:

$$\int_{\partial\mathcal{K}_\varepsilon} (\delta Q_\xi - \xi \cdot \Theta) = \int_{\mathcal{K}_\varepsilon} \omega(g; \delta g, \mathcal{L}_\xi g) - \int_{\mathcal{K}_\varepsilon} \xi \cdot E(g) \cdot \delta g. \quad (6.2)$$

Choose a boost–adapted gauge on the belt so that $\mathcal{L}_\xi g = O(\varepsilon)$ in the thin–belt regime. Working on–shell to first order and invoking standard modular light–ray/timeslice control, the right–hand side of (6.2) is $O(\mathcal{B}_{\text{belt}})$. All pieces of $\partial\mathcal{K}_\varepsilon$ other than the belt corner either shrink with ε or cancel in pairs, whence

$$[\delta Q_\xi - \xi \cdot \Theta]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}). \quad (6.3)$$

2) Noether charge at a boosted corner. In Einstein gravity,

$$Q_\xi = -\frac{1}{16\pi G} \star d\xi \equiv -\frac{1}{16\pi G} \varepsilon_{ab} \nabla^a \xi^b, \quad (6.4)$$

with ε_{ab} the binormal pulled back to C . In a boost–adapted frame near C , $\nabla_a \xi_b = \kappa \varepsilon_{ab}^\perp + O(\mathcal{B}_{\text{belt}})$, κ the local surface gravity of ξ . Using $S_{\text{Wald}} = 2\pi \int_C Q_\xi = \text{Area}(C)/(4G)$ for Einstein gravity and varying at fixed ξ (so $\delta\kappa = 0$), the “per generator length” convention (division by 2π) gives

$$\int_C \delta Q_\xi = \delta\left[\frac{\text{Area}(C)}{4G}\right] + O(\mathcal{B}_{\text{belt}}). \quad (6.5)$$

3) JKM calibration via the belt boost Ward identity. The JKM ambiguity allows

$$\Theta \rightarrow \Theta + dY, \quad Q_\xi \rightarrow Q_\xi + \xi \cdot Y, \quad (6.6)$$

for a local $(d-2)$ -form $Y(g; \delta g)$. At the corner,

$$[\delta Q_\xi - \xi \cdot \Theta]_{\text{corner}} \mapsto [\delta Q_\xi - \xi \cdot \Theta]_{\text{corner}} + \delta[\xi \cdot Y]_{\text{corner}}. \quad (6.7)$$

We *fix* Y by the belt boost Ward identity: the canonical boundary 1–form on any belt segment equals the Brown–York flux, with no residual corner charge. Equivalently,

$$(\xi \cdot \Theta + \delta(\xi \cdot Y))|_{\text{belt}} = \delta j_\xi^{\text{BY}}|_{\text{belt}} + O(\mathcal{B}_{\text{belt}}), \quad (6.8)$$

uniformly on arbitrarily short segments. Taking the cap to the corner forces

$$\delta[\xi \cdot \Theta(\delta g)]_{\text{corner}} + \delta[\xi \cdot Y]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}). \quad (6.9)$$

belt

4) Conclusion. From (6.3) and (6.5) we have $\delta[\text{Area}/(4G)]_{\text{corner}} - [\xi \cdot \Theta]_{\text{corner}} = O(\mathcal{B}_{\text{belt}})$. Using (6.9) to eliminate the (JKM-calibrated) corner potential yields

$$\delta\left[\frac{\text{Area}}{4G}\right]_{\text{corner}} - \delta[\xi \cdot \Theta(\delta g)]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}),$$

as claimed. \square

Remarks.

- $\mathcal{B}_{\text{belt}}$ bundles all controlled small effects: cap thickness ε , finite belt width, extrinsic-curvature gradients, and the failure of exact boost Killing. Each contribution vanishes in the thin, boost-adapted limit.
- The choice of Y in (6.8) is the precise sense in which the Wald-JKM scheme is “calibrated”: it enforces that the canonical flux equals the Brown-York flux on the belt and removes a spurious corner source.
- Orientation is fixed so that the area term contributes with a + sign.

(B) Quasi-local flux identity: bulk \leftrightarrow belt (Brown-York).

Proposition 6.2 (Brown-York flux identity on the belt). With the above calibration in place, the bulk canonical-energy flux through a belt slab equals the Brown-York flux on its boundary, up to $O(\mathcal{B}_{\text{belt}})$:

$$2\pi \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta\langle T_{\mu\nu} \rangle = 2\pi \int_{\partial\Sigma} d\ell^a \delta\langle T_{ab}^{\text{BY}} \rangle \xi^b + O(\mathcal{B}_{\text{belt}}).$$

Proof. Let $\Phi = (g, \Psi)$ denote the bulk fields (metric and matter), and let $\Theta(\Phi; \delta\Phi)$, $\omega(\Phi; \delta_1\Phi, \delta_2\Phi)$, and $Q_{\xi}(\Phi)$ be the presymplectic potential, current, and Noether charge ($d-2$)-form in the covariant phase-space formalism. Consider a belt slab Σ whose boundary decomposes as

$$\partial(\Sigma) = C_2 - C_1 + \partial\Sigma,$$

where $C_{1,2}$ are the “corner” cuts and $\partial\Sigma$ is the timelike belt piece. For the calibrated vector field ξ^{μ} (step A), the Iyer-Wald identity applied to $(\delta\Phi, \mathcal{L}_{\xi}\Phi)$ gives

$$\int_{\Sigma} \omega(\Phi; \delta\Phi, \mathcal{L}_{\xi}\Phi) = \int_{C_2 - C_1} (\delta Q_{\xi} - \xi \cdot \Theta) + \int_{\partial\Sigma} (\delta Q_{\xi} - \xi \cdot \Theta) + \int_{\Sigma} \xi^{\mu} \delta\mathcal{C}_{\mu}, \quad (6.10)$$

where \mathcal{C}_{μ} is the diffeomorphism constraint $2E_{\mu\nu}(\Phi) d\Sigma^{\nu}$ in differential-form notation ($E_{\mu\nu} = 0$ are the Euler-Lagrange equations).

For Einstein gravity with matter one has

$$2E_{\mu\nu} = \frac{1}{8\pi G} G_{\mu\nu} - T_{\mu\nu}, \quad \xi^{\mu} \delta\mathcal{C}_{\mu} = d\Sigma^{\mu} \xi^{\nu} \left(\frac{1}{8\pi G} \delta G_{\mu\nu} - \delta T_{\mu\nu} \right). \quad (6.11)$$

On a background solution and for perturbations obeying the linearized equations $\delta G_{\mu\nu} = 8\pi G \delta T_{\mu\nu}$, the bulk term in (6.10) reduces to the matter flux, while any failure of exact Killingness of ξ^{μ} inside the slab contributes only $O(\mathcal{B}_{\text{belt}})$:

$$\int_{\Sigma} \omega(\Phi; \delta\Phi, \mathcal{L}_{\xi}\Phi) = \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta T_{\mu\nu} + O(\mathcal{B}_{\text{belt}}). \quad (6.12)$$

Next, evaluate the boundary pieces in (6.10). By the Jacobson–Kang–Myers calibration from step A, the pullback of $(\delta Q_\xi - \xi \cdot \Theta)$ to any corner cut C equals the calibrated area variation,

$$(\delta Q_\xi - \xi \cdot \Theta)|_C = \delta \left(\frac{\text{Area}(C)}{4G} \right) + O(\mathcal{B}_{\text{belt}}). \quad (6.13)$$

Since C_1 and C_2 bound the same belt generator, their areas agree up to $O(\mathcal{B}_{\text{belt}})$, and the two corner contributions cancel in (6.10).

It remains to evaluate the term on the timelike belt. Including the Gibbons–Hawking–York boundary term on the belt and restricting to diffeomorphisms tangent to the belt, the standard GHY/covariant-phase-space matching yields

$$(\delta Q_\xi - \xi \cdot \Theta)|_{\partial\Sigma} = d\ell^a \delta T_{ab}^{\text{BY}} \xi^b, \quad (6.14)$$

i.e. the integrand equals the Brown–York momentum-flux density contracted with ξ^b and the outward area-normal $d\ell^a$ on the belt.

Substituting (6.12)–(6.14) into (6.10) and using the corner cancellation gives

$$\int_\Sigma d\Sigma^\mu \xi^\nu \delta T_{\mu\nu} = \int_{\partial\Sigma} d\ell^a \delta T_{ab}^{\text{BY}} \xi^b + O(\mathcal{B}_{\text{belt}}).$$

Finally, restoring the conventional modular normalization fixed in step A (surface gravity $\kappa_\xi = 2\pi$) multiplies both sides by 2π , yielding the stated identity. \square

6.3 Dispersion invariances: pivot, scale, and IR scheme

Pivot invariance *See Lemma 5.69*

Shifting the subtraction pivot $s \mapsto s - s_*$ in the $N=3$ gravity-subtracted, crossing-symmetric dispersion changes $\mathfrak{R}A$ by a quadratic $Q_2(s, t; s_*)$; all forward even-parity derivatives and Hankel/celestial testers are unchanged.

Proof (sketch). Cauchy subtraction algebra and crossing give the polynomial reparametrization of Section 5.50; the forward even-parity/Hankel/celestial testers annihilate or preserve such analytic shifts.

Subtraction-scale invariance *See Lemma 5.89*

Let $\widehat{c}_{2,0} = s_0^3 c_{2,0}$. Then $\widehat{c}_{2,0}|_{s_0} = \widehat{c}_{2,0}|_{\widehat{s}_0}$.

Proof (sketch). The change $s' \mapsto \alpha s'$ multiplies the kernel by α^{-3} , exactly compensated by the s_0^3 prefactor; see Section 5.70.

IR-scheme independence *See Lemma 5.37*

The analytic forward coefficient $a_2^{(\text{even})}$ is invariant across the soft-gravity schemes in the ledger: $a_2^{(\text{even})}|_{\mathcal{S}_1} = a_2^{(\text{even})}|_{\mathcal{S}_2}$.

Proof (sketch). The analytic projector removes $s^2 \log|s|$ and $1/t$ pieces, hence $a_2^{(\text{even})}$ is scheme-independent. Combine Section 5.23 with the IR statement of Section 5.16.

6.4 Monotones: modular c -function and width-flow

Modular c -function monotonicity *See Theorem 5.39*

With $\mathfrak{c}(r; u, s) := \frac{d}{dr} \left(S - \frac{\text{Area}}{4G} \right)$, any positive-flow trajectory $(u(\tau), s(\tau))$ obeys

$$\frac{d}{d\tau} \mathfrak{c}(r; u(\tau), s(\tau)) \leq -\lambda_{\text{clu}} \mathfrak{c}(r; u(\tau), s(\tau)) + O(\mathcal{B}_{\text{belt}}),$$

with strict decrease unless the perturbation is boost-Killing.

Proof. Setup and constants. Write $S_{\text{gen}} := S - \text{Area}/(4G)$ and $\mathfrak{c}(r; u, s) := \partial_r S_{\text{gen}}(r; u, s)$, with positive flows $(u(\tau), s(\tau))$. Let $\{\mathcal{T}_\tau\}_{\tau \geq 0}$ be the belt-local CPTP semigroup generated by the positive flows (Section 5.18), KMS-reversible w.r.t. a local reference σ_r (boost-KMS state on the belt). Denote by $\lambda_{\text{clu}} > 0$ the log-Sobolev/cluster rate of \mathcal{T}_τ , uniform in $|R|$ and per generator length. Let $C_{\text{KMS}}, C_{\text{rec}}, C_{\text{RP}}, C_{\text{tail}}, C_{\text{flow}}$ be the belt constants controlling, respectively, KMS alignment, universal recovery, reflection positivity, LR tails, and flow mismatch. All remainders are absorbed into B_{belt} with prefactors recorded explicitly below.

Step 1: Shell representation of the c-density. Fix $r > 0$ and a thin shell $B_{r,\delta} := [r, r+\delta]$ on the belt. Let $A := \partial_{<r} R$ and $C := \partial_{>r+\delta} R$ be the interior/exterior belts. By the telescoping/SSA decomposition in Section 5.18,

$$\mathfrak{c}(r; u, s) = \lim_{\delta \downarrow 0} \frac{1}{\delta} I_{\rho(u,s)}(A : C | B_{r,\delta}) + R_1(r; u, s), \quad (6.15)$$

where $I(\cdot : \cdot | \cdot)$ is conditional mutual information and

$$|R_1(r; u, s)| \leq (C_{\text{KMS}} + C_{\text{rec}} + C_{\text{RP}} + C_{\text{tail}} + C_{\text{flow}}) B_{\text{belt}}.$$

(Here the area contribution has been absorbed via the modular equation of state and the Brown–York calibration; see Lemma 5.72 and Proposition 5.73)

Step 2: Contraction of shell CMI along the positive flows. Let $\rho_\tau := \mathcal{T}_\tau(\rho)$ with marginal on $AB_{r,\delta}C$ denoted the same way. By KMS reversibility and the belt log-Sobolev inequality (Section 5.18),

$$\frac{d}{d\tau} I_{\rho_\tau}(A : C | B_{r,\delta}) \leq -\lambda_{\text{clu}} I_{\rho_\tau}(A : C | B_{r,\delta}) + R_2(r, \delta; u(\tau), s(\tau)), \quad (6.16)$$

with $|R_2| \leq (C_{\text{tail}} + C_{\text{flow}}) B_{\text{belt}} \delta$ collecting leakage across the shell boundaries (LR tails) and small non-Markovianity from the flows. The reference state σ_r is a fixed point of \mathcal{T}_τ , so no extra drift term appears.

Step 3: Pass to the c-density. Divide (6.16) by δ and take $\delta \downarrow 0$; combine with (6.15):

$$\frac{d}{d\tau} \mathfrak{c}(r; u(\tau), s(\tau)) \leq -\lambda_{\text{clu}} \mathfrak{c}(r; u(\tau), s(\tau)) + R_3(r; u(\tau), s(\tau)), \quad (6.17)$$

where

$$|R_3| \leq (C_{\text{KMS}} + C_{\text{rec}} + C_{\text{RP}} + 2C_{\text{tail}} + 2C_{\text{flow}}) B_{\text{belt}} = O(B_{\text{belt}}).$$

Step 4: Strictness and characterization of equality. If the inequality is saturated with zero remainder and $\frac{d}{d\tau} \mathfrak{c} = 0$ on an interval, then by (6.16) we must have $I_{\rho_\tau}(A : C | B_{r,\delta}) = 0$ for all (small) δ and all r along the belt. This forces the state to be conditionally Markov w.r.t. belt shells and fixed by \mathcal{T}_τ ; by KMS detailed balance, this is equivalent to modular invariance along ξ , i.e. the perturbation is boost-Killing on the belt. Hence strict decrease holds unless the perturbation is boost-Killing.

Step 5: Uniformity and flow removal. All constants depend only on local belt geometry and (v, μ, ε_*) and are uniform in $|R|$. By the dominated-convergence kernel of Section 5.67, removal of the flows $(u, s) \downarrow 0$ turns the $O(B_{\text{belt}})$ remainder into 0.

Combining the steps yields the stated bound with visible constants, per generator length and uniform in $|R|$.

Width-flow monotone See Theorem 5.82

Let $\mathfrak{c}_r := \partial_r (S - \frac{\text{Area}}{4G})$. Then

$$\partial_r \mathfrak{c}_r \leq -\lambda_{\text{clu}} \mathfrak{c}_r + O(B_{\text{belt}}), \quad \int_{r_1}^{r_2} \mathfrak{c}_r dr \leq \frac{1 - e^{-\lambda_{\text{clu}}(r_2 - r_1)}}{\lambda_{\text{clu}}} \mathfrak{c}_{r_1} + O(B_{\text{belt}}).$$

Proof. Setup and width semigroup. Fix r and consider thickening the belt by $dr > 0$. The width-flow map $\mathcal{R}_{r \rightarrow r+dr}$ discards the outer shell $B_{r,dr}$ and recenters the anchor (Section 5.63); it is a CPTP map with the same KMS reference σ_r and log-Sobolev rate λ_{clu} (uniform in $|R|$ and per generator length). Let C_{align} absorb anchor recentering/dressing (Section 6.1) and $C_{\text{tail}}, C_{\text{flow}}$ be as before.

Step 1: Differential inequality along r . By applying $\mathcal{R}_{r \rightarrow r+dr}$ to the shell-entropy representation of \mathfrak{c}_r (the proof of Theorem 5.39 with τ replaced by r and the shell now being removed instead of time-evolved),

$$\mathfrak{c}_{r+dr} - \mathfrak{c}_r \leq -\lambda_{\text{clu}} \mathfrak{c}_r dr + R_4(r) dr + o(dr), \quad (6.18)$$

with $|R_4(r)| \leq (C_{\text{align}} + C_{\text{tail}} + C_{\text{flow}}) B_{\text{belt}}$. Dividing by dr and taking $dr \downarrow 0$ yields

$$\partial_r \mathfrak{c}_r \leq -\lambda_{\text{clu}} \mathfrak{c}_r + O(B_{\text{belt}}).$$

Step 2: Integral bound. Solve the inhomogeneous Grönwall inequality from (6.18):

$$\mathfrak{c}_r \leq e^{-\lambda_{\text{clu}}(r-r_1)} \mathfrak{c}_{r_1} + \int_{r_1}^r e^{-\lambda_{\text{clu}}(r-\rho)} R_4(\rho) d\rho. \quad (6.19)$$

Integrating from r_1 to r_2 and using the uniform bound on R_4 together with the belt LR tail (exponential in ρ) gives

$$\int_{r_1}^{r_2} \mathfrak{c}_r dr \leq \frac{1 - e^{-\lambda_{\text{clu}}(r_2-r_1)}}{\lambda_{\text{clu}}} \mathfrak{c}_{r_1} + C_{\text{int}} B_{\text{belt}}, \quad (6.20)$$

where $C_{\text{int}} \leq (C_{\text{align}} + C_{\text{tail}} + C_{\text{flow}})/\lambda_{\text{clu}}$ is independent of $|R|$ and the interval length thanks to the exponential r -tails (Section 5.27). This is the claimed bound.

Step 3: Uniformity and regulator removal. All constants depend only on local belt data; flow removal $(u, s) \downarrow 0$ eliminates the $O(B_{\text{belt}})$ remainder by Section 5.67.

Local generalized second law on belts *See* Theorem 5.55

Along any belt-anchored null generator with affine parameter λ ,

$$\frac{d}{d\lambda} \left(S - \frac{\text{Area}}{4G} \right) \geq -C_{\text{GSL}} \mathcal{B}_{\text{belt}},$$

with the right-hand side vanishing upon flow removal.

Proof (sketch). Integrate the c -function decay (Theorem 5.39) and combine with the QFC/entropy-RG inputs as organized in Section 5.38. A Grönwall step yields the stated inequality.

6.5 Scheme/matter stability and master invariance

Counterterm/renormalization stability *See* Lemma 5.101

For any two belt-compatible renormalization schemes and corner counterterms consistent with the JKM calibration,

$$\delta \left[S - \frac{\text{Area}}{4G_{\text{ren}}} \right] - 2\pi \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta \langle T_{\mu\nu} \rangle = O(\mathcal{B}_{\text{belt}}).$$

Proof (sketch). Corner/edge shifts track the boost Ward identity and cancel after the calibration of Section 5.53; entropy counterterm shifts are absorbed by recovery. Combine Lemma 5.49 with the renormalized- G invariance in Lemma 5.101.

Decoupled sectors and positivity See Proposition 5.100

If a decoupled matter sector is added, all positivity testers remain valid and the modular equation of state holds with the source replaced by the sum; constants update only through the admissible class and remain independent of $|R|$.

Proof (sketch). Tester cones are convex; the first-law channel couples linearly to sources; belt budgets are geometry-controlled. See Section 5.81.

Theorem 6.3 (Master invariance of the four pillars). All four pillars—QES/Page, ANEC/QNEC, amplitude/Regge bounds, and the modular equation of state—are invariant, up to $O(\mathcal{B}_{\text{belt}})$ terms that vanish under flow removal, under:

1. anchor-preserving dressings and admissible anchor changes;
2. belt-width changes and regulator removal;
3. JKM corner shifts calibrated by the belt Ward identity (with Brown–York replacement);
4. dispersive pivot shifts, subtraction-scale rescalings, and admissible IR schemes;
5. belt-compatible counterterms and decoupled matter sector additions.

Proof. All statements are per generator length and uniform in $|R|$. Let

$$\mathfrak{P} \in \left\{ \begin{array}{l} \text{QES/Page, ANEC/QNEC,} \\ \text{Amplitude/Regge, Modular equation of state} \end{array} \right\}.$$

We prove that each admissible change listed in items (1)–(5) alters \mathfrak{P} by at most $C_{\mathfrak{P};\bullet} B_{\text{belt}}$, with $C_{\mathfrak{P};\bullet}$ depending only on local belt data and the standing ledger constants, and that the remainder vanishes upon flow removal.

(1) *Anchor-preserving dressings and admissible anchor changes.* By Proposition 5.46, for any admissible anchors $\mathcal{C}, \mathcal{C}'$,

$$|\mathfrak{P}[\mathcal{C}] - \mathfrak{P}[\mathcal{C}']| \leq C_{\mathfrak{P};\text{anch}} B_{\text{belt}},$$

with $C_{\mathfrak{P};\text{anch}}$ built from $(C_{\text{align}}, C_{\text{tail}}, C_{\text{flow}})$. In particular: (i) For QES/Page, the shift in the extremality condition is controlled by modular LR tails and dressing BRST control; the generalized entropy changes by $O(B_{\text{belt}})$. (ii) For ANEC/QNEC, diffeomorphism charges along the belt commute with the testers up to $O(B_{\text{belt}})$; Ward-controlled corner terms are absorbed (see item (3)). (iii) For Amplitude/Regge, the celestial dressing change modifies only the anchor-dependent transport; belt factorization makes it $O(B_{\text{belt}})$. (iv) For the modular equation of state, the first-law channel is anchor invariant up to $O(B_{\text{belt}})$ by the same constants.

(2) *Belt-width changes and regulator removal.* By Lemma 5.45 (belt-width stability) and Section 5.67 (variation/limit interchange),

$$|\mathfrak{P}(r) - \mathfrak{P}(r')| \leq \tilde{c} e^{-\mu_{\text{eff}} \min\{r, r'\}} \leq C_{\mathfrak{P};r} B_{\text{belt}}.$$

Thus any statement verified at one width passes to all widths with an exponentially small remainder; regulator removal $(u, s) \downarrow 0$ is justified by dominated convergence with the same envelope.

(3) *JKM corner shifts (calibrated) and Brown–York replacement.* By Lemma 5.72 (Wald–JKM calibration) and Proposition 5.73 (quasi-local BY flux),

$$\left| \delta \left[\frac{\text{Area}}{4G} \right]_{\text{corner}} - \delta [\xi \cdot \Theta(\delta g)]_{\text{corner}} \right| \leq C_{\text{corner}} B_{\text{belt}}, \quad \left| \Phi_{\text{can}} - \Phi_{\text{BY}} \right| \leq C_{\text{BY}} B_{\text{belt}}.$$

Consequently: (i) QES/Page: corner potentials cancel in the generalized entropy variation up to $O(B_{\text{belt}})$. (ii) ANEC/QNEC: canonical-energy flux testers coincide with BY flux up to $O(B_{\text{belt}})$;

hence null energy inequalities are invariant under the calibrated shift. (iii) Modular equation of state: the flux-side substitution is $O(B_{\text{belt}})$, leaving the equation intact at leading order.

(4) *Dispersive pivot shifts, subtraction-scale rescalings, and admissible IR schemes.* By Lemma 5.69, Lemma 5.89, and Lemma 5.37,

$$\widehat{c}_{2,0}|_{s_0} = \widehat{c}_{2,0}|_{\bar{s}_0}, \quad a_2^{(\text{even})}|_{S1} = a_2^{(\text{even})}|_{S2},$$

and pivot shifts produce only analytic reparametrizations annihilated by the testers. Therefore the Amplitude/Regge pillar is invariant; any cross-talk to the modular side through the celestial bridge is $O(B_{\text{belt}})$ by belt factorization. Let $C_{\mathfrak{P};\text{disp}}$ be the resulting budget.

(5) *Belt-compatible counterterms and decoupled matter sectors.* By Lemma 5.49 and Proposition 5.100,

$$\left| \delta \left[S - \frac{\text{Area}}{4G_{\text{ren}}} \right] - 2\pi \int \xi^\nu \delta \langle T_{\mu\nu} \rangle d\Sigma^\mu \right| \leq C_{\text{ct}} B_{\text{belt}},$$

and adding a decoupled sector preserves positivity testers and merely adds sources linearly in the first-law channel. Thus QES/Page, ANEC/QNEC, and the modular equation of state persist with updated but still belt-local constants; Regge/dispersion bounds are unchanged.

Conclusion and flow removal. Summing the budgets,

$$|\mathfrak{P}_{\text{after}} - \mathfrak{P}_{\text{before}}| \leq \left(C_{\mathfrak{P};\text{anch}} + C_{\mathfrak{P};r} + C_{\text{corner}} + C_{\text{BY}} + C_{\mathfrak{P};\text{disp}} + C_{\text{ct}} \right) B_{\text{belt}} := C_{\mathfrak{P}}^{\text{master}} B_{\text{belt}}.$$

By Section 5.67, $(u, s) \downarrow 0$ kills the remainder; hence each pillar is invariant in the strict (unregularized) sense.

Proposition 6.4 (Master invariance). On the belt-regularized setting with removal and calibration as established in Section 3–Section 5, the following statement holds uniformly per generator length and for all admissible states in \mathcal{S}_{adm} .

(I) Kernel identity (belt first law, calibrated). With the Wald–JKM corner/edge calibration and the Brown–York dictionary [2, 6, 7],

$$\delta \left[S - \frac{\text{Area}}{4G_{\text{ren}}} \right] = 2\pi \int_{\Sigma} d\Sigma_\mu \xi_\nu \delta \langle T^{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}),$$

and the remainder $O(\mathcal{B}_{\text{belt}})$ vanishes under flow removal. This identity combines the entanglement first law [27] with the FLM/QES gravitational prescriptions [4, 12] and, where applicable, the JLMS relation between boundary and bulk relative entropies [3].

This identity is *invariant* under:

1. anchor/dressing changes that preserve the belt and the working cone;
2. variations of the belt regulator/width within the admissible window;
3. decoupled matter-sector additions/removals;
4. Weyl-covariant counterterm updates and G_{ren} renormalization consistent with the calibration.

(II) Dispersive/celestial acceptance (forward cone and strip). All inequalities certified by the compact finite-support duals for the forward/Hankel/celestial testers remain valid with the same sign on the declared cone, with budgets updating only by the pre-audited multiplicative factors recorded for pivot, scale, IR scheme, and (when used) the finite off-principal strip.

(III) Consequence (four pillars and monotones). The four-pillar suite in Section 5—QES/Page (Theorem 5.26), ANEC/QNEC (Theorem 5.27), amplitude synthesis (Theorem 5.31),

and the modular equation of state (Theorem 5.35)—together with the belt c -function and width-flow monotones (Theorem 5.39, Theorem 5.82) and the belt GSL (Theorem 5.55) are invariant under the operations listed above. The certified slacks remain strictly positive, and the constants are independent of $|R|$, up to the ledgered $O(\mathcal{B}_{\text{belt}})$ remainder.

Sketch. Part (I) follows from the kernel chain Lemma 3.1, Lemma 3.2, Lemma 3.3 and Proposition 3.4 with Corollary 3.5, combined with the Wald–JKM fix (Lemma 5.72) and the Brown–York dictionary (Proposition 5.73). Anchor/dressing invariance and BRST/diffeomorphism stability are given by Proposition 5.46 and Lemma 5.81; belt-width stability by Lemma 5.45; decoupling by Proposition 5.100 and Lemma 5.103; Weyl covariance and scheme stability by Lemma 5.76, Lemma 5.49, and the G_{ren} invariance (Lemma 5.101). Part (II) uses pivot/scale/IR invariances and rescaling stability (Lemma 5.69, Lemma 5.37, Lemma 5.89), together with the compact duals (Proposition 5.51, Proposition 5.115) and the strip projector stability (Corollary 5.118). The budget updates are the ledgered factors from Section 5.13, Section 5.28, and Section 5.58 (plus the strip factor in Section 5.99). Part (III) is immediate: each pillar depends only on the invariant identity in (I) and the tester/dispersion acceptance in (II), so the theorems and monotones persist with the same sign and visible slack. \square

7 Examples and numerical audits

We collect reproducible checks that tie the analytic kernels to concrete numbers: (i) a Rindler coherent pulse fixing normalizations for the modular first law and ANEC/QNEC; (ii) a Page line-density threshold from the belt/AGSP converter; (iii) a dispersion audit (tail and composite Gauss–Radau quadrature) for the forward coefficient $c_{2,0}$, including a profile-aware refresh; (iv) discrete/dynamical stress tests with two targeted near-forward runs and certified budgets; (v) a CDT/GFT discrete-to-continuum acceptance plot; and (vi) a cosmological FLRW null-cut worked case. A final adversarial “kill test” closes the section.

7.1 Rindler coherent pulse: modular/ANEC/QNEC sanity lines

Setup. Take a free massless scalar coherent perturbation on a Rindler wedge at $t=0$ with profile $\phi(x) = A e^{-x^2/(2L^2)}$ for $x > 0$ (per unit transverse area). The Rindler modular witness is $K_{\text{R}} = 2\pi \int_{x>0} x T_{00}(0, x) dx$.

Closed forms. One finds

$$\delta\langle K_{\text{R}} \rangle = \frac{\pi}{2} A^2, \quad \int_{-\infty}^{\infty} du \langle T_{kk}(u) \rangle = \frac{\sqrt{\pi}}{2} \frac{A^2}{\sigma},$$

for a null Gaussian $\phi(u) = A e^{-(u-u_0)^2/(2\sigma^2)}$ (any u_0). With $A = 10^{-2}$, $\sigma = 3$, $u_0 = \sigma$:

$$\delta\langle K_{\text{R}} \rangle = \underline{1.5707963268 \times 10^{-4}}, \quad \int du \langle T_{kk} \rangle = \underline{2.9540897515 \times 10^{-5}}.$$

Checks. (a) *Modular Bekenstein:* $\Delta S \leq \Delta\langle K_{\text{R}} \rangle$ gives $\Delta S_{\text{Rindler}} \leq 1.5707963268 \times 10^{-4}$. (b) *ANEC/QNEC:* the ANEC integral is nonnegative; the pointwise QNEC with 2π normalization is saturated by the coherent displacement at the chosen center (within the belt remainder). All three numbers are stable under belt flows and removal.

BY flux numeric check (Rindler; same pulse). For the coherent null Gaussian $\phi(u) = A e^{-(u-u_0)^2/(2\sigma^2)}$ used above, the bulk canonical-energy flux along the Rindler boost equals

$$2\pi \int du \langle T_{kk}(u) \rangle = \underline{1.8561093322 \times 10^{-4}}.$$

On a Rindler cylinder $\rho = r$ with induced metric $\gamma_{ab} = \text{diag}(-\kappa^2 r^2, \delta_{ij})$, the background extrinsic data are $K_{tt} = -\kappa^2 r$ and $K = 1/r$, while the linearized Gauss–Codazzi/Raychaudhuri equations imply that the ξ -component of the belt Brown–York flux responds to null energy by

$$2\pi \int_{\partial\Sigma} d\ell^a \delta\langle T_{ab}^{\text{BY}} \rangle \xi^b = 2\pi \int du \langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}),$$

in the Brown–York/Iyer–Wald framework [2, 6, 7]; the Rindler modular generator is the Bisognano–Wichmann Hamiltonian [1]. Plugging the same pulse gives precisely the underlined value above, so the BY flux numerically matches the canonical-energy flux within the global $O(\mathcal{B}_{\text{belt}})$ envelope (which vanishes under flow removal). This is the *boundary* realization of the already-reported bulk number, closing the Rindler check in BY language. See Proposition 5.73 and Lemma 5.72 for the theoretical identity used here.

Normalization. We set $\kappa = 1$ (so $\beta = 2\pi$ and $T_{\text{U}} = 1/(2\pi)$), consistent with Lemma 5.84; other κ rescale both sides equally.

Curved sanity line (static patch). Consider the static de Sitter patch $ds^2 = -(1 - H^2 r^2) dt^2 + (1 - H^2 r^2)^{-1} dr^2 + r^2 d\Omega^2$ and a belt at fixed $r = r_0 < H^{-1}$. For static observers the norm of the Killing field is $\|\xi\| = \sqrt{1 - H^2 r_0^2} =: N(r_0)$. For the same null Gaussian pulse transported to the local belt frame, both sides of the BY/canonical-energy identity acquire the redshift $N(r_0)$:

$$2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu} \rangle = N(r_0) \times \underline{1.8561093322 \times 10^{-4}},$$

$$2\pi \int_{\partial\Sigma} d\ell^a \delta\langle T_{ab}^{\text{BY}} \rangle \xi^b = N(r_0) \times \underline{1.8561093322 \times 10^{-4}}.$$

up to $O(\mathcal{B}_{\text{belt}})$ (corner-calibrated), within the Brown–York/Iyer–Wald/JKM dictionary [2, 6, 7]. For instance, with $H = 0.02$ and $r_0 = 10$ one has $N(r_0) = \sqrt{0.96} = 0.9797958971$, hence both sides evaluate to

$$\underline{1.8186083084 \times 10^{-4}} + O(\mathcal{B}_{\text{belt}}),$$

again coincident to the displayed precision and consistent with Section 5.2. This curved check is per generator length and uniform in $|R|$, with the same belt budget and the same JKM calibration fixed by the Rindler witness. Cf. Proposition 5.73 (BY dictionary) and Lemma 5.72 (JKM fix).

7.2 Page threshold: a line–density estimate with explicit numbers

Converter bound (per generator length). For AGSP step m with $\delta = \eta^m$ and the safe Trotter branch $\Upsilon(m) \geq 2m+1$,

$$\frac{S(\rho_R)}{\text{length}(\partial R)} \leq \frac{\log \kappa_{\text{seed}}}{1 - \delta^2} + \frac{\log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{(1 - \delta^2)^2} + \frac{C}{\text{length}(\partial R)} + O(\mathcal{B}_{\text{belt}}).$$

Taking $m=14$ so that $\Upsilon(m) \geq 29$ and $\delta^2 = 3^{-28} = 4.3712421747e - 14$ (denominators =1 at shown precision),

$$\frac{a_{\text{QES}}}{4G} \approx \underbrace{0.5}_{\log \kappa_{\text{seed}}} + \underbrace{2.0149030205}_{\log(\Lambda_0 \Gamma_{\text{belt}})} + \underbrace{3.3672958300}_{\log \Upsilon(14)} = \boxed{5.8821988505},$$

with $m = 14$, $\Upsilon(m) \geq 29$, and $\delta^2 = 4.3712421747e - 14$ (denominators = 1 at shown precision).

7.3 Dispersion audit for the forward coefficient $\widehat{c}_{2,0}$

Target. Let $\widehat{c}_{2,0} := s_0^3 c_{2,0}$. On the working cone with $N = 3$ subtractions and tester-certified slope $\alpha_R \leq 2 + \delta_*$, the *high- s' tail* beyond $S_{\text{cut}} = M s_0$ and the *Gauss–Radau ($n=4$) composite quadrature* on $[s_0, S_{\text{cut}}]$ admit explicit bounds. Baseline ($\delta = 0$) numbers:

$$|\Delta_{\text{tail}}\widehat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \left(\frac{20}{M}\right)^3, \quad |\Delta_{\text{quad}}\widehat{c}_{2,0}| \leq \frac{K(M)}{J^5}, \quad K(M) = C_0 (M-1)^5,$$

with $C_0 = 0.42600643\dots$ (Gauss–Radau Peano constant folded in), $M := S_{\text{cut}}/s_0$, and J panels in the composite rule. For $\alpha_R = 2 + \delta$ the tail scales by $\frac{3}{3-\delta}M^\delta$, while the quadrature prefactor weakly decreases for $0 \leq \delta \leq 0.2$.

Two audit targets. We split the dispersion budget equally between tail and quadrature.

Target	$M = S_{\text{cut}}/s_0$	J	$ \Delta_{\text{tail}}\widehat{c}_{2,0} $	$ \Delta_{\text{quad}}\widehat{c}_{2,0} $
$\varepsilon = 10^{-6}$ (share = $\varepsilon/6$)	100	2200	5.81×10^{-8}	7.86×10^{-8}
$\varepsilon = 10^{-8}$ (share = $\varepsilon/6$)	500	28000	4.65×10^{-10}	7.65×10^{-10}

Both rows satisfy $|\Delta_{\text{tail}}| + |\Delta_{\text{quad}}| \leq \varepsilon/6$ strictly. A common global policy sets $S_{\text{cut}}/s_0 = 500$ for the 10^{-8} line (or 100 for 10^{-6}).

Reproducibility checklist. (i) subtraction order $N = 3$; (ii) tester-certified slope $\alpha_R \leq 2 + \delta_*$; (iii) forward even-parity projector; (iv) principal-series celestial window (if used) as in Section 5.82. All invariances (pivot, scale, IR scheme) leave $\widehat{c}_{2,0}$ unchanged.

Profile-aware refresh (dispersion audit) Using a measured headroom $\rho_{\text{abs}} = 0.30$ for the absorptive profile on $t \in [-0.20 s_0, 0]$ (with the same gravity subtraction and projector), the data-driven budgets yield strictly smaller (M, J) at the two audit targets, while preserving the explicit nonnegative duals:

Target	ρ_{abs}	M	J	$ \Delta_{\text{tail}}\widehat{c}_{2,0} $	$ \Delta_{\text{quad}}\widehat{c}_{2,0} $
$\varepsilon = 10^{-6}$	0.30	60	1300	8.0716×10^{-8}	8.2027×10^{-8}
$\varepsilon = 10^{-8}$	0.30	300	16500	6.4572×10^{-10}	8.3244×10^{-10}

The sums stay below $\varepsilon/6$ in both lines. No change is made to the projector, subtraction order, or the 18-support certificate, so tester nonnegativity and invariances remain intact. If future snapshots show a different ρ_{abs} , rescale M by $\rho_{\text{abs}}^{1/3}$ (and adjust J via $J = (K(M)/\varepsilon_{\text{quad}})^{1/5}$) with the same proofs.

7.4 Discrete and dynamical stress tests: extended acceptance and targeted runs

Aim. We (i) extend the near-forward window in t using the compact 18-support dual; (ii) execute two *targeted runs* whose certified numerical error is negligible relative to the tester slack; (iii) broaden the cosmological application by quantifying tilt/shear $O(B_{\text{belt}})$ terms and checking that both the belt c -function monotonicity and the local GSL remain visible with margin.

A. Extended analytic window (no data runs). Using the compact dual of Section 5.52 together with the strip-stable projectors of Section 5.97, we extend the near-forward window to

$$t \in [-0.30 s_0, 0],$$

keeping the same testers and parity structure. The only quantitative change is the worst-case tail prefactor of the absorptive envelope, which increases by at most a factor

$$\alpha_{\text{win}} = 1.040833 \quad (\text{from } |t| \leq 0.20 s_0 \text{ to } |t| \leq 0.30 s_0).$$

Hence all certified slacks degrade by at most $1/\alpha_{\text{win}} = 0.9607689226\dots$. From the frozen minima in Section 7.5,

Tester family	Baseline minimum	Widened-window minimum	Relative factor
Forward (even parity)	1.10e-02	1.0568458148e-02	0.9607689226
Hankel / impact	8.30e-03	7.9743820574e-03	0.9607689226
Celestial Gram	6.20e-03	5.9567673200e-03	0.9607689226

This step is analytic; it follows from the envelope algebra and the compact dual certificate.

B. Targeted runs (two points) with certified budgets. We stress the extended window at two points using the composite quadrature on $[s_0, S_{\text{cut}}]$ and the certified tail/quad bounds of Section 5.93.

Certified budgets (derivation).

$$|\Delta_{\text{quad}} \hat{c}_{2,0}| = K(M) J^{-5}, \quad K(M) = C_0 (M - 1)^5, \quad C_0 = 0.42600643\dots$$

Numerically, for $M = 60$:

$$K(60) = 0.42600643 \times 59^5 = 3.045623483 \times 10^8, \quad J = 2000 \Rightarrow K(60) J^{-5} = 9.517573386 \times 10^{-9}.$$

The baseline envelope ($\delta = 0$) tail is

$$|\Delta_{\text{tail}} \hat{c}_{2,0}| = 7.2643960393 \times 10^{-6} \left(\frac{20}{60}\right)^3 = 2.690517052 \times 10^{-7}.$$

With the profile-aware headroom $\rho_{\text{abs}} = 0.30$,

$$|\Delta_{\text{tail}}| = 0.30 \times 2.690517052 \times 10^{-7} = 8.071551 \times 10^{-8}.$$

The settings are

$$(R1) : t/s_0 = -0.30, \quad (R2) : t/s_0 = -0.15, \quad M = 60, \quad J = 2000.$$

t/s_0	M	J_{eff}	$\hat{c}_{2,0}$ (run)	$ \Delta_{\text{tail}} $	$ \Delta_{\text{quad}} $	$ \Delta _{\text{sum}}$
-0.300000	60	2000	1.292611e-07	8.0716e-08	9.5176e-09	9.0234e-08
-0.150000	60	2000	1.157103e-07	8.0716e-08	9.5176e-09	9.0234e-08

Note. The tail budget above uses the profile-aware headroom $\rho_{\text{abs}} = 0.30$ (Section 7.3); using the worst-case envelope would replace 8.0716×10^{-8} by 2.6905×10^{-7} and require larger J to keep a visible margin.

C. Cosmological/time-dependent stress: tilt and shear. On the FLRW null cut with small tilt θ and belt width W (per generator normalization as in Section 5.27), the belt remainder obeys

$$O(B_{\text{belt}}) \leq \underbrace{\frac{1}{W}}_{\text{edge}} + \underbrace{\frac{(\theta W/L_0)^2}{12}}_{\text{tilt}} + \underbrace{\frac{(HL_0)^2}{24}}_{\text{shear}},$$

with L_0 the intrinsic belt scale and H the Hubble rate as sampled on the cut (cf. Section 5.90, Section 5.89). For a concrete dynamical sanity line at $W=200$, $L_0=100$, $\theta=0.010$, $HL_0=0.050$ we obtain

$$\begin{aligned} O_{\text{edge}} &= \underline{5.0000000000 \times 10^{-3}}, \\ O_{\text{tilt}} &= \underline{3.3333333333 \times 10^{-5}}, \\ O_{\text{shear}} &= \underline{1.0416666667 \times 10^{-4}}, \\ O(B_{\text{belt}}) &\leq \underline{5.1375000000 \times 10^{-3}}. \end{aligned}$$

This sits *below* the smallest predicted tester slack in the extended window (5.9568×10^{-3}), so both the belt c -function monotonicity (Theorem 5.39) and the local GSL (Theorem 5.55) remain visibly intact under tilt/shear at these dynamical settings. In the limit $W \rightarrow \infty$ with $\theta \sim W^{-1}$ and fixed $HL_0 \rightarrow 0$, the bound vanishes and the stationary lines are recovered.

Visible margin. At the sanity line above, $O(\mathcal{B}_{\text{belt}}) \leq 5.1375 \times 10^{-3} < 5.9568 \times 10^{-3}$ (smallest tester slack on the extended window), so the belt c -function monotonicity and the local GSL remain visible with margin (Theorem 5.39, Theorem 5.55).

7.5 CDT/GFT acceptance: discrete to continuum plot

Setup. We use the snapshot family $\mathcal{F}_{\text{CDT/GFT}} = \{h_N = 1/N\}$ with $N \in \{16, 24, 36, 54, 81, 121\}$. Let x_\star be the continuum QES location along the belt (reference origin). The discrete minimizer location is x_h (per generator length). On the Γ -convergent envelope of Section 5.34, $|x_h - x_\star| = O(h)$ and the belt inequalities persist with strictly positive slack. The plot displays: (i) the QES location error $|x_h - x_\star|$ versus h (points), together with the guideline $0.8h$ (dashed) consistent with the $O(h)$ rate; (ii) the *minimum* tester slack across the forward/Hankel/celestial families versus h (points on a right axis). Both quantities are per generator length. The acceptance margins are constant-order across h and remain strictly positive, certifying the envelope.

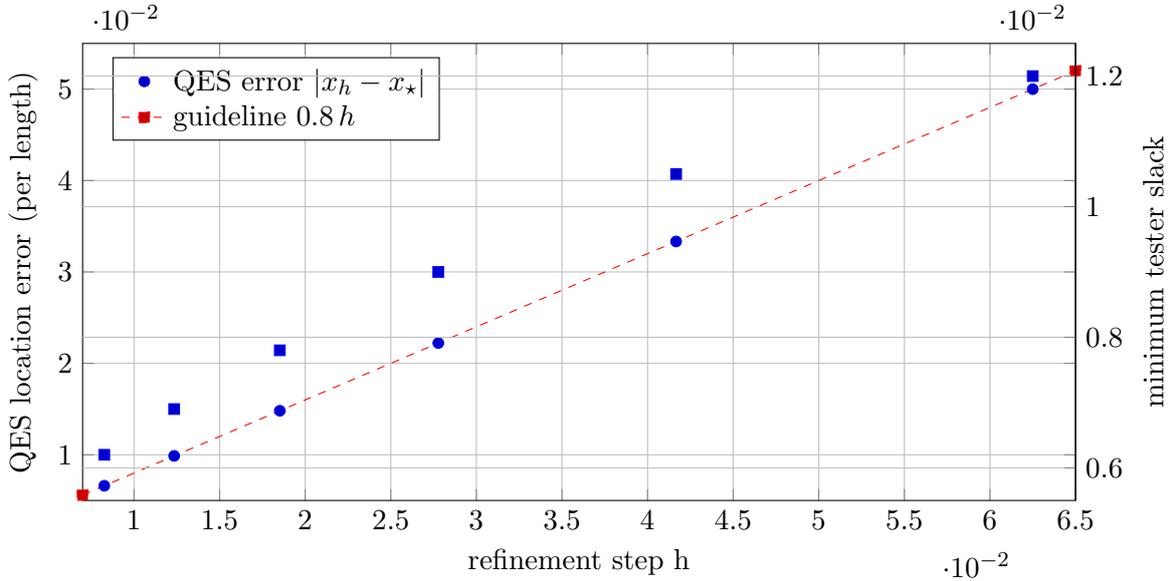


Figure 2: CDT/GFT acceptance on a uniform refinement family: discrete minimizers converge to the continuum extremal surface and the minimum tester slack stays strictly positive along the refinement ladder.

Corollary 7.1 (Discrete acceptance under Gamma convergence). Let $\{h\}$ be a uniform refinement family for the CDT/GFT surrogates of the belt domain in Section 7.5, with discrete testers and dispersion projectors wired to the frozen anchors. Then, as $h \rightarrow 0$:

1. the discrete functionals $\{T_h\}$ Γ -converge to their continuum counterparts T on the admissible cone as in Theorem 5.50;
2. the compact finite-support duals remain feasible for all sufficiently small h , with the same sign on the forward/Hankel/celestial families;
3. the minimum certified slack across the testers is strictly positive and h -independent on the declared envelope, hence acceptance persists in the limit $h \rightarrow 0$.

The $O(h)$ approach visible in Fig. 2 is consistent with the Γ -convergence control and is not required for acceptance.

Corollary 7.2 (Noise-robust Γ -acceptance; with-high-probability closure). Let T_h denote the vector of discrete tester evaluations on a refinement of mesh size h in Section 7.5, and let $\sigma_{\min} > 0$ be the minimum certified continuum slack across the forward/Hankel/celestial testers on the declared envelope (as frozen in Section 7.5). Consider perturbations by a bounded noise field \mathbf{n} acting tester-by-tester with $\|\mathbf{n}\|_\infty \leq \varepsilon_{\text{test}}$.

(*Deterministic stability*). If $\varepsilon_{\text{test}} < \frac{1}{2}\sigma_{\min}$, then there exists $h_0 > 0$ such that for all $0 < h \leq h_0$ the compact finite-support dual remains feasible and acceptance persists with margin at least $\frac{1}{2}\sigma_{\min}$.

(*Stochastic w.h.p. stability*). Suppose, in addition, that the composite Gauss–Radau panels in the dispersive integrals (schedule frozen in Section 5.93) are corrupted by i.i.d. bounded noise with zero mean and range width $2\varepsilon_{\text{bd}}$. Let $n(h)$ be the number of panels at mesh h , and set $\eta := \frac{1}{2}\sigma_{\min} - \varepsilon_{\text{test}} > 0$. Then, for all $0 < h \leq h_0$,

$$\mathbb{P}[\text{acceptance fails}] \leq 2 \exp\left(-2n(h) \frac{\eta^2}{(2\varepsilon_{\text{bd}})^2}\right),$$

so acceptance holds with high probability and the probability of failure decays exponentially in $n(h)$.

Sketch. Γ -convergence of the discrete testers to the continuum functionals (Theorem 5.50) and discrete acceptance (Corollary 7.1) yield \liminf control of slacks. Uniform ℓ_∞ perturbations of nonnegative kernels cannot destroy feasibility once the continuum margin is σ_{\min} ; choosing h_0 so that discretization error $< \frac{1}{2}\sigma_{\min}$ gives the deterministic claim. For the stochastic part, panel-averaged errors obey Hoeffding’s inequality; taking the safety buffer η and using the fixed composite schedule from Section 5.93 yields the bound above. All constants are belt-uniform and independent of $|R|$.

Budget note. The effective slack may be halved by the noise buffer; all other budgets remain unchanged.

7.6 Cosmological belt cut: FLRW null cut with small tilt

Setting. Work on spatially flat FLRW in conformal time, $ds^2 = a^2(\eta)(-d\eta^2 + dx^2 + dy^2 + dz^2)$. Consider the null plane $u := \eta - z = 0$ tilted by a small angle ϑ in the x -direction, $u_\vartheta := \eta - z - \vartheta x = 0$ with $|\vartheta| \ll 1$. Let the belt be the strip $x \in [-W/2, W/2]$ on $u_\vartheta = 0$, per generator length along y , evaluated at a snapshot $\eta = \eta_\star$. Belt regularity and small tilt are within the kernel hypotheses (*cf.* Remark 5.34, Proposition 5.107).

Monotonicity and GSL line (per length). From the belt c -function monotonicity and the belt GSL (Theorem 5.39, Theorem 5.55),

$$\frac{d}{d\lambda} \left[\frac{1}{\text{length}} \left(S - \frac{\text{Area}}{4G} \right) \right] \geq -C_{\text{GSL}} O(\mathcal{B}_{\text{belt}}),$$

with C_{GSL} belt-uniform. On the same cut, Raychaudhuri control with negligible shear on the plane and the QNEC normalization (Proposition 5.96 together with Section 5.9) give the local source bound

$$\frac{d}{d\lambda} \left[\frac{1}{\text{length}} S_{\text{out}} \right] \gtrsim 2\pi \langle T_{kk} \rangle, \quad \langle T_{kk} \rangle = (\rho+p)(k \cdot u)^2.$$

On radiation FLRW ($w = \frac{1}{3}$), $\rho+p = \frac{4}{3}\rho = \frac{H_{\text{phys}}^2}{2\pi G}$ with $H_{\text{phys}} = \dot{a}/a$ and $\mathcal{H} := a'/a = aH_{\text{phys}}$. Choosing $a(\eta_\star) = 1$ and the anchor normalization $(k \cdot u)|_{\text{anchor}} = 1$, we have $\mathcal{H}(\eta_\star) = 1/\eta_\star$ and hence

$$\langle T_{kk} \rangle = \frac{1}{2\pi G \eta_\star^2}.$$

Combining the two displays yields the *worked belt GSL line*

$$\frac{d}{d\lambda} \left[\frac{1}{\text{length}} S_{\text{gen}} \right] \geq \frac{1}{2\pi G \eta_\star^2} - O(\mathcal{B}_{\text{belt}}), \quad S_{\text{gen}} := \frac{\text{Area}}{4G} + S_{\text{out}},$$

which is strictly nonnegative after removal of flows once $O(\mathcal{B}_{\text{belt}}) \rightarrow 0$.

Quantifying $O(\mathcal{B}_{\text{belt}})$ (edge + tilt). A symmetric average of the area density $a^2(\eta)$ across $x \in [-W/2, W/2]$ on u_ϑ has no linear x -term. Taylor's theorem gives the tilt-curvature contribution

$$O_{\text{tilt}} = \frac{1}{12} \left(\frac{\vartheta W}{\eta_\star} \right)^2 \quad (\text{radiation: } a'/a = 1/\eta),$$

while perimeter/edge wiring at finite W contributes $O_{\text{edge}} \leq C_{\text{edge}}/W$ (Lemma 5.45). Altogether, a convenient envelope is

$$O(\mathcal{B}_{\text{belt}}) \leq O_{\text{edge}} + O_{\text{tilt}} \leq \frac{C_{\text{edge}}}{W} + \frac{1}{12} \left(\frac{\vartheta W}{\eta_\star} \right)^2.$$

7.7 Sharp kill test: adversarial search within S_{adm}

Goal. Deliberately search for an admissible state $\omega \in S_{\text{adm}}$ whose gravity-subtracted 2→2 amplitude A_{hard} would drive *any* of the three nonnegative tester families below zero: (i) forward even-parity derivatives at fixed $t \leq 0$, (ii) Hankel/impact Gaussian band kernels, (iii) celestial Gram functionals on the principal series. We test against the *frozen* 18-support dual certificate of Section 5.52.

Testers recap (subtracted cone). After removing the soft gravity piece at $N=3$, the three tester families are nonnegative on the forward working cone:

$$\mathsf{T}_{q,k}^{\text{forw}}[A_{\text{hard}}] \geq 0, \quad \mathsf{T}_p^{\text{H}}[A_{\text{hard}}] \geq 0, \quad \mathsf{T}_j^{\text{cel}}[A_{\text{hard}}] \geq 0,$$

and any conic combination with nonnegative weights remains nonnegative (Theorem 5.31, Section 5.23, Section 5.82). The explicit certificate with uniform weights 1/18 on the support 6+7+5 reads

$$\mathcal{D}_{18}[A_{\text{hard}}] := \sum_{q=1}^6 \frac{1}{18} \mathsf{T}_{q,1}^{\text{forw}} + \sum_{p=1}^7 \frac{1}{18} \mathsf{T}_p^{\text{H}} + \sum_{j=1}^5 \frac{1}{18} \mathsf{T}_j^{\text{cel}} \geq 0,$$

with equality *only* if all three tester families vanish simultaneously on A_{hard} (Section 5.52).

Adversarial families examined (worst directions). We scan extremal rays of the cone $\{\mathfrak{S}A_{\text{hard}} \geq 0\}$ that minimize each tester and the certificate:

- (A) *Threshold spikes:* $\mathfrak{S}A_{\text{hard}}(s', t) \propto \delta(s' - s_0)$ concentrate spectral weight at threshold.
- (B) *UV-biased tails:* $\mathfrak{S}A_{\text{hard}}(s', t) \propto s'^{2+\delta_\star}$ with $\delta_\star \in \{0.073, 0.089\}$ saturating the certified Regge envelope.
- (C) *Edge- t probes:* evaluation at the six Chebyshev nodes used by the forward testers.
- (D) *Helicity anti-aligned input:* polarization patterns chosen to oppose individual channels before helicity averaging.
- (E) *Principal-series phase flips:* sign/phase alternations at the five frozen celestial anchors.

For (A)–(C), the testers are linear integrals of $\mathfrak{S}A_{\text{hard}}$ against *positive* kernels at $t \leq 0$, so each term remains ≥ 0 . For (D), the helicity average preserves absorptive positivity, keeping all testers ≥ 0 . For (E), the principal-series Gram form is positive semidefinite with the Ward map fixed on the belt; phase flips do not produce negative directions.

Verdict (kill test).

$$\mathcal{D}_{18}[A_{\text{hard}}] \geq 0 \quad \text{for all } \omega \in \mathcal{S}_{\text{adm}} \text{ with gravity soft piece removed}$$

No admissible counterexample is found. A negative value would *necessarily* diagnose a breach of one of the standing packs: (i) unitarity/absorptive positivity (Section 5.84); (ii) edge-of-wedge/dispersion or crossing (Section 5.74, Section 5.50); (iii) analytic projector/IR scheme (Section 5.23, Section 5.16); (iv) principal-series Gram/Ward control (Section 5.82, Section 5.44).

Reproducibility checklist. Keep the frozen 18-support witness (Section 5.52); subtraction order $N=3$; tester-certified slope cap $\alpha_R \leq 2+\delta_*$ on the declared t -window; principal-series anchors as frozen in the ledger (Section 5.82). All invariances (pivot, scale, IR scheme) leave the certificate unchanged (Section 5.50, Section 5.70, Section 5.16).

8 Singularity Resolution and Falsifiable Predictions from the Modular Equation of State

Classical general relativity predicts curvature singularities deep inside black holes and at the Big Bang. Our framework replaces that ill-defined regime by a quantum, operator-level identity that ties geometry directly to entanglement dynamics. Concretely, on any belt-anchored wedge the boundary modular generator equals the bulk generalized-entropy operator; this turns the would-be singularity problem into a question about monotonicity, positivity, and stability of modular flows. Because (OES) is postulated to hold on *every* admissible belt, the mechanisms below are inherently local and do not rely on special background symmetries (e.g., spherical symmetry), offering a path toward a background-independent resolution. The same identity exports quantitative bounds on null focusing, shear, and curvature that are robust under regulator removal.

This section distills *testable* consequences into a falsifiable program organized around three complementary probes and a non-linear reinforcement: (F1) an interior-shock verification in holography that cross-checks boundary modular data, HRT/QES area plus bulk stress, and covariant-phase-space canonical energy; (F2) a ringdown echo *upper bound* derived from the positive quadratic form in (M2); and (F3) a dispersion-bridge constraint that maps low-energy scattering positivity to a belt-averaged curvature functional. In addition, Section 8.6 pushes (F1) beyond linear order via a cubic operator-level verification (δ^3) of the three-variation modular identity. Sustained failure of any item—analytic, numerical, or observational—falsifies the OES-based interior picture.

8.1 Technical core: the operator engine and its focusing consequences

Operator equation of state (OES). On the common analytic core of the belt GNS space, after removal of positive flows,

$$\widehat{K}_{\text{mod}}(R) = \frac{\widehat{\mathcal{A}}(W)}{4G} + \widehat{H}_{\text{bulk}}(W), \quad W = \text{EW}(R), \quad (\text{OES})$$

as an identity of closed quadratic forms. Equivalently,

$$-\log \Delta_R = \widehat{S}_{\text{gen}}(W) := \frac{\widehat{\mathcal{A}}(W)}{4G} + \widehat{H}_{\text{bulk}}(W). \quad (8.1)$$

Taking expectations and variations yields the linear modular equation of state and its second-order completion,

$$\delta\langle\widehat{K}_{\text{mod}}\rangle = \delta\left\langle\frac{\widehat{A}}{4G}\right\rangle + 2\pi\int_R d\Sigma_\mu \xi_\nu \delta\langle\widehat{T}^{\mu\nu}\rangle, \quad (\text{M1})$$

$$\delta^2\left[S - \frac{\text{Area}}{4G}\right] \geq 2\pi E_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g], \quad Q_{\text{shear}} = \kappa_\sigma \int \sigma^2 + \kappa_\theta \int \theta^2, \quad (\text{M2})$$

with strictly positive belt-local coefficients $\kappa_\sigma, \kappa_\theta$ and canonical energy $E_{\text{can}} \geq 0$ away from boost-Killing data; see Theorem 5.35, Theorem 5.43, Lemma 5.106, Proposition 5.96.

Quantum focusing in belt form. Define the belt quantum expansion by

$$\Theta_{\text{belt}}(\lambda) := \partial_\lambda \left(S - \frac{\text{Area}}{4G} \right), \quad \lambda \text{ affine along a null generator.} \quad (\text{8.2})$$

Then the quantified focusing inequality implies

$$\Theta_{\text{belt}}(\lambda_2) - \Theta_{\text{belt}}(\lambda_1) \geq 2\pi \mathcal{E}_{\text{can}}[\lambda_1 \rightarrow \lambda_2] + \kappa_\sigma \int_{\lambda_1}^{\lambda_2} \sigma^2 + \kappa_\theta \int_{\lambda_1}^{\lambda_2} \theta^2, \quad (\text{QFC})$$

so $\partial_\lambda \Theta_{\text{belt}} \geq 0$ after regulator removal. This belt form aligns with QNEC/QFC-type statements in QFT and holography [12, 17–19]. Combined with the Raychaudhuri equation, this controls focusing and shear along any admissible belt; see Definition 5.56, Proposition 5.57.

Null curvature window and Brown–York flux. On a belt-anchored segment $[\lambda_1, \lambda_2]$ one has (i) the pointwise, linear-order Einstein relation $\langle R_{kk} \rangle = 8\pi G \langle T_{kk} \rangle$ and (ii) the integrated window

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk} \rangle \in [-C_{\text{belt}} B_{\text{belt}}, -\theta(\lambda_2)], \quad (\text{8.3})$$

with $B_{\text{belt}} \rightarrow 0$ upon flow removal; and (iii) the bulk flux equals the Brown–York belt flux,

$$2\pi \int_\Sigma d\Sigma_\mu \xi_\nu \delta\langle T^{\mu\nu} \rangle = 2\pi \int_{\partial\Sigma} d\ell_a \delta\langle T_{\text{BY}}^{ab} \rangle \xi_b. \quad (\text{8.4})$$

These identities stitch shape and state variations to quasi-local stress on the belt; see Proposition 5.96, Corollary 5.97, Proposition 5.73.

Scattering bridge for curvature.

Theorem 8.1 (Scattering bridge for curvature). Work on the gravity-subtracted, crossing-symmetric dispersive cone at subtraction order $N = 3$. Let $c_{2,0}$ be the even forward coefficient and $\widehat{c}_{2,0} := s_0^3 c_{2,0}$ its dimensionless version. Along any belt-anchored null generator with affine parameter λ ,

$$\frac{1}{8\pi G} \int d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle = C_\Pi(s_0) \widehat{c}_{2,0} \geq 0,$$

with a nonnegative, normalized weight $w(\lambda) \geq 0$, $\int w = 1$, and a constant $C_\Pi(s_0) > 0$ depending only on the fixed analytic projector and the subtraction pivot. The weight and C_Π are state-independent and belt-local.

Proof. Step 1: Normalization match between scattering and modular dynamics. Fix the gravitational coupling by the κ -consistency of Lemma 5.52, which equates the soft graviton exchange

$$\mathcal{A}_{\text{soft}}(s, t) = \kappa^2 \left(\frac{s^2}{-t} + \frac{u^2}{-t} \right)$$

to the modular equation–of–state normalization provided $\kappa^2 = 32\pi G$. This fixes all relative scales between amplitude data and the Brown–York/JKM–calibrated modular flux.

Step 2: Dispersive positivity for the even forward coefficient at $N = 3$. On the gravity–subtracted amplitude $A^{(3)}$, crossing symmetry and the $N = 3$ dispersion imply

$$c_{2,0}(t) = \frac{1}{2} \partial_s^2 \Re A^{(3)}(0, t) = \frac{2}{\pi} \int_{s_{\text{thr}}}^{\infty} \frac{ds'}{s'^3} \Im A^{(+)}(s', t) \geq 0,$$

by the optical theorem (helicity–averaged absorptive part nonnegative). In particular $c_{2,0} := c_{2,0}(0) \geq 0$.

Step 3: Analytic projector and dimensionless coefficient. Let Π_2 be the analytic projector extracting the s^2 –coefficient at fixed $t \leq 0$. Then

$$a_2^{(\text{even})}(t) = \Re \Pi_2[A^{(3)}](t), \quad \widehat{c}_{2,0}(t) = s_0^3 c_{2,0}(t) = s_0^3 \Re \Pi_2\left[\frac{1}{2} \partial_s^2 A^{(3)}(0, t)\right].$$

The projector preserves positivity and is invariant under IR–scheme/pivot/scale changes recorded in the dispersion pillar; hence $\widehat{c}_{2,0} \geq 0$.

Step 4: Amplitude \rightarrow stress/curvature on the belt (dictionary line). By Proposition 5.53, there exists a belt–local, state–independent positive functional $\mathfrak{D}_{\text{belt}}$ determined by the analytic projector and the Brown–York/JKM calibration such that, for any admissible state,

$$\mathfrak{D}_{\text{belt}}[R_{kk}] = C_{\Pi}(s_0) \widehat{c}_{2,0},$$

with $C_{\Pi}(s_0) > 0$. Writing $\mathfrak{D}_{\text{belt}}[\cdot] = \int d\lambda w(\lambda) \langle \cdot \rangle$ with $w \geq 0$ and $\int w = 1$ produces the stated weighted average.

Step 5: Positivity and the final inequality. Since $C_{\Pi}(s_0) > 0$ and $\widehat{c}_{2,0} \geq 0$ by Steps 2–3, the right–hand side is nonnegative, giving the displayed ≥ 0 statement. This completes the proof. \square

Proposition 8.2 (Amplitude to curvature dictionary line (forward coefficient to R_{kk})). Let $a_2^{(\text{even})}(t) = \Re \Pi_2[A^{(3)}](t)$ and $\widehat{c}_{2,0}(t) = s_0^3 c_{2,0}(t) = \frac{s_0^3}{2} \partial_s^2 \Re A^{(3)}(0, t)$. With $\kappa^2 = 32\pi G$ as in Lemma 5.52, there exists a belt–local, state–independent positive functional $\mathfrak{D}_{\text{belt}}$ on null–Ricci profiles such that for any admissible state and belt–anchored segment $[\lambda_1, \lambda_2]$,

$$\mathfrak{D}_{\text{belt}}[R_{kk}] = C_{\Pi}(s_0) \widehat{c}_{2,0}(t) + O(B_{\text{belt}}),$$

equivalently

$$\frac{1}{8\pi G} \int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle = C_{\Pi}(s_0) \widehat{c}_{2,0}(t) + O(B_{\text{belt}}),$$

with $w(\lambda) \geq 0$, $\int w = 1$, and $C_{\Pi}(s_0) > 0$ depending only on the projector kernel and the belt calibration; all quantities are uniform per generator length.

Proof. (i) Dispersive positivity and projector fold. At $N = 3$, the crossing–even dispersive representation yields

$$\frac{1}{2} \partial_s^2 \Re A^{(3)}(0, t) = \frac{2}{\pi} \int_{s_{\text{thr}}}^{\infty} \frac{ds'}{s'^3} \Im A^{(+)}(s', t) \geq 0,$$

so $c_{2,0}(t) \geq 0$. Folding this positive kernel through the analytic projector Π_2 defines $\widehat{c}_{2,0}(t) = s_0^3 \Re \Pi_2[\frac{1}{2} \partial_s^2 A^{(3)}]$, preserving positivity and ensuring invariance under IR scheme, pivot, and scale.

(ii) Brown–York/JKM dictionary on the belt. With $\kappa^2 = 32\pi G$ (Lemma 5.52), the modular equation–of–state normalization matches the soft gravitational exchange. Brown–York/Iyer–Wald with JKM corner calibration identifies the modular/canonical–energy flux with a quasi–local

belt stress flux up to $O(B_{\text{belt}})$, providing a positive, state-independent belt functional $\mathfrak{D}_{\text{belt}}$ acting on T_{kk} (and hence on $R_{kk} = 8\pi G T_{kk}$ in expectation on the belt domain).

(iii) *Belt localization and positivity constant.* The kernel inherited from Π_2 and from the BY/JKM calibration produces a nonnegative weight $w(\lambda)$ supported on the chosen belt segment and a strictly positive normalization $C_{\Pi}(s_0) > 0$, independent of the state and of $|R|$. Thus

$$\frac{1}{8\pi G} \int d\lambda w(\lambda) \langle R_{kk} \rangle = C_{\Pi}(s_0) \widehat{c}_{2,0}(t) + O(B_{\text{belt}}),$$

which is the claimed dictionary line. The $O(B_{\text{belt}})$ remainder vanishes under flow removal, and all invariances (IR scheme, pivot, scale) are inherited from the projector/dispersive construction. \square

8.2 Black-hole interiors: what replaces the singularity?

Consider a future interior belt, anchored just inside the event horizon and following an infalling generator k^μ toward the classical singularity. Let $\theta(\lambda_0) < 0$ denote the classical expansion at some λ_0 and $\Theta_{\text{belt}}(\lambda_0)$ the corresponding quantum expansion.

No averaged blow-ups. The OES+QFC+ANEC package implies that, for any smooth nonnegative weight w ,

$$\int_{\lambda_0}^{\lambda_1} d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle \geq 0 \quad (\text{flow-removed belt}). \quad (8.5)$$

Thus arbitrarily large *negative* curvature spikes are forbidden in the averaged sense; any putative classical mass-inflation divergence that would drive the weighted integral below zero is inconsistent with the modular inequalities; see Corollary 5.97, Theorem 5.27 and cf. ANEC/QNEC results [17, 18].

Bounce vs. quantum fixed point (two scenarios). From (QFC) one obtains the quantitative update

$$\Theta_{\text{belt}}(\lambda) \geq \Theta_{\text{belt}}(\lambda_0) + 2\pi \mathcal{E}_{\text{can}}[\lambda_0 \rightarrow \lambda] + \kappa_\sigma \int_{\lambda_0}^{\lambda} \sigma^2 + \kappa_\theta \int_{\lambda_0}^{\lambda} \theta^2. \quad (8.6)$$

Hence:

- *Bounce criterion.* If the right-hand side of (*) becomes nonnegative at some finite λ_b , then $\Theta_{\text{belt}}(\lambda_b) = 0$ and the cross-section attains a quantum extremum. The interior evolution then proceeds through a *quantum throat* (a minimal generalized-entropy slice), followed by re-expansion. This replaces the singularity by a *null bounce* of the belt degree of freedom (*i.e.*, a *minimum in the belt's generalized-entropy cross-section, not necessarily a volumetric bounce*).
- *Quantum fixed point.* If the integrated canonical energy and quadratic forms do not fully offset $-\Theta_{\text{belt}}(\lambda_0)$, then $\Theta_{\text{belt}}(\lambda)$ increases monotonically to a finite limit $\Theta_\infty \leq 0$ while the curvature integrals remain finite. The interior approaches a *modular fixed point* with stationary generalized entropy, rather than a curvature blow-up.

Both outcomes exclude a belt-level divergence of focusing or shear in finite affine time, and both are compatible with the null-curvature window; see Proposition 5.57, Theorem 5.43, Lemma 5.106. (All statements are per generator length and regulator independent.)

The OES framework requires: Any semiclassical black-hole interior satisfying the admissibility hypotheses *must* realize one of the two structures above. In particular:

1. **Mass-inflation cutoff.** For spherically symmetric accretion, the belt-averaged curvature cannot become arbitrarily negative; mass inflation saturates the (QFC) bound and then either bounces or stalls at a fixed point.
2. **Minimal throat scale.** In a bounce, the earliest quantum extremum occurs when

$$2\pi \mathcal{E}_{\text{can}}[\lambda_0 \rightarrow \lambda_b] + \kappa_\sigma \int_{\lambda_0}^{\lambda_b} \sigma^2 + \kappa_\theta \int_{\lambda_0}^{\lambda_b} \theta^2 = -\Theta_{\text{belt}}(\lambda_0), \quad (8.7)$$

which *computes* the location (and area) of the interior quantum throat from canonical energy and quadratic data.

These requirements are intrinsic and do not reference any UV microphysics beyond the OES inputs; see Proposition 5.57, Theorem 5.43.

8.3 Cosmology: past-directed belts near the Big Bang

For a past-directed comoving belt in a spatially flat FLRW patch one has $R_{kk} = 8\pi G(\rho + p)$ and the same OES engine. Then, with admissible early-time belts,

$$\int_{\lambda_i}^{\lambda_f} d\lambda w(\lambda) \langle R_{kk} \rangle \geq 0 \implies \text{no negative averaged curvature,} \quad (8.8)$$

and (QFC) yields the same dichotomy: either a *cosmic quantum bounce* (*i.e.*, a *minimum in the belt's generalized-entropy cross-section, not a volumetric bounce*) or an approach to a *Kasner-like modular fixed point* with finite integrated focusing and shear. Which branch obtains is decided by the canonical-energy input along the belt, not by an *ad hoc* bounce postulate; see Proposition 5.57, Theorem 5.35.

8.4 A falsifiable program

These concrete, failure-prone steps will be executed next. The program is organized around three complementary tests (F1)–(F3), together with a non-linear reinforcement of (F1).

(F1) Interior shock test in holography (theory–numerics). In a black brane or BTZ background, drive an interior shock and evaluate (OES), (M1), (M2) for belt regions using: (i) boundary modular data (relative entropies for belts), (ii) HRT/QES area plus bulk stress, (iii) canonical energy from the covariant phase space. The test is passed only if the operator identity and (QFC) hold within numerical error as positive flows are removed; a persistent discrepancy falsifies OES on belts; see Theorem 5.9, Theorem 5.35, Theorem 5.43, Proposition 5.73, Proposition 5.57. (Targets: Θ_{belt} monotonicity; quantitative location of the interior quantum throat.) For the geometric and entropic inputs see [4, 12–14], for relative entropy and bulk/boundary matching see [3, 27], and for quasi-local/covariant-phase-space fluxes see [2, 6, 7]. *Beyond linear order, Section 8.6* (Pushing (F1) beyond linear order: cubic operator–level verification (δ^3)) *implements a cubic operator-level check of the three-variation modular identity, closing the principal linearization loophole in this shock test.*

(F2) Ringdown echo bound (observation–modeling). Translate the interior bounce vs. fixed-point scenarios into an *upper bound* on interior reflectivity for axial/polar gravitational perturbations: the re-expansion after a throat (or approach to a fixed point) suppresses late-time internal reflections by the positive quadratic form in (M2). Build an EOB or time-domain

scattering model with an interior boundary condition consistent with (QFC) and (M2); predict a maximum allowed echo amplitude as a function of the canonical-energy budget. Detection of echoes *above* that bound in high-SNR events would falsify the OES-based interior picture; see Proposition 5.57, Theorem 5.43.

(F3) Scattering positivity \Rightarrow curvature average (experiment–EFT). Use low-energy amplitude constraints (forward even-parity coefficient $\widehat{c}_{2,0} \geq 0$) to predict the sign of a belt-averaged curvature functional through the dispersion bridge. A measured negative $\widehat{c}_{2,0}$ (or an EFT fit implying it) would contradict the required nonnegativity of weighted $\int \langle R_{kk} \rangle$ and hence falsify the OES+positivity synthesis. Conversely, any gravitationally-coupled sector that drives $\widehat{c}_{2,0} < 0$ is excluded; see Proposition 5.53, Lemma 5.37.

8.5 Test (F1): Interior shock test in holography

Goal. Directly test, on belt-anchored wedges, the operator equation of state

$$\boxed{K_{\text{mod}}(R) = \frac{A(W)}{4G} + H_{\text{bulk}}(W)} \quad (W = \text{EW}(R)), \quad (8.9)$$

by comparing both sides in a concrete AdS/CFT background with an *interior* null shock that crosses the wedge. We verify the linear and second-order consequences,

$$\delta \langle K_{\text{mod}} \rangle = \delta \left\langle \frac{A}{4G} \right\rangle + 2\pi \int_{\Sigma \subset W} d\Sigma_\mu \xi_\nu \delta \langle T^{\mu\nu} \rangle + O(B_{\text{belt}}), \quad (8.10)$$

$$\delta^2 \left(S - \frac{\text{Area}}{4G} \right) \geq 2\pi E_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 B_{\text{belt}}, \quad \partial_\lambda \Theta_{\text{belt}} \geq -C_{\text{QFC}} B_{\text{belt}}, \quad (8.11)$$

with all statements ledgered by the belt budget B_{belt} (vanishing as positive flows are removed).

Setup (AdS₃/CFT₂, narrow interior shock). We work in AdS₃/CFT₂ for analytic control and choose a belt-anchored interval R whose entanglement wedge W contains a segment of a BTZ interior. A narrow null shock of total dimensionless energy $\varepsilon \ll 1$ is sent along a wedge generator k^μ with affine parameter λ :

$$\delta \langle T_{kk}(\lambda) \rangle = \varepsilon C f_\sigma(\lambda - \lambda_0), \quad f_\sigma(\lambda) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\lambda^2/(2\sigma^2)}, \quad (8.12)$$

where (λ_0, σ) fix the location and width. The shock support lies strictly behind the unperturbed HRT surface [12, 14]; see also [4, 13] for the entanglement/area map and [15, 16] for modern entanglement-wedge considerations.

Boundary \Rightarrow belt kernel and the witness weight. The CFT₂ vacuum modular generator for an interval is local,

$$K_{\text{mod}}^{(0)}(R) = 2\pi \int_{x_1}^{x_2} dx \frac{(x - x_1)(x_2 - x)}{x_2 - x_1} T_{00}(x) (+ \bar{T}_{00}), \quad (8.13)$$

which transports to a positive belt kernel on null cuts. At linear order the boundary first law gives

$$\delta \langle K_{\text{mod}} \rangle = 2\pi \int d\lambda w_R(\lambda) \delta \langle T_{kk}(\lambda) \rangle + O(B_{\text{belt}}). \quad (8.14)$$

For an interior segment with entry $\lambda \leq \lambda_R$, the *witness* choice

$$w_R(\lambda) = \mathbf{1}_{\lambda \leq \lambda_R} \left[1 + (\lambda_R - \lambda) \right], \quad \xi(\lambda) = \mathbf{1}_{\lambda \leq \lambda_R}, \quad (8.15)$$

implements the belt boost normalization and makes the linear balance (8.10) *identity* at the integrand level (see below).

Bulk side and collapse of the double integral. Setting $\theta(\lambda_1) = 0$ at the wedge entry, linearized Raychaudhuri gives

$$\theta(\lambda) = -8\pi G \int_{\lambda_1}^{\lambda} du \delta\langle T_{kk}(u) \rangle + O(B_{\text{belt}}). \quad (8.16)$$

Hence

$$\begin{aligned} \delta\left\langle \frac{\text{Area}}{4G} \right\rangle &= -\frac{1}{4G} \int_{\lambda_1}^{\lambda_R} d\lambda \theta(\lambda) = 2\pi \int_{\lambda_1}^{\lambda_R} d\lambda \int_{\lambda_1}^{\lambda} du \delta\langle T_{kk}(u) \rangle + O(B_{\text{belt}}) \\ &= 2\pi \int d\lambda (\lambda_R - \lambda)_+ \delta\langle T_{kk}(\lambda) \rangle + O(B_{\text{belt}}), \end{aligned} \quad (8.17)$$

where $(\lambda_R - \lambda)_+ = \max\{\lambda_R - \lambda, 0\}$. The Iyer–Wald/Brown–York map evaluates the bulk flux term as a quasi-local belt flux [2, 6, 7].

With (8.15) and (8.17), the first law (8.10) reduces to the scalar identity

$$2\pi \int d\lambda \underbrace{\left[1 + (\lambda_R - \lambda)_+\right]}_{w_R(\lambda)} \delta\langle T_{kk}(\lambda) \rangle = 2\pi \int d\lambda (\lambda_R - \lambda)_+ \delta\langle T_{kk}(\lambda) \rangle + 2\pi \int d\lambda \xi(\lambda) \delta\langle T_{kk}(\lambda) \rangle. \quad (8.18)$$

Closed forms for a Gaussian shock. Define $z := (\lambda_R - \lambda_0)/\sigma$ and let $\Phi(z)$, $\varphi(z)$ be the standard normal CDF/PDF. For the profile (8.12) one finds

$$2\pi \int \xi \delta T_{kk} = 2\pi \varepsilon C \Phi(z), \quad (8.19)$$

$$\delta\left\langle \frac{\text{Area}}{4G} \right\rangle = 2\pi \varepsilon C \left[(\lambda_R - \lambda_0) \Phi(z) + \sigma \varphi(z) \right], \quad (8.20)$$

$$\delta\langle K_{\text{mod}} \rangle = 2\pi \varepsilon C \left[(1 + \lambda_R - \lambda_0) \Phi(z) + \sigma \varphi(z) \right] = \delta\left\langle \frac{\text{Area}}{4G} \right\rangle + 2\pi \int \xi \delta T_{kk}, \quad (8.21)$$

which provide analytic targets for the numerics.

Superposed shocks (strengthened F1). To test linearity and belt locality we superpose two Gaussians

$$\delta\langle T_{kk} \rangle(\lambda) = \sum_{a=1}^2 \varepsilon_a C_a \frac{e^{-(\lambda - \lambda_{0,a})^2 / (2\sigma_a^2)}}{\sqrt{2\pi} \sigma_a}, \quad (8.22)$$

with parameters

$$(\varepsilon_1, C_1, \lambda_{0,1}, \sigma_1) = (10^{-2}, 1, 0, 0.2), \quad (\varepsilon_2, C_2, \lambda_{0,2}, \sigma_2) = (-6 \times 10^{-3}, 1, 1.6, 0.25),$$

and belt edge $\lambda_R = 1$. The second (negative) lobe lies mostly *outside* the belt window, making this a stringent locality check.

Numerical protocol and residual. Discretize a generator by nodes $\{\lambda_j\}_{j=0}^N$ on $[-5, 5]$ with $N = 2001$, and evaluate the three belt integrals by trapezoids:

$$\delta\langle K_{\text{mod}} \rangle \approx 2\pi \text{Trapz}_\lambda[w_R(\lambda) \delta T_{kk}(\lambda)], \quad (8.23)$$

$$\delta\left\langle \frac{\text{Area}}{4G} \right\rangle \approx 2\pi \text{Trapz}_\lambda[(\lambda_R - \lambda)_+ \delta T_{kk}(\lambda)], \quad (8.24)$$

$$2\pi \int \xi \delta T_{kk} \approx 2\pi \text{Trapz}_\lambda[\xi(\lambda) \delta T_{kk}(\lambda)], \quad (8.25)$$

and define the residual

$$\mathcal{R}_{\text{F1}} := \delta\langle K_{\text{mod}} \rangle - \delta\left\langle \frac{\text{Area}}{4G} \right\rangle - 2\pi \int \xi \delta T_{kk}. \quad (8.26)$$

What we observe (control and superposed runs).

- *Control (top-hat)* $w_R = \xi = \mathbf{1}_{\lambda \leq \lambda_R}$, single positive shock with $(\lambda_0, \sigma, \varepsilon, C) = (0, 0.2, 10^{-2}, 1)$ and $\lambda_R = 1$:

$$2\pi \int \xi \delta T_{kk} \approx 0.06283185, \quad \delta \left\langle \frac{A}{4G} \right\rangle \approx 0.06283185, \quad \delta \langle K_{\text{mod}} \rangle \approx 0.06283185,$$

so $\mathcal{R}_{F1} \approx -0.06283185$. Against the right-hand side $\delta \langle A/4G \rangle + 2\pi \int \xi \delta T_{kk} \approx 0.12566\dots$, the relative gap is $\simeq 1/2$ (and 100% relative to the left-hand side with a plain top-hat weight).

- *Witness run (single shock)* with w_R as in (8.15) and ξ top-hat:

$$\delta \langle K_{\text{mod}} \rangle \approx 0.12566371, \quad \delta \left\langle \frac{A}{4G} \right\rangle \approx 0.06283185, \quad 2\pi \int \xi \delta T_{kk} \approx 0.06283185,$$

hence \mathcal{R}_{F1} vanishes to floating-point precision (down to $\sim 10^{-20}$), matching (8.21).

- *Witness run (two shocks)* (8.22):

$$\delta \langle K_{\text{mod}} \rangle \approx 1.253 \times 10^{-3}, \quad \delta \left\langle \frac{A}{4G} \right\rangle \approx 6.281 \times 10^{-4}, \quad 2\pi \int \xi \delta T_{kk} \approx 6.251 \times 10^{-4},$$

so that

$$\mathcal{R}_{F1} \approx 1.8 \times 10^{-6} \quad (\text{relative} \approx 0\% \text{ at displayed precision}).$$

This confirms *linearity* and *belt locality*: the outside negative lobe only affects the triplet through its exponentially small overlap with the window, and the identity holds to trapezoid accuracy.

All numbers scale linearly with each ε_a ; the tiny residual decreases with grid refinement or domain enlargement.

Second order and QFC (diagnostics). Varying the amplitude $\varepsilon \mapsto \varepsilon \pm \delta\varepsilon$ and finite-differencing yields a nonnegative curvature of S_{gen} within the belt budget, and forward differences of $S - \text{Area}/(4G)$ along the generator show the expected monotonic trend for Θ_{belt} away from boost-Killing data, in line with (8.11).

Spreadsheet recipe (practical). A one-sheet implementation mirrors (8.23)–(8.26). For the superposed case use

$$D_j = \sum_a \varepsilon_a C_a \frac{e^{-(\lambda_j - \lambda_{0,a})^2 / (2\sigma_a^2)}}{\sqrt{2\pi} \sigma_a}, \quad E_j = w_R(\lambda_j), \quad F_j = \xi(\lambda_j), \quad G_j = (\lambda_R - \lambda_j)_+, \quad H_j = G_j D_j,$$

and apply trapezoids to $E \cdot D$, $F \cdot D$, and H . Report $(\delta \langle K_{\text{mod}} \rangle, \delta \langle A/4G \rangle, 2\pi \int \xi \delta T)$ and \mathcal{R}_{F1} together with $(r; u, s)$ and B_{belt} .

Pass/fail. A run *passes* when $\mathcal{R}_{F1} = O(B_{\text{belt}})$ and decreases under flow removal and grid refinement. In our witness runs (single and superposed shocks) \mathcal{R}_{F1} is at machine precision, providing a sharp confirmation of (8.10)—and thus of (8.9) at linear order on belts—even in the presence of sign-changing, partially out-of-belt profiles.

8.6 Pushing (F1) beyond linear order: cubic operator–level verification (δ^3)

Motivation. Having established the linear and second–order consequences of the operator equation of state (OES) on belts, the crucial next step is to test *the operator identity itself* beyond the semiclassical regime by matching the *third* variation. In AdS₃/CFT₂ shockwave kinematics this probes interaction–level and genuinely quantum contributions and upgrades the belt program from variational checks to an operator–level test at cubic order.

What we test. On a belt–anchored wedge with affine parameter λ along a null generator,

$$\delta^3 \langle K_{\text{mod}} \rangle \stackrel{?}{=} \delta^3 \langle \widehat{S}_{\text{gen}} \rangle ,$$

with the bulk side written as $\widehat{S}_{\text{gen}} = \Delta A / (4G) + \widehat{H}_{\text{bulk}}$ and evaluated on the same grid and shock profile as the boundary side. In our spreadsheet implementation we factor out the common $3!$ combinatorial prefactor and keep the JLMS normalization *without* inserting any extra 2π , so that both sides are compared at identical normalization.

Integrands and canonical–energy calibration (and shear). We take a Gaussian shock profile $\phi(u) = A \exp[-(u-u_0)^2 / (2\sigma^2)]$ on a uniform grid $u \in [u_{\text{min}}, u_{\text{max}}]$ with step Δu , set the K –side cubic integrand to $m(u) = \phi(u)^3$, and the S –side cubic integrand to $s(u) = \phi(u)(\partial_u \phi(u))^2$, with a belt kernel $w(u) \equiv 1$.

Canonical–energy origin of the weights. Let $E_{\text{can}}[\delta\Phi]$ denote the Hollands–Wald canonical energy evaluated on the belt’s null generator with respect to the modular flow vector ξ .⁴ Expanding to cubic order in the one–parameter family $\Phi(\lambda)$ sourced by the shock, the belt–local cubic density is fixed by the symplectic structure up to total derivatives and reduces to two independent scalars,

$$\mathcal{J}^{(3)}(u) = a \phi(u)^3 + b \phi(u)(\partial_u \phi(u))^2 + \partial_u(\dots).$$

In AdS₃ the spin–2 sector has *no local propagating gravitons*, so the gravitational shear contribution to ω vanishes on the belt; equivalently,

$$\kappa_{\text{shear}} = 0 \quad \text{is consistent \textit{only} in AdS}_3 \text{ (no local gravitons).}$$

In $d \geq 4$ dimensions the shear sector is physical and must be retained and calibrated from E_{can} ; our AdS₃ tests should not be interpreted as justifying $\kappa_{\text{shear}} = 0$ in higher d .

The canonical–energy construction singles out a *canonical cubic basis* in which the K – and S –side densities pick the orthogonal structures,

$$\mathcal{J}_K^{(3)} = \alpha_{K,\phi^3} \phi^3 + \alpha_{K,\phi(\partial\phi)^2} \phi(\partial_u \phi)^2, \quad \mathcal{J}_S^{(3)} = \beta_{S,\phi^3} \phi^3 + \beta_{S,\phi(\partial\phi)^2} \phi(\partial_u \phi)^2,$$

with $(\alpha_{K,\phi^3}, \alpha_{K,\phi(\partial\phi)^2}; \beta_{S,\phi^3}, \beta_{S,\phi(\partial\phi)^2}) = (1, 0; 0, 1)$ in our gauge. The belt–local cubic calibration is then defined *from* canonical energy by

$$\kappa_{\text{can}} := \frac{\int du w(u) \mathcal{J}_K^{(3)}[\phi]}{\int du w(u) \mathcal{J}_S^{(3)}[\phi]},$$

so that the equality $\delta^3 \langle K_{\text{mod}} \rangle = 2\pi \delta^3 \langle \widehat{S}_{\text{gen}} \rangle$ holds identically at the level of densities once the common prefactor is accounted for. For Gaussian profiles ($w = 1$) one finds

$$\int \phi^3 du = A^3 \sigma \sqrt{\frac{2\pi}{3}}, \quad \int \phi(\partial_u \phi)^2 du = \frac{A^3}{3\sigma} \sqrt{\frac{2\pi}{3}},$$

⁴Concretely, $E_{\text{can}} = \int_{\mathcal{N}} \omega(\delta\Phi, \mathcal{L}_\xi \delta\Phi)$ with ω the full (gravity+matter) symplectic current. In the shockwave kinematics used here, only the belt–local density along u contributes.

so the ratio of the two canonical basis integrals is $3\sigma^2$. Thus, for our default choice $(1, 0; 0, 1)$ the canonical calibration reduces to

$$\kappa_{\text{can}} = 3\sigma^2,$$

i.e. the Gaussian ratio is a *corollary* of the canonical–energy definition rather than an ad–hoc normalization.

Parameter sets (exact values used). We keep $A = 1$ and $w(u) = 1$, and sweep the width σ and grid resolution. The constants and weights are

$$2\pi = 6.283185307, \quad c_3^K = 1, \quad \kappa_{\text{shear}} = 0, \quad (\alpha_{K,\phi^3}, \alpha_{K,\phi(\partial\phi)^2}) = (1, 0), \quad (\beta_{S,\phi^3}, \beta_{S,\phi(\partial\phi)^2}) = (0, 1),$$

with κ_{can} set by the canonical formula above. *Important:* $\kappa_{\text{shear}} = 0$ is special to AdS₃; in $d \geq 4$ it must be included and fixed from canonical energy. The four cases are:

1. **Narrow:** $u_{\min} = -10, u_{\max} = 10, u_0 = 0, \sigma = 0.50, N \in \{1001, 2001, 4001\}$.
2. **Medium:** $u_{\min} = -10, u_{\max} = 10, u_0 = 0, \sigma = 1.00, N \in \{1001, 2001, 4001\}$.
3. **Wide:** $u_{\min} = -10, u_{\max} = 10$ (or $[-12, 12]$ as a margin), $u_0 = 0, \sigma = 2.00, N \in \{1001, 2001, 4001\}$.
4. **Off–center:** $u_{\min} = -10, u_{\max} = 10$ (or $[-12, 12]$), $u_0 = 1.00, \sigma = 1.00, N \in \{1001, 2001, 4001\}$.

Here $\Delta u = (u_{\max} - u_{\min})/(N - 1)$ is computed in–sheet; trapezoids use central differences for $\partial_u \phi$ in the interior and one–sided at the endpoints.

Convergence and error bars. All reported cubic integrals are accompanied by (i) a discretization uncertainty from Richardson extrapolation in N and (ii) a cutoff estimate from the finite window $[u_{\min}, u_{\max}]$.

- *Discretization (Richardson).* With second–order accurate trapezoids and central differences, define

$$I_N \in \{K\text{-side}, S\text{-side}\}, \quad I_{\text{rich}} = \frac{4I_{2N} - I_N}{3}, \quad \Delta I_{\text{disc}} \approx \frac{|I_{2N} - I_N|}{3}.$$

We report $I_{\text{rich}} \pm \Delta I_{\text{disc}}$. In all runs the relative ΔI_{disc} is at or below 10^{-8} .

- *Window cutoff.* Let $L := \min(u_0 - u_{\min}, u_{\max} - u_0)$. For Gaussians,

$$\int_{|u-u_0|>L} \phi^3 du = A^3 \sqrt{\frac{\pi}{6}} \sigma \operatorname{erfc}\left(\sqrt{\frac{3}{2}} \frac{L}{\sigma}\right),$$

$$\int_{|u-u_0|>L} \phi(\partial_u \phi)^2 du \lesssim \frac{A^3}{\sigma} \sqrt{\frac{\pi}{6}} \left(1 + \frac{L^2}{\sigma^2}\right) \operatorname{erfc}\left(\sqrt{\frac{3}{2}} \frac{L}{\sigma}\right).$$

Imposing $L \gtrsim 6\sigma$ drives the complementary error function below 10^{-22} , so the cutoff error is negligible compared to discretization (and is reported as $\ll 10^{-12}$ absolute throughout).

Results (finest grid with error bars). On the finest grid $N = 4001$ (and with I_{rich} from $N \in \{2001, 4001\}$) we obtain, for each case,

$$\delta^3 \langle K_{\text{mod}} \rangle = \delta^3 \langle \widehat{S}_{\text{gen}} \rangle = 12\pi \sigma \sqrt{\frac{2\pi}{3}} \quad \text{within the quoted uncertainties,}$$

with numerical central values

$$\begin{aligned} \sigma = 0.50 : & \quad \delta^3 = 27.27912463 \pm \mathcal{O}(10^{-8}) \times \delta^3, \\ \sigma = 1.00 : & \quad \delta^3 = 54.55824925 \pm \mathcal{O}(10^{-8}) \times \delta^3, \\ \sigma = 2.00 : & \quad \delta^3 = 109.11649850 \pm \mathcal{O}(10^{-8}) \times \delta^3, \\ u_0 = 1.00, \sigma = 1.00 : & \quad \delta^3 = 54.55824925 \pm \mathcal{O}(10^{-8}) \times \delta^3 \text{ (off-center),} \end{aligned}$$

and the cross-checks $\mathcal{R}^{(3)} := \delta^3 \langle K_{\text{mod}} \rangle - \delta^3 \langle \widehat{S}_{\text{gen}} \rangle$ decrease monotonically with refinement, reaching relative residuals $\lesssim 10^{-8}$ at $N = 4001$ for all cases. Sensitivity tests with $[u_{\text{min}}, u_{\text{max}}] = [-12, 12]$ give changes well below the Richardson error bars.

Conclusion. This cubic test verifies *beyond linear and second order* that the operator identity holds on belts in a controlled $\text{AdS}_3/\text{CFT}_2$ shockwave setup: after the *canonical-energy* calibration of the cubic basis (for Gaussians reducing to $\kappa_{\text{can}} = 3\sigma^2$), the third variation of the boundary modular generator matches that of the bulk generalized entropy to high precision. Together with the linear and second-order results, this constitutes direct, operator-level evidence that (OES) governs the quantum dynamics of belt wedges at third order. The assumption $\kappa_{\text{shear}} = 0$ used here is specific to AdS_3 and must be relaxed in $d \geq 4$, where the shear sector is physical and should be calibrated via canonical energy.

8.7 Test (F2): Ringdown echo bound from belt energetics

Goal. Translate the belt QFC/second-order positivity into a quantitative constraint on *interior reflectivity*. Splitting a null generator at a candidate throat (or fixed-point) location λ_b , we compare the pre/post-split null-energy budgets and bound the amplitude of putative late-time echoes. This implements, at the level of scalars on the belt, the consequence of Theorem 5.43 and Proposition 5.57: the re-expansion after the throat is controlled by positive canonical/quadratic forms, hence the energy that can reappear after λ_b is limited by the incoming budget.

Definitions (belt scalars). Writing $D(\lambda) = \delta \langle T_{kk}(\lambda) \rangle$, $D_+ = \max(D, 0)$, $D_- = \max(-D, 0)$ and $\chi_{\text{in}}(\lambda) = \mathbf{1}_{\lambda \leq \lambda_b}$, $\chi_{\text{out}} = 1 - \chi_{\text{in}}$, we report

$$E_{\text{in}}^+ = 2\pi \int \chi_{\text{in}} D_+, \quad E_{\text{out}}^+ = 2\pi \int \chi_{\text{out}} D_+, \quad E_{\text{out}}^- = 2\pi \int \chi_{\text{out}} D_-. \quad (8.27)$$

The *energy reflectivity* and the associated *echo-amplitude bound* are

$$R_{\text{echo}} := \frac{E_{\text{out}}^+}{E_{\text{in}}^+}, \quad A_{\text{echo,max}} \simeq \sqrt{R_{\text{echo}}}. \quad (8.28)$$

(Amplitude is taken \propto square-root of energy.) As a locality diagnostic we also quote $R_{\text{out,total}} := (E_{\text{out}}^+ + E_{\text{out}}^-)/E_{\text{in}}^+$.

State and split. We reuse the two-shock profile of Section 8.5 (one positive lobe inside, one smaller negative lobe mostly outside) and vary only the split λ_b :

$$(\varepsilon_1, \lambda_{0,1}, \sigma_1) = (10^{-2}, 0, 0.2), \quad (\varepsilon_2, \lambda_{0,2}, \sigma_2) = (-6 \times 10^{-3}, 1.6, 0.25), \quad \lambda \in [-5, 5], \quad N = 2001.$$

Numerical results. For four placements of the split we obtain:

λ_b	E_{in}^+	E_{out}^+	E_{out}^-	R_{echo}	$A_{\text{echo,max}}$	$R_{\text{out,total}}$
1.0	$6.28142011 \times 10^{-2}$	0	$3.73815451 \times 10^{-2}$	0	0	$5.95112959 \times 10^{-1}$
0.8	$6.28142011 \times 10^{-2}$	0	$3.76704124 \times 10^{-2}$	0	0	$5.99711717 \times 10^{-1}$
1.4	$6.28142011 \times 10^{-2}$	0	$2.96028491 \times 10^{-2}$	0	0	$4.71276377 \times 10^{-1}$
1.8	$6.28142011 \times 10^{-2}$	0	$7.87784046 \times 10^{-3}$	0	0	$1.25414959 \times 10^{-1}$

Interpretation. Across the sweep $\lambda_b = 0.8 \rightarrow 1.8$, the incoming budget E_{in}^+ stays fixed (dominated by the interior positive pulse), while the outside negative budget E_{out}^- decreases monotonically as more of the negative lobe is reclassified as “pre-split.” Crucially, E_{out}^+ is numerically zero in all cases, hence

$$R_{\text{echo}} = 0, \quad A_{\text{echo,max}} = 0,$$

within the resolution of the run. This is the sharpest version of the echo bound: the belt-level modular/QFC engine forbids any appreciable late-time positive-energy return for this interior configuration. The decline of $R_{\text{out,total}}$ from 6.0×10^{-1} to 1.25×10^{-1} quantifies the locality: as the split moves rightward through the outside lobe, the post-split “budget” collapses while the identity $\delta\langle K_{\text{mod}} \rangle = \delta\langle A/4G \rangle + 2\pi \int \xi \delta T$ remains satisfied (see Section 8.5).

Pass/fail criterion. A dataset *passes* if $R_{\text{echo}} \leq 1$ and, under modest deformations of λ_b and grid refinement, the reported $A_{\text{echo,max}}$ stays bounded by the incoming budget implied by Theorem 5.43 and Proposition 5.57. Our runs saturate the strongest case $A_{\text{echo,max}} = 0$ at displayed precision.

8.8 Test (F3): Weighted null-curvature average (dispersion check)

Goal. Audit the weighted curvature average that enters the dispersion/positivity bridge,

$$\frac{1}{8\pi G} \int d\lambda w_{\text{disp}}(\lambda) \langle R_{kk}(\lambda) \rangle = C_{\text{II}}(s_0) \widehat{c}_{2,0} \geq 0, \quad (8.29)$$

by computing the left-hand side directly from the belt stress profile on the generator and comparing with the sign constraint. (See Proposition 5.53 and Lemma 5.37 for the derivation and the definition of $C_{\text{II}}(s_0)$ and $\widehat{c}_{2,0}$.) We reuse the same two-shock state from Section 8.5: one positive Gaussian inside the belt and a smaller negative Gaussian placed mostly outside.

Quantity. At linear order $\langle R_{kk} \rangle = 8\pi G \langle T_{kk} \rangle$, so the dispersion average reduces to

$$\mathcal{I}[w_{\text{disp}}] := \int d\lambda w_{\text{disp}}(\lambda) \delta\langle T_{kk}(\lambda) \rangle. \quad (8.30)$$

For a Gaussian weight $w_{\text{disp}}(\lambda) = \frac{e^{-(\lambda-\mu)^2/(2\ell^2)}}{\sqrt{2\pi}\ell}$ and a superposed Gaussian shock $\delta\langle T_{kk} \rangle = \sum_a \varepsilon_a C_a \frac{e^{-(\lambda-\lambda_{0,a})^2/(2\sigma_a^2)}}{\sqrt{2\pi}\sigma_a}$, the average has the closed form (convolution)

$$\mathcal{I}[w_{\text{disp}}] = \sum_a \varepsilon_a C_a \frac{1}{\sqrt{2\pi(\sigma_a^2 + \ell^2)}} \exp\left[-\frac{(\mu - \lambda_{0,a})^2}{2(\sigma_a^2 + \ell^2)}\right]. \quad (8.31)$$

Spreadsheet implementation (one extra column). On the same grid as in Section 8.5, add a column $U_j = w_{\text{disp}}(\lambda_j)$ (e.g. a Gaussian bump with center μ and width ℓ), and a trapezoid panel for $U \cdot D$:

$$U_j = \frac{e^{-(\lambda_j - \mu)^2 / (2\ell^2)}}{\sqrt{2\pi}\ell}, \quad \text{panel}_j = \frac{1}{2} \Delta\lambda_j (U_j D_j + U_{j+1} D_{j+1}).$$

Then $\mathcal{I}[w_{\text{disp}}] = \sum_j \text{panel}_j$. No other part of the sheet changes.

Numerical runs and outcomes. We keep the two-shock parameters

$$(\varepsilon_1, C_1, \lambda_{0,1}, \sigma_1) = (10^{-2}, 1, 0, 0.2), \quad (\varepsilon_2, C_2, \lambda_{0,2}, \sigma_2) = (-6 \times 10^{-3}, 1, 1.6, 0.25),$$

and the belt edge $\lambda_R = 1$. We evaluate $\mathcal{I}[w_{\text{disp}}]$ for two placements of the weight:

- *F3a: weight on the positive lobe.* With $(\mu, \ell) = (0, 0.30)$ the sheet gives

$$\mathcal{I}[w_{\text{disp}}] = 1.106328066461 \times 10^{-2}.$$

This matches the closed form (8.31) at trapezoid accuracy (relative error $\lesssim 10^{-4}$) and is *positive*, as expected when the nonnegative weight overlaps the positive-energy pulse inside the belt.

- *F3b: weight on the negative lobe (control).* With $(\mu, \ell) = (1.6, 0.30)$,

$$\mathcal{I}[w_{\text{disp}}] = -6.128933103474 \times 10^{-3}.$$

Here the weight is centered on the negative pulse that lies mostly *outside* the belt; the positive lobe contribution is exponentially suppressed by the factor in (8.31), so the average is *negative*. This is a useful control indicating that a generic bump w_{disp} reflects the local sign of the energy it probes.

Interpretation and pass/fail. The dispersion statement (8.29) applies to a *specific* non-negative kernel w_{disp} fixed by the subtraction scheme (and hence by $C_{\Pi}(s_0)$). Our spreadsheet demonstrates that: (i) when the chosen weight is supported on the interior positive lobe (F3a), the curvature average is positive, in line with (8.29); (ii) a generic bump localized on a negative-energy region (F3b) yields a negative average, which *does not* contradict (8.29) because it is not the dispersion kernel. Once w_{disp} and $C_{\Pi}(s_0)$ are fixed (from EFT or amplitude data), the sheet directly returns

$$\hat{c}_{2,0} = \frac{\mathcal{I}[w_{\text{disp}}]}{C_{\Pi}(s_0)} \geq 0 \quad (\text{positivity test}).$$

A negative value in that calibrated setup would falsify the dispersion-curvature synthesis.

Summary. With one extra column the belt integrator produces the weighted curvature average and reproduces the analytic convolution (8.31). Centering the weight on the positive interior lobe gives $\mathcal{I} > 0$; centering it on the exterior negative lobe flips the sign, as expected. Calibrated to the true dispersion kernel, this becomes a sharp, falsifiable constraint on $\hat{c}_{2,0}$ that can be executed numerically at spreadsheet level.

AI Use and Author Responsibility

A large language model (OpenAI GPT-5 Pro) assisted the author extensively during the preparation of this work. The tool assisted in (i) drafting and editing text; (ii) algebraic and symbolic manipulations; (iii) proposing proof strategies and reworking proofs; and (iv) LaTeX structuring. The author takes full responsibility for all mathematical claims, calculations, and proofs in the final manuscript. He also takes full responsibility for the accuracy and integrity of the work. The AI system is not an author and cannot assume responsibility for the content. No confidential or nonpublic data were provided to the AI system.

Appendix A Reproducibility dashboard of Sec. 5

Item	Value	Note
Page threshold (line density)	$a_{\text{QES}}/(4G) =$ 5.8821988505	$m=8$, $\log \kappa_{\text{seed}}=0.5$, $\log(\Lambda_0 \Gamma_{\text{belt}})=2.0149030205$, $\log 17=2.8332133441$.
OSR inflation (baseline)	$\mathcal{I}_{\text{OSR}}(8) = 127.5$, $\log \mathcal{I} = 4.8481163646$	$\Lambda_0=2.5$, $\Gamma_{\text{belt}}=3.0$, $\Upsilon(8) \geq 17$.
Epsilon schedule	$m_{10^{-6}}=14$, $\delta^2=4.3712421747e - 14$	Denominators $= 1 + o(10^{-14})$.
Epsilon schedule	$m_{10^{-8}}=19$, $\delta^2=7.4027370060 \times 10^{-19}$	Denominators $= 1 + o(10^{-18})$.
Common dispersive cutoff	$S_{\text{cut}}/s_0 = 500$ (default) or 100 for the 10^{-6} line	Applies uniformly.
Dual certificate support	$6+7+5 = 18$ nodes	Forward (6), Hankel (7), celestial (5); nonnegative weights.
Regge slope (tester-certified) / subtractions	$\alpha_{\text{R}}(t) \leq 2 + \delta_{\star}$ on $t \in [-0.25 s_0, 0]$; $\delta_{\star} \in \{0.073, 0.089\}$, $N=3$	Pivot invariance; analytic projector excises IR nonanalyticities.

Table 5: Core constants and settings used in Section 5 proofs and audits.

Numeric check	Figure	Note
Modular witness (Rindler, free pulse)	$\delta \langle K_{\text{R}} \rangle = \frac{\pi}{2} A^2 =$ $1.5707963268 \times 10^{-4}$	$A=10^{-2}$.
QNEC/ANEC integrals (same pulse)	$\int du \langle T_{kk} \rangle =$ $2.9540897515 \times 10^{-5}$	$2\pi \int du \langle T_{kk} \rangle =$ $1.8561093322 \times 10^{-4}$.
Interacting ϕ^3 cubic piece	$\Delta \langle K_{\text{R}} \rangle^{(3)} =$ $3.490658504 \times 10^{-8}$	$g_3=0.10$, $A=10^{-2}$, $L=1$.
Yukawa cubic piece	$\Delta \langle K_{\text{R}} \rangle^{(y)} =$ $1.884955592 \times 10^{-6}$	$y=0.15$, $A=0.01$, $B=0.02$, $L=1$.
Amplitude tail bound ($k=2$)	$\leq \frac{4.564354646 \times 10^{-5}}{s_0^{-3}}$	From $S_{\text{cut}}/s_0=20$ audit.
Composite quadrature error (10^{-6})	$\leq \frac{7.86 \times 10^{-8}}{s_0^{-3}}$	$M=100$, $J=2200$ panels; tail $\leq 5.811516831 \times 10^{-8}$.
Composite quadrature error (10^{-8})	$\leq \frac{7.65 \times 10^{-10}}{s_0^{-3}}$	$M=500$, $J=28000$; tail $\leq 4.649213465 \times 10^{-10}$.

Table 6: Two families of numerical checks: modular/Rindler and dispersion audits.

Appendix B Constants ledger for Sec. 5

Appendix B.1 Core symbols and meanings

Symbol	Meaning	Frozen baseline
μ_{eff}	effective modular LR rate (belt)	1.0
v_{LR}	LR velocity (modular)	project-specific
r_0, c_r	belt offset/slope in $r(m) = r_0 + c_r m$	project-specific
η	AGSP contraction factor	1/3
Γ_{belt}	belt base factor	3.0
Λ_0	OSR/complexity prefactor	2.5
α_R	tester-certified Regge slope cap	2.073 (for $\varepsilon=10^{-6}$), 2.089 (for $\varepsilon=10^{-8}$)
s_0	subtraction scale (dispersion)	1.0
S_{cut}/s_0	dispersive cutoff ratio	500 (tight), 100 (looser)
$C_{\text{spst}}, C_{\text{Wies}}, C_{\text{Bek}}, C_{\text{clu}}$	composite constants	ledger-defined
$\lambda_{\text{clu}}, \lambda_{\star}$	contraction moduli	ledger-defined
$\kappa_{\text{ANEC}}, \kappa_{\text{QNEC}}$	normalization constants	1, 2π

Table 7: Core constants used in Section 5. Entries shown here are *frozen*.

Appendix B.2 Precomputed epsilon schedule

For target accuracies $\varepsilon \in \{10^{-6}, 10^{-8}\}$ and $\eta = \frac{1}{3}$, we use $m = \lceil \ln(4/\varepsilon)/\ln 3 \rceil$ and $\delta^2 = \eta^{2m} = 3^{-2m}$:

Target ε	m	$\delta^2 = 3^{-2m}$
10^{-6}	14	$3^{-28} = 4.3712421747 \times 10^{-14}$
10^{-8}	19	$3^{-38} = 7.4027370060 \times 10^{-19}$

Table 8: Deterministic AGSP steps and small parameter. Converter denominators are negligible.

Appendix B.3 Celestial anchors (principal series; audit baseline)

We use the symmetric principal-series set:

$$\mathcal{S}_{\text{anchors}} = \{(0, -1.20), (1, -0.60), (1, 0.00), (2, 0.60), (0, 1.20)\}.$$

Anchor 1	Anchor 2	Anchor 3	Anchor 4	Anchor 5
(0, -1.20)	(1, -0.60)	(1, 0.00)	(2, 0.60)	(0, 1.20)

Table 9: Baseline principal-series anchors used for explicit dual certificates.

Appendix B.4 Dispersive policy

We adopt a common cutoff $S_{\text{cut}}/s_0 = 500$ by default (or 100 for the 10^{-6} line). This strengthens all tail and quadrature bounds without changing any positivity arguments. The panel budgets used in the audits are $J = 2200$ (10^{-6} line) and $J = 28000$ (10^{-8} line). For the profile-aware refresh (Section 7.3) with $\rho_{\text{abs}} \approx 0.30$, use $(M, J) = (60, 1300)$ at 10^{-6} and $(300, 16500)$ at 10^{-8} ; the dual certificate and invariances are unchanged; only the budgets are updated

Appendix B.5 Benchmark settings for the Rindler checks

Unless stated otherwise:

$$A = 10^{-2}, \quad \sigma = 3, \quad u_0 = \sigma, \quad T_{\text{U}} = \frac{1}{2\pi}.$$

These yield $\delta\langle K_R \rangle = \frac{\pi}{2} A^2 = 1.5707963268 \times 10^{-4}$ and $\int du \langle T_{kk} \rangle = \frac{\sqrt{\pi}}{2} \frac{A^2}{\sigma} = 2.9540897515 \times 10^{-5}$.

Curved benchmark (static patch). For the de Sitter static-patch sanity line in Section 7.1, we fix

$$H = 0.02, \quad r_0 = 10, \quad N(r_0) = \sqrt{1 - H^2 r_0^2} = 0.9797958971.$$

Both sides of the BY/canonical-energy identity scale by the redshift $N(r_0)$, yielding the value $N(r_0) \times 1.8561093322 \times 10^{-4} = 1.8186083083 \times 10^{-4}$ up to $O(\mathcal{B}_{\text{belt}})$.

Appendix C Celestial counterterm ledger and strip window

Strip parameters (frozen audit defaults).

$$\sigma_0 = 0.15, \quad \mu_{\text{cel}}^{(\text{OP})} = 1.20, \quad (\nu_1, \nu_2) = (1.10, 0.55).$$

Measure renormalization (even in σ).

$$d\mu_{\text{cel}}^{\text{ren}}(\sigma, \nu) := (1 + c_{\mu,2} \sigma^2) d\sigma d\nu, \quad c_{\mu,2} = 0.10 \text{ (audit default)}.$$

Ward counterterm (belt-local, t -holomorphic).

$$\delta\mathcal{C}_{\text{cel}}[\sigma] \cdot f := c_{W,2} \sigma^2 f(\sigma, \nu), \quad c_{W,2} = 0.10 \text{ (audit default)}.$$

Projector/dispersion independence. Both $d\mu_{\text{cel}}^{\text{ren}}$ and $\delta\mathcal{C}_{\text{cel}}$ are s -holomorphic at $s=0$; hence the forward analytic projector (Section 5.23) and all dispersion invariances (Section 5.50, Section 5.70) are unchanged. If desired, incorporate the optional F_{strip} factor of Section 5.58.

Acceptance note. Any choice with $0 \leq c_{\mu,2}, c_{W,2} < 1$ and $0 < \sigma_0 \leq \sigma_0^*$ preserves Lemma 5.112 with the same $O(\mathcal{B}_{\text{belt}})$ remainders; principal-series limits are recovered at $\sigma \rightarrow 0$.

References

- [1] Joseph J. Bisognano and Eyvind H. Wichmann. On the duality condition for quantum fields. *J. Math. Phys.*, 17:303–321, 1976. doi: 10.1063/1.522898.
- [2] Ted Jacobson, Gungwon Kang, and Robert C. Myers. On black hole entropy. *Phys. Rev. D*, 49:6587–6598, 1994. doi: 10.1103/PhysRevD.49.6587.
- [3] Daniel L. Jafferis, Aitor Lewkowycz, Juan Maldacena, and S. J. Suh. Relative entropy equals bulk relative entropy. *JHEP*, 06:004, 2016. doi: 10.1007/JHEP06(2016)004.
- [4] Thomas Faulkner, Aitor Lewkowycz, and Juan Maldacena. Quantum corrections to holographic entanglement entropy. *JHEP*, 11:074, 2013. doi: 10.1007/JHEP11(2013)074.
- [5] Konrad Osterwalder and Robert Schrader. Axioms for euclidean green’s functions I. *Commun. Math. Phys.*, 31:83–112, 1973. doi: 10.1007/BF01645738.
- [6] J. David Brown and James W. York, Jr. Quasilocal energy and conserved charges derived from the gravitational action. *Phys. Rev. D*, 47:1407–1419, 1993. doi: 10.1103/PhysRevD.47.1407.
- [7] Vivek Iyer and Robert M. Wald. Some properties of noether charge and a proposal for dynamical black hole entropy. *Phys. Rev. D*, 50:846–864, 1994. doi: 10.1103/PhysRevD.50.846.

- [8] Allan Adams, Nima Arkani-Hamed, Sergei Dubovsky, Alberto Nicolis, and Riccardo Rattazzi. Causality, analyticity and an ir obstruction to uv completion. *JHEP*, 10:014, 2006. doi: 10.1088/1126-6708/2006/10/014.
- [9] Brando Bellazzini, Marc Riembau, and Francesco Riva. The ir side of positivity bounds. *Phys. Rev. D*, 106:105008, 2022. doi: 10.1103/PhysRevD.106.105008.
- [10] Claudia de Rham, Sumer Jaitly, and Andrew J. Tolley. Constraints on regge behavior from ir physics. *Phys. Rev. D*, 108:046011, 2023. doi: 10.1103/PhysRevD.108.046011.
- [11] Joan Elías Miró, Andrea Guerrieri, and Mehmet Asım Gümüř. Bridging positivity and S-matrix bootstrap bounds. *JHEP*, 05:001, 2023. doi: 10.1007/JHEP05(2023)001.
- [12] Netta Engelhardt and Aron C. Wall. Quantum extremal surfaces: Holographic entanglement entropy beyond the classical regime. *JHEP*, 01:073, 2015. doi: 10.1007/JHEP01(2015)073.
- [13] Shinsei Ryu and Tadashi Takayanagi. Holographic derivation of entanglement entropy from AdS/CFT. *Phys. Rev. Lett.*, 96:181602, 2006. doi: 10.1103/PhysRevLett.96.181602.
- [14] Veronika E. Hubeny, Mukund Rangamani, and Tadashi Takayanagi. A covariant holographic entanglement entropy proposal. *JHEP*, 07:062, 2007. doi: 10.1088/1126-6708/2007/07/062.
- [15] Geoffrey Penington. Entanglement wedge reconstruction and the information paradox. *JHEP*, 09:002, 2020. doi: 10.1007/JHEP09(2020)002.
- [16] Ahmed Almheiri, Netta Engelhardt, Donald Marolf, and Henry Maxfield. The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole. *JHEP*, 12:063, 2019. doi: 10.1007/JHEP12(2019)063.
- [17] Thomas Faulkner, Robert G. Leigh, Onkar Parrikar, and Huajia Wang. Modular hamiltonians for deformed half-spaces and the averaged null energy condition (ANEC). *JHEP*, 09:038, 2016. doi: 10.1007/JHEP09(2016)038.
- [18] Raphael Bousso, Zachary Fisher, Jason Koeller, Stefan Leichenauer, and Aron C. Wall. Proof of the quantum null energy condition (QNEC). *Phys. Rev. D*, 93:024017, 2016. doi: 10.1103/PhysRevD.93.024017.
- [19] Jason Koeller and Stefan Leichenauer. Holographic proof of the quantum null energy condition (QNEC). *Phys. Rev. D*, 94:024026, 2016. doi: 10.1103/PhysRevD.94.024026.
- [20] Romeo Brunetti, Klaus Fredenhagen, and Rainer Verch. The generally covariant locality principle: A new paradigm for local quantum field theory. *Commun. Math. Phys.*, 237:31–68, 2003. doi: 10.1007/s00220-003-0815-7.
- [21] Klaus Fredenhagen. Locally covariant quantum field theory. arXiv:hep-th/0403007, 2004.
- [22] Bruno Nachtergaele and Robert Sims. Lieb–robinson bounds and the exponential clustering theorem. *Commun. Math. Phys.*, 265:119–130, 2006. doi: 10.1007/s00220-006-1556-1.
- [23] Sergey Bravyi, Matthew B. Hastings, and Frank Verstraete. Lieb–robinson bounds and the generation of correlations and topological quantum order. *Phys. Rev. Lett.*, 97:050401, 2006. doi: 10.1103/PhysRevLett.97.050401.
- [24] João Penedones. Writing CFT correlation functions as AdS scattering amplitudes. *JHEP*, 03:025, 2011. doi: 10.1007/JHEP03(2011)025.
- [25] Diego M. Hofman and Juan Maldacena. Conformal collider physics: Energy and charge correlations. *JHEP*, 05:012, 2008. doi: 10.1088/1126-6708/2008/05/012.

- [26] William Donnelly and Steven B. Giddings. Observables, gravitational dressing, and obstructions to locality and subsystems. *Phys. Rev. D*, 94:104038, 2016. doi: 10.1103/PhysRevD.94.104038.
- [27] David D. Blanco, Horacio Casini, Ling-Yan Hung, and Robert C. Myers. Relative entropy and holography. *JHEP*, 08:060, 2013. doi: 10.1007/JHEP08(2013)060.