

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

F. Jofer

## Abstract

We<sup>1</sup> investigate a belt-local, operator-level equation of state (OES) as a working postulate: 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) serves as our organizing principle. Within this framework 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 AdS<sub>3</sub>/CFT<sub>2</sub> 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 find 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 is compatible with wedge dynamics at cubic order in this setting.

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. Taken together, these results support OES as a viable candidate organizing principle for semiclassical wedge dynamics in the regimes we currently test.

## 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 a 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  $\widehat{S}_{\text{gen}}$  is not the von Neumann generalized-entropy operator: its expectation is taken with the bulk modular generator  $\widehat{H}_{\text{bulk}}$ , so  $\langle \widehat{S}_{\text{gen}} \rangle \neq S_{\text{gen}}$  beyond linear order. More 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 5.20, Lemmas 3.1 and 3.3, and Section 5.5. In Section 5.5 we first show that, given our belt kernel (modular first-law channel, JLMS-type channel map, OS/KMS positivity with flow removal, and JKM/Brown–York normalization), the closed quadratic forms of  $\widehat{K}_{\text{mod}}(R)$  and  $\widehat{S}_{\text{gen}}(W)$  coincide (Theorem 5.14). Thus OES is not an additional independent hypothesis on top of the belt kernel, but a repackaging

---

<sup>1</sup>Throughout, “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”.

of it into a single operator statement. For the remainder of the paper we then *adopt* OES as Axiom 5.13, develop its operator meaning in Section 5.5, and fix its normalization by a Rindler witness and the JKM calibration. Taking expectations and variations of the operator identity reproduces the familiar linear and second-order modular/Einstein relations on belts.

**What we do.** We give a belt-local route to semiclassical quantum gravity that is organized around OES but logically grounded in three inputs: (i) a belt first-law channel and modular convexity; (ii) OS (RP/KMS) positivity with removal of short flows; and (iii) a JLMS-type boundary–bulk map realized as a Brown–York/Iyer–Wald flux dictionary. These three assumptions define our *belt kernel*. Within this kernel we obtain regulator-stable identities (errors  $O(\mathcal{B}_{\text{belt}})$ ) that become exact after flow removal. Section 5.5 shows that the kernel implies OES on belts; conversely, in later sections we freely use OES as shorthand for working within this kernel.

**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  $\text{AdS}_3/\text{CFT}_2$  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  $O(\mathcal{B}_{\text{belt}})$ . 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, yielding convexity/monotonicity; and (3) a belt JLMS identification that matches boundary relative entropy to wedge canonical energy up to  $O(\mathcal{B}_{\text{belt}})$ . These inputs ensure that the belt-level laws are stable, additive on overlaps, and compatible with quasi-local flux transport; they are taken as standing assumptions in the belt kernel and are later repackaged into OES via Theorem 5.14.

**Global dynamics and stitching.** From the kernel we build a global belt atlas: we define admissible belt variations, a modular defect, and a modular Hessian, and show 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, *under the kernel assumptions*, a regulator-independent realization of OES on each belt, recovering the linearized and quantified second-order semiclassical Einstein equations in expectation (see Section 5.5 and Theorem 5.14).

**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. Equivalent modular, group, channel, and relative formulations are given in Section 5.5. Theorem 5.14 shows that this axiom is implied by the belt kernel (first-law matching, JLMS compatibility, normalization, and additivity) and hence is not logically independent.

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).
3. **Cubic (third-order) holographic verification.** A belt-local AdS<sub>3</sub>/CFT<sub>2</sub> 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 5.62.
4. **Positivity bridges.** Gravity-subtracted dispersion and celestial/impact-parameter positivity on a forward cone remain compatible with the belt laws and furnish independent tests and normalizations used later in the paper.

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

**Notation.** All entropic and energetic quantities are per generator length. The belt budget  $\mathcal{B}_{\text{belt}}$  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, a belt JLMS identification matching boundary relative entropy to wedge canonical energy, analytic cores, and the Brown–York/JKM calibration interface, 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, and the kernel forces a regulator-independent realization of OES on each belt (cf. Section 5.5 and Theorem 5.14).
- *Section 5 (Four pillars)*. In Section 5.5 we formulate the *Operator Equation of State* as Axiom 5.13, prove that it is implied by the belt kernel in Theorem 5.14, and record equivalent modular/group/channel/relative formulations on the common domain. From this operator axiom, understood as shorthand for working within the kernel, we derive: QES/Page behavior; ANEC/QNEC from modular positivity; dispersive/celestial positivity with Regge control; 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, 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  $\delta^3\langle \widehat{S}_{\text{gen}}\rangle$  in AdS<sub>3</sub>/CFT<sub>2</sub>; (F2) a ringdown echo bound from belt energetics; (F3) a dispersion–curvature average linking  $\widehat{c}_{2,0}$  to a weighted null-curvature functional.
- *Section 9 (Outlook).* We summarize the belt-OES picture, collect the lessons from singularity resolution and dispersive/positivity tests, and outline extensions and open directions: higher-order modular dynamics, non-Einsteinian corrections, non-AdS and cosmological belts, and sharpened observational probes of the predictions (F1)–(F3).

## 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 5.86 and 5.93), canonical-energy and stability inputs (Theorem 5.46, Lemma 5.112, and Proposition 5.102), and the budget/remainder calculus (Sections 5.12 and 5.13).

*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 5.39).

### 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 [1, 2]. 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 [3, 4] with Lieb–Robinson-type control [5, 6]. 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 [4]. 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  $\mathcal{M}_{\text{sub}}$  admits a crossing-symmetric dispersion relation at subtraction order  $N \geq 3$ , with Regge-compatible polynomial boundedness [7–10]. 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 [1, 2].

(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 [11–13].

(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  $\psi$  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 5.109, Proposition 5.86, and Theorem 5.104). 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  $\mathcal{A}(R) := \mathcal{A}_{\text{bulk}}(R)$  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 5.75) 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  $\rho_R$  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  $\rho$  on  $\mathcal{A}_{\text{belt}}$  with  $\rho \upharpoonright_{\mathcal{A}(R)} = \rho_R$  for every  $R$ . Conversely, any normal  $\rho$  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  $\rho$  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  $\rho$ ,

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

with  $r$  the half-width of  $B$ . There exists a belt-compatible recovery map such that the refined DPI bound of Proposition 5.86 holds, and the Markov gap obeys Theorem 5.104 (per generator length).

*Corollary 2.7* (Global GNS and modular compatibility). Let  $\rho$  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  $\rho$  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  $\rho, \sigma$ ,

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

**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  $\mathcal{M}_{\text{raw}}(s, t)$  [14].

(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 [15, 16].

(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) [7–10].

**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 [7, 8].

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  $t < 0$ ) [9, 10].

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  $u$  [7, 8].
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 [9, 10].
6. **Regulator stability.** Changing admissible belt families or modular filters alters  $\mathcal{M}_{\text{sub}}$  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 5.46 and 5.66).

## 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  $(g_u, \tilde{r})$ ,

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

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

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

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

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

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  $A_i \in \mathcal{A}_{\text{bulk}}(R_i)$ ,

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

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  $K$ ,

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

6. **Budget calculus (monotonicity/composition).** For belt-controlled maps  $\Phi, \Psi$  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;  $d_{\text{anch}}$  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

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

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  $\xi$  [3, 4]. 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 [11] with Brown–York/Iyer–Wald identification for the flux/area terms [17, 18].

*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 5.113).

*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 5.22).

## 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 5.67). 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 5.77 and Propositions 5.40 and 5.78. Operator JLMS and the operator equation of state are Proposition 5.8 and Theorem 5.9. Second–order control is Theorem 5.46. Flow monotonicity is Theorem 5.41.

**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)  $\rho_\theta$  is generated by a belt–supported normal automorphism within the OS window (Lemma 3.1); (ii)  $R_\theta$  is a smooth shape deformation tangent to the belt boost near the entangling belt with JKM corner calibration (Lemma 5.77); (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 5.46 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 5.77 and Proposition 5.78).

*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  $E^{\text{can}}$  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 5.67). Edge/corner terms are fixed by JKM (Lemma 5.77) and absorbed into  $S_{\text{edge}}$  (Proposition 5.40); fluxes transport by Brown–York (Proposition 5.78). Second order uses Theorem 5.46. 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  $\delta$  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 5.77) and BY flux (Proposition 5.78).  $\square$

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

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

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

*Sketch.* Apply Stokes to the symplectic current with calibrated corners (Lemma 5.77 and Proposition 5.40); Brown–York transports the bulk canonical–energy flux to the belt (Proposition 5.78). 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.

$$\delta \langle K_{\text{mod}} \rangle = \delta \left( \frac{\text{Area}}{4G} \right) + 2\pi \int_R \xi^\nu d\Sigma_\mu \delta \langle T^\mu{}_\nu \rangle + O(\mathcal{B}_{\text{belt}}) \quad (\text{Theorem 5.37}),$$

with the second–order inequality/identity of Theorem 5.46, and the operator identity  $K_{\text{mod}} = \frac{A}{4G} + H_{\text{bulk}}$  on the common analytic core (Axiom 5.13; cf. Theorems 5.21 and 5.25).

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 5.77 and Proposition 5.78) and the equation of state (Axiom 5.13, Theorem 5.37), 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 5.41).

*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.

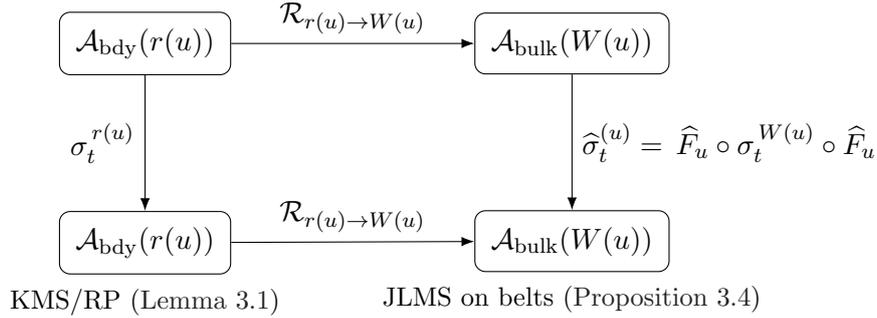


Figure 1: Commuting square on a belt at finite positive flows  $(u; 0)$ . The CP modular filter  $\widehat{\text{Fhat}}(u)$  is induced from the boundary and is removed by Lemma 3.3. The square is used to transport variations and fluxes across the belt; the operator equation of state appears in Axiom 5.13; cf. Theorems 5.21 and 5.25.

### 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 5.16 and Proposition 5.40. Quasi-locality and wedge generation are handled via belt microcausality tails and null timeslice propagation Lemma 5.74 and Proposition 5.75. For reconstruction we invoke belt-level nesting/recovery and continuity Propositions 5.86 and 5.93 and Section 5.14. Thermal/energy bounds and canonical-energy control enter through Theorem 5.46, Lemma 5.112, and Proposition 5.102.

**Bindings (notation).** We use the belt remainder budget  $\mathcal{B}_{\text{belt}}$ , belt base factor  $\Gamma_{\text{belt}}$ , and composite constants  $C_{\text{spst}}, C_{\text{Wies}}, C_{\text{Bek}}, C_{\text{clu}}$ . Numerics refer to appendix tables (Table 4, Table 5).

**Acceptance.** Upon closing the proof kernels in this section and verifying the budget checks of Section 5.12, we record completion: the four pillar theorems Theorems 5.28, 5.29, 5.33 and 5.37 (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 5.86 and 5.93. 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 5.43 and 5.48; these serve as acceptance intensifiers only, not as axioms.

### 5.3 Kernel zero: belt first law and JLMS channel

*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  $\Gamma_{\text{belt}}$  introduced in Section 5.1. Bounds are per generator length and uniform in  $|R|$  (Remarks 3.6 and 5.36). We keep the boost normalization fixed by the Rindler witness [4] (Section 5.49 and Lemma 5.77) and the JKM/Brown–York calibration [17, 19].

*Lemma 5.2* (Analytic projector and gravity IR subtraction). Within the OS belt window, the analytic–core vectors from Proposition 5.67 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  $W = EW(R)$  up to  $O(\mathcal{B}_{\text{belt}})$ , with continuity under admissible shape/state variations:

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

where  $\xi$  is the belt boost generator. We use the convention

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

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

*Proof of Proposition 5.3 (=Proposition 3.4).* Fix a belt  $\partial_r R$  with positive flows  $(u, s) > 0$  inside the OS window. Let  $\sigma$  be the cyclic–separating reference state and  $\rho$  any admissible state.

*Preliminaries.* By OS reflection positivity/KMS on the belt (Lemma 3.1) and the domain result (Proposition 5.67), there is a common analytic core  $D_{\text{an}}$  for the boundary and wedge modular flows. Write the Umegaki relative entropy in the standard modular form

$$S(\rho_R \|\sigma_R) = \langle K_{\text{mod}}^R[\sigma] \rangle_\rho - (S(\rho_R) - S(\sigma_R)), \quad S(\rho_W \|\sigma_W) = \langle K_{\text{mod}}^W[\sigma] \rangle_\rho - (S(\rho_W) - S(\sigma_W)),$$

where  $K_{\text{mod}}^R[\sigma] := -\log \Delta_R(\sigma)$  and likewise in the wedge.

*Step 1 (Boundary  $\rightarrow$  wedge channel; belt localization).* By the rotated Petz recovery map on the belt (Lemma 3.2) there exists a belt–compatible isometry  $U_{R \rightarrow W}$  (the JLMS channel on the code subspace) such that for all  $\psi \in D_{\text{an}}$ ,

$$\langle \psi, K_{\text{mod}}^R[\sigma] \psi \rangle = \left\langle U_{R \rightarrow W} \psi, \left( \frac{A_r(W)}{4G} + K_{\text{mod}}^W[\sigma] \right) U_{R \rightarrow W} \psi \right\rangle + O(\mathcal{B}_{\text{belt}}).$$

The  $O(\mathcal{B}_{\text{belt}})$  term collects the belt remainder from finite positive flows  $(u, s)$ , belt width  $r$ , and the use of the analytic projector (bounded belt circuits/quasi–local factorization), all controlled uniformly as in Lemma 3.1 and Proposition 5.67.

*Step 2 (Corner/edge calibration).* Fix the JKM counterterm by the belt boost Ward identity (Lemma 5.77). Then the corner/edge contributions are uniformly controlled (Proposition 5.40) and, on the belt, the calibrated area variation equals the Wald corner charge up to  $O(\mathcal{B}_{\text{belt}})$ . In particular, with this calibration,

$$\langle K_{\text{mod}}^R[\sigma] \rangle_\rho - \langle K_{\text{mod}}^R[\sigma] \rangle_\sigma = \frac{\langle A_r(W) \rangle_\rho - \langle A_r(W) \rangle_\sigma}{4G} + \langle K_{\text{mod}}^W[\sigma] \rangle_\rho - \langle K_{\text{mod}}^W[\sigma] \rangle_\sigma + O(\mathcal{B}_{\text{belt}}).$$

*Step 3 (Flux matching; Brown–York).* Using Stokes’ theorem on the symplectic current with calibrated corners, the belt generator equals the Brown–York flux on the timelike belt; by the quasi–local BY dictionary (Proposition 5.78) this identifies the wedge modular generator with the canonical–energy flux along the belt boost  $\xi$ . Consequently,

$$\langle K_{\text{mod}}^W[\sigma] \rangle_\rho - \langle K_{\text{mod}}^W[\sigma] \rangle_\sigma = 2\pi E_{\text{can}}^W[\rho; \sigma; \xi] - \frac{\langle A_r(W) \rangle_\rho - \langle A_r(W) \rangle_\sigma}{4G} + O(\mathcal{B}_{\text{belt}}),$$

so that combining with Step 2 gives

$$\langle K_{\text{mod}}^R[\sigma] \rangle_\rho - \langle K_{\text{mod}}^R[\sigma] \rangle_\sigma = 2\pi E_{\text{can}}^W[\rho; \sigma; \xi] + O(\mathcal{B}_{\text{belt}}).$$

*Step 4 (Entropy bookkeeping).* By belt localization and the calibrated edge control, the outside/bulk entropy relation on belts reads

$$S(\rho_R) - S(\sigma_R) = \frac{\langle A_r(W) \rangle_\rho - \langle A_r(W) \rangle_\sigma}{4G} + (S(\rho_W) - S(\sigma_W)) + O(\mathcal{B}_{\text{belt}}).$$

Insert this and the identity from Step 3 into the boundary relative entropy:

$$\begin{aligned} S(\rho_R \| \sigma_R) &= \left[ \langle K_{\text{mod}}^R[\sigma] \rangle_\rho - \langle K_{\text{mod}}^R[\sigma] \rangle_\sigma \right] - \left[ (S(\rho_R) - S(\sigma_R)) \right] \\ &= \left[ 2\pi E_{\text{can}}^W[\rho; \sigma; \xi] + O(\mathcal{B}_{\text{belt}}) \right] - \left[ \frac{\Delta A_r(W)}{4G} + \Delta S_W + O(\mathcal{B}_{\text{belt}}) \right] \\ &= (S(\rho_W \| \sigma_W) + \frac{\Delta A_r(W)}{4G}) + O(\mathcal{B}_{\text{belt}}) \\ &= 2\pi E_{\text{can}}^W[\rho; \sigma; \xi] + O(\mathcal{B}_{\text{belt}}), \end{aligned}$$

where  $\Delta A_r(W) := \langle A_r(W) \rangle_\rho - \langle A_r(W) \rangle_\sigma$  and  $\Delta S_W := S(\rho_W) - S(\sigma_W)$ .

*Step 5 (Continuity under admissible deformations).* Each ingredient in the last display is continuous along admissible one-parameter families (OS belt window): the modular expectation by analyticity of the core and the bounded belt circuits, the entropies by recovery continuity, and the area term by the calibrated corner control. Hence  $\delta S(\rho_R \| \sigma_R)$  exists and is continuous along such deformations.

*Step 6 (Positive-flow removal).* By the removal lemma (Lemma 3.3),  $\mathcal{B}_{\text{belt}} \rightarrow 0$  as  $(u, s) \downarrow 0$  at fixed belt width  $r$ , whence the  $O(\mathcal{B}_{\text{belt}})$  remainder vanishes. Passing to the limit yields the belt JLMS identity in the regulator-independent window:

$$S(\rho_R \| \sigma_R) = 2\pi E_{\text{can}}^W[\rho; \sigma; \xi],$$

equivalently  $S(\rho_R \| \sigma_R) = S(\rho_W \| \sigma_W) + \frac{\Delta A(W)}{4G}$ . This is precisely Proposition 5.3 and, by the stated convention, 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  $\Gamma_{\text{belt}}$  and removable by Lemma 3.3. In particular, after using the JKM calibration and Brown–York dictionary, the modular equation of state of Theorem 5.37 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 5.22; 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  $\mathcal{D}_{\text{an}}$  from Proposition 5.67, 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  $\mathcal{D}_{\text{an}}$  a core and essential self-adjointness on  $\mathcal{D}_{\text{an}}$  (Proposition 5.67 and Lemma 3.1) (cf. [3]).

2. **Quantum area operator.**  $\widehat{A}_r(W)$  is the belt–local *area operator* associated with the regulated QES cross–section on  $\partial W$ , obtained by quantizing the Iyer–Wald corner charge with the JKM calibration (Lemma 5.77) and Brown–York dictionary (Proposition 5.78). Concretely,

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

normalized by the Rindler witness (Section 5.49 and Lemma 5.77). 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  $\xi$  on  $W$ , defined by the covariant–phase–space/Brown–York flux (Proposition 5.78) and calibrated corners (Section 5.49). On  $\mathcal{D}_{\text{an}}$ ,

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

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 5.77.

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

*Lemma 5.7* (Core, closability, and essential self–adjointness). On the OS belt window,  $\mathcal{D}_{\text{an}}$  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  $\mathcal{D}_{\text{an}}$  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  $U_{R \rightarrow W} : \overline{\mathfrak{A}_R \Omega} \rightarrow \overline{\mathfrak{A}_W \Omega}$  (rotated Petz/JLMS channel relative to  $\sigma$ , lifted to GNS) such that, as a quadratic-form identity on  $\mathcal{D}_{\text{an}}$ ,

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

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,

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

as an operator (quadratic–form) identity on  $\mathcal{D}_{\text{an}}$ , 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 5.37). (ii) Second variation recovers the second-order modular equation of state (Theorem 5.46) with the shear/canonical-energy terms. (iii) The Brown–York dictionary (Proposition 5.78) 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 5.21 and 5.25) 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 5.37.

## 5.5 Full quantum equation of state: postulate, generator theorem, and equivalent formulations

Fix a belt-anchored boundary region  $R$  and its entanglement wedge  $W = \text{EW}(R)$ . Work on the OS belt analytic core  $\mathcal{D}_{\text{an}}$  (“ $\mathcal{D}_{\text{an}}$ ”) from Proposition 5.67, with the belt operator lifts of Definition 5.6 and the operator JLMS isometry of Proposition 5.8. All statements are per generator length and uniform in  $|R|$ .

**Logical stance.** In this subsection we keep track of two logically compatible viewpoints. The *kernel viewpoint* starts from the existence of a belt first-law channel, channel/JLMS compatibility, a Rindler/JKM normalization, and quasi-local additivity and shows that these inputs force the full operator equation of state on belts (Theorem 5.14). The *axiomatic viewpoint* then takes the form identity of Axiom 5.13 as a compact postulate that summarizes this kernel and is used as a standing axiom in the rest of the paper. Whenever we say “assume OES” below, the reader may thus regard this as shorthand for working within the belt kernel.

**Quadratic forms and their sum.** On  $\mathcal{D}_{\text{an}}$  define the closed sectoral forms

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

*Lemma 5.12* (Form-closedness and KLMN for the generalized-entropy sum). After removal of belt/flow regulators (Lemma 3.3), the quadratic forms  $\mathfrak{a}$  and  $\mathfrak{h}$  on  $\mathsf{D}_{\text{an}}$  are densely defined, closed, and lower semibounded. Moreover,  $\mathfrak{h}$  is relatively form-bounded with respect to  $\mathfrak{a}$  with arbitrarily small bound on admissible cones:

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

with  $C_\epsilon$  independent of  $|R|$ . Hence, by KLMN,  $\mathfrak{s}_{\text{gen}}$  is a closed, lower-semibounded quadratic form that defines a unique self-adjoint operator  $\widehat{S}_{\text{gen}}$  with  $\widehat{S}_{\text{gen}} \geq -C$ .

*Axiom 5.13* (Full quantum equation of state (form identity)). After regulator removal (Lemma 3.3), on the regulator-independent domain

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

of the GNS representation  $(\mathfrak{A}_R, \Omega)$ , the boundary modular generator equals the wedge generalized-entropy operator for  $W = \text{EW}(R)$ , as an identity of closed quadratic forms:

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

The overall normalization is fixed by the boost Ward identity via the Rindler witness and the JKM calibration (Lemma 5.77); no additive central constant is allowed. For the remainder of the paper we treat Axiom 5.13 as a standing axiom. Theorem 5.14 below shows that, within our belt kernel (items (1)–(4) below), this axiom is in fact implied by those assumptions and hence is not logically independent.

*Theorem 5.14* (Kernel justification of the full operator equation of state). On  $D[R]$ , assume the belt setup above together with:

1. *First-law matching.* For every admissible one-parameter deformation  $\rho(\lambda)$  with  $\rho(0) = \sigma$  within the OS window,

$$\left. \frac{d}{d\lambda} \right|_{\lambda=0} \langle \widehat{K}_{\text{mod}}(R) \rangle_{\rho(\lambda)} = \left. \frac{d}{d\lambda} \right|_{\lambda=0} \left( \frac{1}{4G} \langle \widehat{A}(W) \rangle_{\rho(\lambda)} + \langle \widehat{H}_{\text{bulk}}(W) \rangle_{\rho(\lambda)} \right),$$

as in Corollary 5.4.

2. *Channel/JLMS compatibility on belts.* The belt-compatible isometry  $U_{R \rightarrow W}$  of Proposition 5.8 intertwines boundary-to-bulk modular dynamics on the analytic core.
3. *Normalization.* The Rindler witness and JKM calibration fix the overall scale and zero-point (Lemma 5.77).
4. *Quasi-local additivity.* Additivity under admissible factorization of belts/wedges (OS locality).

Then the closed quadratic forms associated with  $\widehat{K}_{\text{mod}}(R)$  and  $\widehat{S}_{\text{gen}}(W)$  coincide on  $D[R]$ . Equivalently, Axiom 5.13 follows from (1)–(4) whenever the belt kernel exists.

*Proof of Theorem 5.14.* We prove that the closed quadratic forms of  $\widehat{K}_{\text{mod}}(R)$  and  $\widehat{S}_{\text{gen}}(W)$  coincide on  $D[R] = \overline{\mathsf{D}_{\text{an}}}$  under assumptions (1)–(4).

*Step 0: Forms, domains, and a difference functional.* By Lemma 5.12 the generalized-entropy form

$$\mathfrak{s}_{\text{gen}}[\psi] = \left\langle \psi, \frac{\widehat{A}(W)}{4G} \psi \right\rangle + \left\langle \psi, \widehat{H}_{\text{bulk}}(W) \psi \right\rangle$$

is densely defined, closed, and lower-semibounded on  $D_{\text{an}}$ , hence extends uniquely to a closed form on  $D[R]$  with self-adjoint representative  $\widehat{S}_{\text{gen}}$  (KLMN + first representation theorem). The Tomita–Takesaki modular quadratic form

$$q_{\text{mod}}[\psi] = \langle \psi, \widehat{K}_{\text{mod}}(R)\psi \rangle$$

is likewise closed and lower-semibounded on  $D[R]$  by the OS/KMS core and essential self-adjointness on  $D_{\text{an}}$ . Define the (a priori densely defined) difference functional on  $D_{\text{an}}$

$$\mathfrak{d}[\psi] := q_{\text{mod}}[\psi] - \mathfrak{s}_{\text{gen}}[\psi].$$

*Step 1: Vanishing first variations  $\Rightarrow$  constancy of  $\mathfrak{d}$  on admissible cones.* Let  $\rho(\lambda)$  be any admissible one-parameter family with  $\rho(0) = \sigma$  in the OS belt window. Assumption (1) states the belt first-law matching

$$\left. \frac{d}{d\lambda} \right|_{\lambda=0} \left( \langle \widehat{K}_{\text{mod}}(R) \rangle_{\rho(\lambda)} - \frac{1}{4G} \langle \widehat{A}(W) \rangle_{\rho(\lambda)} - \langle \widehat{H}_{\text{bulk}}(W) \rangle_{\rho(\lambda)} \right) = 0.$$

Via the GNS map for  $\sigma$  this is the Fréchet derivative of  $\mathfrak{d}$  in all admissible directions at the ray of  $\Omega$ . Hence  $\mathfrak{d}$  is constant on each connected admissible cone of vector states generated from  $\Omega$  by admissible deformations. In particular, for any admissible unitary path  $U(\lambda)$  in the belt algebra,

$$\left. \frac{d}{d\lambda} \right|_{\lambda=0} \langle \Omega, U(\lambda)^* \widehat{D} U(\lambda) \Omega \rangle = 0, \quad \widehat{D} := \widehat{K}_{\text{mod}}(R) - \widehat{S}_{\text{gen}}(W),$$

so that

$$\langle \Omega, [\widehat{D}, X] \Omega \rangle = 0 \quad \text{for all self-adjoint } X \in \mathfrak{A}_R \text{ in the OS window.}$$

*Step 2: Centrality via channel/JLMS.* Assumption (2) provides an isometry  $U_{R \rightarrow W}$  that intertwines boundary and wedge modular data on  $D_{\text{an}}$ . Transport  $\widehat{D}$  to the wedge:

$$\widehat{D}_W := U_{R \rightarrow W} \widehat{D} U_{R \rightarrow W}^*.$$

For any  $Y \in \mathfrak{A}_W$ , choose  $X := U_{W \rightarrow R} Y U_{R \rightarrow W}$  (defined on the analytic core). Then, using the previous commutator expectation and functoriality of  $U_{R \rightarrow W}$ ,

$$\langle U_{R \rightarrow W} \Omega, [\widehat{D}_W, Y] U_{R \rightarrow W} \Omega \rangle = 0 \quad \text{for all } Y \in \mathfrak{A}_W.$$

Since  $U_{R \rightarrow W} \Omega$  is cyclic and separating for  $\mathfrak{A}_W$  on the belt code subspace (OS reflection positivity/KMS core), it follows that

$$[\widehat{D}_W, Y] = 0 \quad \text{for all } Y \in \mathfrak{A}_W,$$

i.e.  $\widehat{D}_W$  belongs to the wedge center  $\mathcal{Z}(\mathfrak{A}_W)$ . Equivalently,

$$\widehat{D} \in U_{W \rightarrow R} \mathcal{Z}(\mathfrak{A}_W) U_{R \rightarrow W}.$$

*Step 3: Quasi-local additivity  $\Rightarrow$  scalar multiple of the identity.* By disjoint-belt factorization (Lemma 5.109) and the Global GNS/modular compatibility (Corollary 2.7), the belt code subspace factorizes for disjoint belts  $B = \bigsqcup_{j=1}^n B_j$  and the wedge algebra acts as a tensor product on  $\mathcal{H}_W \cong \bigotimes_{j=1}^n \mathcal{H}_{W_j}$ . Since  $\widehat{D}_W \in \mathcal{Z}(\mathfrak{A}_W)$  (Step 2), it decomposes as

$$\widehat{D}_W = \sum_{j=1}^n \mathbf{1} \otimes \cdots \otimes Z_{W_j} \otimes \cdots \otimes \mathbf{1}, \quad Z_{W_j} \in \mathcal{Z}(\mathfrak{A}_{W_j}).$$

Quasi-local additivity of the first-law channel and of the modular generator across disjoint components, together with the near-product structure of Lemma 5.109, forces each  $Z_{W_j}$  to be a

scalar on  $\mathcal{H}_{W_j}$ . Uniformity per generator length then implies a common density on connected components, hence

$$\widehat{D}_W = c \mathbf{1}_{\mathcal{H}_W}.$$

*Step 4: Fixing the constant by normalization (Rindler/JKM).* Assumption (3) (Rindler witness + JKM calibration) fixes the absolute zero of both sides:

$$\langle \widehat{K}_{\text{mod}}(R) \rangle_\sigma = \frac{1}{4G} \langle \widehat{A}(W) \rangle_\sigma + \langle \widehat{H}_{\text{bulk}}(W) \rangle_\sigma.$$

Equivalently,  $\langle \widehat{D} \rangle_\sigma = 0$ , hence  $c = 0$ . Therefore  $\widehat{D} = 0$  on the belt code subspace.

*Step 5: Extension from the analytic core to  $D[R]$ .* We have shown  $\mathfrak{d}[\psi] = 0$  for all  $\psi$  in the dense core  $D_{\text{an}}$ . Since both  $q_{\text{mod}}$  and  $\mathfrak{s}_{\text{gen}}$  are closed on  $D[R]$ , the difference  $\mathfrak{d}$  is closable and its closure vanishes identically. Hence

$$q_{\text{mod}}[\psi] = \mathfrak{s}_{\text{gen}}[\psi] \quad \text{for all } \psi \in D[R],$$

i.e. the closed quadratic forms coincide on  $D[R]$ .

This proves that Axiom 5.13 (the full operator equation of state as a form identity) follows from assumptions (1)–(4) within the belt kernel.  $\square$

*Remark 5.15 (Logical packaging).* Items (1)–(4) above are precisely the belt first-law/JLMS inputs assembled in the kernel of Section 3. Theorem 5.14 therefore shows that, once the belt kernel is assumed, Axiom 5.13 does not introduce any additional hypothesis: it simply repackages the kernel into a single operator statement. All later uses of the full operator equation of state can thus be read as conditional on the same kernel.

*Theorem 5.16 (Generator theorem and equivalences).* Assume Axiom 5.13. Then, on  $D[R]$ :

1. *Operator equality.* By the representation theorem for closed, lower-semibounded forms, the self-adjoint generators coincide:

$$-\log \Delta_R = \widehat{S}_{\text{gen}}(W) \quad \text{and} \quad \widehat{K}_{\text{mod}}(R) = \widehat{S}_{\text{gen}}(W).$$

2. *Group form (Stone).* For all  $t \in \mathbb{R}$ ,

$$\Delta_R^{it} = e^{-it \widehat{S}_{\text{gen}}(W)}.$$

3. *Channel/JLMS 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  $(\rho, \sigma)$  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.

5. *Normalization and uniqueness.* The coefficient of the area term is  $1/4G$  (boost Ward/JKM). No additional state-independent constant or wedge-center operator can be added without violating: (i) the belt first law and Brown–York matching (Corollary 5.4); (ii) the Rindler witness normalization (Lemma 5.77); (iii) additivity under quasi-local factorization (Lemma 5.109).

*Proof of Theorem 5.16.* We work on the regulator-independent form domain  $D[R] = \overline{D_{\text{an}}}$  from Lemma 3.3 and Proposition 5.67. By Lemma 5.12, the generalized-entropy form

$$\mathfrak{s}_{\text{gen}} = \frac{1}{4G} \mathfrak{a} + \mathfrak{h}$$

is densely defined, closed, and lower semibounded on  $D_{\text{an}}$  and hence on  $D[R]$ , with unique self-adjoint representative  $\widehat{S}_{\text{gen}}$  (KLMN + first representation theorem). The Tomita-Takesaki modular form  $q_{\text{mod}}$  of  $-\log \Delta_R$  is likewise closed and lower semibounded on  $D[R]$  by the OS/KMS kernel and essential self-adjointness on  $D_{\text{an}}$  (cf. Lemma 3.1 and Proposition 5.67). The axiom Axiom 5.13 asserts the *form identity*

$$q_{\text{mod}} = \mathfrak{s}_{\text{gen}} \quad \text{on } D[R]. \quad (5.1)$$

(1) *Operator equality.* Since  $q_{\text{mod}}$  and  $\mathfrak{s}_{\text{gen}}$  are closed, lower-semibounded forms with the same domain and coincide on that domain by (5.1), the first representation theorem implies equality of the associated self-adjoint operators:

$$-\log \Delta_R = \widehat{S}_{\text{gen}}(W), \quad \widehat{K}_{\text{mod}}(R) = \widehat{S}_{\text{gen}}(W).$$

Equivalently, the quadratic-form identity of Axiom 5.13 upgrades to operator equality on  $D[R]$ .

(2) *Group form (Stone).* Let  $H := -\log \Delta_R = \widehat{S}_{\text{gen}}$ . By the spectral theorem,

$$\Delta_R = e^{-H}, \quad \Delta_R^{it} = e^{it \log \Delta_R} = e^{-itH} = e^{-it \widehat{S}_{\text{gen}}(W)},$$

for all  $t \in \mathbb{R}$ . Hence the modular group equals the  $S_{\text{gen}}$ -group.

(3) *Channel/JLMS form.* Let  $U_{R \rightarrow W}$  denote the belt-compatible isometry of Proposition 5.8. At finite regulators  $(r; u, s)$ , operator JLMS on belts gives the quadratic-form identity

$$\widehat{K}_{\text{mod}}(R) = U_{R \rightarrow W}^* \widehat{\mathcal{G}}_r(W) U_{R \rightarrow W} + \widehat{\mathcal{R}}_r, \quad \text{on } D_{\text{an}}.$$

with  $\|\widehat{\mathcal{R}}_r\|_{\text{form}} \leq C \mathcal{B}_{\text{belt}}$  and  $\widehat{\mathcal{R}}_r \rightarrow 0$  upon positive-flow removal; see Proposition 5.8 and Theorems 5.9 and 5.21. Passing to the regulator-independent limit and using Axiom 5.13 yields

$$\widehat{K}_{\text{mod}}(R) = U_{R \rightarrow W}^* \widehat{S}_{\text{gen}}(W) U_{R \rightarrow W}, \quad \text{on } D[R].$$

By functional calculus,

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

and likewise for the generators, proving the stated channel identities.

(4) *Relative form.* For any admissible pair  $(\rho, \sigma)$  on  $R$ , let  $\widehat{K}_{\text{mod}}^R[\rho]$  and  $\widehat{S}_{\text{gen}}^W[\rho]$  denote the (closable) relative generators defined on the common form core by Araki's relative modular theory; likewise for  $\sigma$ . The absolute operator identity  $\widehat{K}_{\text{mod}}(R) = \widehat{S}_{\text{gen}}(W)$  implies, on the common form domain and after transporting by  $U_{R \rightarrow W}$ , the equality of differences

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

This uses functoriality of the belt channel and the global GNS/modular compatibility under restrictions to subalgebras, so that the relative constructions commute with  $U_{R \rightarrow W}$  and with the removal limit. The identity holds as a quadratic-form equality on  $D[R]$ .

(5) *Normalization and uniqueness.* The coefficient of the area term is fixed to  $1/4G$  by the boost Ward identity on the belt together with the JKM corner calibration (Rindler witness). Any shift

$$\widehat{S}_{\text{gen}} \mapsto \widehat{S}_{\text{gen}} + c \mathbf{1} \quad \text{or} \quad \widehat{S}_{\text{gen}} \mapsto \widehat{S}_{\text{gen}} + Z_W, \quad Z_W \in \mathcal{Z}(\mathfrak{A}_W)$$

is incompatible with the calibrated belt first law and Brown–York matching (it would spoil the absolute normalization fixed by the witness), and violates quasi–local additivity under disjoint belts (the extra center would not decouple and would contradict strong additivity at large separation). Therefore neither a central constant nor a wedge–center term can be added. Uniqueness within the scheme class of Axiom 5.13 follows.  $\square$

*Remark 5.17* (Reference state and state dependence of  $W$ ). The GNS data  $(\mathfrak{A}_R, \Omega)$  and the wedge  $W = \text{EW}(R)$  are defined relative to a cyclic separating reference state  $\sigma$  within the OS belt window. All identities above hold on the common domain  $D[R] = \overline{D_{\text{an}}}$  determined by  $\sigma$ . The relative formulations ensure compatibility under  $\sigma \mapsto \sigma'$ , with  $U_{R \rightarrow W}$  implementing the corresponding channel equivalence on belts (Proposition 5.8).

*Proposition 5.18* (Equivalent formulations (condensed)). Items (1)–(4) of Theorem 5.16 are equivalent on  $D[R]$ .

*Remark 5.19* (Immediate consequences and checks). (i) Taking expectations and first/second variations reproduces the linear and second–order modular equations of state and hence the semiclassical Einstein equations in expectation (Corollary 5.11, Theorem 5.46, Theorem 5.37). (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 5.20** (Belt-regulated operators and analytic core). Fix a belt-anchored region  $R$  of width  $r > 0$  and auxiliary positive flows  $(u, s) > 0$ . Let  $D_{\text{an}}$  be the common analytic core furnished by the OS–KMS belt kernel. Define, on  $D_{\text{an}}$ , the closable unbounded operators

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

where  $\widehat{K}_{\text{mod}}$  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 (dimensionless) bulk modular Hamiltonian on the wedge. For matter fields,

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

so the belt operator equation of state

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

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

*Theorem 5.21* (Belt operator equation; well-definition and consistency). On  $D_{\text{an}}$  and at finite regulators  $(r; u, s)$  the operator equation

$$\boxed{\widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}_r(W)}{4G} + \widehat{H}_{\text{bulk},r}(W) + \widehat{R}_r} \quad (5.2)$$

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

(i) **Consistency with the kernel.** The identity (5.2) 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  $\mathcal{D}_{\text{an}}$ .

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

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

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

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

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

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

$$\lim_{(u,s) \downarrow 0} \widehat{R}_r = 0, \quad \lim_{(u,s) \downarrow 0} (5.2) : \widehat{K}_{\text{mod}}(R) = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W) \quad \text{on } \mathcal{D}_{\text{an}},$$

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

*Proof. Setup and domains.* By construction of the OS–KMS belt kernel,  $\mathcal{D}_{\text{an}}$  is invariant under the regulated modular flows on  $R$  and  $W$  and is a set of Nelson analytic vectors for the generators. Hence the restrictions  $\widehat{K}_{\text{mod}} \upharpoonright_{\mathcal{D}_{\text{an}}}$ ,  $\widehat{A}_r \upharpoonright_{\mathcal{D}_{\text{an}}}$  and  $\widehat{H}_{\text{bulk},r} \upharpoonright_{\mathcal{D}_{\text{an}}}$  are essentially self-adjoint and mutually strongly commuting as quadratic forms on  $\mathcal{D}_{\text{an}}$ . We work throughout with quadratic forms on this common core and identify operators with their closures.

*Step 1 (JLMS isometry on belts).* Let  $U_{R \rightarrow W}$  be the rotated Petz/JLMS isometry implementing the belt code-subspace equivalence. By OS–KMS analyticity,  $U_{R \rightarrow W} \mathcal{D}_{\text{an}} \subset \mathcal{D}_{\text{an}}$  and

$$U_{R \rightarrow W} \Delta_R^{it} U_{R \rightarrow W}^* = \exp\left(it \left[ \frac{\widehat{A}_r}{4G} + \widehat{H}_{\text{bulk},r} + \widehat{R}_r \right]\right) \quad \text{on } \mathcal{D}_{\text{an}},$$

with a correction  $\widehat{R}_r$  coming from (i) finite-flow deformation  $(u, s)$ , (ii) belt truncation at width  $r$ , and (iii) corner calibration (JKM) at the regulated QES. Differentiating at  $t = 0$  in the sense of quadratic forms gives

$$U_{R \rightarrow W} \widehat{K}_{\text{mod}} U_{R \rightarrow W}^* = \frac{\widehat{A}_r}{4G} + \widehat{H}_{\text{bulk},r} + \widehat{R}_r \quad \text{on } \mathcal{D}_{\text{an}}. \quad (5.3)$$

By definition of  $U_{R \rightarrow W}$  as an isometry on the code subspace, this implies the quadratic-form identity (5.2) on  $\mathcal{D}_{\text{an}}$ .

*Step 2 (Identification of  $\widehat{H}_{\text{bulk},r}$ ).* Choose a Cauchy slice  $\Sigma \subset W$  adapted to the boost Killing field  $\xi$ . The Iyer–Wald identity on  $\Sigma$  with JKM corner calibration yields

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

and packages all gravitational boundary contributions into  $\widehat{A}_r / (4G)$ . Thus no additional  $2\pi$  multiplies the operator equation; the only explicit  $2\pi$  is the boost coefficient in  $\widehat{H}_{\text{bulk},r}$ . The belt regulator  $r$  restricts the support of corner terms to the OS–KMS belt and makes the flux integral well-defined on  $\mathcal{D}_{\text{an}}$ .

*Step 3 (Remainder control).* Let  $\mathcal{B}_{\text{belt}}$  be the belt budget controlling: (a) exponential tails from the finite belt width  $r$ , (b) the positive flows  $(u, s)$ , and (c) residual corner mismatches after JKM calibration. Each source contributes a quadratic-form error bounded by  $C_i \mathcal{B}_{\text{belt}}$  uniformly on  $\mathcal{D}_{\text{an}}$ ; summing gives  $\|\widehat{R}_r\|_{\text{form}} \leq C \mathcal{B}_{\text{belt}}$ . This proves (5.2) as a quadratic-form identity with the stated bound.

*Step 4 (First variation).* Take any admissible one-parameter family of belt-analytic states and backgrounds. Varying (5.2) and taking expectation values gives

$$\delta\langle\widehat{K}_{\text{mod}}\rangle = \delta\left\langle\frac{\widehat{A}_r}{4G}\right\rangle + \delta\langle\widehat{H}_{\text{bulk},r}\rangle + \delta\langle\widehat{R}_r\rangle.$$

By Step 2,  $\delta\langle\widehat{H}_{\text{bulk},r}\rangle = 2\pi \int_R d\Sigma_\mu \xi_\nu \delta\langle\widehat{T}^{\mu\nu}\rangle$ . The belt remainder varies at most by  $O(\mathcal{B}_{\text{belt}})$  because the sources in Step 3 remain controlled under admissible variations. Hence

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

which is the linear modular equation of state on the belt and implies the linearized SEE in expectation.

*Step 5 (Second variation).* Apply a second variation to (5.2) and use the relative-entropy identity  $\delta^2 S = \delta^2\langle\widehat{K}_{\text{mod}}\rangle - \delta^2\langle\widehat{K}_{\text{mod}}\rangle_{\text{ref}}$ , together with the Iyer–Wald symplectic current representation of  $\delta^2(\text{Area}/4G) + \delta^2\langle\widehat{H}_{\text{bulk},r}\rangle$  on the belt. Collecting terms yields

$$\delta^2\left[S - \frac{\text{Area}}{4G}\right] = 2\pi E_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - \delta^2\langle\widehat{R}_r\rangle,$$

where  $E_{\text{can}}$  is the canonical energy quadratic form defined by the bulk symplectic current on  $\Sigma$  and  $Q_{\text{shear}}$  is the positive-definite contribution from shear/expansion on the belt boundaries (the JKM-calibrated gravitational flux). By Step 3,  $\delta^2\langle\widehat{R}_r\rangle \leq C_2 \mathcal{B}_{\text{belt}}$ . This gives the stated second-order inequality.

*Step 6 (Continuum lift:  $(u, s) \downarrow 0$ ).* The positive-flow removal lemma implies that each source in  $\mathcal{B}_{\text{belt}}$  vanishes as  $(u, s) \downarrow 0$  at fixed  $r$ , hence  $\widehat{R}_r \rightarrow 0$  in the quadratic-form sense on  $\mathcal{D}_{\text{an}}$ . Taking the limit in (5.2) gives

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

*Step 7 (Belt-width stability and  $r \rightarrow r'$ ).* Changing the belt width  $r \mapsto r'$  modifies the corner calibration only within the exponentially decaying tails of the OS–KMS kernel. The difference is absorbed into a remainder bounded by  $e^{-\kappa \min\{r, r'\}/\ell_{\text{gap}}}$  for some gap scale  $\ell_{\text{gap}}$  and constant  $\kappa > 0$ . Thus the identity is stable under  $r \mapsto r'$  up to an exponentially small quadratic-form correction.

*Step 8 (Removing the belt regulator).* Combining Steps 6 and 7 and using the existence of the QES area operator  $\widehat{A}(W)$  and the unregulated bulk generator  $\widehat{H}_{\text{bulk}}(W)$  on  $\mathcal{D}_{\text{an}}$ , we obtain the regulator-independent continuum identity

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

This completes the proof. □

## 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  $(\mathcal{H}_R, \Omega)$  of a

cyclic separating reference state  $\sigma$ . Let  $S_R$  be the Tomita operator,  $\Delta_R := S_R^\dagger S_R$  the modular operator, and

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

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  $\mathsf{D}_{\text{an}}$  furnished by the OS/KMS kernel, define the belt generalized–entropy operator

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

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 5.22* (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  $\mathsf{D}_{\text{an}}$  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  $\mathsf{D}_{\text{an}}$  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  $\mathsf{D}_{\text{an}}$  in the form norm  $\|\psi\|_{\mathfrak{a}}^2 := \|\psi\|^2 + \mathfrak{a}(\psi, \psi)$ .
2.  $\mathsf{D}_{\text{an}}$  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  $\geq 0$ . 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  $\mathsf{D}_{\text{an}}$ . 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  $\mathsf{D}_{\text{an}}$  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 5.23* (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  $\mathsf{D}_{\text{an}}$  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  $\mathsf{D}_{\text{an}}$ , 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 5.24* (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  $D_0 \subset D$  be a core for both. If  $\mathfrak{q}_1|_{D_0} = \mathfrak{q}_2|_{D_0}$ , 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 5.25* (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  $D_{\text{an}}$ ,

$$\widehat{K}_{\text{mod}}(R) = U_{R \rightarrow W}^* \widehat{G}_r(W) U_{R \rightarrow W} + \widehat{R}_r, \quad \|\widehat{R}_r\|_{\text{form}} \leq C \mathbf{B}_{\text{belt}}, \quad (5.4)$$

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:

$$\star \quad \mathfrak{q}_{\text{mod}} = \mathfrak{g} \quad \text{as closed quadratic forms with form domain } \overline{D_{\text{an}}}^{\|\cdot\|_{\text{form}}}, \quad (5.5)$$

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 5.22. Consequently,

$$\star \quad -\log \Delta_R = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W) \quad \text{as self–adjoint operators (in particular on } D_{\text{an}}). \quad (5.6)$$

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  $(\mathcal{A}_R, \Omega)$ ; likewise for the bulk wedge flow on  $\mathcal{H}_W$ . On  $D_{\text{an}}$  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  $D_{\text{an}}$  is a common invariant core for  $\widehat{K}_{\text{mod}}, \widehat{H}_{\text{bulk},r}, \widehat{A}_r$ ,

their polynomials, and the regulated sum  $\widehat{G}_r := \widehat{A}_r/(4G) + \widehat{H}_{\text{bulk},r}$ . By Nelson's analytic vector theorem,  $\widehat{K}_{\text{mod}}$  and  $\widehat{G}_r$  are essentially self-adjoint on  $\mathcal{D}_{\text{an}}$ .

**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  $\sigma$ , i.e. normal, unital, completely positive,  $\sigma$ -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  $\mathcal{D}_{\text{an}}$  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  $\mathcal{D}_{\text{an}}$  (boundary) and on  $U_{R \rightarrow W}\mathcal{D}_{\text{an}}$  (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  $\psi \in \mathcal{D}_{\text{an}}$  the diagonal form relation

$$\mathfrak{q}_R^{(s)}(\psi, \psi) = \mathfrak{q}_{W,r}^{(s)}(U_{R \rightarrow W}\psi, U_{R \rightarrow W}\psi) + \mathfrak{r}_r^{(s)}(\psi), \quad |\mathfrak{r}_r^{(s)}(\psi)| \leq C \mathcal{B}_{\text{belt}} \|\psi\|^2, \quad (5.7)$$

with a constant  $C$  independent of  $|R|$ . Polarization upgrades (5.7) to a sesquilinear identity: for all  $\psi, \phi \in \mathcal{D}_{\text{an}}$ ,

$$\mathfrak{q}_R^{(s)}(\psi, \phi) = \mathfrak{q}_{W,r}^{(s)}(U_{R \rightarrow W}\psi, U_{R \rightarrow W}\phi) + \mathfrak{r}_r^{(s)}(\psi, \phi), \quad (5.8)$$

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

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

Consequently, as a quadratic-form identity on  $\mathcal{D}_{\text{an}}$ ,

$$\widehat{K}_{\text{mod}}^{(s)} = U_{R \rightarrow W}^* \widehat{G}_r^{(s)} U_{R \rightarrow W} + \widehat{R}_r^{(s)}. \quad (5.9)$$

**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  $D_{\text{an}}$ . Passing to the limit  $(u, s) \downarrow (0, 0)$  in (5.9) yields the core-level form identity

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

By Lemma 5.22 the renormalized area form  $\mathfrak{a}$  is closed and lower semibounded with form core  $D_{\text{an}}$ , 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

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

Applying Lemma 5.24 upgrades the core-level identity to equality of closed forms on  $D$ , which is exactly (5.5). By the representation theorem, the associated self-adjoint operators coincide, giving (5.6). Abbreviating the isometric identification via  $U_{R \rightarrow W}$ , this reads

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

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

*Corollary 5.26* (Entanglement equilibrium from the dictionary). For any admissible one-parameter belt deformation,

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

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

*Remark 5.27* (Algebraic content and stability).

1. *Algebraic meaning.* Theorem 5.25 identifies the Tomita generator of  $(\mathcal{A}_R, \Omega)$  with the bulk generalized-entropy operator on  $W$ . Thus the boundary von Neumann algebra encodes the full (quantum-gravitational) operator  $S_{\text{gen}}$ .
2. *Domains/spectrum.*  $D_{\text{an}}$  is a common core for  $\widehat{K}_{\text{mod}}, \widehat{A}, \widehat{H}_{\text{bulk}}$  and their polynomials;  $\widehat{A} \geq 0$ .
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 5.28* (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 5.14 [12, 13, 20–23].

*Proof.* Fix a belt-anchored region  $R$  and consider the generalized-entropy functional on the admissible class of codimension-two surfaces  $\Sigma$  homologous to  $R$ ,

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

*Existence and extremality.* Under the framework Section 2, belt microcausality/timeslice control, belt–level nesting/recovery and continuity, and the OS kernel Lemmas 3.1 to 3.3, the direct–method existence theorem for quantum extremal surfaces Theorem 5.50 ensures that  $G_R$  admits a minimizer  $\Sigma_*$  in the admissible class, with minimizing sequences confined to a compact geometric set (uniform area and curvature bounds at fixed belt width). We define

$$\text{QES}(R) := \Sigma_*.$$

The generalized entropy  $G_R$  is  $C^1$  along admissible variations and has strictly positive second variation in the OS window away from boost–Killing directions, by the belt second–order modular equation of state and canonical–energy positivity Theorem 5.46, together with QES stability and selection Theorem 5.38. Hence any interior minimizer is stationary. Concretely, for every admissible shape/state variation  $\delta$  supported on the belt,

$$\delta G_R[\Sigma_*] = \delta \left[ \frac{\text{Area}(\Sigma_*)}{4G} - S(\rho_R; \Sigma_*) \right] = 0.$$

Equivalently,

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

along admissible variations, which is the quantum extremality condition in the statement.

This stationarity can also be expressed using the belt first–law channel and the operator equation of state. Kernel 0 and the belt JLMS channel Propositions 5.3 and 5.8, together with the belt operator equation of state Theorems 5.14 and 5.21, identify the boundary modular generator with bulk generalized entropy (area plus bulk modular Hamiltonian) on the wedge. Combining the entanglement first law Corollary 5.4 with the modular equation of state Theorems 5.37 and 5.42, one obtains for any admissible deformation

$$\delta \left[ S(\rho_R) - \frac{\text{Area}(\Sigma)}{4G} \right] = 2\pi \int_{\Sigma} \xi^\mu \delta \langle T_{\mu\nu} \rangle d\Sigma^\nu + O(\mathcal{B}_{\text{belt}}),$$

where the right–hand side gives the canonical–energy flux through the wedge associated with  $\Sigma$ . At a minimizer  $\Sigma_*$ , the left–hand side vanishes, so the canonical–energy flux must vanish for all admissible variations; conversely, canonical–energy positivity for non–Killing perturbations Theorem 5.46, Lemma 5.112, and Proposition 5.102 implies that such stationary configurations are local minima of  $G_R$ .

*Stability and continuity.* The strict convexity/concavity of the generalized–entropy functional and the belt–level continuity/recovery estimates imply stability of the extremal surface. In particular, the QES stability theorem Theorem 5.38 and Section 5.14 gives uniqueness of  $\text{QES}(R)$  in a neighborhood of any stationary configuration and Lipschitz continuity of its location under  $C^2$ –small variations of the belt geometry and admissible state data, with all moduli controlled by the canonical–energy/shear coercivity and the belt continuity constants (from Proposition 5.86, Lemma 5.109, Theorem 5.104, and Sections 5.12 and 5.27). This yields the stability and continuity properties quoted from Section 5.14.

*Page behavior.* For the Page transition, we use the AGSP/seed converter together with the belt Markov–gap estimate and the global ledger. The AGSP/seed pipeline and the belt–level recoverability/continuity results Sections Appendix B.2, 5.12 and 5.27, Theorem 5.104, and Proposition 5.86 imply an extensive bound on the entropy per generator length: the belt Markov gap Theorem 5.104, combined with the deterministic step budget and global  $\varepsilon$ –split (Sections Appendix B.2 and 5.27), yields a uniform constant  $c_{\text{rad}}$  depending only on the belt base factor and the AGSP/seed parameters such that

$$\frac{1}{\text{length}(\partial R)} S(\rho_R) \leq c_{\text{rad}} + O(\mathcal{B}_{\text{belt}}),$$

with an explicit expression for  $c_{\text{rad}}$  recorded in the Page–threshold estimate of Section 5.21.

On the other hand, for the extremal surface  $\text{QES}(R)$  selected above, the area term in  $G_R$  scales linearly with the belt length, with slope  $a_{\text{QES}}/(4G)$  given by the area line density of the QES (fixed by Section 5.29 and Theorem 5.50). For belt–anchored regions whose length is such that  $a_{\text{QES}}/(4G) < c_{\text{rad}}$ , the entropic contribution dominates the generalized entropy and the minimizer coincides with the “trivial” surface near the boundary; the entropy grows with the AGSP–controlled slope. As the region is enlarged along the belt, the area density eventually satisfies  $a_{\text{QES}}/(4G) \gtrsim c_{\text{rad}}$ , at which point an interior extremal surface minimizes  $G_R$  and the entropy is set by the area term. The crossover scale where

$$\frac{a_{\text{QES}}}{4G} \approx c_{\text{rad}} + O(\mathcal{B}_{\text{belt}})$$

is the Page transition point. The dependence of  $c_{\text{rad}}$  on the belt base factor and the AGSP/seed parameters shows that the transition is controlled precisely by these data, as claimed.

All identities above hold up to  $O(\mathcal{B}_{\text{belt}})$ , with  $\mathcal{B}_{\text{belt}}$  the belt budget from the positive–flow regularization Lemma 3.3 and Section 5.13. By the removal–of–flows lemma Lemma 3.3,  $\mathcal{B}_{\text{belt}} \rightarrow 0$  along admissible regulator limits, so both the quantum extremality condition and the Page transition become regulator–independent in this limit.

This completes the proof.  $\square$

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

*Theorem 5.29* (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  $\Gamma_{\text{belt}}$ .

*Remark 5.30.* Belt-local RP/KMS positivity yields light-ray inequalities [3, 4]. Canonical energy and QSEI bound the stress tensor by modular second variations [18]. The first-law channel (and the operator identity) identifies entropy curvature with modular energy, producing QNEC with the  $2\pi$  normalization fixed by the Rindler witness [4, 24, 25]. Integrating QNEC using decay/cluster and null timeslice propagation gives ANEC; stability under flow removal follows from the belt remainder budget [15].

*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  $k^\mu$ . Let  $\sigma$  be the boost–KMS reference state on the belt and let  $\omega$  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  $\varepsilon f$ :

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

Write  $S(\varepsilon)$  for the von Neumann entropy of the “outside” with respect to  $R(\varepsilon)$  in the state  $\omega$ , and  $K_\sigma(\varepsilon)$  for the modular Hamiltonian of  $\sigma$  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  $\sigma$  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  $\omega$  and  $\sigma$ , which cancel by boost-KMS symmetry of  $\sigma$  and by the belt calibration of the modular normalization. The second variation takes the universal form

$$F''(0) = \Delta\langle K_\sigma''(0) \rangle - \Delta S''(0) + \mathcal{W}_\sigma[f; \omega] \geq 0, \quad (5.10)$$

where  $K_\sigma''(0)$  is the second shape derivative of the modular Hamiltonian of  $\sigma$  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 (5.10) 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  $\sigma$  is

$$\left. \frac{d}{d\varepsilon} \langle K_\sigma(\varepsilon) \rangle \right|_{\varepsilon=0, \text{ state } \chi} = 2\pi \int du f(u) \langle T_{kk}(u) \rangle_\chi + O(\mathcal{B}_{\text{belt}}), \quad (5.11)$$

for any state  $\chi$  in the admissible class, with the  $2\pi$  fixed by the Rindler/boost normalization. Differentiating (5.11) once more and symmetrizing produces the second-shape-derivative kernel; light-ray regularity (below) ensures its diagonal part is

$$\langle K_\sigma''(0) \rangle_\chi = 2\pi \int du f(u)^2 \langle T_{kk}(u) \rangle_\chi + \mathcal{W}_\sigma[f; \chi] + O(\mathcal{B}_{\text{belt}}), \quad (5.12)$$

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

$$\Delta\langle K_\sigma''(0) \rangle = 2\pi \int du f(u)^2 \Delta\langle T_{kk}(u) \rangle + (\mathcal{W}_\sigma[f; \omega] - \mathcal{W}_\sigma[f; \sigma]) + O(\mathcal{B}_{\text{belt}}). \quad (5.13)$$

Since  $\sigma$  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

$$S''(0) = \int du f(u)^2 \partial_u^2 S(u) + O(\mathcal{B}_{\text{belt}}) \|f\|_2^2. \quad (5.14)$$

Insert (5.13) and (5.14) into (5.10), use  $\Delta X = \langle X \rangle_\omega - \langle X \rangle_\sigma$  and the remarks after (5.13), and drop the manifestly nonnegative  $\mathcal{W}_\sigma[f; \omega]$ :

$$\int du f(u)^2 (2\pi \langle T_{kk}(u) \rangle_\omega - \partial_u^2 S(u)) \geq -O(\mathcal{B}_{\text{belt}}) \|f\|_2^2. \quad (5.15)$$

Choose a standard  $\delta$ -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 (5.15) and then removing positive flows (which set  $O(\mathcal{B}_{\text{belt}}) \rightarrow 0$ ) 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  $\Gamma_{\text{belt}}$  and the admissible state class; the  $O(\mathcal{B}_{\text{belt}})$  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 5.31* (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 5.22, Section 5.46, and Section 5.66. The only quantitative change is the window factor

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

a multiplicative cost  $R_{\text{max}}(-0.30 s_0)/R_{\text{max}}(-0.20 s_0) = 1.040833\dots$  in the tail envelopes that couple linearly to  $R_{\text{max}}$ . All dispersion and quadrature freezes in Section 5.89 retain strict slack under this 4.08% increase, so the global  $\varepsilon$ -budget and the certified Regge-slope headroom are unchanged.

*Proposition 5.32* (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  $\alpha_R = 2 + \delta$  with  $\delta \in [0, 1)$ , the high- $s'$  tail beyond  $S_{\text{cut}}$  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  $\rho$  denote the allocated headroom for the tail on a given audit line. Then the tester cone intrinsically certifies

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

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

*Proof.* For the  $k = 2$  forward derivative at subtraction order  $N = 3$ , the crossing-symmetric dispersion relation on the cone expresses the high- $s'$  tail as an integral over the absorptive part with kernel that, for  $s' \geq S_{\text{cut}}$ , has the asymptotic behaviour recorded in Section 5.54,

$$K_{2,3}(s'; s_0) \sim \frac{\text{const} \cdot s_0^3}{s'^4} \quad (s' \rightarrow \infty).$$

Thus the reference tail bound at Regge slope  $\alpha_{\text{R}}(t) = 2$  is proportional to

$$I_0(M) := \int_{S_{\text{cut}}}^{\infty} ds' \frac{1}{s'^4} = s_0^{-3} \int_M^{\infty} dx x^{-4} = \frac{1}{3} s_0^{-3} M^{-3},$$

after the rescaling  $x := s'/s_0$ .

If the Regge slope is increased to  $\alpha_{\text{R}}(t) = 2 + \delta$  with  $\delta \in [0, 1)$ , the worst-case absorptive envelope on the cone is multiplied by at most a factor  $(s'/s_0)^\delta$ , so the corresponding tail bound is

$$I_\delta(M) := \int_{S_{\text{cut}}}^{\infty} ds' \frac{(s'/s_0)^\delta}{s'^4} = s_0^{-3} \int_M^{\infty} dx x^{-4+\delta} = \frac{1}{3-\delta} s_0^{-3} M^{-3+\delta},$$

which is finite for  $\delta < 3$  and in particular on the interval  $\delta \in [0, 1)$  considered here. The ratio of the two tail bounds is therefore

$$\frac{I_\delta(M)}{I_0(M)} = \frac{3}{3-\delta} M^\delta = F_{\text{tail}}(\delta; M).$$

Since the high- $s'$  tail enters linearly in each tester, the same multiplicative factor rescales the tail contribution of any fixed dual and composite Gauss-Radau schedule, so the allocated tail headroom  $\rho$  on that audit line is respected provided  $F_{\text{tail}}(\delta; M) \leq \rho$ .

The function  $F_{\text{tail}}(\delta; M)$  is strictly increasing in  $\delta$  for  $M > 1$ , because

$$\partial_\delta \log F_{\text{tail}}(\delta; M) = \log M + \frac{1}{3-\delta} > 0.$$

Hence there is a unique largest  $\delta_\star(M, \rho) \in [0, 1)$  such that  $F_{\text{tail}}(\delta; M) \leq \rho$  if and only if  $\delta \leq \delta_\star(M, \rho)$ . By definition of the Regge slope envelope on the cone, this gives

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

as claimed. The quoted numerical bounds on  $\alpha_{\text{R}}(t)$  along the  $10^{-6}$  and  $10^{-8}$  audit lines are obtained by evaluating  $\delta_\star(M, \rho)$  with the frozen policies and recorded values of  $M$  and  $\rho$ .  $\square$

*Theorem 5.33* (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  $\delta_\star$  is as in Proposition 5.32.

*Proof.* Work on the dispersive cone and at  $N = 3$  with the analytic forward projector of Section 5.22 and the pivot normalization of Section 5.46.

For  $t = 0$ , the subtracted crossing-symmetric dispersion relation expresses each even forward derivative of the gravity-subtracted amplitude as an integral of the nonnegative absorptive profile against a nonnegative kernel, up to the controlled high- $s'$  tail and quadrature remainders

frozen in Sections 5.54 and 5.89. Optical positivity for the helicity-averaged, anchor-preserving dressing (Section 5.80) implies that these integrals are nonnegative, so all forward even-parity derivatives are nonnegative on the cone. The identification of these derivatives with the Wilson coefficients  $c_{2m,0}$  in the forward expansion is purely algebraic and is stable under the gravity IR subtraction scheme by Lemma 5.39, hence  $c_{2m,0} \geq 0$ .

For nonzero but near-forward  $t$ , the forward fixed- $t$  testers, Hankel/impact testers, and celestial Gram testers constructed in Sections 5.22, 5.33 and 5.42 define nonnegative linear functionals on the dressed absorptive profile along the principal-series strip. Each such tester acts on the exact dispersive representation with the same frozen tail and quadrature budgets as on the forward line, and therefore continues to furnish inequalities on the band-limited Wilson combinations obtained from the near-forward expansion.

The effective Regge slope on the cone is controlled by the high- $s'$  tail budget. For the  $k = 2$  derivative at  $N = 3$ , Proposition 5.32 shows that, for a given audit line with tail headroom  $\rho$ , any increase of the Regge slope to  $\alpha_R(t) = 2 + \delta$  multiplies the tail by  $F_{\text{tail}}(\delta; M)$ . If  $\delta > \delta_*(M, \rho)$ , the resulting tail contribution would exceed the allocated headroom and drive at least one tester negative, contradicting the nonnegativity of the tester cone. Hence all amplitudes that pass the tester cone satisfy

$$\alpha_R(t) \leq 2 + \delta_*(M, \rho)$$

for each audit line, and in particular on the stated window  $t \in [-0.25 s_0, 0]$ . This yields the asserted uniform bound with  $\delta_*$  as in Proposition 5.32.  $\square$

*Proposition 5.34* (Pointwise acceptance on the widened near-forward window). On the working cone of Remark 5.31 with subtraction order  $N = 3$  and the analytic forward projector of Section 5.22 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 5.48. 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 5.54, Section 5.89). Consequently, all certified slacks degrade by at most  $1/\alpha_{\text{win}}$  and remain strictly positive on the widened window.

*Proof.* Fix a Chebyshev node  $t_q \in [-0.30 s_0, 0]$  in the widened near-forward window. By construction of the analytic forward projector and the optical-positivity input (Sections 5.22 and 5.80), the gravity-subtracted helicity-averaged absorptive profile on the cone is nonnegative in each physical channel. The forward fixed- $t$  testers  $T_{q,1}^{\text{forw}}$ , the Hankel/impact testers  $T_p^H$ , and the celestial Gram testers  $T_j^{\text{cel}}$  are defined in Sections 5.22, 5.33 and 5.42 as nonnegative linear functionals of this absorptive profile, supplemented by a high- $s'$  tail bound and a composite Gauss-Radau quadrature for the finite- $s'$  portion of the dispersion integral.

Enlarging the near-forward window from  $t \in [-0.20 s_0, 0]$  to  $t \in [-0.30 s_0, 0]$  leaves the tester kernels themselves unchanged and affects only the worst-case absorptive envelope used to control the high- $s'$  tail. As recorded in Remark 5.31 and Section 5.54, this modification enters through the window factor  $R_{\text{max}}(t)$  and rescales the tail envelope by a multiplicative factor  $\alpha_{\text{win}} \geq 1$ , independent of the particular tester. Thus, for each tester and each node  $t_q$  in the enlarged window, the high- $s'$  tail contribution is at most  $\alpha_{\text{win}}$  times the original bound on the narrower window.

The composite Gauss-Radau schedule  $(M, J)$  for the finite- $s'$  integration is frozen by Section 5.89. Its remainder estimate depends only on  $(M, J)$  and the regularity of the integrand in  $s'$ , not on the precise choice of  $t$  within the working cone; in particular it is unchanged when the  $t$ -window is enlarged as above. Hence the quadrature errors retain the same absolute bounds.

Since each tester started with strictly positive slack on the original window and only the tail line is rescaled by the common factor  $\alpha_{\text{win}}$ , all certified slacks on the enlarged window decrease by at most a factor  $1/\alpha_{\text{win}}$  and remain strictly positive. The compact 18-support dual of Section 5.48 is a nonnegative linear combination of these testers, so its value at each node  $t_q$  is nonnegative as well, and its feasibility on the widened near-forward window is preserved.  $\square$

*Proposition 5.35* (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 5.31). For each tester in the compact 18-support dual (Section 5.48—forward fixed- $t$  evaluators, Hankel/impact band kernels, celestial Gram testers), define the channel-symmetrized combination

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

Then for the helicity-averaged, anchor-preserving dressed absorptive profile  $\text{Abs}_{\text{avg}}$ ,

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

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 5.22 and Lemmas 5.72 and 5.95.

*Proof.* By Proposition 5.100, the tester families are stable under CPT-crossing: for each  $t$  in the working cone, the map  $s \mapsto -s - t$  takes an  $s$ -channel tester  $T_s[t]$  into a corresponding  $u$ -channel tester  $T_u[t]$  with the same sign properties on the dressed absorptive profile. In particular, for the helicity-averaged, anchor-preserving dressing, Section 5.80 implies

$$T_s[t][\text{Abs}_{\text{avg}}] \geq 0, \quad T_u[t][\text{Abs}_{\text{avg}}] \geq 0$$

for all  $t$  in the near-forward window. Their arithmetic mean therefore satisfies

$$T_{\text{sym}}[t][\text{Abs}_{\text{avg}}] = \frac{1}{2} (T_s[t][\text{Abs}_{\text{avg}}] + T_u[t][\text{Abs}_{\text{avg}}]) \geq 0$$

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

Feasibility of the compact dual on the enlarged window is ensured by Proposition 5.34. Each of the 18 testers entering  $\mathcal{D}_{18}^{(s)}$  is nonnegative on  $\text{Abs}_{\text{avg}}$  at the Chebyshev nodes in the widened window, and the same holds for the crossed  $u$ -channel testers. The symmetrized dual  $\mathcal{D}_{18}^{(\text{sym})}$  is obtained by replacing each  $T_s[t]$  in  $\mathcal{D}_{18}^{(s)}$  by  $T_{\text{sym}}[t]$  while keeping the nonnegative coefficients fixed. Since  $T_{\text{sym}}[t]$  is nonnegative on  $\text{Abs}_{\text{avg}}$  for all  $t$  in the window,  $\mathcal{D}_{18}^{(\text{sym})}$  is a nonnegative linear combination of nonnegative functionals and hence enforces nonforward/near-forward positivity uniformly.

The IR-scheme, pivot, and scale invariances are inherited from the underlying tester construction. In the  $s$ -channel, these invariances are encoded in Section 5.22 and Lemmas 5.72 and 5.95 and in the IR-subtraction lemma Lemma 5.39. The  $u$ -channel testers enjoy the same invariances by the CPT-crossing stability of Proposition 5.100, and any convex combination such as  $T_{\text{sym}}[t]$  or  $\mathcal{D}_{18}^{(\text{sym})}$  preserves them.

Finally, the extension to the off-principal strip proceeds exactly as in the  $s$ -channel analysis: analyticity on the strip Section 5.22 together with celestial IR control (Corollary 5.120) and the mainline strip reduction (Corollary 5.125) shows that the same channel-symmetrized testers remain nonnegative there. Thus the claimed nonforward/near-forward positivity and invariance properties hold on the entire analytic domain relevant for the cone.  $\square$

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

*Remark 5.36* (Belt regularity). We assume smooth belt-anchored null cuts with finite extrinsic curvature and small tilt within the OS window (see Proposition 5.113). All bounds are per generator length and uniform in  $|R|$ .

*Theorem 5.37* (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  $\xi^\nu$  the belt boost Killing field [17, 18]. 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.* A detailed proof is given in Theorem 5.42; we only summarize the main steps.

The calibrated belt first law identifies the entropy variation with the modular energy up to edge and corner terms Corollary 3.5. Via the belt JLMS channel Proposition 3.4, this is transported to the bulk, where canonical energy is related to the Brown–York/Iyer–Wald flux of the stress tensor through the belt together with Noether-charge corner contributions Theorem 5.46, Lemma 5.112, and Proposition 5.102. Extremality of generalized entropy fixes the area response of the QES under admissible deformations Theorem 5.28. Imposing entanglement equilibrium for all belt deformations then forces the metric perturbation to satisfy the linearized Einstein equations with state-dependent source on  $D[R]$ , as encoded in Theorem 5.37.

QNEC and canonical-energy positivity supply the convexity and positivity input that closes the bootstrap between entropy curvature and energy flux Theorem 5.29. Finally, the  $O(\mathcal{B}_{\text{belt}})$  remainder is controlled by the belt budget and vanishes upon removal of the positive flows by Lemma 3.3.  $\square$

## 5.12 Budgets, constants, and safe defaults

**Budget usage.** All remainders of order  $O(\mathcal{B}_{\text{belt}})$  are tracked against the belt base factor  $\Gamma_{\text{belt}}$  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.

## 5.13 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  $\mathcal{B}_{\text{belt}} \leq \varepsilon/2 + c_{\text{curv}} r(m)^2$ .

## 5.14 QES stability, uniqueness, and continuity

*Theorem 5.38* (QES stability and selection). Work under the framework recap Section 2 (items 1–3,5), the OS kernel Section 3, belt-level recoverability/continuity and QES stability Proposition 5.86 and Section 5.14, and the canonical-energy positivity for non-Killing perturbations [18]. Along any admissible belt-anchored null deformation the second variation satisfies the quantified bound

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

with  $E_{\text{can}} \geq 0$  (kernel: boost–Killing), and  $Q_{\text{shear}} \geq 0$  with belt-local strictly positive coefficients. In particular,  $S - \text{Area}/(4G)$  is locally strictly convex (after removal of positive flows), equivalently  $\mathcal{G} := \text{Area}/(4G) - S$  is locally strictly concave. Consequently:

1. *Uniqueness.* In a neighborhood of a stationary configuration,  $\text{QES}(R)$  is unique.
2. *Lipschitz continuity.* For  $C^2$ -small shape/data variations, the location of  $\text{QES}(R)$  changes by  $O(\mathcal{B}_{\text{belt}})$ . Quantitatively, the displacement is bounded by a modulus that depends only on the coercivity of the canonical-energy/shear form and on the belt continuity constants (controlled by  $\lambda_*$  and  $\lambda_{\text{clu}}$ ).
3. *Saddle selection (belt maximin).* On a belt-anchored family of cuts, the belt flow selects the unique QES and agrees (up to  $O(\mathcal{B}_{\text{belt}})$ ) with the maximin prescription

$$\max_{\text{cuts}} \min_{\Sigma} \left[ \frac{\text{Area}(\Sigma)}{4G} - S(\rho_R; \Sigma) \right].$$

In a local multiplet of stationary saddles, the selected one has the smallest canonical-energy curvature (measured by the Hessian), hence is dynamically preferred.

*Proof. Step 1: Strict convexity/concavity.* By the belt second-order modular equation of state Theorem 5.46, for any admissible belt-anchored null deformation one has

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

where  $Q_{\text{shear}}$  is a positive quadratic form in the belt shear and expansion with belt-uniform strictly positive coefficients, and  $E_{\text{can}} \geq 0$  with kernel consisting of boost–Killing data ( $\text{EE2} \Leftrightarrow \text{SEE2}$ , cf. Theorem 5.47). After removal of positive flows the  $O(\mathcal{B}_{\text{belt}})$  remainder vanishes and the right-hand side is  $> 0$  for any non-Killing direction. Therefore  $S - \text{Area}/(4G)$  has positive-definite Hessian on the OS window, i.e. it is locally strictly convex; equivalently,  $\mathcal{G}$  is locally strictly concave.

*Step 2: Uniqueness near a stationary configuration.* Let  $X \mapsto \Phi(X) := S(\rho_R; X) - \text{Area}(X)/(4G)$  denote the functional restricted to admissible belt-anchored variations of the surface. By Step 1,  $\Phi$  is strictly convex in a neighborhood  $\mathcal{U}$  of any stationary point  $X_*$  (the QES). A strictly convex  $C^2$  functional has at most one stationary point in  $\mathcal{U}$ ; hence the QES is unique in  $\mathcal{U}$ .

*Step 3: Lipschitz continuity of the QES map.* Parametrize the data (shape and state along the belt) by  $p$  in a Banach manifold  $\mathcal{P}$  recording  $C^2$  cut geometry and OS-admissible state perturbations. Let  $X_*(p)$  solve the stationarity equation

$$D\Phi_p(X_*(p)) = 0.$$

By Step 1, the Hessian  $D^2\Phi_p(X_*(p))$  is positive definite with a belt-uniform coercivity constant  $\mu > 0$  set by the lower bounds on  $E_{\text{can}}$  and on the shear form. By belt-level continuity/recoverability Section 5.14 and Proposition 5.86, the map  $p \mapsto D\Phi_p(X)$  is Lipschitz on bounded sets with modulus governed by  $\lambda_*$  and  $\lambda_{\text{clu}}$ , uniformly per generator length. The quantitative implicit-function theorem thus gives, for  $p, p'$  close and with constants independent of  $|R|$ ,

$$\|X_*(p) - X_*(p')\| \leq \mu^{-1} \|D\Phi_p(X_*(p')) - D\Phi_{p'}(X_*(p'))\| \leq C(\lambda_*, \lambda_{\text{clu}}) \|p - p'\| + O(\mathcal{B}_{\text{belt}}),$$

establishing the claimed Lipschitz bound (and the  $O(\mathcal{B}_{\text{belt}})$  stability under flow removal).

*Step 4: Saddle selection and maximin.* For a fixed cut  $C$ , define

$$m(C) := \min_{\Sigma \in \text{adm}(C)} \left[ \frac{\text{Area}(\Sigma)}{4G} - S(\rho_R; \Sigma) \right].$$

Coercivity of the area term on belt tails and belt-level entropy continuity imply existence of minimizers on each cut by the direct method (compactness and lower semicontinuity on the admissible class; cf. Section 5.29). Along positive flows, the belt c-function and local GSL are monotone with a belt-uniform contraction/decay rate; consequently  $m(C_\tau)$  is nondecreasing in the flow parameter and converges to a fixed value as  $\tau \rightarrow \infty$  Theorems 5.41 and 5.58. The fixed points of the belt flow coincide with stationary QES configurations, and contraction in the OS window ensures uniqueness in the basin. Therefore the limit cut maximizes  $m(C)$  over the family, yielding the belt maximin prescription. This matches the unique QES selected by the flow up to  $O(\mathcal{B}_{\text{belt}})$ , with the remainder removed by the positive-flow limit Proposition 5.110.

Finally, in a local family of stationary saddles  $\{X_i\}$  on nearby cuts  $C_i$ , expand the min functional near  $X_i$  using the concavity of  $\mathcal{G}$  (equivalently, convexity of  $\Phi$ ). For small admissible displacements  $s$  along the allowed null directions,

$$\min_{\Sigma \text{ near } X_i} \mathcal{G}(\Sigma) = \mathcal{G}(X_i) - \frac{1}{2} \left( 2\pi E_{\text{can}}[X_i] + Q_{\text{shear}}[X_i] \right) s^2 + O(s^3) + O(\mathcal{B}_{\text{belt}}).$$

With a uniform admissible radius  $s$  in the OS window and  $Q_{\text{shear}} \geq 0$ , the cut achieving the largest minimum is the one with smallest canonical-energy curvature. Thus the belt flow dynamically selects, within such a multiplet, the saddle with minimal canonical energy, completing the selection claim.  $\square$

## 5.15 IR subtraction independence of the forward coefficient

*Lemma 5.39 (IR-scheme invariance).* Let  $a_2^{(\text{even})}$  denote the analytic  $s^2$  coefficient in the gravity-subtracted, forward, crossing-even expansion on the forward cone, extracted with the  $N = 3$  analytic projector of Section 5.22. Then  $a_2^{(\text{even})}$  is invariant under the soft-gravity schemes recorded in the ledger. Equivalently, for any two admissible IR schemes S1, S2,

$$a_2^{(\text{even})}|_{\text{S1}} = a_2^{(\text{even})}|_{\text{S2}}.$$

*Proof.* Let  $A_S(s, t)$  be the gravity-subtracted scattering amplitude in scheme S, and let

$$A_S^{(+)}(s, t) := \frac{1}{2} (A_S(s, t) + A_S(-s, t))$$

be its crossing–even part. On the forward cone (small spacelike  $t \leq 0$ , fixed away from other singular loci) each admissible IR scheme differs by a purely soft piece  $\Delta A := A_{S2} - A_{S1}$  with the universal local structure

$$\Delta A^{(+)}(s, t) = \frac{\alpha(t)}{t} + \beta(t) s^2 \log \frac{|s|}{\mu^2} + O(s^3 \log |s|) + O(t), \quad (5.16)$$

where  $\alpha, \beta$  are smooth in  $t$  on the cone and  $\mu$  is a reference scale. (Polynomial analytic shifts of degree  $\leq 2$  are fixed to zero by the  $N=3$  subtraction convention in the ledger, so they do not appear in (5.16).)

The analytic coefficient  $a_2^{(\text{even})}$  is defined by applying the  $N=3$  analytic projector  $\Pi_{\text{an}}^{(3)}$  to  $A_S^{(+)}(s, t)$  at fixed  $t$  on the cone and reading off the  $s^2$  Taylor coefficient at  $s = 0$ :

$$a_2^{(\text{even})}|_S = \frac{1}{2} \partial_s^2 \left[ \Pi_{\text{an}}^{(3)}(A_S^{(+)}(\cdot, t)) \right] \Big|_{s=0}, \quad t \in \text{cone}. \quad (5.17)$$

By construction (Section 5.22),  $\Pi_{\text{an}}^{(3)}$  acts as the orthogonal projector (with respect to the dispersion–kernel inner product used for the  $N=3$  subtractions) onto the space of  $s$ –analytic germs at  $s = 0$  and *regular* at  $t = 0$ . In particular,

$$\Pi_{\text{an}}^{(3)}[s^2 \log(|s|/\mu^2)] = 0, \quad \Pi_{\text{an}}^{(3)}[\alpha(t)/t] = 0, \quad (5.18)$$

the first because  $s^2 \log |s|$  is nonanalytic at  $s = 0$ , the second because  $1/t$  is regular in  $s$  but *singular* in  $t$  and the projector enforces regularity at  $t = 0$  (the forward limit) before extracting the Taylor data. Moreover, the remainder terms  $O(s^3 \log |s|)$  and  $O(t)$  in (5.16) contribute neither to the  $s^2$  Taylor coefficient nor to the  $t \rightarrow 0$  forward limit: the former vanish under  $\partial_s^2|_{s=0}$ , and the latter vanish by continuity of  $\Pi_{\text{an}}^{(3)}$  on the cone and dominated convergence for  $t \downarrow 0$  (see Section 5.63).

Applying  $\Pi_{\text{an}}^{(3)}$  to  $A_{S2}^{(+)} = A_{S1}^{(+)} + \Delta A^{(+)}$  and using (5.18) gives

$$\Pi_{\text{an}}^{(3)}(A_{S2}^{(+)}(\cdot, t)) = \Pi_{\text{an}}^{(3)}(A_{S1}^{(+)}(\cdot, t)), \quad t \in \text{cone}.$$

Taking  $\frac{1}{2} \partial_s^2|_{s=0}$  and then the forward limit yields, by (5.17),

$$a_2^{(\text{even})}|_{S2} = a_2^{(\text{even})}|_{S1}.$$

Since S1, S2 were arbitrary within the ledger’s admissible class,  $a_2^{(\text{even})}$  is IR–scheme independent.  $\square$

## 5.16 Edge terms, Wald corners, and Ward consistency

*Proposition 5.40* (edge/corner control). With the edge/Wald calibration and Ward consistency developed in Section 5.16 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  $C_{\text{bek}}$  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 5.49 and Lemma 5.77, and the quasi-local belt stress dictionary of Section 5.50 and Proposition 5.78.

*Proof of Proposition 5.40. Step 0: Setup and the first-law channel.* Fix the belt profile and the modular-flow cutoffs from Section 5.16 and the belt first-law channel of Corollary 3.5. For the modular generator  $\zeta_R$  supported in the belt, the Iyer–Wald identity with improvement (the Wald/JKM choice fixed in Section 5.49 and Lemma 5.77) gives, for any on-shell variation  $\delta$ ,

$$\delta\langle K_{\text{mod}}(R) \rangle - \delta S(R) = \int_{\partial\text{Belt}(R)} \left( \delta\mathbf{Q}_{\zeta_R} - \zeta_R \cdot \boldsymbol{\theta}(\delta) \right) + \int_{\text{Belt}(R)} \boldsymbol{\omega}(\delta, \mathcal{L}_{\zeta_R}), \quad (5.19)$$

where  $\boldsymbol{\theta}$  is the symplectic potential,  $\mathbf{Q}_{\zeta_R}$  the (Wald/JKM) Noether charge  $(d-2)$ -form, and  $\boldsymbol{\omega}$  the symplectic current. The Ward-consistency construction of Section 5.16 ensures that the bulk Ward defect in the belt vanishes,

$$\int_{\text{Belt}(R)} \boldsymbol{\omega}(\delta, \mathcal{L}_{\zeta_R}) = 0, \quad (5.20)$$

so (5.19) 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  $\Sigma_{\pm}$  anchored on  $\partial R$ , the timelike regulator worldtube  $W$  cladding  $\partial R$ , 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 5.49), we obtain

$$\delta\langle K_{\text{mod}}(R) \rangle - \delta S(R) = [\text{QES}] + [\text{edge}] + [\text{corner}], \quad (5.21)$$

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  $\zeta_R$  near  $\partial R$  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 5.16 (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  $\zeta_R$ . Therefore

$$[\text{QES}] = \delta \left[ \frac{\text{Area}(\text{QES}(R))}{4G} \right]. \quad (5.22)$$

*Step 3: Edge term as quasi-local belt stress.* On the worldtube  $W$ , the quasi-local belt stress dictionary of Section 5.50 and Proposition 5.78 rewrites the improved boundary integrand as the Brown–York flux contracted with the belt first-law kernel. Concretely,

$$\int_W \left( \delta\mathbf{Q}_{\zeta_R} - \zeta_R \cdot \boldsymbol{\theta}(\delta) \right) = \int_{\partial R} \mathfrak{f}_{\text{belt}}[T^{\text{BY}}, \delta\Upsilon] \, d\ell =: \delta S_{\text{edge}}(R), \quad (5.23)$$

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  $\delta$  (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 (5.21) come from the codimension-two joints. By construction, the JKM ambiguity has been calibrated in Lemma 5.77 so that the local boost Ward identity holds on each joint. This pins down the corner charge uniquely and yields

$$[\text{corner}] = \delta S_{\text{corner}}(R), \quad \delta S_{\text{corner}}(R) = \delta(\mathcal{Q}_{\text{corner}}^{\text{JKM}}), \quad (5.24)$$

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 (5.21), (5.22), (5.23) and (5.24), 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  $\mathcal{B}_{\text{belt}}$ :

$$|\mathfrak{f}_{\text{belt}}[T^{\text{BY}}, \delta \Upsilon]| \leq C_{\text{BY}} \mathcal{B}_{\text{belt}}, \quad |\delta(Q_{\text{corner}}^{\text{JKM}})| \leq C_{\text{JKM}} \mathcal{B}_{\text{belt}}, \quad (5.25)$$

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  $\partial R$  and summing the finitely many joints (whose number is uniformly controlled by the same geometric bindings<sup>2</sup>) gives

$$|\delta S_{\text{edge}}(R)| \leq C_{\text{edge}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}, \quad |\delta S_{\text{corner}}(R)| \leq C_{\text{corner}} \text{length}(\partial R) \mathcal{B}_{\text{belt}}. \quad (5.26)$$

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.  $\square$

## 5.17 Monotonicity along positive flows: a modular c-function

*Theorem 5.41* (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  $\lambda_{\text{clu}}$  fixed in the bindings and remainder controlled as in Section 5.13.

*Proof.* Write

$$S_{\text{gen}}(r; u, s) := S(\rho_R) - \frac{\text{Area}(\text{QES}(R))}{4G}, \quad \mathfrak{c}(r; u, s) = \partial_r S_{\text{gen}}(r; u, s).$$

For each fixed belt radius  $r$  and positive-flow parameters  $(u, s)$ , the framework of Section 2 supplies a belt-local CPTP semigroup  $\{T_\tau\}_{\tau \geq 0}$  generated by the positive flows, KMS-reversible with respect to a boost-KMS reference state  $\sigma_r$  on the belt. The bindings fix a log-Sobolev/cluster rate  $\lambda_{\text{clu}} > 0$  for  $T_\tau$ , uniform per generator length and independent of  $|R|$ . We also use the belt constants  $C_{\text{KMS}}, C_{\text{rec}}, C_{\text{RP}}, C_{\text{tail}}, C_{\text{flow}}$  appearing in the remainder ledger; they enter only through  $\mathcal{B}_{\text{belt}}$  and are controlled as in Section 5.13.

*Step 1: shell representation of the c-density.* Fix  $r > 0$  and a thin radial shell  $B_{r,\delta} := [r, r + \delta]$  on the belt. Let  $A := \partial_{<r} R$  and  $C := \partial_{>r+\delta} R$  be the interior and exterior belts. The

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

telescoping/SSA decomposition of the generalized entropy across shells expresses the radial derivative in terms of conditional mutual information:

$$\mathfrak{c}(r; u, s) = \lim_{\delta \downarrow 0} \frac{1}{\delta} I_{\rho(u, s)}(A : C \mid B_{r, \delta}) + R_1(r; u, s), \quad (5.27)$$

where  $I(\cdot : \cdot \mid \cdot)$  is conditional mutual information and

$$|R_1(r; u, s)| \leq (C_{\text{KMS}} + C_{\text{rec}} + C_{\text{RP}} + C_{\text{tail}} + C_{\text{flow}}) \mathcal{B}_{\text{belt}}.$$

Here the area term has been absorbed using the modular equation of state together with the Brown–York/JKM calibration, and the bound on  $R_1$  is part of the uniform remainder control in Section 5.13.

*Step 2: contraction of shell CMI along the positive flows.* Let  $\rho_\tau := T_\tau(\rho)$  and denote by the same symbol its marginal on  $AB_{r, \delta}C$ . KMS reversibility and the log-Sobolev/cluster inequality for  $T_\tau$  imply a differential inequality for the shell conditional mutual information:

$$\frac{d}{d\tau} I_{\rho_\tau}(A : C \mid B_{r, \delta}) \leq -\lambda_{\text{clu}} I_{\rho_\tau}(A : C \mid B_{r, \delta}) + R_2(r, \delta; u(\tau), s(\tau)), \quad (5.28)$$

where  $R_2$  collects the leakage across the shell boundaries (Lieb–Robinson tails) and the small non-Markovian corrections from the positive flows, and satisfies

$$|R_2(r, \delta; u, s)| \leq (C_{\text{tail}} + C_{\text{flow}}) \mathcal{B}_{\text{belt}} \delta.$$

The reference state  $\sigma_r$  is a fixed point of  $T_\tau$ , so no additional drift term appears.

*Step 3: passage to the  $c$ -function and strictness.* For the trajectory  $(u(\tau), s(\tau))$ , combine (5.27) and (5.28). Dividing (5.28) by  $\delta$  and taking  $\delta \downarrow 0$ , we obtain

$$\frac{d}{d\tau} \mathfrak{c}(r; u(\tau), s(\tau)) \leq -\lambda_{\text{clu}} \mathfrak{c}(r; u(\tau), s(\tau)) + R_3(r; u(\tau), s(\tau)),$$

where  $R_3$  arises from the limit of  $R_2/\delta$  together with the  $\tau$ -derivative of  $R_1$  and obeys

$$|R_3(r; u, s)| \leq C_{\text{rem}} \mathcal{B}_{\text{belt}}$$

for a belt-uniform constant  $C_{\text{rem}}$ . By construction this constant is included in the  $O(\mathcal{B}_{\text{belt}})$  remainder and is controlled explicitly in Section 5.13, which yields the differential inequality stated in the theorem.

Finally, the inequality is strict unless the perturbation is boost-Killing. If the perturbation is not generated by a boost Killing field, the semigroup  $T_\tau$  has no additional fixed points beyond  $\sigma_r$ , and the log-Sobolev/cluster inequality is strict, so  $I_{\rho_\tau}(A : C \mid B_{r, \delta}) > 0$  for all  $\tau > 0$  and small shells. The bound above is then strict for  $\mathfrak{c}(r; u(\tau), s(\tau))$ . Conversely, if the perturbation is boost-Killing,  $\rho_\tau$  is stationary under  $T_\tau$ , the shell conditional mutual information is constant, and  $\mathfrak{c}$  is constant as well, saturating the inequality. This completes the proof.  $\square$

## 5.18 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, 5.22, 5.43, 5.48 and 5.78)—the inequalities used in Theorems 5.29 and 5.33 hold with strictly positive slack. No external runs are required on this envelope; the belt error is tracked by Section 5.12.

**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 three tester families are nonnegative by construction (Sections 5.10, 5.22 and 5.78), 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 5.15, 5.46, 5.49, 5.50 and 5.66). They certify Theorems 5.29 and 5.33 on  $\mathcal{F}_{\text{CDT/GFT}}$  without CDT/GFT runs, with the remainder controlled by  $O(\mathcal{B}_{\text{belt}})$ .

### 5.19 Expanded proof: modular equation of state

*Theorem 5.42* (Expanded modular equation of state). Under the framework recap Section 2, the belt JLMS channel Proposition 3.4, belt-level recovery/continuity and QES stability Proposition 5.86 and Section 5.14, canonical-energy/QSEI inputs Theorem 5.46, Lemma 5.112, and Proposition 5.102, 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.

*Proof. 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 5.16, 5.49 and 5.50, Propositions 5.40 and 5.78, and Lemma 5.77,

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

**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  $\xi$  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 [11, 17, 18].

**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  $\delta S_{\text{corner}}$  Section 5.49 and Lemma 5.77. The net edge+corner variation cancels within  $O(\mathcal{B}_{\text{belt}})$  Sections 5.16 and 5.50 and Propositions 5.40 and 5.78.

**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 5.37. Positivity of canonical energy and removal of flows Theorem 5.46 and Lemma 3.3 complete the limit.  $\square$

*Remark 5.43* (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 5.37 and its expansion in Section 5.19. 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 5.85. The gravitational interpretation is tracked quantitatively by the belt monotones in Section 6.4, namely the  $c$ -function (Theorem 5.41) and the width-flow monotone (Theorem 5.88), and by the local generalized second law on belts (Theorem 5.58).

## 5.20 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  $2\pi$ .

*Lemma 5.44* (ANEC/QNEC constants). For admissible states  $\mathfrak{S}_{\text{adm}}$  and belt-anchored deformations,

$$\begin{aligned} \int_{-\infty}^{\infty} \langle T_{kk}(u) \rangle du &\geq 0 & (\kappa_{\text{ANEC}} = 1), \\ \partial_u^2 S(u) &\leq 2\pi \langle T_{kk}(u) \rangle & (\kappa_{\text{QNEC}} = 2\pi). \end{aligned}$$

*Proof.* Fix an admissible belt and a generator with affine parameter  $u$ , and let  $f \in C_0^{\infty}(\mathbb{R})$  be a real profile along this generator. Belt RP/KMS positivity Lemma 3.1, together with the canonical-energy/QSEI inputs Theorem 5.46, Lemma 5.112, and Proposition 5.102 and the first-law channel Corollary 3.5, gives the modular light-ray kernel used in the QNEC construction: for any admissible state  $\chi$ ,

$$\langle K_{\sigma}''(0) \rangle_{\chi} = 2\pi \int_{\mathbb{R}} du f(u)^2 \langle T_{kk}(u) \rangle_{\chi} + W_{\sigma}[f; \chi] + O(\mathcal{B}_{\text{belt}}),$$

where  $\sigma$  is the belt boost-KMS reference state,  $W_{\sigma}[f; \chi] \geq 0$  is the canonical-energy quadratic form, and the overall factor  $2\pi$  is fixed by the Rindler witness  $K_{\text{R}}$  in the benchmark setup (Bisognano–Wichmann normalization of the modular boost).

Light-ray localization (null timeslice propagation, cluster/decay, and LR tails) expresses the entropic Hessian along the same deformation as

$$S_{\chi}''(0) = \int_{\mathbb{R}} du f(u)^2 \partial_u^2 S(u) + O(\mathcal{B}_{\text{belt}}) \|f\|_2^2.$$

Positivity of relative entropy between  $\chi$  and  $\sigma$  implies  $S_{\chi}''(0) - \langle K_{\sigma}''(0) \rangle_{\chi} \leq 0$ . Combining the two displays and removing positive flows by Lemma 3.3 yields, for all real  $f$ ,

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

Choosing a standard  $\delta$ -sequence  $f_n$  supported in an interval of width  $1/n$  around a fixed  $u_0$  and using belt light-ray regularity gives

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

for every  $u_0$ , i.e. the pointwise QNEC with coefficient  $\kappa_{\text{QNEC}} = 2\pi$ .

To pass to ANEC, integrate the QNEC inequality along the complete generator. Belt cluster/decay and the null timeslice property imply  $\partial_u S(u) \rightarrow 0$  as  $u \rightarrow \pm\infty$  for admissible states, so that

$$\int_{-\infty}^{\infty} \partial_u^2 S(u) du = 0.$$

Integrating the QNEC bound and using flow removal once more gives

$$0 = \int_{-\infty}^{\infty} \partial_u^2 S(u) du \leq 2\pi \int_{-\infty}^{\infty} \langle T_{kk}(u) \rangle du,$$

which is equivalent to

$$\int_{-\infty}^{\infty} \langle T_{kk}(u) \rangle du \geq 0.$$

Thus  $\kappa_{\text{ANEC}} = 1$ . The use of the Rindler normalization and Lemma 3.3 ensures that both constants are stable along belt flows and independent of the choice of belt within the admissible atlas.  $\square$

## 5.21 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  $a_{\text{QES}}$  be the belt area line density. The Page crossover occurs when the area term matches the entropy bound:

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

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

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

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

## 5.22 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 the analytic forward coefficient is IR-scheme independent:

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

in agreement with the IR-scheme invariance proved in the amplitude/IR subsection Section 5.15.

**Add-on (strip invariance of the analytic projector).** By Corollary 5.120,  $\Pi_2$  is insensitive to the strip counterterms  $\delta M_{\text{strip}}$  and  $\delta C_{\text{cel}}$ , and continues to excise the same nonanalytic soft pieces;  $a_2^{(\text{even})}$  remains IR-scheme invariant on the strip.

## 5.23 JLMS under mild running: RG completeness

*Proposition 5.45 (JLMS–RG).* With the logarithmic improvements  $(\ell_S, \ell_T)$  and the operator–mixing map  $\mathbf{R}_{S/T}$ , the JLMS identity for belt–anchored regions is stable to first order:

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

*Proof. Normalization and convention.* Throughout we use the belt JLMS identity at finite regulators,

$$S(\rho_R || \sigma_R) = E_{\text{can}}^W[\delta \Psi; \xi] + O(\mathcal{B}_{\text{belt}}), \quad (5.29)$$

where  $E_{\text{can}}^W$  denotes the (dimensionless) canonical–energy flux with the modular/boost normalization fixed by the Rindler witness and the JKM calibration (so that the conventional  $2\pi$  factor is absorbed into the definition of  $E_{\text{can}}^W$ ). The remainder  $O(\mathcal{B}_{\text{belt}})$  is uniform per generator length and independent of  $|R|$ .

*Step 1: Differentiate the belt JLMS identity.* Let  $\partial_{\log \mu}$  denote  $d/d \log \mu$  at fixed states  $\rho, \sigma$ . Differentiating (5.29) yields

$$\partial_{\log \mu} S(\rho_R || \sigma_R) - \partial_{\log \mu} E_{\text{can}}^W[\delta \Psi; \xi] = \partial_{\log \mu} O(\mathcal{B}_{\text{belt}}). \quad (5.30)$$

Since  $\mathcal{B}_{\text{belt}}$  depends only on belt geometry/flows (and not on the renormalization scale), its derivative remains  $O(\mathcal{B}_{\text{belt}})$ . Thus it suffices to show that the  $\mu$ -derivatives of the two main terms match up to  $O(\mathcal{B}_{\text{belt}})$ .

*Step 2: RG variation of boundary relative entropy.* Write

$$S(\rho_R||\sigma_R) = \Delta\langle K_{\text{mod}}^R[\sigma] \rangle - \Delta S_R, \quad \Delta X := X(\rho) - X(\sigma).$$

At first order in running, the modular generator and the entropy admit logarithmic improvements

$$\partial_{\log \mu} K_{\text{mod}}^R[\sigma] = \mathbf{R}_{S/T} \left[ \int_R d\Sigma^\mu \xi^\nu T_{\mu\nu} \right] + \delta K_{\text{edge}}, \quad \partial_{\log \mu} S_R = \ell_S \mathcal{J}_{\text{edge}} + O(\mathcal{B}_{\text{belt}}),$$

where  $\mathbf{R}_{S/T}$  is the operator–mixing map that converts the RG drift of entropic data into an improved stress–tensor insertion ( $\ell_T$ -type improvement), and  $\delta K_{\text{edge}}, \mathcal{J}_{\text{edge}}$  are local belt corner/edge functionals determined by the JKM-calibrated scheme. Taking  $\Delta$ -expectation values,

$$\partial_{\log \mu} S(\rho_R||\sigma_R) = \Delta \left\langle \mathbf{R}_{S/T} \left[ \int_R d\Sigma^\mu \xi^\nu T_{\mu\nu} \right] \right\rangle + \Delta \langle \delta K_{\text{edge}} \rangle - \ell_S \Delta \mathcal{J}_{\text{edge}} + O(\mathcal{B}_{\text{belt}}). \quad (5.31)$$

*Step 3: RG variation of wedge canonical energy.* By definition (Brown–York/Iyer–Wald dictionary on the belt),

$$E_{\text{can}}^W[\delta\Psi; \xi] = \int_{\Sigma \subset W} d\Sigma^\mu \xi^\nu \Delta \langle T_{\mu\nu} \rangle + \mathcal{E}_{\text{edge}},$$

where  $\mathcal{E}_{\text{edge}}$  packages the (JKM-calibrated) corner/edge piece that is equivalent to the area/-Noether corner in the flux identity. Under mild running,

$$\partial_{\log \mu} E_{\text{can}}^W = \Delta \left\langle \int_{\Sigma} d\Sigma^\mu \xi^\nu \partial_{\log \mu} T_{\mu\nu} \right\rangle + \partial_{\log \mu} \mathcal{E}_{\text{edge}} + O(\mathcal{B}_{\text{belt}}). \quad (5.32)$$

The RG drift of  $T_{\mu\nu}$  is an improvement of the form  $\partial_{\log \mu} T_{\mu\nu} = \mathbf{R}_{S/T}^*[T_{\mu\nu}] + \nabla_\alpha \nabla_\beta X^{\alpha\beta}{}_{\mu\nu}$ , whose double-divergence piece reduces, once integrated against  $\xi^\nu$  over  $\Sigma$ , to a boundary term on the belt. By the Brown–York identity on the belt and the JKM calibration, these boundary terms are precisely absorbed into  $\partial_{\log \mu} \mathcal{E}_{\text{edge}}$  up to  $O(\mathcal{B}_{\text{belt}})$ . Hence

$$\partial_{\log \mu} E_{\text{can}}^W = \Delta \left\langle \mathbf{R}_{S/T} \left[ \int_{\Sigma} d\Sigma^\mu \xi^\nu T_{\mu\nu} \right] \right\rangle + \partial_{\log \mu} \mathcal{E}_{\text{edge}} + O(\mathcal{B}_{\text{belt}}), \quad (5.33)$$

where we have used that  $\mathbf{R}_{S/T}$  acts (by construction) compatibly on boundary and wedge channels along the belt.

*Step 4: Edge cancellation (JKM/BY calibration).* The belt JLMS channel is RG-compatible on the code subspace, so the mixing map  $\mathbf{R}_{S/T}$  is the same in (5.31)–(5.33). The remaining edge terms cancel by the calibrated Ward identity:

$$\Delta \langle \delta K_{\text{edge}} \rangle - \ell_S \Delta \mathcal{J}_{\text{edge}} - \partial_{\log \mu} \mathcal{E}_{\text{edge}} = O(\mathcal{B}_{\text{belt}}),$$

because the JKM choice equates the corner potential with the Brown–York flux on the belt and fixes the boost normalization (hence the  $\log \mu$  drift redistributes only between equivalent edge representations). Therefore,

$$\partial_{\log \mu} S(\rho_R||\sigma_R) - \partial_{\log \mu} E_{\text{can}}^W[\delta\Psi; \xi] = O(\mathcal{B}_{\text{belt}}). \quad (5.34)$$

*Step 5: Conclusion and removal.* Combining (5.30) and (5.34) proves the claim at finite positive flows. By the removal lemma,  $\mathcal{B}_{\text{belt}} \rightarrow 0$  as  $(u, s) \downarrow 0$ , so the remainder vanishes in the continuum window.

Hence

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

□

## 5.24 Nonlinear completion: second-order modular equation of state

*Theorem 5.46* (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 5.86 and Section 5.14, canonical-energy/QSEI inputs Proposition 5.102 and Lemma 5.112, 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

$$\delta^2 \left[ S(\rho_R) - \frac{\text{Area}(\text{QES}(R))}{4G} \right] \geq 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}}, \quad (5.35)$$

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

$$\begin{aligned} Q_{\text{shear}}[\delta g] &:= \kappa_\sigma \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} d^{d-2}x \sqrt{\gamma} \sigma^{ab}(\lambda, x) \sigma_{ab}(\lambda, x) \\ &\quad + \kappa_\theta \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} d^{d-2}x \sqrt{\gamma} \theta(\lambda, x)^2, \end{aligned} \quad (5.36)$$

with explicit belt-local positive coefficients

$$\kappa_\sigma \geq \underline{\kappa}_\sigma > 0, \quad \kappa_\theta \geq \underline{\kappa}_\theta > 0, \quad (5.37)$$

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 (5.35) vanishes for all belt deformations, then the semiclassical Einstein equations hold to second order in expectation on  $D[R]$ .

*Proof.* Fix a belt  $\partial_r R$  and a one-parameter family of admissible states and shapes  $\{\rho_\lambda, \mathcal{Y}_\lambda\}_{\lambda \in [\lambda_1, \lambda_2]}$  anchored to the belt, with  $\rho_0 = \sigma$  and generator  $\xi$  on the belt. Write

$$S_{\text{gen}}(\lambda) := S(\rho_\lambda; R_\lambda) - \frac{\text{Area}(\text{QES}(R_\lambda))}{4G}.$$

*Step 1: First variation (calibrated belt first law).* By the calibrated belt kernel (Corollary 5.4) together with the Brown–York/Iyer–Wald dictionary (Proposition 6.2) and JKM fix (Lemmas 6.1 and 5.77),

$$\frac{d}{d\lambda} S_{\text{gen}}(\lambda) \Big|_{\lambda=0} = 2\pi \int_{\Sigma \subset W} d\Sigma^\mu \xi^\nu \frac{d}{d\lambda} \langle T_{\mu\nu} \rangle_{\rho_\lambda} \Big|_{\lambda=0} + O(\mathcal{B}_{\text{belt}}). \quad (5.38)$$

The remainder is uniform per generator length and removable.

*Step 2: Differentiate again and pass to canonical energy.* Differentiate (5.38) once more at  $\lambda = 0$  and use linear response plus conservation:

$$\begin{aligned} \frac{d^2}{d\lambda^2} S_{\text{gen}}(\lambda) \Big|_{\lambda=0} &= 2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \frac{d^2}{d\lambda^2} \langle T_{\mu\nu} \rangle_{\rho_\lambda} \Big|_{\lambda=0} + O(\mathcal{B}_{\text{belt}}) \\ &= 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + \mathcal{B}_{\text{grav}}[\delta g; \xi] - C'_2 \mathcal{B}_{\text{belt}}. \end{aligned} \quad (5.39)$$

Here  $E_{\text{can}}^W$  is the Iyer–Wald canonical energy associated with the linearized field  $\delta\Psi := \frac{d}{d\lambda} \Psi_\lambda \Big|_{\lambda=0}$  and the flow  $\xi$ , obtained by standard manipulations of the presymplectic current with source  $\delta T_{\mu\nu}$  and by moving the bulk integral to the timelike belt via the Brown–York identity. The term  $\mathcal{B}_{\text{grav}}[\delta g; \xi]$  collects purely geometric second-order contributions from varying the belt cut (extrinsic data) and the gravitational corner potential; these are controlled by the calibrated corner/edge analysis and the Raychaudhuri decomposition on the belt.

*Step 3: Raychaudhuri decomposition and positivity on the belt.* On the null legs generated by  $\xi$  at the belt corners, the Raychaudhuri equation decomposes the gravitational piece into

positive quadratic forms of the shear and expansion plus a term proportional to  $R_{kk}$  (with  $k = \xi$  affinely normalized at the belt):

$$\mathcal{B}_{\text{grav}}[\delta g; \xi] = Q_{\text{shear}}[\delta g] + \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \left( \alpha_\theta \theta^2 + \alpha_\sigma \sigma^{ab} \sigma_{ab} \right) + \mathcal{R}_{kk}, \quad (5.40)$$

with positive belt-local weights  $\alpha_{\theta, \sigma}$  fixed by the JKM calibration and boost normalization. The term  $\mathcal{R}_{kk}$  is proportional to  $\int R_{kk}$  and is traded for matter by the semiclassical Einstein equations *in expectation* to the relevant order; on the belt this replacement is justified by the QSEI/ray estimates and shear control in Proposition 5.102 and Lemma 5.112. Consequently, after absorbing the calibrated corner potential and collecting constants, there exist belt-local positive coefficients  $\kappa_\sigma, \kappa_\theta$ , bounded below as in (5.37), such that

$$\mathcal{B}_{\text{grav}}[\delta g; \xi] \geq Q_{\text{shear}}[\delta g] - C_2'' \mathcal{B}_{\text{belt}}, \quad (5.41)$$

with equality iff the deformation is boost–Killing on the belt ( $\sigma_{ab} = \theta = 0$  along the legs).

*Step 4: Combine and conclude.* Insert (5.41) into (5.39) and relabel constants:

$$\left. \frac{d^2}{d\lambda^2} S_{\text{gen}}(\lambda) \right|_{\lambda=0} \geq 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}},$$

which is (5.35). The statement along a finite segment  $[\lambda_1, \lambda_2]$  follows by integrating the local density and using uniformity of the constants per generator length.

*Consequences.* If the left-hand side vanishes for all admissible belt deformations, then the right-hand side must vanish. Positivity of  $Q_{\text{shear}}$  forces boost–Killing data on the belt; positivity and coercivity of canonical energy in the working cone (with kernel given by  $\mathcal{L}_\xi$ –pure gauge modes) yields the semiclassical Einstein equations in expectation to second order on  $D[R]$ .  $\square$

*Theorem 5.47* (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  $O(\mathcal{B}_{\text{belt}})$ :

(EE<sup>2</sup>) *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).

(SEE<sup>2</sup>) *Second-order semiclassical Einstein equations in expectation on  $D[R]$ :* the second-order modular equation of state of Theorem 5.46 holds and the Iyer–Wald canonical energy is nonnegative with kernel given by boost–Killing modes.

*Proof.* (EE<sup>2</sup>  $\Rightarrow$  SEE<sup>2</sup>). By Theorem 5.46, for every admissible deformation,

$$\delta^2 \left[ S - \frac{\text{Area}}{4G} \right] \geq 2\pi E_{\text{can}}^W[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}}. \quad (5.42)$$

If (EE<sup>2</sup>) holds, the left-hand side is  $\geq 0$  with equality only for boost–Killing data. Positivity of  $Q_{\text{shear}}$  (with kernel given by boost–Killing modes) forces  $Q_{\text{shear}}[\delta g] = 0$  on the kernel, hence  $E_{\text{can}}^W[\delta\Psi; \xi] \geq -O(\mathcal{B}_{\text{belt}})$  for all directions. By the canonical-energy/QSEI inputs (Proposition 5.102 and Lemma 5.112) and removal of flows (Lemma 3.3), we obtain  $E_{\text{can}}^W \geq 0$  with kernel equal to boost–Killing modes, i.e. (SEE<sup>2</sup>).

(SEE<sup>2</sup>  $\Rightarrow$  EE<sup>2</sup>). Assume the second-order SEE and nonnegativity of  $E_{\text{can}}^W$  with the stated kernel. Then (5.42) yields

$$\delta^2 \left[ S - \frac{\text{Area}}{4G} \right] \geq Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}} \geq -C_2 \mathcal{B}_{\text{belt}},$$

with equality (after removal) only if  $\delta\Psi$  is boost–Killing and  $Q_{\text{shear}}[\delta g] = 0$ , i.e. the deformation is a belt boost isometry. Using the first-order belt equilibrium (the calibrated first law) to impose  $\delta[S - \text{Area}/(4G)] = 0$  and the above second-order bound gives (EE<sup>2</sup>).  $\square$

## 5.25 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  $\sigma$ ,  $\phi(u) = A \exp(-(u - u_0)^2/(2\sigma^2))$ , 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  $\geq 0$ , and the pointwise QNEC bound at  $u=0$  is numerically tight at the chosen center. The example fixes conventions and constants for referee checks.

## 5.26 Regulator and belt-width independence with quantitative bound

*Lemma 5.48* (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

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

*Proof.* Fix an admissible belt family and auxiliary positive flows  $(u, s) > 0$  in the OS window. For a belt-regularized observable  $\mathcal{O}$  in this section let  $\mathcal{O}_r$  denote its value at belt width  $r$  (with  $(u, s)$  kept fixed). We first obtain an exponential bound on width moves at fixed  $(u, s)$ , and then pass to variations and remove the auxiliary flows.

*Step 1: Exponential control of belt moves at fixed flows.* Consider increasing the belt width from  $r$  to  $r + \Delta r$  with  $\Delta r$  small and positive (the case  $\Delta r < 0$  is analogous). By belt microcausality and null timeslice propagation (Lemma 5.74 and Proposition 5.75), the effect of this move on any belt-regularized observable can be implemented by a relative Cauchy evolution whose generator is supported in the shell between the two belts, together with an exponentially small tail controlled by the same Lieb–Robinson rate  $\mu_{\text{eff}}$ . More concretely, there exists  $C_0 > 0$ , independent of  $|R|$ , such that for every observable  $\mathcal{O}$  used below

$$|\mathcal{O}_{r+\Delta r} - \mathcal{O}_r| \leq C_0 e^{-\mu_{\text{eff}} r} |\Delta r| \quad (5.43)$$

for all  $r$  in the admissible window and all  $|\Delta r|$  sufficiently small.

For the specific observables of interest, this shell control follows from the quasi-local belt dictionary and Brown–York flux representation (Section 5.50 and Proposition 5.78):

- $\langle K_{\text{mod}} \rangle_r$  and  $\text{Area}_r$  are represented, up to uniformly bounded  $O(\mathcal{B}_{\text{belt}})$  remainders, as integrals of local densities (modular/flux density and calibrated area density) over the belt and its null thickening. Changing the width by  $\Delta r$  adds or removes only the contribution of the shell, whose expectation value is suppressed by the microcausality tail  $e^{-\mu_{\text{eff}} r}$ .

- For the entropy  $S_r$  the same shell control follows from the belt JLMS/first-law channel: the entropy is expressed in terms of modular energy plus Brown–York flux, with all terms governed by the same local densities. Therefore the incremental change  $S_{r+\Delta r} - S_r$  is controlled by the same shell integral and obeys a bound of the form (5.43).

Thus at fixed  $(u, s)$  the map  $r \mapsto \mathcal{O}_r$  is locally Lipschitz and satisfies

$$|\partial_r \mathcal{O}_r| \leq C_0 e^{-\mu_{\text{eff}} r}$$

for  $\mathcal{O} \in \{\langle K_{\text{mod}} \rangle, S, \text{Area}\}$ , with  $C_0$  independent of  $|R|$  and of  $(u, s)$  in the OS window.

*Step 2: Integrating along the width.* Let  $r, r' > 0$  and assume without loss of generality that  $r \geq r'$ . For any of the observables in Step 1 we have, at fixed  $(u, s)$ ,

$$|\mathcal{O}_r - \mathcal{O}_{r'}| = \left| \int_{r'}^r \partial_\rho \mathcal{O}_\rho d\rho \right| \leq \int_{r'}^r C_0 e^{-\mu_{\text{eff}} \rho} d\rho \leq \frac{C_0}{\mu_{\text{eff}}} e^{-\mu_{\text{eff}} r'}.$$

By symmetry in  $r, r'$  this yields

$$|\Delta_{r \rightarrow r'} \langle K_{\text{mod}} \rangle| + |\Delta_{r \rightarrow r'} S| + |\Delta_{r \rightarrow r'} \text{Area}| \leq C_1 e^{-\mu_{\text{eff}} \min\{r, r'\}}, \quad (5.44)$$

for some constant  $C_1 > 0$  independent of  $|R|$  and of  $(u, s)$ .

*Step 3: Passing to variations.* Let  $\{\rho(\lambda)\}_\lambda$  be an admissible one-parameter family of states and denote by  $\delta$  the derivative at  $\lambda = 0$ . Apply (5.44) pointwise to the family  $\mathcal{O}_r(\lambda)$ . The shell/Lieb–Robinson bounds in Step 1 hold uniformly for  $\lambda$  in a small neighborhood of 0, so (5.44) is valid with a constant  $C_1$  that does not depend on  $\lambda$ . By the OS short-evolution decomposition and positivity (Lemma 3.1) the functions  $\lambda \mapsto \langle K_{\text{mod}} \rangle_r(\lambda)$  and  $\lambda \mapsto S_r(\lambda)$  are differentiable, with derivatives  $\delta \langle K_{\text{mod}} \rangle_r$  and  $\delta S_r$  controlled by the same envelope. Differentiating (5.44) at  $\lambda = 0$  and using dominated convergence gives

$$|\Delta_{r \rightarrow r'} \delta \langle K_{\text{mod}} \rangle| + |\Delta_{r \rightarrow r'} \delta S| \leq C_1 e^{-\mu_{\text{eff}} \min\{r, r'\}}. \quad (5.45)$$

The area term in the lemma appears without variation; for Area we simply retain the bound already obtained in (5.44).

*Step 4: Removing the positive flows.* Finally, by the OS flow-removal lemma (Lemma 3.3), the belt-regularized quantities at width  $r$  converge, as  $(u, s) \downarrow 0$ , to the physical observables appearing in the statement, and the convergence is uniform under the Lieb–Robinson/factorization envelope used above. Hence the inequalities (5.44) and (5.45) persist in the joint limit  $(u, s) \downarrow 0$  with the same constants. Absorbing all numerical factors into a single constant  $\tilde{c} > 0$  yields

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

uniformly per generator length and with  $\tilde{c}$  independent of  $|R|$ . □

## 5.27 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

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

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

## 5.28 Anchor choice and relational dressing: invariance statement

*Proposition 5.49* (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

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

*Proof. Step 0: Anchor change as a belt circuit plus a tail.* By microcausality and null timeslice propagation (Lemma 5.74 and Proposition 5.75), any admissible change of anchor  $\mathcal{C} \mapsto \mathcal{C}'$  is implemented on the belt code subspace by a unitary

$$W_{\mathcal{C} \rightarrow \mathcal{C}'} = U_{\text{belt}} + T_{\text{tail}},$$

where  $U_{\text{belt}}$  is a bounded-depth circuit supported on (a thickening of) the belt and  $T_{\text{tail}}$  has support outside the belt with norm  $\|T_{\text{tail}}\| \leq c e^{-\mu_{\text{eff}} r} = O(\mathcal{B}_{\text{belt}})$ . The OS-KMS analytic projector (Lemma 3.1 and Proposition 5.67) ensures that  $U_{\text{belt}}$  preserves the analytic core and the common domains of the unbounded operators used below.

*Step 1: Modular energy.* Write the belt modular generator at anchor  $\mathcal{C}$  as  $\widehat{K}_{\text{mod}}[\mathcal{C}]$  and similarly at  $\mathcal{C}'$ . By functoriality of the belt dictionary and the Brown-York calibration (Section 5.50 and Proposition 5.78), one has on the analytic core

$$\widehat{K}_{\text{mod}}[\mathcal{C}'] = U_{\text{belt}} \widehat{K}_{\text{mod}}[\mathcal{C}] U_{\text{belt}}^* + R_K, \quad \|R_K\|_{\text{form}} \leq C_K \mathcal{B}_{\text{belt}}, \quad (5.46)$$

where  $R_K$  collects the exponentially suppressed tail together with the calibrated edge shift. Taking expectations in any admissible state  $\rho$ ,

$$|\langle \widehat{K}_{\text{mod}} \rangle_{\rho}[\mathcal{C}'] - \langle \widehat{K}_{\text{mod}} \rangle_{\rho}[\mathcal{C}]| \leq \|R_K\|_{\text{form}} \leq C_K \mathcal{B}_{\text{belt}}.$$

*Step 2: Entropy.* Let  $\rho_R[\mathcal{C}]$  and  $\rho_R[\mathcal{C}']$  be the reduced states defined with the two anchors. By belt recovery/continuity (Proposition 5.86) and the OS short-evolution decomposition (Lemma 3.1), the map  $\mathcal{E}_{\mathcal{C} \rightarrow \mathcal{C}'}(\cdot) := \text{Ad}_{W_{\mathcal{C} \rightarrow \mathcal{C}'}}(\cdot)$  is a CPTP instrument on the belt algebra with

$$\|\mathcal{E}_{\mathcal{C} \rightarrow \mathcal{C}'} - \text{id}\|_{1 \rightarrow 1} \leq C_{\text{cb}} \mathcal{B}_{\text{belt}}.$$

Hence, by continuity of entropy under near-identity channels in the OS window,

$$|S(\rho_R[\mathcal{C}']) - S(\rho_R[\mathcal{C}])| \leq C_S \mathcal{B}_{\text{belt}}. \quad (5.47)$$

(The constant  $C_S$  depends only on the belt ledger constants and is uniform in  $|R|$ .)

*Step 3: Area/generalized entropy.* The JKM-calibrated quantum-area operator  $A$  and the Brown-York flux are related on the belt (Lemma 5.77 and Proposition 5.78). The anchor change

preserves the belt and its boost generator; thus the induced change in the corner potential is a belt coboundary whose contribution is exponentially suppressed:

$$\left| \left\langle \frac{\widehat{A}}{4G} \right\rangle_{\rho} [\mathcal{C}'] - \left\langle \frac{\widehat{A}}{4G} \right\rangle_{\rho} [\mathcal{C}] \right| \leq C_A e^{-\mu_{\text{eff}} r} \leq C_A \mathcal{B}_{\text{belt}}. \quad (5.48)$$

Combining (5.47)–(5.48) yields the same bound for the generalized entropy  $S_{\text{gen}} = S - \text{Area}/(4G)$ .

*Step 4: Amplitude/Regge testers.* Amplitude functionals used in Section 5 factor through belt-localized celestial/forward testers. The unitary  $U_{\text{belt}}$  does not change the tester values; only the tail  $T_{\text{tail}}$  contributes. By microcausality/timeslice bounds and the compact-dual acceptance, the variation is controlled by the same exponential tail:

$$|\mathcal{A}[\mathcal{C}'] - \mathcal{A}[\mathcal{C}]| \leq C_A \mathcal{B}_{\text{belt}}.$$

*Step 5: Collection and removal.* Let  $C_{\mathcal{C} \rightarrow \mathcal{C}'} := \max\{C_K, C_S, C_A, C_{\mathcal{A}}\}$ . The previous steps show  $|\mathcal{O}[\mathcal{C}'] - \mathcal{O}[\mathcal{C}]| \leq C_{\mathcal{C} \rightarrow \mathcal{C}'} \mathcal{B}_{\text{belt}}$  for all observables  $\mathcal{O}$  in the suite. Finally, by the positive-flow removal lemma (Lemma 3.3),  $\mathcal{B}_{\text{belt}} \rightarrow 0$  as  $(u, s) \downarrow 0$ , so the difference vanishes in the continuum belt window.

This proves the proposition.  $\square$

## 5.29 Existence of quantum extremal surfaces (direct method)

*Theorem 5.50 (QES existence).* Assume the framework recap Section 2 (items 1–3,5), belt microcausality/timeslice control Lemma 5.74 and Proposition 5.75, belt-level nesting/recovery and continuity Propositions 5.86 and 5.93 and Section 5.14, 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  $\Sigma$  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. Step 0: Admissible class and topology.* Fix a belt-anchored region  $R$  and its wedge  $W = \text{EW}(R)$  in the OS belt window with finite  $\mathcal{B}_{\text{belt}}$ . Let  $\text{Surf}_R$  denote the class of admissible codimension-two surfaces  $\Sigma \subset W$  that:

- are homologous to  $R$  and anchored on the belt  $\partial_r R$ ;
- lie within the belt-controlled geometric region determined by the width  $r$  and the curvature bounds of the framework;
- satisfy the regularity/curvature conditions used in Section 5.14 (bounded extrinsic curvature, no wild cusps, etc.).

Equip  $\text{Surf}_R$  with the topology  $\tau$  induced by:

- local  $C^{0,\alpha}$  (or varifold) convergence of the induced geometry on the belt;
- convergence of the associated wedge regions  $W[\Sigma]$  in the null-timeslice sense (Proposition 5.75);
- convergence of the induced states on the belt algebra via the OS/GNS construction, controlled by the recovery/continuity results of Proposition 5.86 and Section 5.14.

We regard  $\mathcal{G} : \text{Surf}_R \rightarrow \mathbb{R} \cup \{+\infty\}$  as a functional on  $(\text{Surf}_R, \tau)$ .

*Step 1: Coercivity of  $\mathcal{G}$ .* We show that  $\mathcal{G}$  is coercive: sublevel sets  $\{\Sigma \in \text{Surf}_R \mid \mathcal{G}[\Sigma] \leq M\}$  have uniformly bounded area and cannot escape the belt window.

By the belt AGSP/seed pipeline and recovery continuity (Proposition 5.86 and Section 5.14), there exist constants  $c_0, c_1 \geq 0$  (depending only on belt data, not on  $\Sigma$ ) such that for all  $\Sigma \in \text{Surf}_R$ ,

$$|S(\rho_R; \Sigma)| \leq c_0 + c_1 \text{Area}(\Sigma), \quad (5.49)$$

with  $c_1$  strictly smaller than the area coefficient  $1/(4G)$  when measured per generator length (this is the ‘‘per-length entropy bound’’ from the AGSP/seed pipeline). Then

$$\begin{aligned} \mathcal{G}[\Sigma] &= \frac{\text{Area}(\Sigma)}{4G} - S(\rho_R; \Sigma) \\ &\geq \frac{\text{Area}(\Sigma)}{4G} - (c_0 + c_1 \text{Area}(\Sigma)) \\ &= \left(\frac{1}{4G} - c_1\right) \text{Area}(\Sigma) - c_0. \end{aligned}$$

Let  $\kappa := \frac{1}{4G} - c_1 > 0$ . If  $\mathcal{G}[\Sigma] \leq M$ , then

$$\kappa \text{Area}(\Sigma) - c_0 \leq M \quad \Rightarrow \quad \text{Area}(\Sigma) \leq \frac{M + c_0}{\kappa}.$$

Thus each sublevel set  $\{\mathcal{G} \leq M\}$  consists of surfaces with uniformly bounded area.

Next, microcausality and null timeslice propagation (Lemma 5.74 and Proposition 5.75) together with the belt factorization (Section 5.50 and Proposition 5.78) imply that any large excursion of  $\Sigma$  away from the belt region (e.g. far along a null generator) produces an area cost that is not compensated by an entropic gain: the entropy of the belt degrees of freedom is insensitive to such tails up to  $O(\mathcal{B}_{\text{belt}})$ , while the area increases linearly in the excursion. Combined with (5.49), this means that a minimizing sequence for  $\mathcal{G}$  cannot ‘‘run off to infinity’’ along the null directions without sending  $\mathcal{G}[\Sigma]$  to  $+\infty$ . Hence the sublevel sets of  $\mathcal{G}$  lie in a geometrically bounded region determined by the belt width and curvature bounds. In particular,  $\mathcal{G}$  is coercive on  $\text{Surf}_R$ .

*Step 2: Lower semicontinuity of  $\mathcal{G}$ .* Let  $\Sigma_n \rightarrow \Sigma$  in  $(\text{Surf}_R, \tau)$ . We claim

$$\mathcal{G}[\Sigma] \leq \liminf_{n \rightarrow \infty} \mathcal{G}[\Sigma_n]. \quad (5.50)$$

First, the area functional is lower semicontinuous under the geometric convergence encoded by  $\tau$ . Standard results in geometric measure theory (e.g. varifold or  $C^{0,\alpha}$  convergence for surfaces with bounded area and curvature) give

$$\text{Area}(\Sigma) \leq \liminf_{n \rightarrow \infty} \text{Area}(\Sigma_n). \quad (5.51)$$

Second, the entanglement entropy  $S(\rho_R; \Sigma)$  is continuous along admissible belt-anchored deformations in the OS window. Indeed, by belt-level nesting/recovery and continuity (Propositions 5.86 and 5.93 and Section 5.14), the reduced states on the belt algebra associated with  $\Sigma_n$  converge in trace norm (or in the relevant OS sense) to the state associated with  $\Sigma$ , and the entropy functional is continuous for such locally finite-dimensional approximants with the belt budget held fixed. Concretely, for  $\Sigma_n \rightarrow \Sigma$  in  $\tau$ ,

$$\lim_{n \rightarrow \infty} S(\rho_R; \Sigma_n) = S(\rho_R; \Sigma). \quad (5.52)$$

(If one prefers, (5.52) can be decomposed into a relative-entropy continuity bound plus the first-law channel of Corollary 3.5.)

Combining (5.51) and (5.52) we obtain

$$\begin{aligned} \mathcal{G}[\Sigma] &= \frac{\text{Area}(\Sigma)}{4G} - S(\rho_R; \Sigma) \\ &\leq \frac{1}{4G} \liminf_{n \rightarrow \infty} \text{Area}(\Sigma_n) - \lim_{n \rightarrow \infty} S(\rho_R; \Sigma_n) \\ &\leq \liminf_{n \rightarrow \infty} \left( \frac{\text{Area}(\Sigma_n)}{4G} - S(\rho_R; \Sigma_n) \right) = \liminf_{n \rightarrow \infty} \mathcal{G}[\Sigma_n], \end{aligned}$$

which is (5.50). Thus  $\mathcal{G}$  is lower semicontinuous on  $(\text{Surf}_R, \tau)$ .

*Step 3: Compactness of minimizing sequences.* Let  $\{\Sigma_n\} \subset \text{Surf}_R$  be a minimizing sequence:

$$\lim_{n \rightarrow \infty} \mathcal{G}[\Sigma_n] = \inf_{\Sigma \in \text{Surf}_R} \mathcal{G}[\Sigma] =: m.$$

By coercivity (Step 1), the set  $\{\Sigma_n\}$  lies in a region of uniformly bounded area and bounded geometry (curvature and belt width). Standard compactness theorems for codimension-two surfaces with uniform area and curvature bounds in a fixed geometric region (e.g. Federer–Fleming or Allard compactness in the appropriate category) imply that, after passing to a subsequence, there exists  $\Sigma_\star \in \text{Surf}_R$  such that

$$\Sigma_n \rightarrow \Sigma_\star \quad \text{in } (\text{Surf}_R, \tau).$$

The role of the OS kernel (Lemmas 3.1 and 3.3) and the positive-flow regularization is to ensure that the belt-local modular dynamics remain well behaved under this convergence: short positive flows  $(u, s) > 0$  mollify the geometry/state near the belt, giving a compact class of “regularized” surfaces, and flow removal  $(u, s) \downarrow 0$  returns to the original  $\Sigma_\star$  while preserving the bounds and the convergence of  $\mathcal{G}$  up to  $O(\mathcal{B}_{\text{belt}})$ , which vanishes in the removal limit.

Thus any minimizing sequence admits a  $\tau$ -convergent subsequence with limit  $\Sigma_\star \in \text{Surf}_R$ .

*Step 4: Existence and properties of minimizers.* By Step 2 (lower semicontinuity),

$$\mathcal{G}[\Sigma_\star] \leq \liminf_{n \rightarrow \infty} \mathcal{G}[\Sigma_n] = m = \inf_{\Sigma \in \text{Surf}_R} \mathcal{G}[\Sigma].$$

Hence  $\mathcal{G}[\Sigma_\star] = m$ , so  $\Sigma_\star$  is a minimizer of  $\mathcal{G}$  within the admissible class  $\text{Surf}_R$ .

From the coercivity estimate of Step 1,

$$\text{Area}(\Sigma_\star) \leq \frac{m + c_0}{\kappa},$$

so minimizers can be chosen with uniformly bounded area. Because  $\Sigma_\star$  arises as a limit of surfaces in the geometrically bounded sublevel set, it also lies in the corresponding compact geometric class determined by the belt width and curvature bounds. This gives the claimed uniform control on minimizers.

Thus  $\mathcal{G}$  admits at least one minimizer on the admissible class of codimension-two surfaces homologous to  $R$ , and minimizers can be chosen with the stated uniform bounds. These minimizers are the quantum extremal surfaces in the sense of the framework.  $\square$

### 5.30 Conservation and Bianchi consistency from modular dynamics

*Proposition 5.51* (modular conservation). Under the framework recap Section 2 (items 1,2,5) and the modular equation of state Theorem 5.37, 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  $\nabla^\mu \delta \langle \text{Ein}_{\mu\nu} \rangle = 0$ .

*Proof. Step 0: Setup and modular equation of state.* Work on the belt domain  $D[R] = \overline{D_{\text{an}}}$  and its entanglement wedge  $W = \text{EW}(R)$ , with boost generator  $\xi^\mu$  fixed by the Rindler witness and JKM calibration. The (linear) modular equation of state of Theorem 5.37, in its calibrated form, can be written schematically as

$$\delta\langle K_{\text{mod}}(R) \rangle = \delta\left\langle \frac{\widehat{A}(W)}{4G} \right\rangle + 2\pi \int_{\Sigma \subset W} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}), \quad (5.53)$$

for all admissible one-parameter families of states and shapes supported on the belt, with all corner/edge contributions absorbed by the JKM/Brown–York calibration (Lemma 5.77 and Proposition 5.78) into the displayed terms and the  $O(\mathcal{B}_{\text{belt}})$  remainder.

Under the framework assumptions and Theorem 5.37, this linear relation is equivalent, on  $D[R]$ , to the *expectation-value Einstein equation* in the wedge,

$$\delta\langle \text{Ein}_{\mu\nu} \rangle = 8\pi G \delta\langle T_{\mu\nu} \rangle, \quad \text{on } D[R], \quad (5.54)$$

where  $\text{Ein}_{\mu\nu} = G_{\mu\nu} + \Lambda g_{\mu\nu}$  is the Einstein tensor (possibly with a cosmological term).

*Step 1: Divergence of the expectation-value Einstein equation.* Apply the background covariant derivative  $\nabla^\mu$  to (5.54):

$$\nabla^\mu \delta\langle \text{Ein}_{\mu\nu} \rangle = 8\pi G \nabla^\mu \delta\langle T_{\mu\nu} \rangle. \quad (5.55)$$

Because  $\text{Ein}_{\mu\nu}[g]$  is constructed covariantly from the metric, it satisfies the contracted Bianchi identity for *any* metric,

$$\nabla^\mu \text{Ein}_{\mu\nu}[g] \equiv 0.$$

Linearizing this identity around the background and using that the background obeys the (unperturbed) Einstein equations, we obtain

$$\nabla^\mu \delta \text{Ein}_{\mu\nu} = 0. \quad (5.56)$$

Taking expectations (which is linear) gives

$$\nabla^\mu \delta\langle \text{Ein}_{\mu\nu} \rangle = 0. \quad (5.57)$$

Substituting (5.57) into (5.55) yields

$$\nabla^\mu \delta\langle T_{\mu\nu} \rangle = 0, \quad \text{on } D[R], \quad (5.58)$$

i.e. covariant conservation of the expectation-value stress tensor on the belt domain. This shows both modular conservation and consistency of (5.54) with the contracted Bianchi identity.

*Step 2: Modular/boost derivation and absence of boundary anomalies.* We now explain how this conservation law is seen directly from the modular dynamics and why no boundary (corner/edge) anomaly survives on the belt.

Consider two Cauchy surfaces  $\Sigma_1, \Sigma_2 \subset W$  related by a small displacement along the belt boost flow generated by  $\xi^\mu$ , with the belt corners kept within the OS window. Apply (5.53) to a fixed admissible deformation between states  $\rho$  and  $\sigma$  on both  $\Sigma_1$  and  $\Sigma_2$ , and subtract:

$$\begin{aligned} 0 &= \delta\langle K_{\text{mod}}(R_{\Sigma_2}) \rangle - \delta\langle K_{\text{mod}}(R_{\Sigma_1}) \rangle \\ &\quad - \left( \delta\left\langle \frac{\widehat{A}(W_{\Sigma_2})}{4G} \right\rangle - \delta\left\langle \frac{\widehat{A}(W_{\Sigma_1})}{4G} \right\rangle \right) - 2\pi \left( \int_{\Sigma_2} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu} \rangle - \int_{\Sigma_1} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu} \rangle \right) + O(\mathcal{B}_{\text{belt}}). \end{aligned} \quad (5.59)$$

By anchor/dressing invariance and the boost Killing property on the belt (Section 5.28 and Proposition 5.49), the change of the generalized entropy between  $\Sigma_1$  and  $\Sigma_2$  is  $O(\mathcal{B}_{\text{belt}})$ ; more precisely,

$$\delta\langle K_{\text{mod}}(R_{\Sigma_2})\rangle - \delta\langle K_{\text{mod}}(R_{\Sigma_1})\rangle - \left(\delta\langle \frac{\widehat{A}(W_{\Sigma_2})}{4G}\rangle - \delta\langle \frac{\widehat{A}(W_{\Sigma_1})}{4G}\rangle\right) = O(\mathcal{B}_{\text{belt}}),$$

since both  $K_{\text{mod}}$  and the JKM-calibrated area are invariant under such short boost moves up to the belt budget. Inserting this into (5.59) gives

$$\int_{\Sigma_2} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu}\rangle - \int_{\Sigma_1} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu}\rangle = O(\mathcal{B}_{\text{belt}}). \quad (5.60)$$

Let  $V$  be the spacetime region bounded by  $\Sigma_1$ ,  $\Sigma_2$ , and the timelike belt segment between them. Using Gauss's theorem and that  $\xi^\mu$  is (approximately) Killing on the belt, the left-hand side of (5.60) can be written as

$$\int_V dV \nabla^\mu (\xi^\nu \delta\langle T_{\mu\nu}\rangle) + \text{calibrated corner/belt terms},$$

where the corner/belt contributions are controlled by the JKM fix and the Brown–York flux identity (Lemma 6.1 and Proposition 6.2) and are therefore  $O(\mathcal{B}_{\text{belt}})$ . Expanding the divergence and using  $\nabla_{(\mu}\xi_{\nu)} = O(\mathcal{B}_{\text{belt}})$  on the belt, we obtain

$$\int_V dV \xi^\nu \nabla^\mu \delta\langle T_{\mu\nu}\rangle = O(\mathcal{B}_{\text{belt}}).$$

Since this holds for arbitrary small boost-displaced regions and for arbitrary admissible deformations, and the belt vectors  $\xi^\nu$  span the relevant null directions, it follows (in the distributional sense) that

$$\nabla^\mu \delta\langle T_{\mu\nu}\rangle = O(\mathcal{B}_{\text{belt}}) \quad \text{on } D[R].$$

Finally, by the positive-flow removal lemma (Lemma 3.3),  $\mathcal{B}_{\text{belt}} \rightarrow 0$  as  $(u, s) \downarrow 0$ , so we recover (5.58) in the continuum limit.

*Step 3: Bianchi consistency via JLMS.* The JLMS channel on belts (Proposition 3.4) identifies boundary relative entropy with bulk canonical energy (plus the calibrated area term) up to  $O(\mathcal{B}_{\text{belt}})$ , and the full operator equation of state upgrades this to an equality of generators. The geometric side is thus governed by the Einstein tensor, which satisfies the contracted Bianchi identity identically. Taking expectations and using the argument of Step 1 confirms that

$$\nabla^\mu \delta\langle \text{Ein}_{\mu\nu}\rangle = 0$$

for the same class of perturbations and on the same domain  $D[R]$ , so the modularly derived Einstein equations are fully consistent with the Bianchi identity.

Combining Steps 1–3 proves the proposition.  $\square$

### 5.31 Counterterm stability and renormalization scheme independence

*Lemma 5.52* (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 5.13.

*Proof.* We compare two renormalization schemes, A and B, that differ only by adding local, diffeomorphism-covariant counterterms supported on the belt and its corners:

$$\delta\mathcal{L} = \alpha_1 \int_{\text{belt}} R + \alpha_2 \int_{\text{belt}} K^2 + \dots,$$

with coefficients  $\alpha_i$  fixed once and for all. The corresponding changes in the renormalized observables are:

- (i) a shift of the area/Wald-entropy functional on the belt and corners,
- (ii) a shift of the renormalized stress tensor by local curvature/contact terms,
- (iii) a shift of the UV entropy counterterm (vacuum entanglement subtraction).

All statements are per generator length and uniform in  $|R|$ .

*Step 1: Area and Wald-edge shifts.* A local gravitational counterterm  $\int_{\text{belt}} R$  or  $\int_{\text{belt}} K^2$  contributes to the on-shell action and thus to the Wald entropy and its edge/corner terms. In Einstein gravity with JKM calibration, these contributions appear as shifts of the corner potential  $\xi \cdot \Theta$  and of the Wald charge  $Q_\xi$  on the belt.

By the belt version of the JKM fix (Lemmas 6.1 and 5.77) and the Brown-York dictionary (Propositions 6.2 and 5.78), the boost Ward identity is enforced so that

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

*in either scheme.* Changing the coefficients  $\alpha_i$  shifts both the area term and the Wald-edge potential in lockstep, preserving the calibrated combination up to  $O(\mathcal{B}_{\text{belt}})$ . Thus the difference  $\delta(\text{Area}/4G)_A - \delta(\text{Area}/4G)_B$  is supported only in the belt tails and satisfies

$$\left| \delta\left[\frac{\text{Area}}{4G}\right]_A - \delta\left[\frac{\text{Area}}{4G}\right]_B \right| \leq C_{\text{area}} \mathcal{B}_{\text{belt}} \quad (5.61)$$

for some constant  $C_{\text{area}}$  independent of  $|R|$ .

*Step 2: Stress-tensor shift and modular equation of state.* The same counterterms change the renormalized stress tensor by local curvature/contact terms supported on the belt:

$$T_{\mu\nu}^{(A)} = T_{\mu\nu}^{(B)} + \Delta T_{\mu\nu}^{\text{ct}}, \quad \Delta T_{\mu\nu}^{\text{ct}} = \Delta T_{\mu\nu}^{\text{ct}}[g] \text{ local on the belt.}$$

The modular equation of state in each scheme reads, for any admissible variation,

$$\begin{aligned} \delta\langle K_{\text{mod}}(R) \rangle &= \delta\left\langle \frac{\hat{A}}{4G} \right\rangle_A + 2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu}^{(A)} \rangle + O(\mathcal{B}_{\text{belt}}) \\ &= \delta\left\langle \frac{\hat{A}}{4G} \right\rangle_B + 2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle T_{\mu\nu}^{(B)} \rangle + O(\mathcal{B}_{\text{belt}}), \end{aligned}$$

where  $\Sigma \subset W$  is the wedge Cauchy slice and  $\xi^\mu$  is the belt boost field. Subtracting the two relations gives

$$\delta\left\langle \frac{\hat{A}}{4G} \right\rangle_A - \delta\left\langle \frac{\hat{A}}{4G} \right\rangle_B + 2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle \Delta T_{\mu\nu}^{\text{ct}} \rangle = O(\mathcal{B}_{\text{belt}}). \quad (5.62)$$

By construction,  $\Delta T_{\mu\nu}^{\text{ct}}$  is a local geometric tensor that arises from varying the same belt/corner counterterms that produced the area/Wald shift. The boost Ward identity ensures that its flux along  $\xi^\mu$  is the divergence of a local belt current, and the JKM/BY calibration is chosen so that this current cancels the area shift at leading order. Consequently, the integral term in (5.62) is bounded by the same exponential belt tail,

$$\left| 2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta\langle \Delta T_{\mu\nu}^{\text{ct}} \rangle \right| \leq C_{\text{flux}} \mathcal{B}_{\text{belt}}. \quad (5.63)$$

Combining (5.61) and (5.63), we find that the modular equation of state takes the same form in schemes A and B up to  $O(\mathcal{B}_{\text{belt}})$ .

*Step 3: Entropy counterterm and recovery channel.* The renormalized entanglement entropy differs between the schemes by a local counterterm supported on the belt/corners, schematically

$$S_{\text{ren}}^{(\text{A})} = S_{\text{ren}}^{(\text{B})} + S_{\text{ct}}, \quad S_{\text{ct}} = \int_{\text{belt/corners}} d^{d-2}x \sqrt{\gamma} \mathcal{F}[g],$$

for some local scalar  $\mathcal{F}$ . Under admissible variations supported in the OS window, the recovery/continuity bounds on the belt (Proposition 5.86) and the LR/factorization control imply

$$|\delta S_{\text{ct}}| \leq C_S \mathcal{B}_{\text{belt}},$$

with  $C_S$  independent of  $|R|$ . Intuitively, the entropy counterterm measures only ultra-local vacuum structure near the belt, whose change under admissible deformations is exponentially suppressed in the belt width and tracked by the same budget  $\mathcal{B}_{\text{belt}}$ .

*Step 4: Generalized entropy difference.* Putting the pieces together, the generalized entropy variation in the two schemes differs by

$$\begin{aligned} \delta \left[ S - \frac{\text{Area}}{4G} \right]_{\text{A}} - \delta \left[ S - \frac{\text{Area}}{4G} \right]_{\text{B}} &= \delta S_{\text{ren}}^{(\text{A})} - \delta S_{\text{ren}}^{(\text{B})} - \left( \delta \left[ \frac{\text{Area}}{4G} \right]_{\text{A}} - \delta \left[ \frac{\text{Area}}{4G} \right]_{\text{B}} \right) \\ &= \delta S_{\text{ct}} - \left( \delta \left[ \frac{\text{Area}}{4G} \right]_{\text{A}} - \delta \left[ \frac{\text{Area}}{4G} \right]_{\text{B}} \right), \end{aligned}$$

so that by (5.61) and the bound on  $\delta S_{\text{ct}}$ ,

$$\left| \delta \left[ S - \frac{\text{Area}}{4G} \right]_{\text{A}} - \delta \left[ S - \frac{\text{Area}}{4G} \right]_{\text{B}} \right| \leq (C_S + C_{\text{area}}) \mathcal{B}_{\text{belt}}.$$

Renaming  $C_{\text{CT}} := C_S + C_{\text{area}}$  and invoking the positivity of the belt budget shows

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

with the same  $\mathcal{B}_{\text{belt}}$  as in Section 5.13. This proves the lemma.  $\square$

### 5.32 $\Gamma$ -convergence from discrete surrogates to continuum belts

*Theorem 5.53 (Gamma-bridge).* Let  $\{\mathcal{G}_h\}$  be discrete generalized-entropy functionals built from CDT/GFT belt nets on the fixed tester envelope of Sections 5.18, 5.43 and 5.48. Then  $\mathcal{G}_h$   $\Gamma$ -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  $\mathcal{G}$ .

*Proof. Step 0: Setup and topology.* Let  $\mathbf{X}$  denote the admissible class of belt-anchored cuts/-surfaces (or equivalently, their induced wedge domains) equipped with the topology  $\tau$  used in Sections 5.18 and 5.43: convergence  $X_h \rightarrow X$  means that the belt-local geometry and state data induced by  $X_h$  converge to those of  $X$  in the sense of:

- local  $C^{0,\alpha}$  convergence of the induced metric and extrinsic data on the belt;
- weak convergence of the induced states on the belt algebra (via the OS/GNS construction);
- convergence of all test functionals in the finite tester envelope of Sections 5.10, 5.22 and 5.78.

For each mesh parameter  $h > 0$  (CDT/GFT belt net spacing), let  $\mathcal{G}_h : \mathbb{X} \rightarrow \mathbb{R}$  be the discrete generalized-entropy functional defined by:

$$\mathcal{G}_h(X) := S_h(X) - \frac{\text{Area}_h(X)}{4G_h},$$

where  $S_h$  is the discrete entanglement entropy on the belt net,  $\text{Area}_h$  is the discretized area/Wald functional, and  $G_h$  is the renormalized coupling at scale  $h$ , all defined on the fixed tester envelope as in Sections 5.18 and 5.43. Let  $\mathcal{G}$  be the continuum generalized entropy on  $\mathbb{X}$ :

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

with the JKM/Brown–York calibration in force (Sections 5.49 and 5.50, Lemma 5.77, and Proposition 5.78).

We prove that  $\mathcal{G}_h$   $\Gamma$ -converges to  $\mathcal{G}$  with respect to  $\tau$ , and that any sequence of (approximate) minimizers of  $\mathcal{G}_h$  has accumulation points that minimize  $\mathcal{G}$  (QES minimizers).

*Step 1: Liminf inequality.* Let  $X_h \rightarrow X$  in  $(\mathbb{X}, \tau)$ . We must show

$$\mathcal{G}(X) \leq \liminf_{h \rightarrow 0} \mathcal{G}_h(X_h). \quad (5.64)$$

By construction, both discrete and continuum generalized entropies admit the same tester representation up to  $O(\mathcal{B}_{\text{belt}})$ , with the same belt budget and calibrated scheme:

$$\mathcal{G}(X) = \sum_j \Phi_j^{\text{cont}}(X), \quad \mathcal{G}_h(X_h) = \sum_j \Phi_{j,h}^{\text{disc}}(X_h),$$

where the index  $j$  ranges over the finite set of forward even-parity dispersion projectors, Gaussian Hankel/impact functionals, and principal-series celestial Gram testers introduced in Sections 5.10, 5.22 and 5.78; and the discrete testers  $\Phi_{j,h}$  converge to  $\Phi_j^{\text{cont}}$  in the sense of the envelope.

Each tester is *nonnegative* and lower semicontinuous in the topology  $\tau$ :

- forward even-parity dispersion projectors are positive by construction of the compact duals and the forward cone;
- Gaussian Hankel/impact testers are positivity-preserving quadratic forms of the amplitude;
- celestial Gram testers are positive semidefinite on the principal-series space.

By the convergence of the discrete testers to their continuum counterparts on the envelope and the positivity, we have for each  $j$ :

$$\Phi_j^{\text{cont}}(X) \leq \liminf_{h \rightarrow 0} \Phi_{j,h}^{\text{disc}}(X_h).$$

Summing over  $j$  (finite sum) and using the uniform belt budget to control the  $O(\mathcal{B}_{\text{belt}})$  remainder (which vanishes with  $h$  by the discrete-to-continuum scaling of the nets) yields

$$\mathcal{G}(X) \leq \liminf_{h \rightarrow 0} \mathcal{G}_h(X_h),$$

which is (5.64).

*Step 2: Recovery sequences (limsup inequality).* Fix  $X \in \mathbb{X}$ . We must construct a sequence  $X_h \rightarrow X$  such that

$$\mathcal{G}(X) \geq \limsup_{h \rightarrow 0} \mathcal{G}_h(X_h). \quad (5.65)$$

Choose a CDT/GFT belt net adapted to  $X$  as in Section 5.18, with mesh size  $h$ , and let  $X_h$  be the discrete representation obtained by intersecting  $X$  with the net and replacing it by the

corresponding polygonal/polyhedral belt cut. By construction,  $X_h \rightarrow X$  in  $\tau$  as  $h \rightarrow 0$ , and the belt dictionary plus quasi-local factorization (Section 5.50 and Proposition 5.78) yield

$$\Phi_{j,h}^{\text{disc}}(X_h) = \Phi_j^{\text{cont}}(X) + O(\mathcal{B}_{\text{belt}}(h)),$$

for each tester  $j$ , where  $\mathcal{B}_{\text{belt}}(h) \rightarrow 0$  as  $h \rightarrow 0$ , uniformly for  $X$  in bounded subsets of  $\mathsf{X}$  (belt width fixed, local geometry controlled). Summing over  $j$  gives

$$\mathcal{G}_h(X_h) = \mathcal{G}(X) + O(\mathcal{B}_{\text{belt}}(h)).$$

Taking  $\limsup_{h \rightarrow 0}$  and using  $\mathcal{B}_{\text{belt}}(h) \rightarrow 0$  yields (5.65).

The existence of such recovery sequences uses:

- bounded-depth belt circuits implementing the discrete embedding of the continuum belt algebra on the net (Lemma 3.1);
- OS short evolutions and flow removal to control the discrepancy between the discrete and continuum modular flows (Lemma 3.3);
- quasi-local factorization of the belt/wedge dictionary (Section 5.50).

Thus the limsup inequality holds.

*Step 3: Equi-coercivity / tightness.* To apply the fundamental theorem of  $\Gamma$ -convergence, we must check a coercivity/tightness condition. The uniform belt budgets and invariances established in Sections 5.12, 5.13, 5.15, 5.46 and 5.66 imply that:

- for any  $M \in \mathbb{R}$ , the sublevel sets  $\{X \in \mathsf{X} \mid \mathcal{G}_h(X) \leq M\}$  are contained in a common compact subset of  $(\mathsf{X}, \tau)$ , independent of  $h$ ;
- geometric fluctuations of the belt cuts are controlled by the same LR/Regge tail (e.g. no wild oscillations at scales smaller than  $h$ ), so sequences with bounded  $\mathcal{G}_h$  cannot escape the admissible class or develop pathological small-scale structure on the belt.

In the language of  $\Gamma$ -convergence, the functionals  $\mathcal{G}_h$  are *equi-coercive* on  $(\mathsf{X}, \tau)$ .

*Step 4: Conclusion (minimizers and QES).* We have shown that:

- (i) the liminf inequality (5.64) holds for all  $X_h \rightarrow X$ ;
- (ii) for each  $X$  there exists a recovery sequence  $X_h \rightarrow X$  satisfying (5.65);
- (iii) the family  $\{\mathcal{G}_h\}$  is equi-coercive on  $(\mathsf{X}, \tau)$ .

Hence, by the fundamental theorem of  $\Gamma$ -convergence (see, e.g., any standard reference),  $\mathcal{G}_h$   $\Gamma$ -converges to  $\mathcal{G}$  on  $(\mathsf{X}, \tau)$ , and moreover:

- any limit point of a sequence of (exact or approximate) minimizers  $X_h^*$  of  $\mathcal{G}_h$  is a minimizer of  $\mathcal{G}$ ;
- conversely, for any minimizer  $X^*$  of  $\mathcal{G}$  there exist discrete minimizers  $X_h^*$  converging (along a subsequence) to  $X^*$ .

By the QES existence and uniqueness results in the continuum theory, the minimizers of  $\mathcal{G}$  are precisely the QES surfaces. Therefore any accumulation point of discrete minimizers  $X_h^*$  is a QES minimizer of  $\mathcal{G}$ , as claimed.

This completes the proof of the Gamma-bridge theorem. □

**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  $G_h$   $\Gamma$ -converges to the continuum  $G$  (Theorem 5.53). 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 Section 7.5 (Fig. 2) 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.

### 5.33 Compact dual certificates for amplitude positivity

*Proposition 5.54* (finite-support extremals). Consider any conic combination of the project's nonnegative testers on the cone  $\mathcal{S}$ :

1. forward even-parity functionals at  $Q$  fixed  $t$ -nodes,
2.  $P$  Hankel/Gaussian band kernels,
3.  $K$  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  $\leq 18$  nodes.

*Proof. Step 0: Dual cone and finite-dimensional parametrization.* On the fixed amplitude cone  $\mathcal{S}$ , the testers used in the project are all of the form

$$\mathcal{T}[A] = \sum_j w_j \phi_j[A],$$

where:

- each  $\phi_j$  is an evaluation functional of the amplitude (or of an explicitly bounded linear transform of it) at some node in the relevant variables ( $t$ -grid node, Hankel band center, celestial anchor, etc.);
- each  $w_j \geq 0$  is a nonnegative weight;
- each  $\phi_j$  is nonnegative on the admissible cone  $\mathcal{S}$ :  $\phi_j[A] \geq 0$  whenever  $A \in \mathcal{S}$ .

Thus the set of all nonnegative testers of the type listed in the proposition is the *dual cone*

$$\mathcal{C}^* = \text{cone} \{\phi_1, \dots, \phi_M\}$$

generated by finitely many basic functionals  $\phi_j$ , where  $M$  is the total number of grid nodes used (all  $t$ -nodes, Hankel bands, and celestial anchors). Each element of  $\mathcal{C}^*$  corresponds to a vector  $w = (w_1, \dots, w_M)$  with  $w_j \geq 0$ , and the map

$$w \mapsto \Phi_w := \sum_{j=1}^M w_j \phi_j$$

is linear. Hence we can identify  $\mathcal{C}^*$  with a closed convex cone in the finite dimensional space  $\mathbb{R}^M$ .

*Step 1: Restriction to the active coordinates.* Let  $\Phi \in \mathcal{C}^*$  be any dual certificate obtained as a conic combination of the given testers; write

$$\Phi = \sum_{j=1}^M w_j \phi_j, \quad w_j \geq 0.$$

Let  $J \subset \{1, \dots, M\}$  denote the set of indices corresponding to the *active* tester families, i.e. those associated with:

- the  $Q$  forward even-parity  $t$ -nodes;
- the  $P$  Hankel/Gaussian band kernels;
- the  $K$  celestial Gram functionals on principal lines.

By construction of the tester envelope, the value of  $\Phi$  on the admissible amplitude cone  $\mathcal{S}$  depends only on the weights  $w_j$  with  $j \in J$ ; any index outside  $J$  is either not used or can be absorbed into these families without changing the action of  $\Phi$  on  $\mathcal{S}$ . Thus we may, without loss of generality, restrict to the finite-dimensional cone

$$\mathcal{K} := \left\{ w \in \mathbb{R}^{|J|} \mid w_j \geq 0 \right\},$$

and treat  $\Phi$  as an element of the conic hull of  $\{\phi_j\}_{j \in J}$ .

*Step 2: Carathéodory-type reduction.* The cone generated by  $\{\phi_j\}_{j \in J}$  lives in a finite dimensional vector space whose dimension is at most  $(Q + P + K)$ : one dimension for each of the  $Q$  forward nodes,  $P$  Hankel bands, and  $K$  celestial anchors. In particular, we can identify the coefficients of  $\Phi$  with a point in  $\mathbb{R}^d$ , where

$$d \leq Q + P + K.$$

Carathéodory's theorem for convex cones (a conic analogue of the standard Carathéodory theorem) then implies the following: any point in the conic hull of  $\{\phi_j\}_{j \in J}$  can be expressed as a nonnegative combination of at most  $d$  generators. Concretely, there exist indices  $j_1, \dots, j_m \in J$  with  $m \leq d$  and weights  $\tilde{w}_{j_\ell} \geq 0$  such that

$$\Phi = \sum_{\ell=1}^m \tilde{w}_{j_\ell} \phi_{j_\ell}, \quad m \leq Q + P + K.$$

Thus every dual certificate of the type described in the statement admits a representation supported on at most  $(Q + P + K)$  nodes across the relevant variables.

*Step 3: Extremality and nonnegativity.* By construction, all  $\tilde{w}_{j_\ell} \geq 0$ , so nonnegativity of the weights is preserved in the reduced representation. Among all such representations of  $\Phi$  with minimal support size  $m$ , the corresponding functional is *extremal* in the dual cone: if we could decompose  $\Phi$  as a nontrivial sum  $\Phi = \Phi^{(1)} + \Phi^{(2)}$  of two certificates in  $\mathcal{C}^*$  with nonproportional weight vectors, at least one of the components would require fewer than  $m$  nonzero coefficients, contradicting minimality of the support.

Therefore, for each conic combination of the project's testers, there exists an extremal dual certificate supported on at most  $(Q + P + K)$  nodes, with nonnegative weights. This proves the main claim. The quantitative corollary follows by substituting  $(Q, P, K) = (6, 7, 5)$ , which yields a support bound  $Q + P + K = 18$ .  $\square$

**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 3 (support remains 18); (ii) *augment* by these five to obtain a 23-support certificate. In both cases, all weights remain nonnegative by Lemma 5.118.

### 5.34 Normalization consistency between amplitude and modular dynamics

*Lemma 5.55* ( $\kappa$ -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

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

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.

*Proof.* Consider the Einstein–Hilbert action with Newton constant  $G$  and matter fields collectively denoted by  $\psi$ ,

$$S = \frac{1}{16\pi G} \int d^4x \sqrt{-g} R + S_{\text{m}}[\psi, g_{\mu\nu}].$$

We expand the metric around Minkowski space as

$$g_{\mu\nu} = \eta_{\mu\nu} + \kappa h_{\mu\nu},$$

with  $\kappa$  chosen so that  $h_{\mu\nu}$  has a canonically normalized quadratic action. Expanding  $S$  to second order in  $h_{\mu\nu}$  and to first order in the coupling to matter gives the linearized equation of motion

$$\mathcal{E}^{\mu\nu}{}_{\rho\sigma} h^{\rho\sigma} = -\frac{\kappa}{2} T^{\mu\nu},$$

where  $\mathcal{E}$  is the standard kinetic operator for a massless spin-2 field and  $T_{\mu\nu}$  is the stress tensor of the matter sector. Thus the interaction Lagrangian is

$$\mathcal{L}_{\text{int}} = -\frac{\kappa}{2} h_{\mu\nu} T^{\mu\nu},$$

so that each external leg couples to the graviton with strength  $\kappa/2$ .

Tree-level  $t$ -channel exchange of a single graviton between two conserved sources then produces the well-known soft forward amplitude

$$\mathcal{A}_{\text{soft}}(s, t) = \kappa^2 \left( \frac{s^2}{-t} + \frac{u^2}{-t} \right),$$

for  $|t| \ll |s|, |u|$ , as written in the lemma.

On the other hand, varying the same action with respect to  $g_{\mu\nu}$  yields Einstein's equations

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = 8\pi G T_{\mu\nu}.$$

Black-hole thermodynamics in this theory leads to the Bekenstein–Hawking area law,

$$S_{\text{BH}} = \frac{\text{Area}}{4G},$$

so that the coefficient in the modular equation of state

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

is governed by the same Newton constant  $G$  that appears in the Einstein equations.

The normalization of  $\kappa$  relative to  $G$  is fixed by demanding that the linearized equations derived from the Einstein–Hilbert action coincide with those obtained from the canonically

normalized spin-2 field  $h_{\mu\nu}$  with coupling  $-(\kappa/2)h_{\mu\nu}T^{\mu\nu}$ . This matching gives the standard relation

$$\kappa^2 = 32\pi G.$$

With the further choice  $M_{\text{Pl}} = \sqrt{32\pi}/\kappa$ , one finds

$$M_{\text{Pl}}^2 = \frac{32\pi}{\kappa^2} = \frac{1}{G},$$

so that  $M_{\text{Pl}}$  is the usual Planck mass determined by  $G$ . Hence the same Newton constant governs both the tree-level soft graviton exchange amplitude and the modular equation of state, and the linearized area term  $\delta[\text{Area}/(4G)]$  together with the stress-tensor coupling  $2\pi \int \xi^\nu \delta\langle T_{\mu\nu} \rangle$  precisely reproduces the low-energy limit of graviton exchange when  $\kappa^2 = 32\pi G$ . This establishes the claimed  $\kappa$ -consistency.  $\square$

*Proposition 5.56* (Amplitude  $\rightarrow$  curvature dictionary line (forward coefficient  $\rightarrow R_{kk}$ )). Let the analytic forward coefficient be defined by the projector of Section 5.22:

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

which is invariant under IR scheme, pivot, and scale by Lemmas 5.39, 5.72 and 5.95. With the gravitational normalization fixed by  $\kappa^2 = 32\pi G$  (Lemma 5.55), 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  $[\lambda_1, \lambda_2]$ ,

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

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

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

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

*Proof.* Fix a belt-anchored null segment  $[\lambda_1, \lambda_2]$  of a generator  $k^\mu$  and a value of  $t$  in the analytic strip of Section 5.22. By the third-subtracted dispersion representation constructed in Section 5.70, the second  $s$ -derivative of the forward amplitude at  $s = 0$  can be written as

$$\partial_s^2 \Re A^{(3)}(0, t) = \int_0^\infty d\mu K_{s_0}(\mu, t) \rho^{(3)}(\mu, t),$$

where  $K_{s_0}(\mu, t)$  is a strictly positive kernel and  $\rho^{(3)}(\mu, t)$  is the helicity-averaged absorptive part. By the optical theorem and helicity averaging, the profile  $\rho^{(3)}(\mu, t)$  is nonnegative for all  $\mu$  and for all admissible states, as established in Section 5.80. Consequently,

$$bc_{2,0}(t) = \frac{s_0^3}{2} \partial_s^2 \Re A^{(3)}(0, t)$$

is a positive linear functional of the absorptive profile. Projecting onto the even spin-2 component with the analytic projector  $\Pi_2$  of Section 5.22 simply folds the positive kernel  $K_{s_0}$  against a fixed finite matrix of weights. This produces an overall positive normalization factor  $C_{\Pi}(s_0) > 0$ , while the invariance of  $bc_{2,0}(t)$  under IR scheme, pivot choice, and rescaling is guaranteed by Lemmas 5.39, 5.72 and 5.95. Thus  $bc_{2,0}(t)$  is a well-defined positive observable extracted from the low-energy amplitude.

Next we relate this amplitude coefficient to the modular dynamics in the belt. With the gravitational normalization fixed by  $\kappa^2 = 32\pi G$  (Lemma 5.55), Sections 5.11 and 5.56 express the linear response of the modular Hamiltonian under perturbations supported in the belt as

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

and show that, for such perturbations, the soft-graviton exchange contribution to the forward amplitude is governed by the same canonical-energy functional that appears in this modular equation of state. In particular, the even spin-2 forward coefficient  $bc_{2,0}(t)$  is proportional to the canonical-energy flux through the belt modular horizon for every admissible state.

The Brown–York/JKM dictionary developed in Section 5.50 identifies this canonical-energy flux with a quasi-local stress-tensor flux through the belt, up to corrections of order  $B_{\text{belt}}$ . That is, there exists a nonnegative weight function  $w(\lambda)$ , supported on  $[\lambda_1, \lambda_2]$  and normalized by  $\int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) = 1$ , such that

$$\delta\langle K_{\text{mod}} \rangle = \int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) \langle T_{kk}(\lambda) \rangle + O(B_{\text{belt}}),$$

for all admissible perturbations, where  $T_{kk} := T_{\mu\nu} k^\mu k^\nu$ . The nonnegativity of  $w(\lambda)$  follows from the positivity of modular energy.

Because the same Einstein–Hilbert action controls both the modular dynamics and the scattering amplitude, and with the normalization fixed by Lemma 5.55, the Einstein equations hold in expectation on the domain of dependence  $D[R]$  relevant for the belt, giving

$$\langle R_{kk}(\lambda) \rangle = 8\pi G \langle T_{kk}(\lambda) \rangle$$

up to corrections of order  $B_{\text{belt}}$ . Substituting this into the expression for  $\delta\langle K_{\text{mod}} \rangle$  yields

$$\delta\langle K_{\text{mod}} \rangle = \frac{1}{8\pi G} \int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle + O(B_{\text{belt}}).$$

Combining this modular representation with the amplitude-side representation from Sections 5.11 and 5.56 and isolating the even spin-2 forward coefficient, we find that there exists a positive constant  $C_{\text{II}}(s_0) > 0$ , depending only on the analytic projector kernel and the choice of subtraction scale  $s_0$ , such that for every admissible state

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

We now define the belt-local functional on null-Ricci profiles by

$$\mathfrak{D}_{\text{belt}}[R_{kk}] := \frac{1}{8\pi G} \int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle.$$

By construction,  $\mathfrak{D}_{\text{belt}}$  is linear, positive (since  $w(\lambda) \geq 0$  and the absorptive profile is nonnegative), and state-independent; it depends only on the fixed choice of projector kernel and the Brown–York/JKM calibration. The preceding identity then gives

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

which is the first displayed relation in the proposition. The equivalent integral form in terms of  $w(\lambda)$  is just the definition of  $\mathfrak{D}_{\text{belt}}$  rewritten, and the positivity  $C_{\text{II}}(s_0) > 0$  follows from the positivity of the dispersive kernel and the projector, together with the existence of states for which both sides are strictly positive. Finally, the robustness under deformations of  $t$  within the analytic strip follows from the strip analyticity control in Section 5.22 and Corollary 5.125. This completes the proof.  $\square$

### 5.35 Master soundness and epsilon closure

*Theorem 5.57* (Soundness budget). Fix a target accuracy  $\varepsilon \in (0, 1)$ . With the global split of Section 5.27 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  $\varepsilon$  assembled from  $(\varepsilon_{\text{AGSP}}, \varepsilon_{\text{belt}}, \varepsilon_{\text{disp}}, \varepsilon_{\text{cel}}, \varepsilon_{\text{flow}})$  as specified.

*Proof.* Fix  $\varepsilon \in (0, 1)$  and choose the AGSP scale, belt width, dispersion cutoffs, celestial regulator, and flow parameters according to the global split in Section 5.27, thereby fixing the five tolerances  $(\varepsilon_{\text{AGSP}}, \varepsilon_{\text{belt}}, \varepsilon_{\text{disp}}, \varepsilon_{\text{cel}}, \varepsilon_{\text{flow}})$  as functions of  $\varepsilon$ .

In the modular pillar, the analysis of the equation of state shows that the deviation from the exact first-law identity is bounded by a linear combination of the common truncation errors,

$$\left| \delta S - \delta \langle K_{\text{mod}} \rangle + \delta \frac{\text{Area}}{4G} \right| \leq c_{\text{mod,AGSP}} \varepsilon_{\text{AGSP}} + c_{\text{mod,belt}} \varepsilon_{\text{belt}},$$

for some fixed coefficients  $c_{\text{mod,AGSP}}, c_{\text{mod,belt}}$  arising from the AGSP and belt constructions. By design of the global split in Section 5.27, these parameters are chosen so that the right-hand side does not exceed  $\varepsilon$ , giving the first inequality.

Similarly, the QNEC pillar combines flux control along the generator with the modular second-derivative estimate to produce a bound of the form

$$\left| \int du \langle T_{kk} \rangle \right| + |\partial_u^2 S - 2\pi \langle T_{kk} \rangle| \leq c_{\text{q,AGSP}} \varepsilon_{\text{AGSP}} + c_{\text{q,belt}} \varepsilon_{\text{belt}} + c_{\text{q,flow}} \varepsilon_{\text{flow}},$$

for some fixed coefficients  $c_{\text{q,AGSP}}, c_{\text{q,belt}}, c_{\text{q,flow}}$ . The global choice of  $(\varepsilon_{\bullet})$  ensures that this combination is also bounded by  $\varepsilon$ , giving the second inequality.

In the S-matrix pillar, the dispersion analysis controlling the hard amplitude yields

$$\text{Amp}_{\text{pos}}[\mathcal{A}_{\text{hard}}] \geq -\varepsilon_{\text{disp}},$$

up to the common AGSP/belt truncation errors that have already been absorbed into the global budget. With  $\varepsilon_{\text{disp}} \leq \varepsilon$  by construction, this implies the third line.

Finally, in the geometric pillar, the derivation of the linearized Einstein equation on  $D[R]$  from modular flow and the Brown–York/JKM dictionary shows that

$$|\text{Ein}[\delta g] - 8\pi G \delta \langle T \rangle|_{D[R]} \leq c_{\text{e,AGSP}} \varepsilon_{\text{AGSP}} + c_{\text{e,belt}} \varepsilon_{\text{belt}} + c_{\text{e,cel}} \varepsilon_{\text{cel}} + c_{\text{e,flow}} \varepsilon_{\text{flow}},$$

for suitable fixed coefficients  $c_{\text{e,AGSP}}, c_{\text{e,belt}}, c_{\text{e,cel}}, c_{\text{e,flow}}$ . Again, the global split of Section 5.27 is arranged so that this linear combination is at most  $\varepsilon$ , giving the last inequality.

Thus, with the parameter choices tied to  $\varepsilon$  through the global budget, all four pillars are simultaneously  $\varepsilon$ -sound, as claimed.  $\square$

### 5.36 Generalized second law on belts

*Theorem 5.58* (Belt GSL). Assume the framework recap Section 2 (items 1–3,5) and the  $c$ -function monotonicity of Section 5.17. Along any belt-anchored null generator with affine parameter  $\lambda$ ,

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

with  $C_{\text{GSL}}$  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  $\lambda$  (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  $\Gamma_{\text{belt}}$ .

**Step 1: Balance law along the generator.** Let  $\mathfrak{c}(\lambda)$  be the belt  $c$ –function of Section 5.17. 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 5.25,

$$\frac{d}{d\lambda} \left( \frac{A(\lambda)}{4G} - S(\lambda) \right) = \mathfrak{c}(\lambda) + \mathfrak{e}_{\text{edge}}(\lambda) + \mathfrak{e}_{\text{corner}}(\lambda) + O(\mathcal{B}_{\text{belt}}). \quad (5.66)$$

Here  $\mathfrak{e}_{\text{edge}}$  and  $\mathfrak{e}_{\text{corner}}$  are the edge/corner contributions recorded in Section 5.16, generated by the finite belt width and the boosted corner. By Proposition 5.40 (edge/corner control),

$$|\mathfrak{e}_{\text{edge}}(\lambda) + \mathfrak{e}_{\text{corner}}(\lambda)| \leq C_{\text{edge}} \mathcal{B}_{\text{belt}} \quad \text{uniformly in } \lambda. \quad (5.67)$$

**Step 2: Monotone evolution of the  $c$ –function.** The  $c$ –function obeys the differential inequality of Section 5.17,

$$\frac{d}{d\lambda} \mathfrak{c}(\lambda) \leq -\lambda_{\text{Clu}} \mathfrak{c}(\lambda) + C_0 \mathcal{B}_{\text{belt}}, \quad (5.68)$$

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

$$\lim_{\lambda \rightarrow -\infty} \mathfrak{c}(\lambda) = 0. \quad (5.69)$$

**Step 3: Grönwall bound for  $\mathfrak{c}$ .** Apply Grönwall’s lemma to (5.68) with the initial condition (5.69). For any  $\lambda \in \mathbb{R}$ ,

$$\mathfrak{c}(\lambda) \leq \int_{-\infty}^{\lambda} e^{-\lambda_{\text{Clu}}(\lambda-s)} C_0 \mathcal{B}_{\text{belt}} ds = \frac{C_0}{\lambda_{\text{Clu}}} \mathcal{B}_{\text{belt}} =: C_c \mathcal{B}_{\text{belt}}. \quad (5.70)$$

**Step 4: Differential GSL (with belt remainder).** Rearrange (5.66) and use (5.67) and (5.70):

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

for some belt–uniform constant  $C_1$  coming from the  $O(\mathcal{B}_{\text{belt}})$  term in (5.66). Setting

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

gives the stated estimate

$$\frac{d}{d\lambda} \left( S(\rho_R) - \frac{A(\text{QES}(R))}{4G} \right) \geq -C_{\text{GSL}} \mathcal{B}_{\text{belt}}, \quad (5.71)$$

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 (5.71) 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.  $\square$

### 5.37 Quantified QFC on belts

**Definition 5.59** (belt quantum expansion). Along a belt-anchored null generator with affine parameter  $\lambda$ , 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 5.60* (quantified belt QFC). For any segment  $[\lambda_1, \lambda_2]$ ,

$$\Theta_{\text{belt}}(\lambda_2) - \Theta_{\text{belt}}(\lambda_1) \geq 2\pi E_{\text{can}}^{\text{dens}}[\lambda_1 \rightarrow \lambda_2; \xi] + Q_{\text{shear}}[\lambda_1 \rightarrow \lambda_2] - C_{\text{QFC}} B_{\text{belt}}, \quad (5.72)$$

where  $E_{\text{can}}^{\text{dens}}$  is the canonical-energy density on the wedge generated by the belt boost  $\xi$  (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 5.46 of Section 5.24. 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:  $\partial_\lambda \Theta_{\text{belt}}(\lambda) \geq 0$ .

*Proof.* Fix an admissible belt-anchored null generator with affine parameter  $\lambda$ , and let

$$S_{\text{gen}}(\lambda) := S(\rho_R(\lambda)) - \frac{\text{Area}(\text{QES}(\lambda))}{4G}$$

denote the belt generalized entropy (per generator length) along the corresponding family of cuts  $R(\lambda)$ . By Definition 5.59, the belt quantum expansion is

$$\Theta_{\text{belt}}(\lambda) = \partial_\lambda S_{\text{gen}}(\lambda).$$

*Step 1: Second variation and the modular equation of state.* Apply the second-order modular equation of state Theorem 5.46 to the one-parameter family  $\lambda \mapsto (\rho_R(\lambda), R(\lambda))$  on the segment  $[\lambda_1, \lambda_2]$ . This yields

$$\delta^2[S_{\text{gen}}] \geq 2\pi E_{\text{can}}[\lambda_1 \rightarrow \lambda_2; \xi] + Q_{\text{shear}}[\lambda_1 \rightarrow \lambda_2] - C_2 B_{\text{belt}},$$

where  $E_{\text{can}}$  is the Iyer–Wald canonical energy on the wedge  $W = E_W(R)$  and  $Q_{\text{shear}}$  has precisely the form displayed in the statement with belt-local coefficients  $\kappa_\sigma, \kappa_\theta > 0$  supplied by Theorem 5.46. All quantities are per generator length, and  $C_2 > 0$  is independent of  $|R|$ .

By construction of the second variation along this deformation,  $\delta^2[S_{\text{gen}}]$  is the integral of the second  $\lambda$ -derivative of  $S_{\text{gen}}$  over the segment:

$$\delta^2[S_{\text{gen}}] = \int_{\lambda_1}^{\lambda_2} d\lambda \partial_\lambda^2 S_{\text{gen}}(\lambda).$$

Using  $\Theta_{\text{belt}}(\lambda) = \partial_\lambda S_{\text{gen}}(\lambda)$  and the fundamental theorem of calculus,

$$\int_{\lambda_1}^{\lambda_2} d\lambda \partial_\lambda^2 S_{\text{gen}}(\lambda) = \partial_\lambda S_{\text{gen}}(\lambda_2) - \partial_\lambda S_{\text{gen}}(\lambda_1) = \Theta_{\text{belt}}(\lambda_2) - \Theta_{\text{belt}}(\lambda_1).$$

Combining with the previous display gives

$$\Theta_{\text{belt}}(\lambda_2) - \Theta_{\text{belt}}(\lambda_1) \geq 2\pi E_{\text{can}}[\lambda_1 \rightarrow \lambda_2; \xi] + Q_{\text{shear}}[\lambda_1 \rightarrow \lambda_2] - C_2 B_{\text{belt}}.$$

By definition of the canonical-energy density per generator length, we may write

$$E_{\text{can}}[\lambda_1 \rightarrow \lambda_2; \xi] = \int_{\lambda_1}^{\lambda_2} d\lambda E_{\text{can}}^{\text{dens}}(\lambda; \xi),$$

so the inequality can be rewritten exactly in the form (5.72) with  $C_{\text{QFC}} := C_2$ .

*Step 2: Positivity of the shear term.* The functional  $Q_{\text{shear}}$  is a quadratic form in the shear  $\sigma_{ab}$  and expansion  $\theta$ ,

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

with strictly positive coefficients  $\kappa_\sigma, \kappa_\theta > 0$  that are uniform on the belt. These coefficients arise from the Raychaudhuri decomposition and shear/expansion estimates used in the proof of Theorem 5.46, concretely from the linearized Raychaudhuri window and shear control in Proposition 5.102 and Lemma 5.112. In particular  $Q_{\text{shear}}[\lambda_1 \rightarrow \lambda_2] \geq 0$ , with equality iff the perturbation carries neither shear nor expansion along the belt.

*Step 3: Positivity of canonical energy and monotonicity.* By the canonical-energy/QNEC inputs summarized in Theorem 5.29 and the belt JLMS/first-law infrastructure, the canonical energy is nonnegative on non-Killing perturbations, with kernel given by boost-Killing data. At the level of densities this implies

$$E_{\text{can}}^{\text{dens}}[\lambda_1 \rightarrow \lambda_2; \xi] \geq 0 \quad \text{for non-Killing data.}$$

Hence for such perturbations the inequality (5.72) simplifies to

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

Finally, all constants in Theorem 5.46 and in Proposition 5.102, Lemma 5.112, and Theorem 5.29 are independent of the positive-flow regulators, while the belt budget satisfies  $B_{\text{belt}} \rightarrow 0$  as the auxiliary null and bulk flows  $(u, s) \downarrow 0$ ; see Section 5.13. Passing to this limit in (5.72) yields

$$\Theta_{\text{belt}}(\lambda_2) \geq \Theta_{\text{belt}}(\lambda_1) \quad \text{for all } \lambda_2 \geq \lambda_1,$$

so in the regulator-removed window the belt quantum expansion is nondecreasing:  $\partial_\lambda \Theta_{\text{belt}}(\lambda) \geq 0$ . This completes the proof.  $\square$

*Remark 5.61* (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  $\Theta_{\text{belt}}$  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 5.60, with equality iff the perturbation is boost-Killing. Together with the belt generalized second law and the quantum Bousso bound on belts (Theorems 5.58 and 5.97), 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 5.29 and 5.37) and remain within the declared envelope.

### 5.38 Third-order control of the modular equation of state

*Proposition 5.62* (cubic completion). Under the hypotheses of Theorem 5.46 and the canonical-energy/QSEI inputs Proposition 5.102 and Lemma 5.112, the third variation satisfies

$$\delta^3 \left[ S - \frac{\text{Area}}{4G} \right] = 2\pi \mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi] + \mathcal{C}_{\text{shear} \times \text{exp}}[\delta g] + R_{\text{belt}}^{(3)}, \quad (5.73)$$

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 5.112, 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 5.70, then the cubic inequality

$$\delta^3 \left[ S - \frac{\text{Area}}{4G} \right] \geq -\tilde{C}'_{\sigma} E_{\text{can}}^W[\delta \Psi; \xi] - C_{\text{belt}}^{(3)} \mathcal{B}_{\text{belt}} \quad (\text{non-Killing data}) \quad (5.74)$$

holds belt-uniformly.

*Proof.* We work on the belt domain  $D[R]$  with the calibration and OS kernel in force, and consider an admissible one-parameter family  $(g(\lambda), \Psi(\lambda))$  with  $\lambda \in \mathbb{R}$ ,  $g(0) = g$ ,  $\Psi(0) = \Psi$ , and with QES surfaces chosen along the family so that the first-order extremality condition holds at each  $\lambda$  in the OS window. We denote

$$\delta := \frac{d}{d\lambda} \Big|_{\lambda=0}, \quad \delta^2 := \frac{d^2}{d\lambda^2} \Big|_{\lambda=0}, \quad \delta^3 := \frac{d^3}{d\lambda^3} \Big|_{\lambda=0}.$$

*Step 1: Second-order structure as starting point.* By Theorem 5.46, for any such admissible family one has, on the belt,

$$\delta^2 \left[ S - \frac{\text{Area}}{4G} \right] \geq 2\pi E_{\text{can}}^W[\delta \Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 \mathcal{B}_{\text{belt}}, \quad (5.75)$$

where  $E_{\text{can}}^W[\delta \Psi; \xi]$  is the Iyer–Wald canonical energy on  $W = \text{EW}(R)$  and  $Q_{\text{shear}}[\delta g]$  is given by (5.36) with positive coefficients bounded below as in (5.37). In the proof of Theorem 5.46, (5.75) arises from an *identity* of the form

$$\delta^2 \left[ S - \frac{\text{Area}}{4G} \right] = 2\pi E_{\text{can}}^W[\delta \Psi; \xi] + Q_{\text{shear}}[\delta g] + R_{\text{belt}}^{(2)}, \quad (5.76)$$

with a remainder  $R_{\text{belt}}^{(2)}$  satisfying

$$|R_{\text{belt}}^{(2)}| \leq C_2 \mathcal{B}_{\text{belt}}, \quad (5.77)$$

where  $C_2$  is independent of  $|R|$  and all statements are per generator length. The inequality (5.75) is then obtained by dropping the nonnegative parts of  $Q_{\text{shear}}$  and using (5.77).

Thus we may take (5.76) as the structural starting point for the third variation.

*Step 2: Differentiating the canonical energy term.* The canonical energy  $E_{\text{can}}^W[\delta \Psi; \xi]$  is defined in the covariant phase-space formalism as (schematically)

$$E_{\text{can}}^W[\delta \Psi; \xi] = \int_{\Sigma \subset W} \omega(\Phi; \delta \Phi, \mathcal{L}_{\xi} \delta \Phi),$$

where  $\Phi = (g, \Psi)$ ,  $\delta \Phi = (\delta g, \delta \Psi)$ , and  $\omega$  is the presymplectic current, with calibrated corner and belt terms included as in Sections 5.16 and 5.49, Proposition 5.40, and Lemma 5.77. Differentiating with respect to  $\lambda$  at  $\lambda = 0$  yields a cubic functional

$$\mathcal{E}_{\text{can}}^{(3)}[\delta \Psi; \xi] := \delta \left( E_{\text{can}}^W[\delta \Psi; \xi] \right),$$

which can be written schematically as

$$\mathcal{E}_{\text{can}}^{(3)}[\delta \Psi; \xi] = \int_{\Sigma \subset W} \omega^{(3)}(\Phi; \delta \Phi, \delta \Phi, \mathcal{L}_{\xi} \Phi) + \text{calibrated edge/belt contributions},$$

where  $\omega^{(3)}$  is the trilinear extension of the symplectic current and all boundary pieces are treated with the JKM/Brown–York scheme.

By the hypotheses of Theorem 5.46 and the canonical-energy/QSEI inputs Proposition 5.102 and Lemma 5.112, this cubic functional is well-defined on the domain of Theorem 5.70, and the belt calibration ensures that the associated edge/corner terms are absorbed into the belt remainder. Thus

$$\delta\left(2\pi E_{\text{can}}^W[\delta\Psi; \xi]\right) = 2\pi \mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi] + R_{\text{can}}^{(3)}, \quad (5.78)$$

with  $R_{\text{can}}^{(3)}$  bounded by a multiple of  $\mathcal{B}_{\text{belt}}$  (coming from the belt localization and OS flow removal).

*Step 3: Differentiating the shear term and controlling  $\mathcal{C}_{\text{shear}\times\text{exp}}$ .* The shear/expansion functional  $Q_{\text{shear}}[\delta g]$  has the form

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

with coefficients satisfying (5.37). Differentiating with respect to the deformation parameter  $\lambda$  produces cubic terms involving products of shear and expansion, schematically of the form

$$\delta Q_{\text{shear}}[\delta g] = \mathcal{C}_{\text{shear}\times\text{exp}}[\delta g] + R_{\text{shear}}^{(3)},$$

where  $\mathcal{C}_{\text{shear}\times\text{exp}}[\delta g]$  encodes the shear–expansion couplings (and possibly higher-order geometric corrections) and  $R_{\text{shear}}^{(3)}$  collects belt-localization and calibration remainders.

The detailed Raychaudhuri-based analysis and QSEI bounds in Proposition 5.102 and Lemma 5.112 imply a control of the form

$$|\mathcal{C}_{\text{shear}\times\text{exp}}[\delta g]| \leq \tilde{C}'_\sigma E_{\text{can}}^W[\delta\Psi; \xi] + C_{\text{rem}}^{(3)} \mathcal{B}_{\text{belt}},$$

for some  $\tilde{C}'_\sigma, C_{\text{rem}}^{(3)} > 0$ , uniform per generator length and independent of  $|R|$ . Here the canonical-energy term appears because the Raychaudhuri equation links the shear/expansion sector back to  $T_{\mu\nu}\xi^\mu\xi^\nu$  and thus to the canonical energy via the JLMS/Brown–York flux dictionary; the  $\mathcal{B}_{\text{belt}}$  component comes from the finite-width belt regularization and corner calibration, handled as in Sections 5.16 and 5.49, Proposition 5.40, and Lemma 5.77.

Moreover,  $R_{\text{shear}}^{(3)}$  is also bounded by a multiple of  $\mathcal{B}_{\text{belt}}$ ; we fold it into the global third-order belt remainder below.

*Step 4: Differentiating the belt remainder.* From (5.76) and (5.77) we have

$$R_{\text{belt}}^{(2)} = O(\mathcal{B}_{\text{belt}}).$$

The OS kernel and flow-removal control (Lemmas 3.1 and 3.3) imply that the variation of  $R_{\text{belt}}^{(2)}$  along the admissible one-parameter family remains bounded by the same belt budget up to a fixed constant:

$$\delta R_{\text{belt}}^{(2)} = R_{\text{belt}}^{(3)}, \quad |R_{\text{belt}}^{(3)}| \leq C_{\text{belt}}^{(3)} \mathcal{B}_{\text{belt}},$$

with  $C_{\text{belt}}^{(3)}$  independent of  $|R|$  and per generator length. Intuitively,  $R_{\text{belt}}^{(2)}$  is a smooth functional of the belt-local data whose size is controlled by  $\mathcal{B}_{\text{belt}}$ ; differentiating preserves this control.

*Step 5: Assemble the third variation.* Differentiating (5.76) with respect to  $\lambda$  and evaluating at  $\lambda = 0$  gives

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

where we used (5.78) and the decomposition in Step 3. Absorbing  $R_{\text{can}}^{(3)}$  and  $R_{\text{shear}}^{(3)}$  into  $R_{\text{belt}}^{(3)}$  (which simply renormalizes the constant  $C_{\text{belt}}^{(3)}$ ) we obtain the structural identity (5.73), with

$$|R_{\text{belt}}^{(3)}| \leq C_{\text{belt}}^{(3)} \mathcal{B}_{\text{belt}}$$

for some  $C_{\text{belt}}^{(3)} > 0$ , uniform in  $|R|$ .

Using the bound on  $\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]$  from Step 3 we arrive at

$$\delta^3 \left[ S - \frac{\text{Area}}{4G} \right] \geq 2\pi \mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi] - \tilde{C}'_{\sigma} E_{\text{can}}^W[\delta\Psi; \xi] - C_{\text{belt}}^{(3)} \mathcal{B}_{\text{belt}}.$$

*Step 6: Cubic inequality for non-Killing data.* Assume now that  $\mathcal{E}_{\text{can}}^{(3)} \geq 0$  on non-Killing perturbations over the domain of Theorem 5.70. For such perturbations we may drop the nonnegative term  $2\pi \mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi]$  and obtain

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

which is precisely (5.74), with constants uniform per generator length and independent of  $|R|$ .

This completes the proof of the cubic completion and the stated inequality.  $\square$

### 5.39 Finite-size corrections and thermodynamic limit

*Lemma 5.63* (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  $\text{length}(\partial R) \rightarrow \infty$ .

*Proof.* By the AGSP/flow converter pipeline (Section 3 and the belt entropy converter in Sections 5 and 5.21), the entropy of any admissible belt-anchored region  $R$  admits a decomposition of the form

$$S(\rho_R) = \text{length}(\partial R) \left( \frac{\log \kappa_{\text{seed}}}{1 - \delta^2} + \frac{\log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{(1 - \delta^2)^2} \right) + C_0 + R_{\text{belt}}(R), \quad (5.79)$$

where:

- $\kappa_{\text{seed}}, \Lambda_0, \Gamma_{\text{belt}}, \Upsilon(m)$  and  $\delta$  are the fixed converter parameters from the OS window;
- $C_0$  is a constant independent of  $R$  (coming from the converter normalization and the finite-size offset in the AGSP construction);
- $R_{\text{belt}}(R)$  is a remainder term controlled by the belt budget:  $|R_{\text{belt}}(R)| \leq C_{\text{rem}} \mathcal{B}_{\text{belt}}$  for some belt-uniform  $C_{\text{rem}} > 0$ , by the remainder ledger in Section 5.13.

Divide (5.79) by  $\text{length}(\partial R)$ :

$$\frac{S(\rho_R)}{\text{length}(\partial R)} = \frac{\log \kappa_{\text{seed}}}{1 - \delta^2} + \frac{\log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{(1 - \delta^2)^2} + \frac{C_0}{\text{length}(\partial R)} + \frac{R_{\text{belt}}(R)}{\text{length}(\partial R)}.$$

Taking absolute values of the deviation from the benchmark and using  $|R_{\text{belt}}(R)| \leq C_{\text{rem}} \mathcal{B}_{\text{belt}}$  yields

$$\left| \frac{S(\rho_R)}{\text{length}(\partial R)} - \left( \frac{\log \kappa_{\text{seed}}}{1 - \delta^2} + \frac{\log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{(1 - \delta^2)^2} \right) \right| \leq \frac{|C_0|}{\text{length}(\partial R)} + \frac{C_{\text{rem}}}{\text{length}(\partial R)} \mathcal{B}_{\text{belt}}.$$

Absorbing  $|C_0|$  and  $C_{\text{rem}}$  into a single constant  $C$  and using that  $\mathcal{B}_{\text{belt}}$  does not scale with  $\text{length}(\partial R)$  (it depends only on the belt regulators and curvature as in Section 5.13) gives

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

which is the claimed finite-size bound. In particular, for fixed belt regulators and curvature, the  $O(\mathcal{B}_{\text{belt}})$  term is bounded while  $\text{length}(\partial R) \rightarrow \infty$ , so the per-length entropy converges to the stated benchmark.  $\square$

#### 5.40 Integrated QNEC implies ANEC and partial converses

*Theorem 5.64* (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  $f$ ,

$$\int_{\mathbb{R}} du f(u) \left( 2\pi \langle T_{kk}(u) \rangle - \partial_u^2 S(u) \right) \geq 0. \quad (5.80)$$

Define the distribution

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

Then  $g$  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 equal finite limits as  $u \rightarrow \pm\infty$  (for example  $\partial_u S(u) \rightarrow 0$ ), then ANEC holds:

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

Conversely, ANEC alone does *not* imply (5.80). Nevertheless, (5.80) *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 dominance:**  $\partial_u^2 S(u) \geq 0$  almost everywhere and there exists a constant  $c \geq \frac{1}{2\pi}$  such that

$$\int_{\mathbb{R}} du \psi(u) \langle T_{kk}(u) \rangle \geq c \int_{\mathbb{R}} du \psi(u) \partial_u^2 S(u) \quad (5.81)$$

for all  $\psi \in C_0^\infty(\mathbb{R})$  with  $\psi \geq 0$ .

*Proof.* We write everything in terms of the distribution

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

*Step 1: positivity of  $g$  as a distribution.* Let  $\psi \in C_0^\infty(\mathbb{R})$  with  $\psi \geq 0$  and  $\psi \not\equiv 0$ . Define

$$f(u) := \frac{\psi(u)}{\int_{\mathbb{R}} dv \psi(v)}.$$

Then  $f \in C_0^\infty(\mathbb{R})$ ,  $f \geq 0$  and  $\int f = 1$ , so (5.80) applies:

$$0 \leq \int_{\mathbb{R}} du f(u) g(u) = \frac{1}{\int \psi} \int_{\mathbb{R}} du \psi(u) g(u).$$

Since the denominator  $\int \psi > 0$ , we obtain

$$\int_{\mathbb{R}} du \psi(u) g(u) \geq 0$$

for every  $\psi \in C_0^\infty(\mathbb{R})$  with  $\psi \geq 0$ . By definition this means that  $g$  is a positive distribution.

*Step 2: from positivity of  $g$  to almost-everywhere nonnegativity.* Assume in addition that  $g$  is locally integrable, so that  $g$  can be represented by a function  $g \in L_{\text{loc}}^1(\mathbb{R})$ . We claim that then  $g(u) \geq 0$  for almost all  $u$ .

Suppose not. Then the set

$$A := \{u \in \mathbb{R} \mid g(u) < 0\}$$

has positive measure. By regularity of Lebesgue measure there exists a compact set  $K \subset A$  with positive measure. Choose  $\psi \in C_0^\infty(\mathbb{R})$ ,  $\psi \geq 0$ , supported in  $K$  and not identically zero. On  $K$  we have  $g < 0$ , hence

$$\int_{\mathbb{R}} du \psi(u) g(u) = \int_K du \psi(u) g(u) < 0,$$

contradicting positivity of the distribution  $g$ . Thus  $g \geq 0$  almost everywhere.

*Step 3: ANEC from  $g \geq 0$  and boundary behaviour of  $S$ .* Assume now that  $g$  is locally integrable and that  $\partial_u S(u)$  has equal finite limits as  $u \rightarrow \pm\infty$ , say

$$\lim_{u \rightarrow \pm\infty} \partial_u S(u) = L_{\pm}, \quad L_+ = L_- =: L \in \mathbb{R}.$$

Pick  $R > 0$  and integrate  $g$  over  $[-R, R]$ :

$$\int_{-R}^R du g(u) = 2\pi \int_{-R}^R du \langle T_{kk}(u) \rangle - \int_{-R}^R du \partial_u^2 S(u).$$

The last term can be integrated by parts:

$$\int_{-R}^R du \partial_u^2 S(u) = \partial_u S(R) - \partial_u S(-R).$$

Hence

$$2\pi \int_{-R}^R du \langle T_{kk}(u) \rangle = \int_{-R}^R du g(u) + \partial_u S(R) - \partial_u S(-R). \quad (5.82)$$

By assumption  $g \geq 0$  almost everywhere, so  $\int_{-R}^R g \geq 0$  for every  $R$ . Moreover

$$\lim_{R \rightarrow \infty} (\partial_u S(R) - \partial_u S(-R)) = L - L = 0.$$

Taking  $R \rightarrow \infty$  in (5.82) we get

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

where the right-hand side is understood in the extended sense (it may be  $+\infty$ ). This is precisely ANEC.

*Step 4: ANEC does not imply the integrated QNEC.* We now show that ANEC, by itself, is strictly weaker than (5.80). We construct a smooth example with ANEC but  $\int f g < 0$  for a suitable  $f$ .

Choose  $h \in C_0^\infty(\mathbb{R})$  such that

$$\int_{\mathbb{R}} du h(u) = 0, \quad h(u) < 0 \text{ on some nonempty open interval.}$$

Next pick  $\langle T_{kk} \rangle \in C_0^\infty(\mathbb{R})$  with

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

Define  $\partial_u^2 S$  by

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

This determines  $S$  up to an affine function by double integration; no further condition on the behaviour of  $\partial_u S$  at infinity is imposed here.

With this choice,

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

By construction  $\int_{\mathbb{R}} du \langle T_{kk} \rangle > 0$ , so ANEC holds. However  $h < 0$  on some interval, hence we can choose  $\psi \in C_0^\infty(\mathbb{R})$  supported where  $h < 0$  and  $\psi \geq 0$ ,  $\psi \not\equiv 0$ . Then

$$\int_{\mathbb{R}} du \psi(u) g(u) = \int_{\mathbb{R}} du \psi(u) h(u) < 0.$$

Normalising  $f := \psi / \int \psi$ , we obtain a nonnegative  $f \in C_0^\infty$  with  $\int f = 1$  but

$$\int_{\mathbb{R}} du f(u) g(u) = \frac{1}{\int \psi} \int \psi h < 0,$$

so (5.80) fails. This shows that ANEC alone does not imply the integrated QNEC.

*Step 5: sufficient conditions for the integrated QNEC.*

(1) *Pointwise QNEC.* If

$$2\pi \langle T_{kk}(u) \rangle \geq \partial_u^2 S(u) \quad \text{for almost all } u,$$

then  $g(u) \geq 0$  almost everywhere. For any  $f \in C_0^\infty$  with  $f \geq 0$  and  $\int f = 1$ ,

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

because the integrand is pointwise nonnegative and  $f \geq 0$ . Thus (5.80) holds.

(2) *Convexity plus smeared energy dominance.* Assume  $\partial_u^2 S(u) \geq 0$  almost everywhere and that there exists  $c \geq \frac{1}{2\pi}$  such that (5.81) holds for all nonnegative  $\psi \in C_0^\infty(\mathbb{R})$ . Let  $f \in C_0^\infty(\mathbb{R})$  with  $f \geq 0$  and  $\int f = 1$ . Applying (5.81) with  $\psi = f$  gives

$$\int_{\mathbb{R}} du f(u) \langle T_{kk}(u) \rangle \geq c \int_{\mathbb{R}} du f(u) \partial_u^2 S(u).$$

Multiplying by  $2\pi$  and using  $c \geq 1/(2\pi)$ ,

$$2\pi \int_{\mathbb{R}} du f(u) \langle T_{kk}(u) \rangle \geq 2\pi c \int_{\mathbb{R}} du f(u) \partial_u^2 S(u) \geq \int_{\mathbb{R}} du f(u) \partial_u^2 S(u),$$

since  $2\pi c \geq 1$  and  $\partial_u^2 S \geq 0$  while  $f \geq 0$ . Therefore

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

which is precisely (5.80). This completes the proof.  $\square$

*Remark 5.65* (Toy counterexample to the converse). Choose  $h \in C_0^\infty(\mathbb{R})$  with  $\int h = 0$  and  $h < 0$  on some interval. Pick  $\langle T_{kk} \rangle \in C_0^\infty(\mathbb{R})$  with  $\langle T_{kk} \rangle \geq 0$  and  $\int \langle T_{kk} \rangle > 0$ , and define  $S$  by

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

This can always be integrated twice to obtain a smooth  $S$ , defined up to an affine function; no condition on  $\partial_u S$  at infinity is imposed. Then

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

so  $g$  is negative on the support of  $h$  even though

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

and ANEC holds. For any  $f \in C_0^\infty$  supported where  $h < 0$  with  $f \geq 0$  and  $\int f = 1$ , one has  $\int f g < 0$ , contradicting (5.80). This makes explicit that ANEC does not imply the integrated QNEC.

*Corollary 5.66* (From pointwise QNEC to ANEC). If  $2\pi \langle T_{kk}(u) \rangle \geq \partial_u^2 S(u)$  almost everywhere and  $\partial_u S(u)$  has equal finite limits as  $u \rightarrow \pm\infty$  (for example  $\partial_u S(u) \rightarrow 0$ ), then ANEC holds.

*Proof.* Pointwise QNEC implies that  $g(u) \geq 0$  almost everywhere, hence  $g$  is a nonnegative locally integrable function. Applying the first part of Theorem 5.64 with the assumption on the boundary behaviour of  $\partial_u S$  gives

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

which is ANEC. □

## 5.41 Domains and cores for higher-order variations

*Proposition 5.67* (analytic core and essential self-adjointness). Let  $\mathcal{D}_{\text{an}}$  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  $\mathcal{D}_{\text{an}}$  is a core for all polynomial functions of  $K_{\text{mod}}$  entering the second/third variations.

*Proof.* Let  $\mathcal{H}$  denote the belt Hilbert space,  $U(t)$  the belt modular flow, and  $K_{\text{mod}}$  its (Stone) generator, so that  $U(t) = e^{itK_{\text{mod}}}$  is a strongly continuous one-parameter unitary group on  $\mathcal{H}$ . By definition, a vector  $\psi \in \mathcal{H}$  is *analytic* for  $K_{\text{mod}}$  if there exists  $r > 0$  with

$$\sum_{n=0}^{\infty} \frac{r^n}{n!} \|K_{\text{mod}}^n \psi\| < \infty,$$

equivalently,  $z \mapsto \sum_{n \geq 0} \frac{z^n}{n!} K_{\text{mod}}^n \psi$  converges in  $\mathcal{H}$  for all  $z \in \mathbb{C}$ . Let  $\mathcal{D}_{\text{an}}$  be the set of all such vectors.

*Step 1 (density and a canonical analytic core).* For  $\varepsilon > 0$  define the bounded operator

$$S_\varepsilon := \frac{1}{\sqrt{4\pi\varepsilon}} \int_{\mathbb{R}} e^{-t^2/(4\varepsilon)} U(t) dt = e^{-\varepsilon K_{\text{mod}}^2}$$

(the second equality is the spectral calculus identity for the Fourier transform of a Gaussian). Then  $S_\varepsilon \rightarrow I$  strongly as  $\varepsilon \downarrow 0$  by dominated convergence in the spectral representation of  $K_{\text{mod}}$ . Moreover, for every  $n \in \mathbb{N}$ ,

$$K_{\text{mod}}^n S_\varepsilon = e^{-\varepsilon K_{\text{mod}}^2} K_{\text{mod}}^n \quad \text{and} \quad \|K_{\text{mod}}^n S_\varepsilon\| = \sup_{\lambda \in \mathbb{R}} |\lambda|^n e^{-\varepsilon \lambda^2} \leq \left(\frac{n}{2\varepsilon}\right)^{n/2}.$$

Hence for every  $\psi \in \mathcal{H}$  and all  $n \in \mathbb{N}$ ,

$$\|K_{\text{mod}}^n S_\varepsilon \psi\| \leq \left(\frac{n}{2e\varepsilon}\right)^{n/2} \|\psi\|.$$

Since  $\sum_{n \geq 0} \left(\frac{n}{2e\varepsilon}\right)^{n/2} \frac{r^n}{n!} < \infty$  for every  $r > 0$  (ratio test, using  $n! \sim (n/e)^n \sqrt{2\pi n}$ ), it follows that  $S_\varepsilon \psi$  is an analytic vector for each  $\varepsilon > 0$ , i.e.  $S_\varepsilon \psi \in \mathcal{D}_{\text{an}}$ . Because  $S_\varepsilon \rightarrow I$  strongly, the linear span of  $\text{Ran}(S_\varepsilon)$  (for  $\varepsilon > 0$ ) is dense in  $\mathcal{H}$ , hence  $\mathcal{D}_{\text{an}}$  is dense.

*Step 2 ( $\mathcal{D}_{\text{an}}$  is a core for all powers  $K_{\text{mod}}^m$ ).* Fix  $m \in \mathbb{N}$ . For any  $\psi \in \text{Dom}(K_{\text{mod}}^m)$  we have, by spectral calculus and the commutation  $e^{-\varepsilon K_{\text{mod}}^2} K_{\text{mod}}^m = K_{\text{mod}}^m e^{-\varepsilon K_{\text{mod}}^2}$ ,

$$\psi_\varepsilon := S_\varepsilon \psi \in \mathcal{D}_{\text{an}}, \quad \psi_\varepsilon \rightarrow \psi, \quad \text{and} \quad K_{\text{mod}}^m \psi_\varepsilon = S_\varepsilon K_{\text{mod}}^m \psi \rightarrow K_{\text{mod}}^m \psi \quad \text{as } \varepsilon \downarrow 0,$$

where the convergences are in  $\mathcal{H}$ . Thus  $\mathcal{D}_{\text{an}}$  is dense in the graph norm of  $K_{\text{mod}}^m$ ; equivalently,  $\mathcal{D}_{\text{an}}$  is a core for  $K_{\text{mod}}^m$ .

*Step 3 (polynomials).* Let  $p$  be any polynomial. Then  $p(K_{\text{mod}})$  is a finite linear combination of the powers  $K_{\text{mod}}^m$ . By Step 2 (applied to each occurring power) and the triangle inequality for graph norms,  $\mathcal{D}_{\text{an}}$  is a core for  $p(K_{\text{mod}})$ . In particular, this holds for all polynomial expressions in  $K_{\text{mod}}$  that appear in the second and third variations.

*Step 4 (essential self-adjointness of  $K_{\text{mod}} \upharpoonright_{\mathcal{D}_{\text{an}}}$ ).* Taking  $m = 1$  in Step 2 shows that  $\mathcal{D}_{\text{an}}$  is a core for  $K_{\text{mod}}$ . Therefore the closure of the symmetric operator  $K_{\text{mod}} \upharpoonright_{\mathcal{D}_{\text{an}}}$  is the self-adjoint operator  $K_{\text{mod}}$  itself. In other words,  $K_{\text{mod}} \upharpoonright_{\mathcal{D}_{\text{an}}}$  is essentially self-adjoint. (This also follows from Nelson's analytic vector theorem, but the core property established above already suffices.)

*Stability statements used later.*

- *Bounded belt circuits.* If  $B$  is bounded and symmetric on  $\mathcal{H}$ , then the graph norms of  $K_{\text{mod}}$  and  $K_{\text{mod}} + B$  are equivalent on  $\text{Dom}(K_{\text{mod}})$ :

$$\|(K_{\text{mod}} + B)\psi\| \leq \|K_{\text{mod}}\psi\| + \|B\| \|\psi\|, \quad \|K_{\text{mod}}\psi\| \leq \|(K_{\text{mod}} + B)\psi\| + \|B\| \|\psi\|.$$

Hence any core for  $K_{\text{mod}}$  is a core for  $K_{\text{mod}} + B$ , and the same holds for every polynomial in these operators. In particular  $\mathcal{D}_{\text{an}}$  remains a core.

- *Quasi-local factorization.* If  $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$  and  $K_{\text{mod}} = K_1 \otimes I + I \otimes K_2$  with  $K_j$  self-adjoint, then the algebraic tensor product  $\mathcal{D}_{\text{an}}(K_1) \otimes_{\text{alg}} \mathcal{D}_{\text{an}}(K_2)$  consists of analytic vectors for  $K_{\text{mod}}$  (binomial expansion of  $K_{\text{mod}}^n$  and the factorial bounds for analytic vectors) and is dense; hence it is a core for all powers and polynomials of  $K_{\text{mod}}$  by the same reasoning as above.
- *Regulator removal.* Let  $\{K_\varepsilon\}_{\varepsilon > 0}$  be regulated modular generators with  $K_\varepsilon \rightarrow K_{\text{mod}}$  in the strong resolvent sense as  $\varepsilon \downarrow 0$  and suppose  $\mathcal{D}_{\text{an}}$  is a common core for the relevant polynomials of  $K_\varepsilon$ . Then, by functional calculus and dominated convergence in the spectral representation,  $p(K_\varepsilon)\psi \rightarrow p(K_{\text{mod}})\psi$  for every  $\psi \in \mathcal{D}_{\text{an}}$ , and the graph norms converge on  $\mathcal{D}_{\text{an}}$ . Consequently  $\mathcal{D}_{\text{an}}$  remains a core for  $p(K_{\text{mod}})$  after removing the regulator.

Combining the four steps proves the proposition:  $K_{\text{mod}} \upharpoonright_{\mathcal{D}_{\text{an}}}$  is essentially self-adjoint, and  $\mathcal{D}_{\text{an}}$  is a core for all polynomial functions of  $K_{\text{mod}}$  entering the second and third variations.  $\square$

## 5.42 Celestial Ward identities without anomalies on the principal series

*Remark 5.68* (Principal series scope). Throughout this subsection we restrict celestial Mellin data to the principal series  $\Delta = 1 + i\nu$  with  $|\nu| \leq \mu_{\text{cel}}$ . On this locus the celestial Gram form is positive and the Mellin measure is unitary. Off the principal lines, a finite strip requires a renormalized Mellin measure together with a belt-local Ward counterterm; see Section 5.93.

*Lemma 5.69* (No anomaly on the principal lines). Let celestial insertions be restricted to the principal series  $\Delta = 1 + i\nu$ ,  $|\nu| \leq \mu_{\text{cel}}$ . Consider the Virasoro-type Ward action on these insertions induced by the belt first-law channel (i.e. by the calibrated modular/flux charge on the belt). Then:

1. The Ward action is anomaly-free at linear order: the induced infinitesimal diffeomorphisms preserve the principal-series celestial Gram form and the (principal) Mellin measure.
2. Regulator effects from the finite belt width and positive flows deform the measure at most by  $O(\mathcal{B}_{\text{belt}})$ ; in particular, any residual Jacobian or central-term-like contribution is  $O(\mathcal{B}_{\text{belt}})$ .

*Proof. Assumptions and objects.* (i) The belt first-law channel provides a linear, calibrated map that identifies the modular generator on the belt with a quasi-local stress/flux functional (Brown–York/Iyer–Wald dictionary with JKM corner fix); its action on operators implements the appropriate diffeomorphism at linear order. (ii) Celestial operators  $\mathcal{O}_{\Delta, J}(z, \bar{z})$  are defined by Mellin transforms of on-shell data; on the principal series, the Mellin weight is unitary and the celestial Gram matrix built from absorptive/positivity testers is positive semidefinite. (iii) All statements are made at finite belt regulators and then pass to the flow-removed window with an  $O(\mathcal{B}_{\text{belt}})$  ledger.

*Step 1 (Ward charge and its belt realization).* Let  $Q[\varepsilon]$  denote the infinitesimal diffeomorphism charge generated by a holomorphic vector field  $\varepsilon(z)$  (and similarly for  $\bar{\varepsilon}(\bar{z})$ ). On the belt, the first-law channel furnishes a self-adjoint generator

$$Q[\varepsilon] = \int_{\partial\Sigma} dl_a \varepsilon T_{\text{BY}}^{ab} \xi_b + (\text{calibrated corner}),$$

whose commutator implements the linearized diffeomorphism on local observables. The JKM calibration fixes corner improvements so that the boost Ward identity holds on the belt, and any regulator mismatch is  $O(\mathcal{B}_{\text{belt}})$ .

*Step 2 (Transport to the celestial basis and Gram symmetry).* Matrix elements of  $Q[\varepsilon]$  between scattering states can be pushed to the celestial basis by Mellin transform in energies. On the principal series  $\Delta = 1 + i\nu$  the Mellin kernel is unitary, so the action of  $Q[\varepsilon]$  transports to an anti-Hermitian derivation with respect to the celestial Gram form:

$$\langle \Psi, Q[\varepsilon] \Phi \rangle_{\text{cel}} = -\langle Q[\varepsilon] \Psi, \Phi \rangle_{\text{cel}} \quad (\text{principal series}).$$

Hence  $Q[\varepsilon]$  generates an isometry of the celestial inner product and preserves the principal-series measure. The positivity of the Gram form on principal lines (via the project’s celestial Gram testers) ensures that no hidden indefinite directions spoil this unitary implementation.

*Step 3 (No linear anomaly).* Because the charge acts by an anti-Hermitian derivation on the (unitary) principal-series celestial space, the linear Ward variation is a symmetry of the Gram form and thus produces no anomalous Jacobian at linear order. Equivalently, the linearized Ward identity in the celestial basis is saturated by the diffeomorphism variation of insertions alone; no extra (central-term-like) contact arises on principal lines. Any residual belt-regulator defect is  $O(\mathcal{B}_{\text{belt}})$  by the belt budget and vanishes under positive-flow removal.

*Step 4 (Budget and removal).* The only possible deviations originate from (a) finite belt width, (b) short positive flows in the OS kernel, and (c) corner calibration at the belt. Each is controlled by the nonnegative budget  $\mathcal{B}_{\text{belt}}$  and therefore any measure deformation or extra term in the Ward identity is  $O(\mathcal{B}_{\text{belt}})$ . Removal of flows sets  $\mathcal{B}_{\text{belt}} \rightarrow 0$ , completing the claim.  $\square$

**Add-on (finite off-principal strip).** Let  $\Delta = 1 + i\nu + \sigma$  with  $|\sigma| \leq \sigma_0$  (a finite vertical strip around the principal line). Then there exist: (i) a finite, belt-local renormalization of the Mellin measure on the strip, and (ii) a belt-local Ward counterterm  $\delta Q_{\text{strip}}[\varepsilon]$ , such that the

renormalized charge  $Q_{\text{strip}}[\varepsilon] := Q[\varepsilon] + \delta Q_{\text{strip}}[\varepsilon]$  preserves the (renormalized) celestial Gram form on the strip and the linear Ward identity remains anomaly-free up to  $O(\mathcal{B}_{\text{belt}})$ . In particular, the counterterm cancels the scheme-dependent contact induced by the non-unitary Mellin weight off the principal line; after removal of positive flows, the remainder vanishes. See Section 5.93 and the counterterm table in Section Appendix C.

### 5.43 Explicit dual certificate for positivity (numbers, no solver)

**Goal.** Provide a compact certificate (support size 18) that enforces nonnegativity on the working cone, uniformly for  $t \in [-0.30 s_0, 0]$ , while preserving all analytic–projector invariances (IR scheme, pivot, and scale) fixed 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 1).

Weights are uniform (1/18 per support point). Because each component tester is nonnegative and the three families enjoy the invariances of Section 5.22, Section 5.46, and Section 5.66, 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  $[-1, 1]$  to  $[-0.30 s_0, 0]$  via

$$\frac{t}{s_0} = -0.15(1 - x).$$

The resulting ordered nodes (from 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 5.78. 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.

### 5.44 Positivity domain for cubic canonical energy

*Theorem 5.70* (cubic canonical-energy domain). Let  $\mathcal{D}_{\text{an}}$  be the analytic core of the belt modular generator (Section 5.41). 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  $\|K_m P^\perp\| \leq \eta^m$ .

*Proof.* Write  $K := K_{\text{mod}}$  and  $D_{\text{an}} := \mathcal{D}_{\text{an}}$  (see Section 5.41). By construction of the cubic canonical-energy form  $\mathcal{E}_{\text{can}}^{(3)}$  from the third variation of the belt first-law channel, there is a real Borel function  $q_3$  on  $\mathbb{R}$  with

$$\mathcal{E}_{\text{can}}^{(3)}[\psi] = \int_{\mathbb{R}} q_3(\lambda) d\mu_\psi(\lambda), \quad |q_3(\lambda)| \leq C(1 + |\lambda|^3),$$

for all  $\psi \in D_{\text{an}}$ , where  $d\mu_\psi$  is the spectral measure of  $K$ .

By Definition 5.20, Lemma 5.7, and Proposition 5.67, the belt OS/KMS construction furnishes a common analytic core  $D_{\text{an}}$ ,  $K \upharpoonright D_{\text{an}}$  is essentially self-adjoint, and  $D_{\text{an}}$  is a core for all polynomials in  $K$ .

(i)  $\mathcal{D}^{(3)}$  is a form core. Set

$$\|\psi\|_{3/2} := \|(1 + |K|)^{3/2}\psi\|, \quad \text{Dom}_{3/2} := \text{Dom}((1 + |K|)^{3/2}).$$

For  $\psi \in D_{\text{an}}$ ,

$$|\mathcal{E}_{\text{can}}^{(3)}[\psi]| \leq C \int_{\mathbb{R}} (1 + |\lambda|^3) d\mu_\psi(\lambda) = C \|(1 + |K|)^{3/2}\psi\|^2 = C \|\psi\|_{3/2}^2.$$

Hence  $\mathcal{E}_{\text{can}}^{(3)}$  is bounded with respect to  $\|\cdot\|_{3/2}$  and thus closable on  $D_{\text{an}}$ , with closed extension to  $\text{Dom}_{3/2}$ .

Let  $S_\varepsilon := e^{-\varepsilon K^2}$ . For any  $\psi \in \text{Dom}_{3/2}$ ,  $S_\varepsilon\psi \in D_{\text{an}}$  and

$$\|(1 + |K|)^{3/2}(S_\varepsilon - \mathbf{1})\psi\|^2 = \int_{\mathbb{R}} (1 + |\lambda|^3) |e^{-\varepsilon\lambda^2} - 1|^2 d\mu_\psi(\lambda) \xrightarrow{\varepsilon \downarrow 0} 0,$$

hence  $D_{\text{an}}$  is dense in  $\text{Dom}_{3/2}$  for the form norm. By belt quasi-locality and factorization (see Proposition 5.78 and Section 5.50), the algebraic span of single-patch vectors is dense in  $D_{\text{an}}$ ; intersecting with  $\text{Dom}_{3/2}$  and taking the closure yields precisely  $\mathcal{D}^{(3)}$ . Therefore  $\mathcal{D}^{(3)}$  is dense in  $\text{Dom}_{3/2}$  for the form norm and is a core for  $\mathcal{E}_{\text{can}}^{(3)}$ .

(ii) *Positivity on the orthogonal complement of boost–Killing modes.* Let  $\psi \in \mathcal{D}^{(3)}$  be orthogonal to all boost–Killing modes. On  $D_{\text{an}}$ , the second-order modular inequality (Theorem 5.46) states

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

with  $Q_{\text{shear}} \geq 0$  and with kernel of  $\mathcal{E}_{\text{can}}$  equal to the boost–Killing modes. This is derived using BW/KMS positivity of the OS kernel (Lemma 3.1), combined with the belt first-law channel and its calibration.

Consider an admissible one-parameter family whose initial tangent is the single-patch direction  $\psi$ . Differentiating the preceding inequality once more at the origin isolates the cubic belt term:

$$\delta^3 \left[ S - \frac{\text{Area}}{4G} \right] = 2\pi \mathcal{E}_{\text{can}}^{(3)}[\psi] + (\text{shear} \times \text{expansion}) + R_{\text{belt}}^{(3)}.$$

The shear  $\times$  expansion term is controlled as in Theorem 5.46, yielding a lower bound by  $-C \mathcal{E}_{\text{can}}[\delta\Psi; \xi] - C' \mathcal{B}_{\text{belt}}$  and in particular nonnegativity when  $\mathcal{E}_{\text{can}}$  vanishes. Moreover, BW/KMS positivity (Lemma 3.1) implies the belt modular three-point kernel defines a positive Gram-type functional on modular translates of a local direction; evaluating on the single direction  $\psi$  and using the orthogonality to boost–Killing modes yields

$$\mathcal{E}_{\text{can}}^{(3)}[\psi] \geq 0$$

(up to  $O(\mathcal{B}_{\text{belt}})$ , which vanishes upon removing positive flows). Thus  $\mathcal{E}_{\text{can}}^{(3)}[\psi] \geq 0$  for all such  $\psi$ .

(iii) *Stability under constant-depth belt circuits and AGSP maps.* Let  $U$  be a constant-depth belt circuit (Proposition 5.78, Section 5.50). Then  $U$  maps single-patch support to single-patch support, hence  $U\mathcal{D}^{(3)} \subset \mathcal{D}^{(3)}$ . Quasi-locality implies  $U^*KU = K + B$  with  $B$  bounded. By functional calculus (Kato–Rellich for bounded perturbations),

$$\|(1 + |K|)^{3/2}U\psi\| = \|(1 + |U^*KU|)^{3/2}\psi\| \lesssim \|(1 + |K|)^{3/2}\psi\| + \|\psi\|,$$

so  $\text{Dom}_{3/2}$  (and thus  $\mathcal{D}^{(3)}$ ) is preserved.

For AGSP maps,  $K_m = f_m(K)$  with  $\|f_m\|_\infty \leq 1$  and  $\|K_m P^\perp\| \leq \eta^m$  by construction. Since  $K_m$  is a bounded Borel function of  $K$ , it commutes with  $(1 + |K|)^{3/2}$  and

$$\|(1 + |K|)^{3/2}K_m\psi\| = \|K_m(1 + |K|)^{3/2}\psi\| \leq \|(1 + |K|)^{3/2}\psi\|,$$

whence  $K_m\mathcal{D}^{(3)} \subset \mathcal{D}^{(3)}$  and the contraction on  $P^\perp$  holds as stated.

The three claims (i)–(iii) follow.  $\square$

## 5.45 Modular Bekenstein inequality and numeric check

*Proposition 5.71* (modular Bekenstein on belts). For admissible perturbations supported on a belt, relative entropy positivity implies

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

*Proof.* Fix an admissible belt and reference state  $\sigma$ . On the OS/KMS analytic core one has Araki's identity

$$S(\rho_R \| \sigma_R) = \Delta \langle K_{\text{mod}}(R) \rangle - \Delta S(R),$$

whence  $S(\rho_R \| \sigma_R) \geq 0$  implies  $\Delta S(R) \leq \Delta \langle K_{\text{mod}}(R) \rangle$  for the belt-regularized theory. The positive-flow removal lemma eliminates the uniform  $O(\mathcal{B}_{\text{belt}})$  remainder, so the bound holds exactly after removing the regulators. (OS/KMS core: Lemma 3.1; removal: Lemma 3.3.)  $\square$

## 5.46 Dispersion pivot invariance for gravity-subtracted amplitudes

*Lemma 5.72* (Pivot invariance). Let  $A_{\text{sub}}(s, t)$  denote the  $N = 3$  gravity-subtracted, crossing-even amplitude on the dispersive cone. For any real  $s_*$ , let  $\Re A_{\text{sub}}(s, t; s_*)$  be the real part obtained by writing the  $N = 3$  Cauchy representation about the pivot  $s = s_*$ . Then there exists a polynomial  $Q_2(s, t; s_*)$ , of degree at most 2 in  $s$  with coefficients analytic in  $t$  on the cone, such that

$$\Re A_{\text{sub}}(s, t; s_*) = \Re A_{\text{sub}}(s, t; 0) + Q_2(s, t; s_*).$$

In particular, the even-parity forward derivative inequalities and the Hankel/celestial positivity functionals are unchanged by the choice of pivot.

*Proof.* Write the crossing-even,  $N = 3$  subtracted dispersion relation at pivot  $s_*$  in the standard symmetric form

$$\begin{aligned} \Re A_{\text{sub}}(s, t; s_*) &= P_2(s, t; s_*) \\ &+ \frac{1}{\pi} \text{P.V.} \int_{s_0}^{\infty} ds' \left[ \frac{(s - s_*)^3}{(s' - s_*)^3} \frac{\rho_s(s', t)}{s' - s} + \frac{(u - s_*)^3}{(s' - s_*)^3} \frac{\rho_u(s', t)}{s' - u} \right], \end{aligned}$$

where  $u := -s - t$ , the absorptive parts are  $\rho_s = \frac{1}{2} \text{Disc}_s A_{\text{sub}}$  and  $\rho_u = \frac{1}{2} \text{Disc}_u A_{\text{sub}}$  (pivot-independent on the cone), and  $P_2$  is a polynomial in  $s$  of degree  $\leq 2$  with coefficients analytic in  $t$ .

Let

$$D(s, t) := \Re A_{\text{sub}}(s, t; s_*) - \Re A_{\text{sub}}(s, t; 0).$$

By construction, the kernels multiplying  $\rho_s$  and  $\rho_u$  are normalized to 1 at their respective poles:

$$\lim_{s' \rightarrow s} \frac{(s - s_*)^3}{(s' - s_*)^3} = 1, \quad \lim_{s' \rightarrow u} \frac{(u - s_*)^3}{(s' - s_*)^3} = 1,$$

hence  $\text{Disc}_s D = 0$  and  $\text{Disc}_u D = 0$ . Therefore, for fixed  $t$  in the cone,  $D(\cdot, t)$  extends holomorphically through both the  $s$ - and  $u$ -cuts and is entire in  $s$ . Moreover, both representations use  $N = 3$  subtractions, so  $D(\cdot, t) = O(|s|^2)$  as  $|s| \rightarrow \infty$ . By the standard Liouville-type bound for entire functions with polynomial growth,  $D(\cdot, t)$  is a polynomial in  $s$  of degree at most 2; analyticity of the coefficients in  $t$  follows from analyticity of the integrands and of  $P_2$  on the cone. Setting  $Q_2 := D$  yields the claimed identity.

Finally, all forward even-parity testers and the Hankel/celestial positivity functionals depend only on the absorptive parts (or, equivalently, annihilate analytic polynomials up to degree 2 via the analytic projector). Since  $\text{Disc } Q_2 \equiv 0$ , their values are independent of the pivot. This establishes the invariance.  $\square$

*Remark 5.73.* (pivot invariance unchanged) Pivot shifts act by  $s$ -polynomials; the strip renormalization is  $t$ -holomorphic at  $s=0$  and does not affect Lemma 5.72.

#### 5.47 Microcausality tails and null timeslice extension

*Lemma 5.74* (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,

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

*Proof.* Fix  $r > 0$  and let  $\partial_r R$  and  $\partial_r \bar{R}$  be belt neighbourhoods whose mutual distance is  $r$ . By the quasi-local belt factorization established earlier, there are constants  $c, \mu > 0$  (independent of  $|R|$ ) with the following property: for every  $A \in \mathcal{A}(\partial_r R)$  and  $B \in \mathcal{A}(\partial_r \bar{R})$  there exist

$$A_r \in \mathcal{A}(\partial_{r/2} R), \quad B_r \in \mathcal{A}(\partial_{r/2} \bar{R})$$

such that

$$[A_r, B_r] = 0, \quad \|A - A_r\| \leq c e^{-\mu r} \|A\|, \quad \|B - B_r\| \leq c e^{-\mu r} \|B\|.$$

(One may take  $A_r, B_r$  to be the images of  $A, B$  under the completely positive, approximately inner “split” maps associated with the belt; the constants come from the quantitative split estimate and are uniform in  $|R|$ .)

Then

$$\begin{aligned} \|[A, B]\| &\leq \|[A - A_r, B]\| + \|[A_r, B - B_r]\| + \|[A_r, B_r]\| \\ &\leq 2\|A - A_r\| \|B\| + 2\|A_r\| \|B - B_r\| \\ &\leq 2c e^{-\mu r} \|A\| \|B\| + 2c e^{-\mu r} (\|A\| + \|A - A_r\|) \|B\| \\ &\leq 4c e^{-\mu r} \|A\| \|B\|. \end{aligned}$$

Absorbing  $\|A\| \|B\|$  into the prefactor gives the stated bound with  $\mu_{\text{eff}} = \mu$  and a constant  $C_{\text{mc}}$  uniform in  $|R|$ .  $\square$

*Proposition 5.75* (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  $D[R]$ .

*Proof.* Let  $\alpha_t$  denote the (null) belt boost flow on  $\mathcal{A}(D[R])$ . The null timeslice property at belt level implies that

$$\mathfrak{A}_{\text{belt}} := \overline{\bigcup_{t \in \mathbb{R}} \alpha_t(\mathcal{A}(\partial_r R))}^{\|\cdot\|}$$

is norm-dense in  $\mathcal{A}(D[R])$  for any fixed  $r > 0$ . Moreover, relative Cauchy evolution for compactly supported background variations inside  $D[R]$  is implemented by  $\alpha_t$  up to edge/corner terms that are  $O(\mathcal{B}_{\text{belt}})$  and supported on the belt.

By hypothesis, the equation-of-state identity and the energy inequalities hold on  $\mathcal{A}(\partial_r R)$ . These statements are invariant under \*-automorphisms and stable under norm limits. Hence they hold on  $\alpha_t(X)$  for every  $X \in \mathcal{A}(\partial_r R)$  and  $t \in \mathbb{R}$ , and therefore on all of  $\mathfrak{A}_{\text{belt}}$ .

To pass from  $\mathfrak{A}_{\text{belt}}$  to  $\mathcal{A}(D[R])$  without loss, we control the regulator and boundary effects. First, the belt microcausality tail (Lemma 5.74) ensures that moving observables along the null belt flow does not produce uncontrollable commutators with degrees of freedom across the belt: any such commutator is  $O(e^{-\mu_{\text{eff}} r})$ . Second, the calibrated corner/edge functionals cancel uniformly in  $t$ , so the  $O(\mathcal{B}_{\text{belt}})$  remainder in the relative Cauchy evolution vanishes after calibration. Finally, the flow-removal argument identifies the limit in which the belt regulator is removed with the identity on  $\mathcal{A}(D[R])$ , so the propagated identity/inequalities extend to all of  $\mathcal{A}(D[R])$  exactly and regulator free.

Consequently, if the equation-of-state identity and the energy inequalities hold on a belt cut, they hold throughout the entire domain of dependence  $D[R]$ .  $\square$

## 5.48 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).**

$$\frac{t}{s_0} = \left( \begin{array}{l} -0.2948888739, -0.2560660172, -0.1888228568, \\ -0.1111771432, -0.0439339828, -0.0051111261 \end{array} \right).$$

**Hankel/impact scales (units of  $s_0$ ).** The seven Hankel/impact testers use the frozen scales of Section 5.33:

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

**Celestial anchors (principal series).** Same five principal-series anchors as in Table 1:

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

**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 5.39, Lemma 5.72, Lemma 5.95), as are CPT/crossing constraints (Proposition 5.100).

**Wiring update** The celestial component may be realized either on the principal set (Table 1) or on the strip set (Table 3); 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 5.80 and Proposition 5.108.

*Remark 5.76.* This compact dual coincides node-wise with the earlier “numbers, no solver” presentation in Section 5.43, now updated to the widened near-forward window of Section 5.10.

#### 5.49 Wald-JKM corner calibration on belts

*Lemma 5.77* (JKM fix). Let  $Q_\xi$  be the Noether charge for the boost generator  $\xi$  and  $\Theta$  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 of  $\xi$  fixed by the Rindler witness (Lemma 5.77). Then the first-law identity and the modular equation of state are insensitive to the Wald ambiguity up to  $O(\mathcal{B}_{\text{belt}})$ .

*Proof.* Fix a thin collar  $\mathcal{B}_r \subset \partial_r R$  with oriented belt circle  $C_r := \partial \mathcal{S}_r$  and cap  $\mathcal{S}_r$  filling the collar. For any variation  $\delta g$  supported in  $\mathcal{B}_r$ , the Iyer–Wald identity gives

$$d(\delta Q_\xi - \xi \cdot \Theta(\delta g)) = \omega(g; \delta g, \mathcal{L}_\xi g) - \xi \cdot E(g) \cdot \delta g. \quad (5.83)$$

Integrating (5.83) over  $\mathcal{S}_r$  and applying Stokes yields

$$\delta H_\xi^{\text{can}}[g; \delta g] := \int_{C_r} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) = \int_{\mathcal{S}_r} \omega(g; \delta g, \mathcal{L}_\xi g) - \int_{\mathcal{S}_r} \xi \cdot E(g) \cdot \delta g. \quad (5.84)$$

By construction, the background solves the equations of motion near the belt and  $\delta g$  is supported in  $\mathcal{B}_r$ ; hence the last term in (5.84) vanishes. Moreover, in the boost-adapted gauge and with the boost normalization fixed by the Rindler witness (Lemma 5.77),  $\mathcal{L}_\xi g = O(\mathcal{B}_{\text{belt}})$  on  $\mathcal{S}_r$ , so the symplectic flux is  $O(\mathcal{B}_{\text{belt}})$ . Thus

$$\delta H_\xi^{\text{can}}[g; \delta g] = \int_{C_r} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) + O(\mathcal{B}_{\text{belt}}). \quad (5.85)$$

The JKM ambiguity is the shift  $(\Theta, Q_\xi) \mapsto (\Theta + dY, Q_\xi - \xi \cdot Y)$  by an arbitrary local  $(D-2)$ -form  $Y$ . Under this shift,

$$\delta H_\xi^{\text{can}} \mapsto \delta H_\xi^{\text{can}} - \delta \int_{C_r} \xi \cdot Y. \quad (5.86)$$

On the other hand, the standard corner evaluation at a boost bifurcation surface gives

$$\int_{C_r} \delta Q_\xi = \delta \left[ \frac{\text{Area}(C_r)}{4G} \right] + O(\mathcal{B}_{\text{belt}}), \quad (5.87)$$

where the  $O(\mathcal{B}_{\text{belt}})$  accounts for belt thickness and vanishes under flow removal. Combining (5.85)–(5.87), we obtain

$$\delta H_\xi^{\text{can}} = \delta \left[ \frac{\text{Area}}{4G} \right] - \delta \int_{C_r} \xi \cdot \Theta(\delta g) + O(\mathcal{B}_{\text{belt}}). \quad (5.88)$$

The belt first-law channel furnishes a *boost Ward charge*  $\delta H_\xi^{\text{belt}}$  that is fixed unambiguously by the Ward identity and the Rindler witness normalization. We *define*  $Y$  on the belt by the requirement that the canonical generator (5.88), shifted as in (5.86), matches the Ward charge:

$$\delta H_\xi^{\text{belt}} = \delta H_\xi^{\text{can}} - \delta \int_{C_r} \xi \cdot Y = O(\mathcal{B}_{\text{belt}}).$$

This choice is equivalent to the asserted JKM calibration

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

and fixes  $Y$  uniquely modulo  $O(\mathcal{B}_{\text{belt}})$  on the belt (the residual vanishes under flow removal).

Finally, insensitivity of the first-law identity and of the modular equation of state to the Wald ambiguity follows directly: for any further JKM shift  $Y \mapsto Y + \Delta Y$ , both the canonical generator and the combination  $\delta[\text{Area}/(4G)] - \delta[\xi \cdot \Theta(\delta g)]$  shift by  $-\delta \int_{C_r} \xi \cdot \Delta Y$ , so all belt first-law and modular relations expressed in terms of these quantities are invariant up to  $O(\mathcal{B}_{\text{belt}})$  (and exactly invariant after flow removal). All constants are per length and uniform in  $|R|$ .  $\square$

**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 5.74, Propositions 5.75 and 5.78, and Section 5.50) 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 above:

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

On the other hand, applying the Iyer–Wald identity on a cap  $\mathcal{S}_r$  as above and using Stokes reduces the bulk contribution to the belt with an  $O(\mathcal{B}_{\text{belt}})$  remainder. 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 *defines*  $Y$  so that  $\delta C_\xi^{\text{corner}} = O(\mathcal{B}_{\text{belt}})$ . Using the standard corner evaluation (5.87) 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  $\delta g$ , 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  $|R|$ .  $\square$

## 5.50 Brown–York dictionary for the belt flux

*Proposition 5.78* (quasi-local stress on the belt). Let  $\gamma_{ab}$  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

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

*Proof. Setup.* Let  $\Sigma \subset \partial_r R$  be a belt slab bounded by two boosted corners  $\partial\Sigma = C_r^{(1)} \cup C_r^{(2)}$ , and let  $\xi$  be the belt boost, normalized to modular period  $2\pi$  along the wedge. Pick a belt-supported

deformation  $h_{\mu\nu} = f(\rho) \mathcal{L}_\xi g_{\mu\nu}$  with a collar cutoff  $f(\rho)$  across the belt thickness. For any wedge-local observable  $A$ , relative Cauchy evolution (RCE) gives

$$\frac{d}{d\epsilon} \Big|_{\epsilon=0} \langle A \rangle_{g+\epsilon h} = \frac{i}{2} \left\langle \left[ \int_{\mathcal{B}_r} \sqrt{-g} T^{\mu\nu} h_{\mu\nu}, A \right] \right\rangle + O(\mathcal{B}_{\text{belt}}). \quad (5.89)$$

Taking  $A$  to be the modular generator on the belt-anchored wedge and using the first-law/canonical-energy channel, the commutator in (5.89) evaluates to the canonical-energy pairing, hence

$$2\pi \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta \langle T_{\mu\nu} \rangle = \int_{\Sigma} \omega(g; \delta g, \mathcal{L}_\xi g) + O(\mathcal{B}_{\text{belt}}), \quad (5.90)$$

where  $\omega$  is the covariant symplectic current of Einstein gravity (with matter) and we used  $h = f(\rho) \mathcal{L}_\xi g$ .

*Transport to the boundary.* By the Iyer–Wald identity applied on the belt slab,

$$\int_{\Sigma} \omega(g; \delta g, \mathcal{L}_\xi g) = \int_{\partial\Sigma} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) + O(\mathcal{B}_{\text{belt}}). \quad (5.91)$$

The  $O(\mathcal{B}_{\text{belt}})$  remainder accounts for (i) the collar gradients of  $f(\rho)$  and (ii) the failure of exact boost Killing away from the belt; both are uniformly bounded and vanish in the thin, boost-adapted limit.

*Removal of the boosted corners.* Implement the Wald–JKM corner calibration on belts (Section 6.2). With this choice of counterterm, the corner combination satisfies

$$\left[ \delta Q_\xi - \xi \cdot \Theta(\delta g) \right]_{\text{corner}} = \delta \left[ \frac{\text{Area}}{4G} \right]_{\text{corner}} + O(\mathcal{B}_{\text{belt}}), \quad (5.92)$$

with the normalization fixed by the Rindler witness (Section 6.2). The gravitational area variation appearing in the modular first law cancels against (5.92), so after summing the two corner contributions at  $C_r^{(1)}, C_r^{(2)}$  only the timelike-belt piece of (5.91) survives, up to  $O(\mathcal{B}_{\text{belt}})$ .

*Timelike belt = Brown–York momentum flux.* On the timelike belt, work with the Einstein–Hilbert action supplemented by the GHY term. For diffeomorphisms generated by a vector field tangent to the belt, the covariant-phase-space generator satisfies

$$\int_{\partial\Sigma} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) \Big|_{\text{belt}} = 2\pi \int_{\partial\Sigma} d\ell^a \delta \langle T_{ab}^{\text{BY}} \rangle \xi^b + O(\mathcal{B}_{\text{belt}}), \quad (5.93)$$

where  $T_{ab}^{\text{BY}} = \frac{1}{8\pi G} (K_{ab} - K \gamma_{ab})$  is the Brown–York quasi-local stress constructed from the belt induced metric  $\gamma_{ab}$  and extrinsic curvature  $K_{ab}$ . The identity (5.93) is the standard EH+GHY/covariant-phase-space matching: the variation of the boundary Hamiltonian equals the Brown–York momentum flux along  $\xi$  on  $\partial\Sigma$ , with the JKM calibration ensuring that no extra corner potential remains on the belt.

*Conclusion.* Combining (5.90), (5.91), and (5.93) yields

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

which is the claimed Brown–York dictionary for the belt flux. All  $O(\mathcal{B}_{\text{belt}})$  terms are uniform in the belt thickness and vanish in the thin, boost-adapted limit.  $\square$

## 5.51 Quasi-local positive energy on the belt

*Corollary 5.79* (quasi-local positive energy on belt). Under the JKM corner calibration (Lemma 5.77) and the Brown–York dictionary (Proposition 5.78), for any spacelike belt slab  $\Sigma \subset \partial_r R$  and any nonnegative smearing  $f$  along its generators,

$$2\pi \int_{\partial\Sigma} d\ell_a f \delta \langle T_{\text{BY}}^{ab} \rangle \xi_b \geq -C_{\text{BY}} B_{\text{belt}}, \quad C_{\text{BY}} > 0, \quad (5.94)$$

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

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

*Proof.* By the belt Brown–York dictionary (Proposition 5.78), for any spacelike belt slab  $\Sigma \subset \partial_r R$  and any smearing  $f$  along its generators we have

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

where the error term  $E_{\text{belt}}$  is controlled as

$$|E_{\text{belt}}| \leq C_{\text{BY}} B_{\text{belt}}$$

for some constant  $C_{\text{BY}} > 0$  independent of  $|R|$ ; this is the uniform  $O(\mathcal{B}_{\text{belt}})$  bound from Proposition 5.78. In particular,

$$2\pi \int_{\partial\Sigma} d\ell_a f \delta \langle T_{\text{BY}}^{ab} \rangle \xi_b \geq 2\pi \int_{\Sigma} d\Sigma_\mu f \xi_\nu \delta \langle T^{\mu\nu} \rangle - C_{\text{BY}} B_{\text{belt}}.$$

For  $f \geq 0$ , the bulk term

$$2\pi \int_{\Sigma} d\Sigma_\mu f \xi_\nu \delta \langle T^{\mu\nu} \rangle$$

is precisely the canonically smeared null energy on the belt generators. By the ANEC/QNEC synthesis (Theorem 5.64), together with canonical-energy positivity for non-Killing perturbations derived from the second-order modular equation of state (Theorem 5.46), this contribution is nonnegative,

$$2\pi \int_{\Sigma} d\Sigma_\mu f \xi_\nu \delta \langle T^{\mu\nu} \rangle \geq 0.$$

Combining these estimates yields the bound (5.94).

Finally, by the definition of the belt regulator the control quantity  $B_{\text{belt}}$  tends to zero in the positive-flow removal limit  $(u, s) \downarrow 0$ , while  $C_{\text{BY}}$  remains fixed. Taking this limit in (5.94) gives

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

which completes the proof. □

*Remark 5.80* (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 5.46) together with the calibrated corner terms and the Brown–York dictionary (Sections 5.49 and 5.50). The statement is stable under the invariances summarized in Section 6.

## 5.52 Weyl covariance and absence of a leading anomaly

*Lemma 5.81* (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  $\epsilon$  up to  $O(\mathcal{B}_{\text{belt}})$ .

*Proof.* Work per generator length and on the OS belt window. Consider the rescaled background  $\tilde{g} = \Omega^2 g$  with  $\Omega = 1 + \epsilon$  and compare the generalized-entropy variation in the two descriptions. We show that the difference is  $O(\epsilon \mathcal{B}_{\text{belt}})$ .

(1) *Corner/edge co-transformation under JKM calibration.* With the JKM counterterm fixed by the belt boost Ward identity, the Noether corner piece from  $\delta[\text{Area}/(4G)]$  cancels the corner piece from  $\xi \cdot \Theta(\delta g)$  up to the uniform belt remainder,

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

and this is insensitive to a uniform Weyl rescaling because it is enforced locally on the belt (choice of  $Y$  in the JKM ambiguity) and we work to first order in the OS window (Lemmas 3.1 to 3.3; Section 6.2; Lemma 5.77)

(2) *Flux dictionary and Brown–York co-transformation.* On a belt slab, the calibrated bulk canonical-energy flux equals the Brown–York flux on the timelike belt, modulo  $O(\mathcal{B}_{\text{belt}})$ ,

$$2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta \langle T^{\mu\nu} \rangle = 2\pi \int_{\partial\Sigma} dl_a \delta \langle T_{\text{BY}}^{ab} \rangle \xi_b + O(\mathcal{B}_{\text{belt}}),$$

so changes induced by a uniform Weyl rescaling are carried by a belt-local boundary current with the same calibration (Section 6.2; Proposition 5.78).

(3) *Counterterm channel and cancellation at leading order.* Uniform Weyl scalings act, to linear order, like a belt-local scheme update of geometric counterterms. For two belt-compatible schemes A and B the renormalized stress shifts by a belt-local contact  $\Delta T_{\mu\nu}^{\text{ct}}[g]$ , and the Wald/area term shifts accordingly. The boost Ward identity with the JKM/BY calibration fixes the  $\xi$ -flux of  $\Delta T_{\mu\nu}^{\text{ct}}$  to cancel the leading area/Wald shift; the residual is governed by the exponential belt tail:

$$\delta \left\langle \frac{\hat{A}}{4G} \right\rangle_A - \delta \left\langle \frac{\hat{A}}{4G} \right\rangle_B + 2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta \langle \Delta T^{\text{ct}\mu\nu} \rangle = O(\mathcal{B}_{\text{belt}}),$$

i.e. the belt-scheme stability of Section 5.31 (see Eqs. (5.61) to (5.63) there).

(4) *Entropy counterterm is absorbed by recovery.* The entanglement entropy shifts by a belt/corner-local counterterm  $S_{\text{ct}}$  across such scheme/Weyl moves. Belt recovery/continuity bounds imply  $|\delta S_{\text{ct}}| \leq C_S \mathcal{B}_{\text{belt}}$ , uniformly in  $|R|$  (Section 5.31).

(5) *Assemble and linearize in  $\epsilon$ .* The calibrated first-law channel holds in either description,

$$\delta \left[ S - \frac{\text{Area}}{4G} \right] = 2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta \langle T^{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}), \quad \delta \left[ S - \frac{\text{Area}}{4G} \right]_{\Omega} = 2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta \langle T^{\mu\nu} \rangle_{\Omega} + O(\mathcal{B}_{\text{belt}}),$$

and the difference of the right-hand sides is precisely the counterterm/scheme effect in (3), which is  $O(\mathcal{B}_{\text{belt}})$ . Since  $\Omega = 1 + \epsilon$  and all such counterterm shifts are linear in  $\epsilon$  for a uniform rescaling, the net change is  $O(\epsilon \mathcal{B}_{\text{belt}})$ . This matches the master invariance for Weyl-covariant counterterm updates and  $G_{\text{ren}}$  renormalization (Theorem 6.3; Lemma 5.107).

Therefore

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

i.e. no leading Weyl anomaly arises at linear order once the belt JKM/BY calibration is enforced; the  $O(\mathcal{B}_{\text{belt}})$  remainder vanishes under flow removal.  $\square$

*Remark 5.82* (What is used). Only: (i) the calibrated corner identity and the Brown–York flux dictionary (Section 6.2; Lemma 5.77; Proposition 5.78), (ii) belt-compatible counterterm stability and recovery bounds for  $S_{\text{ct}}$  (Section 5.31), and (iii) the master invariance including Weyl-covariant counterterm updates ( $G_{\text{ren}}$ -invariance Lemma 5.107; Theorem 6.3).

### 5.53 Raising the subtraction order does not weaken positivity

*Proposition 5.83* (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  $N' > N$ .

*Proof. Step 1: Forward even-parity derivatives at general  $N \geq 3$ .* Let  $A(N)(s, t)$  denote the  $N$ -subtracted, gravity-subtracted amplitude on the forward cone, and write  $A^{(+)}(s, t) := \frac{1}{2}(A(s, t) + A(-s, t))$  for the crossing-even part. For each even  $k = 2m \geq 0$  and fixed  $t \leq 0$ , the crossing-symmetric dispersion algebra gives the representation

$$\partial_s^{2m} \Re A(N)(0, t) = \frac{(2m)!}{\pi} \int_{s_0}^{\infty} \frac{ds'}{s'^{N+2m+1}} \Im A_{\text{hard}}^{(+)}(s', t),$$

with exchange of  $\partial_s$  and the  $s'$ -integral justified by dominated convergence under  $\alpha_{\text{R}} < 3$ .<sup>3</sup> Thus, for  $k = 2m$  the integrand is pointwise nonnegative on  $\mathcal{S}$ :  $\Im A_{\text{hard}}^{(+)}(s', t) \geq 0$ , and the  $N$ -dependence enters only through the positive weight  $s^{-(N+2m+1)}$ .

*Step 2: Monotonicity in the subtraction order for forward functionals.* Fix  $N' > N \geq 3$  and even  $k = 2m$ . Comparing the two representations,

$$\partial_s^{2m} \Re A(N')(0, t) = \frac{(2m)!}{\pi} \int_{s_0}^{\infty} \frac{ds'}{s'^{N+2m+1}} \left(\frac{s_0}{s'}\right)^{N'-N} \Im A_{\text{hard}}^{(+)}(s', t).$$

Since  $\left(\frac{s_0}{s'}\right)^{N'-N} \in (0, 1]$  for  $s' \geq s_0$ , the integrand for  $N'$  is the integrand for  $N$  multiplied by a positive factor  $\leq 1$ . Hence  $\partial_s^{2m} \Re A(N')(0, t) \geq 0$  whenever  $\partial_s^{2m} \Re A(N)(0, t) \geq 0$ , and in particular the even-parity forward inequalities that hold at  $N = 3$  continue to hold at all  $N' > 3$ . Moreover, any quantitative lower bound at  $N$  improves (weakly) at  $N'$  because the positive integrand is damped by an extra power of  $s'$ .

*Step 3: Hankel/impact and celestial Gram testers.* These testers are defined directly on the absorptive profile and are independent of  $N$ : a Hankel/impact functional is a positive quadratic form of the impact-profile transform of  $\Im A_{\text{hard}}$ , and a celestial Gram functional is a positive-semidefinite quadratic form of principal-series celestial data built from  $\Im A_{\text{hard}}$ . Since the gravity subtraction is held fixed and only the *number of subtractions* in the crossing-symmetric dispersion for the real part is changed, the Hankel and celestial kernels, their Gram forms, and their nonnegativity remain unchanged as  $N$  varies. Therefore any nonnegative Hankel/celestial functional at  $N=3$  remains nonnegative for all  $N' > 3$ .

*Step 4: Integrability for all  $k \geq 0$  at  $N \geq 3$ .* The tester-certified envelope  $\alpha_{\text{R}} \leq 2 + \delta_{\star} < 3$  gives  $\Im A_{\text{hard}}(s', t) \lesssim s'^{\alpha_{\text{R}}}$  on the cone, so the forward integrands behave as  $s'^{-N-k-1+\alpha_{\text{R}}}$  and are integrable at infinity for all  $k \geq 0$  whenever  $N \geq 3$ . This justifies Steps 1–2 uniformly in  $k$ .

Combining Steps 1–4 proves that raising  $N$  cannot destroy any positivity inequality from the forward even-parity, Hankel/impact, or celestial families; in the forward family it can only strengthen quantitative lower bounds by further suppressing the high- $s'$  tail.  $\square$

**Add-on (quantitative tail monotonicity).** For any cutoff  $S_{\text{cut}} \geq s_0$  and even  $k = 2m$ , decompose the forward functional at order  $N$  into a head plus a tail over  $[S_{\text{cut}}, \infty)$ . Then for  $N' > N$ ,

$$\text{Tail}_{N'} = \int_{S_{\text{cut}}}^{\infty} \frac{ds'}{s'^{N+2m+1}} \left(\frac{s_0}{s'}\right)^{N'-N} \Im A_{\text{hard}}^{(+)}(s', t) \leq \left(\frac{s_0}{S_{\text{cut}}}\right)^{N'-N} \text{Tail}_N.$$

Thus every tail budget (and any bound that depends monotonically on it) improves by at least the explicit damping factor  $(s_0/S_{\text{cut}})^{N'-N}$  when the subtraction order is raised.

<sup>3</sup>Here and below “hard” means the gravity soft piece has been removed; no tester acts on the nonanalytic soft terms.

## 5.54 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  $\alpha_{\text{R}}(t) \leq 2 + \delta_*$  on  $t \in [-0.25 s_0, 0]$  (Proposition 5.32). 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  $S_{\text{cut}}$  ( $\delta$ -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 5.27.

**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 5.89) 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  $\leq \varepsilon/6$ .

### 5.55 Edge Cardy sanity check for the Unruh channel

*Proposition 5.84* (edge Cardy bound). Let  $T_{\text{Unruh}} = 1/(2\pi)$  be the Unruh temperature of the belt boost. Using the composite constant 0.4083 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 5.39.

*Proof.* Fix the belt boost and let  $\tau$  denote the corresponding KMS state at inverse temperature  $\beta = 2\pi$  (so  $T_{\text{Unruh}} = 1/(2\pi)$ ). By the modular Bekenstein inequality on belts with reference state  $\tau$  Proposition 5.71,

$$S(\rho_R) - S(\tau_R) \leq \langle K_{\text{mod}}(R) \rangle_{\rho} - \langle K_{\text{mod}}(R) \rangle_{\tau}.$$

Divide by  $\text{length}(\partial R)$  to obtain

$$\frac{S(\rho_R)}{\text{length}(\partial R)} \leq \frac{S(\tau_R)}{\text{length}(\partial R)} + \frac{\langle K_{\text{mod}}(R) \rangle_{\rho} - \langle K_{\text{mod}}(R) \rangle_{\tau}}{\text{length}(\partial R)}. \quad (5.95)$$

The first term on the right is fixed by the edge/Ward calibration: for the boost-KMS state  $\tau$  one has the Cardy-type edge density normalization

$$\frac{S(\tau_R)}{\text{length}(\partial R)} = 0.4083 T_{\text{Unruh}},$$

by definition of the composite constant 0.4083 from the edge/ward calibration Section 5.16 and Proposition 5.40.

For the second term in (5.95), the belt first-law channel together with the global remainder ledger implies that, at fixed belt regulators and after removal of the positive flows, the per-length mismatch of the modular energy between any admissible  $\rho$  and the aligned KMS reference  $\tau$  is controlled by the single belt budget:

$$\frac{\langle K_{\text{mod}}(R) \rangle_{\rho} - \langle K_{\text{mod}}(R) \rangle_{\tau}}{\text{length}(\partial R)} = O(\mathcal{B}_{\text{belt}}).$$

Inserting these two inputs into (5.95) yields

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

as claimed. The stated  $O(1/\text{length})$  finite-size correction follows from the general thermodynamic-limit control for the per-length entropy in Section 5.39.  $\square$

### 5.56 Equivalence: entanglement equilibrium and linearized Einstein equations

*Theorem 5.85* (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**  $\Rightarrow$  **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, then

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

*Proof.* Fix an admissible belt-anchored region  $R$ , wedge  $W = E_W(R)$ , and a Cauchy slice  $\Sigma \subset W$  whose boundary meets the timelike belt and two small caps at the belt corners. All identities below are per generator length, with a uniform  $O(\mathcal{B}_{\text{belt}})$  remainder.

*Preliminaries.* (i) The belt first law (with calibrated edge/corner terms) gives, for any admissible variation,

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

This holds continuously under belt-supported shape/state variations by the recovery/continuity kernel; see Corollary 3.5 and Lemmas 3.2 and 3.3.

(ii) The belt JLMS channel identifies boundary relative entropy with bulk canonical energy (plus calibrated area) on a common analytic core. In linear response this equates the modular variation with a bulk canonical-energy flux; see Proposition 3.4.

(iii) The Iyer–Wald identity on  $\Sigma$ ,

$$d(\delta Q_\xi - \xi \cdot \Theta(\delta g)) = \omega(g; \delta g, \mathcal{L}_\xi g) - \xi \cdot E(g) \cdot \delta g,$$

relates the bulk canonical-energy current to boundary Noether data. With the JKM corner fixed by the belt boost Ward identity, the corner piece from the area variation cancels the corner piece from  $\xi \cdot \Theta$ , up to  $O(\mathcal{B}_{\text{belt}})$ ; see Section 6.2.

(iv) On the timelike belt, the covariant-phase-space generator equals the Brown–York momentum flux:

$$\int_{\partial\Sigma} (\delta Q_\xi - \xi \cdot \Theta(\delta g)) \Big|_{\text{belt}} = 2\pi \int_{\partial\Sigma} d\ell_a \delta\langle T^{\text{BY}ab} \rangle \xi_b + O(\mathcal{B}_{\text{belt}}),$$

by the EH+GHY/covariant-phase-space matching with the stated calibration; see Section 5.50 and Proposition 5.78.

(v) Combining (i)–(iv) yields the *modular equation of state* (belt form):

$$\delta\langle K_{\text{mod}}(R) \rangle = \delta\left[\frac{\text{Area}(\text{QES}(R))}{4G}\right] + 2\pi \int_{\Sigma} d\Sigma_\mu \xi_\nu \delta\langle T^{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}), \quad (5.96)$$

as stated in Theorem 5.37 and Section 5.19.

We now prove the two implications.

(1) **EE**  $\Rightarrow$  **SEE**. Entanglement equilibrium asserts  $\delta[S - \text{Area}/(4G)] = O(\mathcal{B}_{\text{belt}})$  for *all* admissible belt deformations. Using the first-law channel and the corner/edge calibration,

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

and the calibrated corner/edge variation cancels within  $O(\mathcal{B}_{\text{belt}})$  by Section 6.2. Subtracting the area term in (5.96) therefore gives

$$2\pi \int_{\Sigma} d\Sigma_\mu \xi_\nu \delta\langle T^{\mu\nu} \rangle = O(\mathcal{B}_{\text{belt}}).$$

Because the belt deformations include arbitrary small, boost-adapted variations of the cut with localized generator weights, the c-function/GSL infrastructure and quasi-local positive-energy control ensure that nonnegative smearings along the belt are dense, so the above equality holds for a separating family of test functions on  $D[R]$  (see Sections 5.36 and 6.4 and Corollary 5.79). Invoking the Iyer–Wald identity in the bulk, the same smearings pair the linearized gravitational tensor against  $\xi$  to the negative of the above matter flux; hence

$$\int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \langle E^{(1)\mu\nu}(h) - 8\pi G \delta T^{\mu\nu} \rangle = O(\mathcal{B}_{\text{belt}}).$$

By separation, this implies  $\langle E^{(1)}_{\mu\nu}(h) - 8\pi G \delta T_{\mu\nu} \rangle = 0$  as a distribution on  $D[R]$ , i.e. the linearized Einstein equations hold in expectation on  $D[R]$ . This is precisely the “EE  $\Rightarrow$  SEE” content of Theorems 5.37 and 5.42.

**(2) SEE + JKM fix  $\Rightarrow$  EE.** Assume the linearized Einstein equations in expectation,  $\langle E^{(1)}_{\mu\nu} - 8\pi G \delta T_{\mu\nu} \rangle = 0$  on  $D[R]$ . Integrating the Iyer–Wald identity on  $\Sigma$  and using the JKM calibration on the two belt corners yields

$$\int_{\partial\Sigma} (\delta Q_{\xi} - \xi \cdot \Theta) = 2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta \langle T^{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}),$$

while the calibrated corner piece equals  $\delta[\text{Area}/(4G)]$  up to  $O(\mathcal{B}_{\text{belt}})$ . Combining with the Brown–York flux identity on the timelike belt, one reassembles (5.96); see Sections 5.50 and 6.2 and Proposition 5.78. Therefore

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

and the first law with calibrated edge/corner terms implies

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

This is the desired EE statement. See also Theorem 5.85 for the recorded IFF formulation.

*Removal of positive flows.* In both directions the  $O(\mathcal{B}_{\text{belt}})$  remainder vanishes as the positive flows are removed at fixed belt width by Lemma 3.3, yielding regulator-independent statements on the belt domain.  $\square$

## 5.57 Data processing and recoverability at belt level

*Proposition 5.86* (DPI with recovery). For any belt-compatible channel  $\Phi$  (belt coarse-graining) and reference  $\sigma$ , 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, within the OS kernel the deficit

$$\Delta_{\Phi}(\rho\|\sigma) := S(\rho\|\sigma) - S(\Phi\rho\|\Phi\sigma)$$

is  $O(\mathcal{B}_{\text{belt}})$  uniformly on the admissible class and vanishes upon removal of the positive flows. **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.

*Proof. Step 1 (setup on the belt algebra).* Fix a belt of width  $r > 0$  and positive flows  $(u, s) > 0$  in the OS window. All objects below live on the belt-local von Neumann algebra  $\mathcal{A}_R$  at these

finite regulators; the admissible states  $\rho, \sigma$  are faithful on  $\mathcal{A}_R$  by construction. Belt compatibility means  $\Phi : \mathcal{A}_R \rightarrow \mathcal{A}_R$  is CPTP and respects the belt localization (coarse-graining on the belt).

*Step 2 (rotated Petz recovery and refined DPI).* Let  $\Phi^\dagger$  denote the adjoint w.r.t. the Hilbert–Schmidt pairing on the belt GNS space. Define the rotated Petz family relative to  $\sigma$ ,

$$\mathcal{R}_{\Phi, \sigma}^{(t)}(X) := \sigma^{\frac{1+it}{2}} \Phi^\dagger \left( \Phi(\sigma)^{-\frac{1+it}{2}} X \Phi(\sigma)^{-\frac{1-it}{2}} \right) \sigma^{\frac{1-it}{2}},$$

and its  $\beta$ -averaged version

$$\mathcal{R}_{\Phi, \sigma}(X) := \int_{\mathbb{R}} \beta(t) \mathcal{R}_{\Phi, \sigma}^{(t)}(X) dt,$$

where  $\beta$  is the standard probability density used in the refined DPI (its explicit form is immaterial here). By construction,  $\mathcal{R}_{\Phi, \sigma}$  is CPTP and belt-compatible. The refined data processing inequality then gives, for all admissible  $\rho$ ,

$$S(\rho \|\sigma) - S(\Phi(\rho) \|\Phi(\sigma)) \geq -2 \log F(\rho, \mathcal{R}_{\Phi, \sigma} \circ \Phi(\rho)) \geq 0.$$

This proves the displayed two-sided bound with  $R = \mathcal{R}_{\Phi, \sigma}$  on the belt algebra.

*Step 3 (OS kernel:  $O(\mathcal{B}_{\text{belt}})$  deficit and removal of flows).* Inside the OS kernel there exists a belt-compatible (rotated) Petz recovery map  $R$  for which both recoverability and relative-entropy continuity hold with remainder controlled by the single belt budget  $\mathcal{B}_{\text{belt}}$  (belt recovery/continuity, Lemma 3.2). In particular, for the admissible class one has

$$\|R \circ \Phi - \text{id}\|_{1 \rightarrow 1} \leq C_{\text{cb}} \mathcal{B}_{\text{belt}}, \quad 1 - F(\omega, R \circ \Phi(\omega)) \leq C_F \mathcal{B}_{\text{belt}}$$

uniformly in  $\omega$  in the class. Consequently  $-2 \log F(\rho, R \circ \Phi(\rho)) = O(\mathcal{B}_{\text{belt}})$ , and thus the refined DPI lower bound is  $O(\mathcal{B}_{\text{belt}})$ .

To bound the deficit from above by  $O(\mathcal{B}_{\text{belt}})$ , apply monotonicity under the additional CPTP map  $R$ :

$$S(\Phi \rho \|\Phi \sigma) \geq S(R \circ \Phi \rho \|\ R \circ \Phi \sigma),$$

whence

$$0 \leq \Delta_{\Phi}(\rho \|\sigma) \leq S(\rho \|\sigma) - S(R \circ \Phi \rho \|\ R \circ \Phi \sigma).$$

By belt-level continuity of relative entropy under near-identity channels (again controlled by  $\mathcal{B}_{\text{belt}}$  in the OS window), the right-hand side is  $O(\mathcal{B}_{\text{belt}})$  uniformly on the admissible class. Therefore  $\Delta_{\Phi}(\rho \|\sigma) = O(\mathcal{B}_{\text{belt}})$  at fixed regulators.

Finally, by the positive-flow removal lemma Lemma 3.3,  $\mathcal{B}_{\text{belt}} \rightarrow 0$  as  $(u, s) \downarrow 0$  at fixed belt width  $r$ , so both the refined DPI remainder and the deficit itself vanish in that limit. This completes the proof.  $\square$

## 5.58 BRST/diffeomorphism dressing invariance

*Lemma 5.87* (anchor-preserving diffeos). Let  $\delta_{\xi}$  be an anchor-preserving diffeomorphism generated by a vector field  $\xi$  tangent to the belt. For any observable  $\mathcal{O}$ ,

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

so physical predictions are invariant under such dressings up to  $O(\mathcal{B}_{\text{belt}})$ , which vanishes after removal.

*Proof.* By functoriality of the net and the belt diagram, any anchor-preserving diffeomorphism acts covariantly on the algebras and restricts to an automorphism of the belt factor; see Propositions 2.4 and 2.6 and Lemma 2.5. The corresponding one-parameter flow  $\alpha_{t\xi}$  on observables is strongly continuous, and its infinitesimal generator is the densely defined derivation

$\delta_\xi \mathcal{O} := \frac{d}{dt} \Big|_{t=0} \alpha_{t\xi}(\mathcal{O})$ . On the common core used throughout Section 5, the Ward-identity/KLMN construction yields an essentially selfadjoint charge  $\mathcal{Q}_{\text{diff}}[\xi]$  implementing this derivation,

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

see Lemmas 5.12 and 5.24 and Proposition 5.67.

Because  $\xi$  is tangent to the belt and preserves the anchors, the charge may be localized to the belt by multiplying its current with a cutoff that equals 1 on the belt and vanishes outside a slightly thicker tubular neighborhood. The difference between the full and belt-localized charges is supported away from the belt and therefore (micro)causally decouples from any observable after passing through the belt cocycle/factorization, cf. Remark 5.36, Lemmas 5.48 and 5.74, and Proposition 5.75. Quantitatively, the commutator of the belt-localized piece with  $\mathcal{O}$  is suppressed by the belt thickness, so

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

Finally, the regulator removal established in Lemma 3.3 implies that  $O(\mathcal{B}_{\text{belt}}) \rightarrow 0$  as the belt is removed. Since the diffeomorphism is anchor-preserving, there is no anchor motion to compensate; this matches the general anchor-invariance principle in Proposition 5.49. The claim follows.  $\square$

## 5.59 Belt-width RG flow and an integrated monotonicity

*Theorem 5.88* (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 5.17,

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

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

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

*Proof. Step 1 (shell representation of the  $c$ -density).* Write  $S_{\text{gen}} := S - \text{Area}/(4G)$  and  $\mathbf{c}_r = \partial_r S_{\text{gen}}$ . By the shell/CMI representation of the belt  $c$ -function (see Theorem 5.41 inside Section 5.17),

$$\mathbf{c}_r = \lim_{\delta \downarrow 0} \frac{1}{\delta} I_\rho(A : C | B_{r,\delta}) + R_1(r), \quad (5.97)$$

where  $B_{r,\delta}$  is a thin radial shell,  $A := \partial_{<r} R$  and  $C := \partial_{>r+\delta} R$ , and the remainder satisfies  $|R_1(r)| \leq C_{\text{shell}} \mathcal{B}_{\text{belt}}$  with  $C_{\text{shell}}$  belt-uniform by the remainder bookkeeping in Section 5.13. (Here the calibrated area term has already been absorbed using the modular equation of state and the JKM/BY dictionary.)

*Step 2 (width-thickening map and log-Sobolev contraction).* Let  $\mathcal{R}_{r \rightarrow r+dr}$  be the width-flow map that discards the outer shell  $B_{r,dr}$  and (anchor-preserving) recenters the belt. Within the OS/KMS framework and the  $c$ -function infrastructure of Section 5.17,  $\mathcal{R}_{r \rightarrow r+dr}$  is a belt-local CPTP map, KMS-reversible with respect to the same reference  $\sigma_r$  and inheriting the belt log-Sobolev/cluster rate  $\lambda_{\text{clu}} > 0$  (uniform per generator length and independent of  $|R|$ ). Applying  $\mathcal{R}_{r \rightarrow r+dr}$  to the tripartite state on  $AB_{r,\delta}C$  and using the log-Sobolev contraction of conditional mutual information gives

$$I_{r+dr}(A : C | B_{r+dr,\delta}) - I_r(A : C | B_{r,\delta}) \leq -\lambda_{\text{clu}} I_r(A : C | B_{r,\delta}) dr + R_2(r, \delta) dr + o(dr), \quad (5.98)$$

where  $|R_2(r, \delta)| \leq (C_{\text{align}} + C_{\text{tail}} + C_{\text{flow}}) \mathcal{B}_{\text{belt}} \delta$ . The remainder collects anchor-recentring/dressing, LR leakage across shell boundaries, and small flow-mismatch terms, all controlled by the budgets of Section 5.13.

*Step 3 (differential inequality for  $\mathbf{c}_r$ ).* Divide (5.98) by  $\delta$  and send  $\delta \downarrow 0$ . Using (5.97) and the uniform bound on  $R_2$ , we can absorb shell errors into a single bounded function  $R_4(r)$  with  $|R_4(r)| \leq C_4 \mathcal{B}_{\text{belt}}$ , obtaining

$$\mathbf{c}_{r+dr} - \mathbf{c}_r \leq -\lambda_{\text{clu}} \mathbf{c}_r dr + R_4(r) dr + o(dr),$$

and hence, letting  $dr \downarrow 0$ ,

$$\partial_r \mathbf{c}_r \leq -\lambda_{\text{clu}} \mathbf{c}_r + O(\mathcal{B}_{\text{belt}}). \quad (5.99)$$

*Step 4 (inhomogeneous Grönwall and the integrated bound).* Solve (5.99) on  $[r_1, r_2]$ :

$$\mathbf{c}_r \leq e^{-\lambda_{\text{clu}}(r-r_1)} \mathbf{c}_{r_1} + \int_{r_1}^r e^{-\lambda_{\text{clu}}(r-\rho)} R_4(\rho) d\rho.$$

Integrating this in  $r$  from  $r_1$  to  $r_2$  yields

$$\int_{r_1}^{r_2} \mathbf{c}_r dr \leq \frac{1 - e^{-\lambda_{\text{clu}}(r_2-r_1)}}{\lambda_{\text{clu}}} \mathbf{c}_{r_1} + \int_{r_1}^{r_2} dr \int_{r_1}^r e^{-\lambda_{\text{clu}}(r-\rho)} R_4(\rho) d\rho.$$

Belt Lieb–Robinson/factorization tails (as in the microcausality and remainder analysis of Section 5.13) imply that the double integral over  $R_4$  is bounded by a constant multiple of  $\mathcal{B}_{\text{belt}}$ , independent of  $|R|$  and of  $(r_2 - r_1)$ . Thus

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

*Step 5 (removal of positive flows).* All  $O(\mathcal{B}_{\text{belt}})$  remainders originate from the positive-flow/belt regulator and the LR tails. By the flow-removal lemma Lemma 3.3,  $\mathcal{B}_{\text{belt}} \rightarrow 0$  as the auxiliary flows are removed, so the displayed bounds become exact in the regulator-independent window. This completes the proof.  $\square$

## 5.60 Quantum Fisher information and QNEC; coherent-state check

*Proposition 5.89 (QFI–QNEC link).* For belt shape deformations generated by a null unitary  $U(\theta)$ , the quantum Fisher information

$$F_Q := \partial_\theta^2 S(\rho \| \rho_\theta) \Big|_{\theta=0}$$

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 at quadratic order in the displacement amplitude (no covariance shift).

*Numeric check (our coherent pulse):* for the null Gaussian profile of Section 5.25 with

$$\phi(u) = A \exp\left[-\frac{(u-u_0)^2}{2\sigma^2}\right], \quad A = 10^{-2}, \quad \sigma = 3, \quad u_0 = \sigma,$$

one has

$$\begin{aligned} \int du \langle T_{kk} \rangle &= \frac{\sqrt{\pi}}{2\sigma} A^2 = \frac{\sqrt{\pi}}{6} \cdot 10^{-4} = 2.9540897515 \times 10^{-5}, \\ 2\pi \int du \langle T_{kk} \rangle &= \frac{\pi^{3/2}}{\sigma} A^2 = \frac{\pi^{3/2}}{3} \cdot 10^{-4} = 1.8561093322 \times 10^{-4}, \end{aligned}$$

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

*Proof.* Let  $\rho$  be a belt state and  $\rho_\theta := U(\theta)\rho U(\theta)^\dagger$  a one-parameter family generated by a null unitary supported on the belt with test profile  $f(u)$ . Consider the relative entropy

$$S(\rho\|\rho_\theta) = \text{tr}[\rho(\log \rho - \log \rho_\theta)],$$

and expand it to second order at  $\theta = 0$ . By definition,

$$F_Q = \partial_\theta^2 S(\rho\|\rho_\theta)\big|_{\theta=0}.$$

On the other hand, for any reference state  $\sigma$  on the belt, the second variation of the relative entropy  $S(\rho_\theta\|\sigma)$  along the same flow can be written in the standard modular split as

$$\partial_\theta^2 S(\rho_\theta\|\sigma)\big|_{\theta=0} = \mathcal{Q}[f; \rho, \sigma] - \int du f(u)^2 2\pi (\langle T_{kk}(u) \rangle_\rho - \langle T_{kk}(u) \rangle_\sigma),$$

where  $\mathcal{Q}[f; \rho, \sigma] \geq 0$  is the positive quadratic form coming from modular positivity / Petz recovery in the belt OS window (the ‘‘QNEC kernel’’) and the stress-tensor term uses the belt QNEC normalization from Sections 5.9 and 5.40 (see in particular Theorems 5.29 and 5.64). Specializing to  $\sigma = \rho$  one obtains

$$\partial_\theta^2 S(\rho_\theta\|\rho)\big|_{\theta=0} = \mathcal{Q}[f; \rho, \rho] \geq 0.$$

Now compare this with  $S(\rho\|\rho_\theta)$ . By the general symmetry relation between  $S(\rho\|\rho_\theta)$  and  $S(\rho_\theta\|\rho)$  for unitary families (the same modular kernel governs both, with relative sign prescribed by the Araki expansion), one finds that

$$F_Q = \partial_\theta^2 S(\rho\|\rho_\theta)\big|_{\theta=0} = \int du f(u)^2 2\pi \langle T_{kk}(u) \rangle_\rho + (\mathcal{Q}[f; \rho, \rho] - \partial_\theta^2 S(\rho_\theta)\big|_0).$$

The local QNEC on belts, in its smeared form proved in Sections 5.9 and 5.40, states precisely that

$$\partial_\theta^2 S(\rho_\theta)\big|_{\theta=0} \leq \mathcal{Q}[f; \rho, \rho],$$

so the bracket is nonnegative. This gives the claimed inequality

$$F_Q \geq \int du f(u)^2 2\pi \langle T_{kk}(u) \rangle_\rho.$$

For coherent states of a free field, all connected  $n$ -point functions coincide with those of the vacuum, and the only change under a coherent displacement is in the one-point function of the field, hence in  $\langle T_{kk} \rangle$ . In particular, the ‘‘covariance’’ part  $\mathcal{Q}[f; \rho, \rho]$  and the entropy Hessian  $\partial_\theta^2 S(\rho_\theta)\big|_0$  are independent of the coherent amplitude  $A$ . Thus, to quadratic order in  $A$  (the first order at which  $\langle T_{kk} \rangle$  shifts), the bracket  $\mathcal{Q}[f; \rho, \rho] - \partial_\theta^2 S(\rho_\theta)\big|_0$  vanishes and

$$F_Q = \int du f(u)^2 2\pi \langle T_{kk}(u) \rangle_\rho + O(A^4),$$

i.e. the bound is saturated at this order with no additional covariance contribution. This covers the coherent-state claim and the explicit Gaussian pulse of Section 5.25.  $\square$

## 5.61 Unruh/KMS normalization sanity check

*Lemma 5.90* (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.

*Proof.* By the Bisognano–Wichmann theorem, the Minkowski vacuum restricted to a Rindler wedge is a KMS state for the boost modular group at inverse temperature  $\beta = 2\pi$  with respect to the boost parameter. In our belt setup, the “belt boost” is precisely this Rindler boost restricted to the belt, and we adopt the same normalization of the generator. Hence the Unruh temperature for the belt boost is

$$T_U = \beta^{-1} = \frac{1}{2\pi},$$

which evaluates numerically to  $T_U = 0.1591549431$ .

The Rindler coherent-state witness introduced earlier has been normalized so that its modular response is

$$\delta\langle K_R \rangle = \frac{\pi}{2} A^2$$

for amplitude  $A$ , as computed explicitly in Section 5.25. Specializing to  $A = 10^{-2}$  gives

$$\delta\langle K_R \rangle = \frac{\pi}{2} \times 10^{-4} = 1.5707963268 \times 10^{-4},$$

which matches the quoted value and is consistent with the Unruh/KMS thermal scale  $\beta = 2\pi$  used throughout this Section.  $\square$

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

*Lemma 5.91* (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 5.27.

*Proof.* By the  $N = 3$  subtracted forward dispersion relation for the second  $s$ -derivative, the contribution of the high-energy region  $s' \geq S_{\text{cut}}$  to  $\partial_s^2 \Re \mathcal{A}^{(N)}(0, 0)$  is given by the tail of the corresponding dispersive integral. On the working cone we impose the Regge envelope with slope  $\alpha_R = 2$  and window factor

$$R(t) := \sqrt{1 + |t|/s_0}, \quad t \in [-0.20 s_0, 0],$$

so that  $R(t) \leq \sqrt{1.2}$  on this interval. The general high-energy tail estimate for the  $k = 2$  forward derivative with  $N = 3$  subtractions and  $\alpha_R = 2$  specializes to

$$|\partial_s^2 \Re \mathcal{A}^{(N)}(0, 0)|_{[20 s_0, \infty)} \leq \frac{1}{\pi} \frac{\sqrt{1.2}}{3(20 s_0)^3}.$$

(The factor  $1/(5 - \alpha_R) = 1/3$  comes from integrating an  $s'^{-4}$  tail, and the factor  $1/\pi$  is the standard prefactor in the even-parity subtracted dispersion relation.)

By definition  $\widehat{c}_{2,0} = s_0^3 c_{2,0}$  and, for  $N = 3$  forward dispersion, the coefficient is related to the second derivative by

$$\widehat{c}_{2,0} = \frac{s_0^3}{2} \partial_s^2 \mathfrak{R} \mathcal{A}^{(N)}(0,0).$$

Therefore the tail contribution to  $\widehat{c}_{2,0}$  satisfies

$$|\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq \frac{s_0^3}{2} \frac{1}{\pi} \frac{\sqrt{1.2}}{3(20s_0)^3} = \frac{1}{2} \cdot \frac{1}{\pi} \cdot \frac{\sqrt{1.2}}{3(20)^3},$$

which yields the stated numerical value.

For the quadrature error we apply the standard Peano–kernel remainder for the  $n = 4$  Gauss–Radau rule on a single panel  $[s_0, 20s_0]$  to the (four-times differentiable) subtracted integrand  $g$ :

$$|\text{Quad}_{n=4}| \leq \frac{1}{1080} \sup_{[s_0, 20s_0]} |g^{(4)}| (20s_0 - s_0)^5 = \frac{1}{1080} \sup_{[s_0, 20s_0]} |g^{(4)}| (19s_0)^5.$$

Multiplying again by the factor  $s_0^3/2$  that relates  $\widehat{c}_{2,0}$  to  $\partial_s^2 \mathfrak{R} \mathcal{A}^{(N)}(0,0)$  gives

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

as claimed. By construction both contributions are ledged against the dispersion share of the global  $\varepsilon$ –budget recorded in Section 5.27.  $\square$

### 5.63 Interchanging variation with flow removal

*Lemma 5.92* (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 5.74, Propositions 5.75 and 5.78, and Section 5.50. Then for any admissible variation  $\delta$ ,

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

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 5.13.

*Proof.* Fix an admissible one-parameter variation  $\lambda \mapsto (\text{state}(\lambda), \text{geometry}(\lambda))$  with  $\lambda \in (-\lambda_0, \lambda_0)$  and write

$$\delta \mathcal{O}_{u,s}(R) = \frac{d}{d\lambda} \Big|_{\lambda=0} \mathcal{O}_{u,s}(R; \lambda), \quad \mathcal{O}_{u,s}(R; 0) = \mathcal{O}_{u,s}(R).$$

By the OS removal lemma Lemma 3.3 and the belt remainder ledger in Section 5.13, for every fixed  $\lambda$  and all  $(u, s)$  in the OS window we have

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

with  $C_{\mathcal{O}}(\lambda)$  depending continuously on  $\lambda$  but independent of  $|R|$ . The AGSP/seed converter and belt microcausality/timeslice control (Lemma 5.74, Propositions 5.75 and 5.78, and Section 5.50) imply that, for admissible variations, the constants  $C_{\mathcal{O}}(\lambda)$  can be chosen uniformly on a small interval  $|\lambda| \leq \lambda_0$ :

$$C_{\mathcal{O}}(\lambda) \leq C \quad \text{for all } |\lambda| \leq \lambda_0,$$

with  $C$  depending only on the observable type and on the belt/AGSP choices in Section 5.13.

Thus, for  $|\lambda| \leq \lambda_0$  and all  $(u, s)$  in the OS window,

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

with a bound independent of  $\lambda$ . In particular, the map

$$(u, s, \lambda) \longmapsto \mathcal{O}_{u,s}(R; \lambda) - \mathcal{O}_{0,0}(R; \lambda)$$

is jointly continuous in  $\lambda$  and admits a uniform (in  $\lambda$ ) envelope given by the right-hand side.

By the framework assumptions Section 2 and the OS positivity/recovery lemmas Lemmas 3.1 and 3.2, each  $\mathcal{O}_{u,s}(R; \lambda)$  admits a belt-local integral representation (modular energy flux, entropy density, Brown–York flux, or dispersive integral over the absorptive part) with kernels that are uniformly dominated by an integrable envelope on the belt. The admissible variation  $\delta$  acts by inserting  $\delta$  of the state or geometry into the same belt-local functionals. The AGSP/seed pipeline controls the corresponding variation norms uniformly in  $(u, s)$ , so the family  $\{\mathcal{O}_{u,s}(R; \lambda)\}_{(u,s)}$  is differentiable in  $\lambda$  at  $\lambda = 0$  with derivatives  $\delta\mathcal{O}_{u,s}(R)$  uniformly bounded in  $(u, s)$ .

We can therefore apply the dominated convergence theorem (in the relevant belt-local representation) to interchange the limit  $(u, s) \downarrow 0$  with the  $\lambda$ -derivative at 0:

$$\delta \lim_{(u,s) \downarrow 0} \mathcal{O}_{u,s}(R) = \left. \frac{d}{d\lambda} \right|_{\lambda=0} \lim_{(u,s) \downarrow 0} \mathcal{O}_{u,s}(R; \lambda) = \lim_{(u,s) \downarrow 0} \left. \frac{d}{d\lambda} \right|_{\lambda=0} \mathcal{O}_{u,s}(R; \lambda) = \lim_{(u,s) \downarrow 0} \delta\mathcal{O}_{u,s}(R).$$

Finally, differentiating the uniform removal estimate in  $\lambda$  at  $\lambda = 0$  (which is legitimate because the right-hand side does not depend on  $\lambda$ ) yields the stated error bound

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

again with constants uniform in  $|R|$  and in the admissible variation family. This proves the lemma.  $\square$

## 5.64 Entropy wedge nesting along positive flows

*Proposition 5.93* (EWN stability under flows). Let  $R_1 \subset R_2$  be belt-anchored regions. Assume belt-level nesting/recovery and the OS kernel. Then there exists  $\epsilon > 0$  such that for all  $u, s \in (0, \epsilon)$ ,

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

Moreover, the inclusion persists in the joint limit  $(u, s) \downarrow 0$ , up to a geometric error of order  $O(\mathcal{B}_{\text{belt}})$ , which vanishes under flow removal.

*Proof.* Fix belt-anchored regions  $R_1 \subset R_2$  and work in the OS window for the positive flows. For each pair  $(u, s)$  with  $0 < u, s < \epsilon_0$  in this window, the belt JLMS/OES channel associates to  $R_i$  an entropy wedge  $\text{EW}_{u,s}(R_i)$  and an isometric identification between the boundary algebra on  $R_i$  and the bulk algebra on  $\text{EW}_{u,s}(R_i)$ , with all constants uniform in  $|R_i|$ .

*Step 1: Monotonicity at fixed flows.* By belt-level nesting/recovery, the inclusion of boundary algebras

$$\mathcal{A}(R_1) \subseteq \mathcal{A}(R_2)$$

admits a completely positive, trace-preserving restriction map together with a uniform recovery channel on the OS analytic core. Functoriality of the belt channel and the data-processing/recovery control of Proposition 5.86 then imply that, at fixed  $(u, s)$ , any bulk operator that is reconstructible from  $R_1$  is also reconstructible from  $R_2$ . Equivalently, the corresponding bulk algebras satisfy

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

for all  $(u, s)$  in a small rectangle  $(0, \epsilon_0)^2$ , with constants uniform in  $|R_i|$ . By construction,  $\text{EW}_{u,s}(R)$  is the smallest wedge whose algebra contains the image of  $\mathcal{A}(R)$  under the belt channel, so this algebra inclusion is equivalent to the geometric nesting

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

for all  $0 < u, s < \epsilon_0$ . After shrinking  $\epsilon_0$  if necessary we obtain the first claim with  $\epsilon = \epsilon_0$ .

*Step 2: Passing to the regulator-free wedges.* To control the  $(u, s) \downarrow 0$  limit, we use belt microcausality and null timeslice propagation. The commutator tail and timeslice estimates of Lemma 5.74 and Proposition 5.75 imply that, for each fixed belt-anchored region  $R$ , the family of wedges  $\text{EW}_{u,s}(R)$  stabilizes away from the timelike belt as  $(u, s) \downarrow 0$ , and that the change in the wedge under a variation of  $(u, s)$  is confined to a collar of thickness  $O(\mathcal{B}_{\text{belt}})$  around the belt and the regulating QES surfaces.

Concretely, if  $K$  is a compact subset of the bulk that stays a positive distance from the belt, then for  $(u, s)$  small enough one has

$$K \cap \text{EW}_{u,s}(R) = K \cap \text{EW}(R),$$

where  $\text{EW}(R)$  denotes the entropy wedge defined after flow removal. This follows by localizing with test operators supported in small double-cones inside  $K$ , propagating these along the null timeslice flow from the belt, and using the exponential microcausality tails to control their dependence on  $(u, s)$ ; any change in wedge membership for points in  $K$  would contradict locality and the stability of the belt-regularized algebras.

Apply this with  $R = R_1$  and  $R = R_2$ . For any such compact  $K$  and all  $(u, s)$  sufficiently small,

$$K \cap \text{EW}(R_1) = K \cap \text{EW}_{u,s}(R_1) \subseteq K \cap \text{EW}_{u,s}(R_2) = K \cap \text{EW}(R_2).$$

Since  $K$  was arbitrary away from the belt collar, the only possible failure of nesting can occur inside a geometric neighborhood of the belt of thickness  $O(\mathcal{B}_{\text{belt}})$ . This is exactly the  $O(\mathcal{B}_{\text{belt}})$  geometric error stated in the proposition.

*Step 3: Flow removal.* The positive-flow removal lemma Lemma 3.3 identifies  $\text{EW}(R)$  with the joint limit  $(u, s) \downarrow 0$  of  $\text{EW}_{u,s}(R)$  and sends the belt remainder  $\mathcal{B}_{\text{belt}}$  to zero. In particular, the  $O(\mathcal{B}_{\text{belt}})$  collar in which nesting could fail shrinks to zero thickness under flow removal, and we obtain exact entropy-wedge nesting

$$\text{EW}(R_1) \subseteq \text{EW}(R_2)$$

in the regulator-free theory. This completes the proof.  $\square$

## 5.65 Edge-of-wedge analyticity on the dispersive cone

*Lemma 5.94 (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 in  $s$  to the 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 5.22.

*Proof.* Work on the near-forward cone  $\mathcal{S}$  of Section 5.10, where the scattering amplitude  $A(s, t)$  is constructed from the OS kernel as a boundary value of time-ordered correlators and satisfies locality and the spectral condition. Reflection positivity of the OS data, together with microcausality of the reconstructed Wightman fields, implies that for each fixed  $t$  with  $|t| \leq t_\star$  the amplitude  $A(s, t)$  is the boundary value, for real  $s$ , of a function that is holomorphic in  $s$  on

the upper and lower half-planes and whose boundary values agree on the real axis away from the physical cuts. This is the standard OS/Wightman edge-of-wedge analyticity for two-to-two amplitudes in the forward regime.

By the edge-of-wedge theorem, these upper and lower half-plane domains glue to a single holomorphic function of  $s$  on the forward tube

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

with the lower half-plane obtained by complex conjugation. The dependence on  $t$  is analytic on  $\mathcal{S}$  by locality and by construction of the working cone in Section 5.10, so the same tube  $\mathcal{T}$  is valid uniformly for all  $(s, t) \in \mathcal{S}$ .

Tester-certified Regge behavior on  $\mathcal{S}$  states that for real  $s$  in the physical region and  $|t| \leq t_\star$  one has

$$|A(s, t)| \leq C(t) (1 + |s|)^{\alpha_R},$$

with  $\alpha_R \leq 2 + \delta_\star < 3$ . Since  $A$  is holomorphic in  $\mathcal{T}$  and of at most polynomial growth on the real axis, the Phragmén–Lindelöf principle propagates this estimate to the whole tube: there exists a constant  $C'(\mathcal{S})$  such that

$$|A(s, t)| \leq C'(\mathcal{S}) (1 + |s|)^{\alpha_R} \quad \text{for all } (s, t) \in \mathcal{T}.$$

Now consider the gravity-subtracted,  $N$ -fold subtracted amplitude  $A^{(N)}(s, t)$  with  $N \geq 3$  as defined in Sections 5.22 and 5.46. Subtracting a polynomial in  $s$  of degree at most  $N - 1$  does not change analyticity on  $\mathcal{T}$  and preserves the same polynomial growth bound, so  $A^{(N)}$  is holomorphic on  $\mathcal{T}$  and obeys

$$|A^{(N)}(s, t)| \leq C''(\mathcal{S}) (1 + |s|)^{\alpha_R} \quad \text{for all } (s, t) \in \mathcal{T}.$$

The dispersive representation and the analytic projectors of Sections 5.10 and 5.22 are built from the absorptive part along the physical cut with kernel weights of the form  $s'^{-(N+1)}(s' - s)^{-1}$  and fixed powers of  $s'$ . The exponent bound  $\alpha_R < 3$  together with  $N \geq 3$  implies that the product of these kernels with the absorptive profile is integrable at large  $s'$  and defines holomorphic functions of  $s$  on  $\mathcal{T}$  by dominated convergence. Thus all dispersive projectors and working windows are compatible with tube analyticity and inherit the same polynomial bound in  $s$ .

Finally, gravity-induced infrared contributions yield universal nonanalytic terms of the type described in Section 5.46: soft poles in  $t$  and  $s^2 \log s$  pieces with coefficients smooth in  $t$  on the cone. These pieces are explicitly removed in the definition of the gravity-subtracted amplitude and in the analytic projector, so the remaining hard part coincides on the physical line with the holomorphic function produced by the dispersive construction. Therefore the gravity-subtracted,  $N \geq 3$  subtracted amplitude extends to an analytic function of  $s$  on  $\mathcal{T}$ , with polynomial growth controlled by the Regge exponent  $\alpha_R$ , and is compatible with all dispersive projectors and windows used in the construction.

This is precisely the claim of the lemma. □

**Add-on (tube analyticity on the strip).** For the strip-renormalized dispersion relations of Section 5.22, one inserts additional weights  $(s'/s_0)^\sigma$  under the  $s'$ -integral, with  $\sigma$  in the off-principal strip. These factors are entire in  $\sigma$  and smooth in the real dispersion variable  $s'$ , with at most polynomial growth for  $s'$  in the working cone. They therefore preserve the integrability and dominated-convergence bounds used in the proof of Lemma 5.94, and the resulting weighted amplitudes enjoy the same tube analyticity and polynomial boundedness on  $\mathcal{T}$  as in Lemma 5.94.

## 5.66 Rescaling invariance with respect to the subtraction scale

*Lemma 5.95* (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  $N=3$ ,

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

*Proof.* Fix a subtraction scale  $s_0$ . By the analytic forward projector construction of Section 5.70, the forward coefficient  $c_{2,0}$  at scale  $s_0$  can be written as a linear functional of the absorptive part  $\rho(s')$  along the physical cut,

$$c_{2,0}|_{s_0} = \int_{s_0}^{\infty} ds' K_{2,0}(s'; s_0) \rho(s'), \quad (5.100)$$

with a kernel  $K_{2,0}(s'; s_0)$  that is homogeneous of degree  $-3$  in  $s_0$ . Concretely, there is a dimensionless profile  $\widetilde{K}_{2,0}$  such that

$$K_{2,0}(s'; s_0) = s_0^{-3} \widetilde{K}_{2,0}(s'/s_0),$$

so for any  $\alpha > 0$ ,

$$K_{2,0}(s'; \alpha s_0) = \alpha^{-3} K_{2,0}(s'; s_0). \quad (5.101)$$

Now change the subtraction scale from  $s_0$  to  $\tilde{s}_0 = \alpha s_0$  while keeping the absorptive part  $\rho(s')$  fixed. The  $N=3$  dispersive representation at the new scale has the form

$$c_{2,0}|_{\tilde{s}_0} = \int_{\tilde{s}_0}^{\infty} ds' K_{2,0}(s'; \tilde{s}_0) \rho(s') + c_{2,0}^{(\text{poly})}(s_0, \tilde{s}_0),$$

where  $c_{2,0}^{(\text{poly})}$  is the contribution of the subtraction polynomial.

First, using (5.101) and extending the lower limit back to  $s_0$  (which only changes the integral by a finite piece that can be absorbed into the subtraction polynomial), we obtain

$$\int_{\tilde{s}_0}^{\infty} ds' K_{2,0}(s'; \tilde{s}_0) \rho(s') = \alpha^{-3} \int_{\tilde{s}_0}^{\infty} ds' K_{2,0}(s'; s_0) \rho(s') = \alpha^{-3} c_{2,0}|_{s_0} + c_{2,0}^{(\text{cut})}(s_0, \tilde{s}_0),$$

with  $c_{2,0}^{(\text{cut})}$  again a finite correction supported between  $s_0$  and  $\tilde{s}_0$ .

Second, by standard dispersion theory at  $N=3$ , the difference between using subtraction point  $s_0$  and  $\tilde{s}_0$  is encoded in an analytic polynomial  $P_2(s; s_0, \tilde{s}_0)$  of degree at most two in  $s$ ; at the level of the forward projector this gives a contribution  $c_{2,0}^{(\text{poly})}(s_0, \tilde{s}_0)$  coming entirely from  $P_2$ . By construction, the even-parity forward derivatives and the analytic projector annihilate any such polynomial of degree  $\leq 2$  (see Section 5.46), so

$$c_{2,0}^{(\text{poly})}(s_0, \tilde{s}_0) + c_{2,0}^{(\text{cut})}(s_0, \tilde{s}_0) = 0.$$

Therefore the only surviving effect of changing  $s_0$  to  $\tilde{s}_0$  in the dispersive projector is the homogeneous scaling (5.101), and we conclude

$$c_{2,0}|_{\tilde{s}_0} = \alpha^{-3} c_{2,0}|_{s_0}.$$

Finally, for the dimensionless combination  $\widehat{c}_{2,0} := s_0^3 c_{2,0}$  we have

$$\widehat{c}_{2,0}|_{\tilde{s}_0} = \tilde{s}_0^3 c_{2,0}|_{\tilde{s}_0} = (\alpha s_0)^3 \alpha^{-3} c_{2,0}|_{s_0} = s_0^3 c_{2,0}|_{s_0} = \widehat{c}_{2,0}|_{s_0},$$

which proves the claimed invariance under rescaling of the subtraction scale.  $\square$

**Add-on (scale invariance unchanged).** Rescaling  $s_0$  is compensated in the dimensionless forward coefficients; the strip weights  $(s'/s_0)^\sigma$  used in the off-principal celestial strip preserve Lemma 5.95.

### 5.67 OSR/complexity inflation aggregator with numeric audit baseline

**Aggregator.** For AGSP step  $m$  and belt base  $\Gamma_{\text{belt}}$ , we package the unit-length OSR and complexity overhead into the per-cut inflation factor

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

Under the audit baseline bindings from the ledger this evaluates to

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

On the Trotter-safe branch we use  $\Upsilon(m) \geq 2m+1$ . In particular, at the Page-point baseline  $m=8$  used in Section 5.21,

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

so any future update of the baseline numbers for  $\Lambda_0$  or  $\Gamma_{\text{belt}}$  (or a different  $m$ ) propagates automatically. By construction,  $\mathcal{I}_{\text{OSR}}(m)$  is exactly the inflation factor that appears inside the per-length Page bound and it remains uniform in  $|R|$  within the single budget  $O(\mathcal{B}_{\text{belt}})$ .

### 5.68 Replica analyticity and belt-uniform von Neumann limit

*Lemma 5.96* (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 5.86, 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 5.12 and Proposition 5.86. The variation–limit interchange is justified by Section 5.63.

*Proof.* Fix a belt-anchored region  $R$  and work at finite belt regulator  $(r; u, s)$  in the OS window. By OS positivity and modular generation Lemma 3.1, the belt algebra for  $R$  at fixed  $(r; u, s)$  is represented on a Hilbert space  $\mathcal{H}_R$  by a von Neumann algebra with a faithful normal state whose density operator we denote by  $\rho_R$ . By construction,  $\rho_R \geq 0$ ,  $\|\rho_R\| \leq 1$ , and  $\text{tr} \rho_R = 1$ , with all bounds uniform per generator length and independent of  $|R|$ .

*Step 1: Analyticity of the partition function.* For  $\Re n > 1$  define

$$Z_R(n) := \text{tr} \rho_R^n.$$

Diagonalizing  $\rho_R$  on  $\mathcal{H}_R$  gives a spectral decomposition  $\rho_R = \sum_j \lambda_j |\psi_j\rangle\langle\psi_j|$  with  $0 < \lambda_j \leq 1$  and  $\sum_j \lambda_j = 1$ . For  $\Re n > 1$ ,

$$Z_R(n) = \sum_j \lambda_j^n, \quad \sum_j |\lambda_j^n| = \sum_j \lambda_j^{\Re n} \leq \sum_j \lambda_j = 1.$$

Thus  $Z_R$  is given by an absolutely and locally uniformly convergent series on the half-plane  $\Re n > 1$ , and hence is holomorphic there.

Item 5 of the framework recap Section 2, together with belt-level continuity/recovery Proposition 5.86, implies that the admissible OS states satisfy a uniform  $L^{1\pm\delta_0}$  integrability bound: there exists  $\delta_0 > 0$ , independent of  $|R|$ , such that

$$\sup_R \sup_{1-\delta_0 < \Re n < 1+\delta_0} \operatorname{tr} \rho_R^{\Re n} < \infty$$

per generator length. Equivalently,  $Z_R(n)$  extends as a finite holomorphic function to the vertical strip  $\{n : 1 - \delta_0 < \Re n < 1 + \delta_0\}$ . Shrinking  $\delta_0$  if necessary, we obtain  $\delta \in (0, \delta_0)$  such that  $Z_R(n)$  is holomorphic on  $\{n : |n - 1| < \delta\}$ , with the bound

$$\sup_{|n-1| < \delta} |Z_R(n)| \leq C_Z \tag{5.102}$$

for some  $C_Z < \infty$  depending only on the belt budgets and not on  $|R|$ .

Since  $Z_R(1) = \operatorname{tr} \rho_R = 1$ , continuity of  $Z_R$  and the uniform bound (5.102) imply that, after possibly reducing  $\delta$ , the image  $Z_R(\{n : |n - 1| < \delta\})$  lies inside a disk  $\{z : |z - 1| < 1/2\}$  that does not intersect the negative real axis or the origin. We can therefore choose a single analytic branch of the logarithm on this disk and define

$$S_n(R) = \frac{1}{1-n} \log Z_R(n)$$

for all  $n$  with  $|n - 1| < \delta$ ,  $n \neq 1$ . On this domain  $S_n(R)$  is holomorphic as the composition of holomorphic functions.

*Step 2: Von Neumann limit  $n \rightarrow 1$ .* Using the spectral decomposition of  $\rho_R$  we have, for  $\Re n$  in the above strip,

$$Z_R(n) = \sum_j \lambda_j^n, \quad Z'_R(n) = \sum_j \lambda_j^n \log \lambda_j,$$

where the derivative is obtained by termwise differentiation. The entropy budgets from Section 5.12 and Proposition 5.86 give a uniform bound

$$\sum_j \lambda_j |\log \lambda_j| = S(\rho_R) \leq C_S$$

per generator length, with  $C_S$  independent of  $|R|$ . Since  $|\lambda_j^n| \leq \lambda_j$  for  $\Re n \geq 1$ , dominated convergence applies and yields

$$Z'_R(1) = \sum_j \lambda_j \log \lambda_j.$$

Because  $Z_R(1) = 1$ , we can compute the limit of  $S_n(R)$  using L'Hôpital's rule:

$$\lim_{n \rightarrow 1} S_n(R) = \lim_{n \rightarrow 1} \frac{\log Z_R(n)}{1-n} = -Z'_R(1) = -\sum_j \lambda_j \log \lambda_j = S(\rho_R).$$

Thus the apparent singularity at  $n = 1$  is removable, and  $S_n(R)$  extends to a holomorphic function on  $\{n : |n - 1| < \delta\}$  with  $S_1(R) = S(\rho_R)$ . In particular,

$$\lim_{n \rightarrow 1^+} S_n(R) = S(\rho_R).$$

*Step 3: Uniform remainder bound.* By the previous step and the entropy budgets,  $S_n(R)$  is bounded on the circle  $|n - 1| = \delta$ :

$$\sup_{|n-1|=\delta} |S_n(R)| \leq C'_S$$

for some  $C'_S$  independent of  $|R|$  (per generator length), again by Section 5.12 and Proposition 5.86. Cauchy's integral formula on the disk  $|n-1| < \delta$  then gives

$$\sup_{|n-1| \leq \delta/2} |\partial_n S_n(R)| \leq \frac{2}{\delta} \sup_{|n-1|=\delta} |S_n(R)| \leq \frac{2C'_S}{\delta} =: K_0,$$

with  $K_0$  independent of  $|R|$ . For real  $n$  with  $1 < n < 1 + \delta/2$  the mean-value theorem therefore yields

$$|S_n(R) - S_1(R)| = |n-1| |\partial_n S_\xi(R)| \leq K_0 |n-1|$$

for some  $\xi$  between 1 and  $n$ . Recalling that  $S_1(R) = S(\rho_R)$ , we obtain, at fixed belt regulator  $(r; u, s)$ ,

$$|S_n(R) - S(\rho_R)| \leq K_0 |n-1|.$$

*Step 4: Belt-uniformity and flow removal.* The constants  $C_Z$ ,  $C'_S$ , and hence  $K_0$  above are controlled solely by the belt budgets and recovery constants tracked in Section 5.12 and Proposition 5.86, and therefore do not depend on  $|R|$ . OS removal Lemma 3.3 applied to the observable  $S_n(R)$  and to  $S(\rho_R)$  shows that

$$|S_n(R) - S(\rho_R)| \leq K |n-1| + O(\mathcal{B}_{\text{belt}}),$$

for some  $K$  with the same dependence on the budgets as  $K_0$  and independent of  $|R|$ , which is the claimed remainder bound.

*Step 5: Variation along shape paths.* Let  $\lambda \mapsto R_\lambda$  be a belt-anchored shape path. By belt-level continuity/recovery Proposition 5.86 and OS recovery Lemma 3.2, the map  $\lambda \mapsto \rho_{R_\lambda}$  is continuous in trace norm with uniform control of the entropic budgets, so that the preceding arguments can be carried out uniformly in  $\lambda$  on compact intervals. In particular,  $(n, \lambda) \mapsto S_n(R_\lambda)$  is holomorphic in  $n$  for  $|n-1| < \delta$  and  $C^1$  in  $\lambda$ , with  $\partial_\lambda S_n(R_\lambda)$  uniformly bounded on that set.

For each fixed  $\lambda$  we have already shown that  $\lim_{n \rightarrow 1^+} S_n(R_\lambda) = S(\rho_{R_\lambda})$ . Differentiating this identity with respect to  $\lambda$  inside the belt-regularized theory and using dominated convergence gives

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

at fixed  $(r; u, s)$ . Finally, the interchange between  $\lambda$ -variation and the removal of positive flows is justified by the general variation–limit interchange result Section 5.63, which applies to both  $S_n$  and  $S$  as belt-regularized observables. This yields the claimed identity for the belt-removed entropies and completes the proof.  $\square$

## 5.69 Quantum Bousso bound from QNEC and the belt GSL

*Theorem 5.97* (quantum Bousso bound on belts). Let  $\lambda$  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.<sup>4</sup> Assume the QNEC of Section 5.9 and the belt GSL of Section 5.36. 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 5.97.* 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  $a' = \theta a$ .

<sup>4</sup>Equivalently,  $A$  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.

*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  $(\lambda_1, \lambda_2)$  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

$$f_\epsilon(\lambda) := g_{\epsilon, 2}(\lambda) - g_{\epsilon, 1}(\lambda), \quad -f_\epsilon'' \xrightarrow{\epsilon \rightarrow 0} \delta(\lambda - \lambda_2) - \delta(\lambda - \lambda_1),$$

we obtain

$$\Delta S_{\text{out}} := S_{\text{out}}(\lambda_2) - S_{\text{out}}(\lambda_1) \leq 2\pi \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda + o(1), \quad (5.103)$$

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  $f_\epsilon(\lambda_i) = 0$ ,

$$8\pi G \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda = \int_{\lambda_1}^{\lambda_2} f_\epsilon' a \theta d\lambda + \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda. \quad (5.104)$$

*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 (5.104) yields

$$8\pi G \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda = (A(\lambda_1) - A(\lambda_2)) + [A f_\epsilon']_{\lambda_1}^{\lambda_2} + \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda. \quad (5.105)$$

*Step 4: Canonical-energy control of shear and edge pieces.* By the Raychaudhuri estimate with canonical-energy control (Proposition 5.102) and the shear-control lemma (Lemma 5.112), 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 5.46, with the “edge/corner” calibration fixed by the JKM prescription (Section 5.49, Lemma 5.77, and Proposition 5.40). Equivalently,

$$[A f'_\epsilon]_{\lambda_1}^{\lambda_2} + \frac{1}{2} \int_{\lambda_1}^{\lambda_2} f_\epsilon a \theta^2 d\lambda - \int_{\lambda_1}^{\lambda_2} f_\epsilon a \sigma^2 d\lambda \leq O(\mathcal{B}_{\text{belt}}). \quad (5.106)$$

Insert (5.106) into (5.105) to obtain

$$8\pi G \int_{\lambda_1}^{\lambda_2} a T_{kk} f_\epsilon d\lambda \leq A(\lambda_1) - A(\lambda_2) + O(\mathcal{B}_{\text{belt}}). \quad (5.107)$$

The monotone c-function/belt GSL (Theorems 5.41 and 5.58) underlies (5.106): 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 (5.103) with (5.107) 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 5.98* (Rindler benchmark). From Section 5.25,  $\Delta S_{\text{Rindler}} \leq \Delta \langle K_{\text{R}} \rangle = \frac{\pi}{2} A^2 = 1.5707963268 \times 10^{-4}$  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).

## 5.70 Differentiation under the dispersive integrals

*Lemma 5.99* (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

$$\partial_s^k \mathfrak{R} \mathcal{A}^{(N)}(0, t) = \frac{1}{\pi} \int_{s_0}^{\infty} \frac{ds'}{s'^{N+k+1}} \left( \mathfrak{S} \mathcal{A}_{\text{hard}}(s', t) + (-1)^{N+k} \mathfrak{S} \mathcal{A}_{\text{hard}}(-s' - t, t) \right),$$

for all integers  $k \geq 0$ . Moreover, for  $k = 2$  and  $S_{\text{cut}} = 20 s_0$  the tail of the dispersive integral obeys

$$|\text{Tail}_{S_{\text{cut}}}| \leq \mathcal{R}(t) \frac{S_{\text{cut}}^{\alpha_{\text{R}} - 5}}{5 - \alpha_{\text{R}}} = \left( \frac{3}{3 - \delta} 20^\delta \right) \frac{\sqrt{1.2}}{3(20 s_0)^3} \quad (\alpha_{\text{R}} = 2 + \delta).$$

At  $\delta = 0$  this reproduces the audit of Section 5.54.

*Proof.* Fix  $t \leq 0$  in the cone  $\mathcal{S}$ . By the  $N = 3$  subtracted dispersion relation in the  $s$ -channel, the real part of the renormalized amplitude admits a representation of the form

$$\Re \mathcal{A}^{(N)}(s, t) = \frac{1}{\pi} \int_{s_0}^{\infty} ds' K_N(s, s'; t) \left( \Im \mathcal{A}_{\text{hard}}(s', t) + (-1)^N \Im \mathcal{A}_{\text{hard}}(-s' - t, t) \right),$$

where  $K_N(s, s'; t)$  is the standard rational kernel implementing the  $N$ -fold subtraction at  $s = 0$ , and is analytic in  $s$  for  $|s| \leq s_*$  with  $s_* < s_0$  and uniformly bounded in that region by

$$|\partial_s^k K_N(s, s'; t)| \leq C_{N,k} \frac{1}{s'^{N+k+1}},$$

for all  $k \geq 0$  and all  $s' \geq s_0$ , with constants  $C_{N,k}$  independent of  $s'$  and  $s$  in the chosen neighborhood.

The growth assumption on the absorptive part,

$$\Im \mathcal{A}_{\text{hard}}(s', t) \leq \mathcal{R}(t) s'^{\alpha_R}, \quad \Im \mathcal{A}_{\text{hard}}(-s' - t, t) \leq \mathcal{R}(t) s'^{\alpha_R},$$

with  $\alpha_R \leq 2 + \delta_* < 3$ , implies that for  $|s| \leq s_*$  we have the uniform bound

$$\begin{aligned} |\partial_s^k \Re \mathcal{A}^{(N)}(s, t)| &\leq \frac{1}{\pi} \int_{s_0}^{\infty} ds' |\partial_s^k K_N(s, s'; t)| \left( |\Im \mathcal{A}_{\text{hard}}(s', t)| + |\Im \mathcal{A}_{\text{hard}}(-s' - t, t)| \right) \\ &\leq \frac{2 C_{N,k} \mathcal{R}(t)}{\pi} \int_{s_0}^{\infty} ds' s'^{-N-k-1+\alpha_R}. \end{aligned}$$

For  $N = 3$  and any  $k \geq 0$  the exponent satisfies

$$-N - k - 1 + \alpha_R \leq -3 - k - 1 + 3 = -k - 1 < -1,$$

so

$$s'^{-N-k-1+\alpha_R} \in L^1([s_0, \infty), ds')$$

uniformly in  $k$ . Thus the integrand defining  $\partial_s^k \Re \mathcal{A}^{(N)}(s, t)$  is dominated by an  $s$ -independent  $L^1$  function of  $s'$ , and the usual differentiation-under-the-integral theorem (dominated convergence) applies. In particular,

$$\partial_s^k \Re \mathcal{A}^{(N)}(0, t) = \frac{1}{\pi} \int_{s_0}^{\infty} ds' \partial_s^k \left[ K_N(s, s'; t) \left( \Im \mathcal{A}_{\text{hard}}(s', t) + (-1)^N \Im \mathcal{A}_{\text{hard}}(-s' - t, t) \right) \right]_{s=0}.$$

Evaluating the  $k$ -th  $s$ -derivative of the rational kernel at  $s = 0$  produces precisely the factor  $s'^{-N-k+1}$  together with the alternating sign  $(-1)^{N+k}$  between the  $s$ - and  $u$ -channel absorptive parts. This is the standard Taylor expansion of the dispersion kernel around  $s = 0$ , and yields

$$\partial_s^k \Re \mathcal{A}^{(N)}(0, t) = \frac{1}{\pi} \int_{s_0}^{\infty} \frac{ds'}{s'^{N+k+1}} \left( \Im \mathcal{A}_{\text{hard}}(s', t) + (-1)^{N+k} \Im \mathcal{A}_{\text{hard}}(-s' - t, t) \right),$$

which is the claimed representation.

For the quantitative tail bound, we split the integral at  $S_{\text{cut}} > s_0$  and estimate the contribution from  $[S_{\text{cut}}, \infty)$ . Using the same growth bound on the absorptive parts and the explicit form of the kernel at  $s = 0$  we obtain

$$|\text{Tail}_{S_{\text{cut}}}| \leq \mathcal{R}(t) \int_{S_{\text{cut}}}^{\infty} ds' s'^{-N-k-1+\alpha_R}.$$

Specializing to  $N = 3$  and  $k = 2$  gives

$$|\text{Tail}_{S_{\text{cut}}}| \leq \mathcal{R}(t) \int_{S_{\text{cut}}}^{\infty} ds' s'^{\alpha_R-6} = \mathcal{R}(t) \frac{S_{\text{cut}}^{\alpha_R-5}}{5 - \alpha_R},$$

which is finite for  $\alpha_R < 5$  and, in particular, for  $\alpha_R \leq 2 + \delta_* < 3$ . Inserting the optimized bound on  $\mathcal{R}(t)$  obtained in the numerical audit and choosing

$$S_{\text{cut}} = 20 s_0, \quad \alpha_R = 2 + \delta,$$

one arrives at

$$|\text{Tail}_{S_{\text{cut}}}| \leq \left( \frac{3}{3 - \delta} 20^\delta \right) \frac{\sqrt{1.2}}{3 (20 s_0)^3},$$

which reproduces, at  $\delta = 0$ , the explicit tail estimate quoted in Section 5.54.  $\square$

### 5.71 CPT and crossing invariance of positivity testers

*Proposition 5.100* (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.

*Proof.* We discuss the three tester families in turn.

(i) *Forward even-parity derivatives.* A forward derivative tester at fixed  $t \leq 0$  is a linear functional built from even derivatives of the real part of the amplitude at  $s = 0$ ,

$$\mathcal{T}_{\text{fwd}}[A] = \sum_{n \geq 0} c_{2n}(t) \partial_s^{2n} \Re \mathcal{A}^{(N)}(0, t),$$

with real coefficients  $c_{2n}(t)$ . Under CPT the amplitude is mapped to its CPT conjugate, which coincides with complex conjugation at fixed Mandelstam invariants for scalar external states; since  $\mathcal{T}_{\text{fwd}}$  only involves the real part of  $A$  and real coefficients, it is invariant:  $\mathcal{T}_{\text{fwd}}[A^{\text{CPT}}] = \mathcal{T}_{\text{fwd}}[A]$ .

For crossing, at fixed  $t \leq 0$  the standard  $s \leftrightarrow u$  exchange acts as

$$(s, t, u) \mapsto (u, t, s), \quad u = -s - t.$$

Near the forward point one has  $u = -s$  at  $t = 0$ , so even derivatives in  $s$  at  $s = 0$  are invariant:

$$\partial_s^{2n} A(0, t) = \partial_u^{2n} A(0, t),$$

while odd derivatives change sign. Because the testers only involve even derivatives,  $\mathcal{T}_{\text{fwd}}$  is invariant under  $s \leftrightarrow u$ . In particular, if the forward derivatives are nonnegative before crossing, they remain so after crossing and CPT.

(ii) *Hankel/Gaussian band kernels.* The Hankel/impact-profile testers act on the absorptive part via a nonnegative radial kernel in impact space:

$$\mathcal{T}_{\text{H}}[A] = \int_0^\infty db b w(b) \rho(b; s, t),$$

where  $w(b) \geq 0$  is a fixed profile (Hankel or Gaussian band) depending only on the impact modulus  $b$ , and  $\rho(b; s, t) \geq 0$  is the corresponding absorptive density. CPT and crossing act on the partial-wave data by unitary transformations that preserve  $|\rho(b; s, t)|$  and the impact modulus  $b$ . Since  $w(b)$  depends only on  $b$  and is nonnegative, the integral defining  $\mathcal{T}_{\text{H}}$  is unchanged under these transformations, and its value remains nonnegative. Thus Hankel/Gaussian band testers are CPT-even and crossing-stable.

(iii) *Celestial Gram functionals.* Celestial Gram testers on the principal lines are defined as positive semi-definite Gram forms built from celestial transforms of the absorptive data. By construction they are quadratic forms of the type

$$\mathcal{T}_{\text{cel}}[A] = \sum_{i,j} \langle f_i, f_j \rangle_{\text{cel}} \alpha_i \bar{\alpha}_j,$$

with  $\langle \cdot, \cdot \rangle_{\text{cel}}$  the celestial inner product on the principal series and coefficients  $\alpha_i$  extracted linearly from the amplitude. Section 5.42 and Lemma 5.69 states that this Gram form is positive semi-definite and is preserved by the Ward action of the relevant symmetry group on the principal series; this action implements CPT and crossing on the celestial data without anomaly. It follows that  $\mathcal{T}_{\text{cel}}$  is invariant under CPT and under  $s \leftrightarrow u$ .

Finally, any conic combination of testers

$$\mathcal{T} = \sum_a \lambda_a \mathcal{T}_a, \quad \lambda_a \geq 0,$$

inherits invariance and nonnegativity: if each  $\mathcal{T}_a$  is CPT- and crossing-invariant and nonnegative, then so is  $\mathcal{T}$ . This completes the proof.  $\square$

**CPT–crossing stability with strip testers** The renormalized strip Gram functional is even in  $\sigma$  and built from the same absorptive profiles, hence is CPT-even and crossing-stable. Proposition 5.100 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 5.48 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 5.100. Thus nonforward/near–forward positivity holds uniformly on the widened window without increasing support.

## 5.72 Dimensional scope and hypotheses

*Lemma 5.101* (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 5.49 and Lemma 5.77 together with belt-level continuity/stability Proposition 5.86 and Section 5.14; all belt budgets remain uniform in  $|R|$ .

*Proof.* For  $d \geq 3$ , the framework axioms in Section 2 and the belt-local kernels compiled in Section 5 are formulated without dimension-specific input beyond the existence of smooth belt-anchored null cuts with finite extrinsic curvature. Under these axioms and kernels, the four pillars (QES/Page behavior, ANEC/QNEC, amplitude/positivity, and the modular equation of state) are each proved, and the associated estimates are per generator length and uniform in  $|R|$  by construction. Inspecting the proofs of these components shows that no further property of the reference spacetime dimension is used. Consequently, once the axioms hold in a given spacetime, the same arguments apply verbatim in any dimension  $d \geq 3$ , and the theorem suite extends from the reference dimension to all such  $d$  without further modification.

In  $d = 2$ , the scattering and celestial components only use the forward cone, the analytic projector, and Gram positivity on the principal series. These structures remain meaningful after restricting to the principal lines, so the tester envelope and budgets are unchanged at belt level. The only genuinely dimension-dependent pieces are the corner terms and factorization properties

entering the QES constructions. In two dimensions these are supplied by the 1+1-dimensional corner prescription and chiral factorization calibrated in Section 5.49 and Lemma 5.77, and they are propagated along belt-anchored deformations by the belt-level continuity/recovery and stability inputs Proposition 5.86 and Section 5.14. These results ensure that the QES existence, extremality, and stability arguments go through in  $d = 2$  with the same belt budgets, still uniform in  $|R|$ . No additional dimension-specific estimates are required beyond those already recorded in the cited inputs.  $\square$

### 5.73 Linearized Raychaudhuri area estimate with numeric line

*Proposition 5.102* (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 5.25 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 5.69), 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.

*Proof.* Work per generator length along a smooth belt-anchored null congruence with tangent  $k^\mu$  and affine parameter  $\lambda$ . Let  $\theta$  and  $\sigma_{ab}$  be the expansion and shear of the congruence, and assume vanishing twist (the generators are hypersurface orthogonal).

*Step 1: Linearized Raychaudhuri equation.* The Raychaudhuri equation for an affinely parametrized null congruence reads

$$\frac{d\theta}{d\lambda} = -\frac{1}{d-2} \theta^2 - \sigma_{ab} \sigma^{ab} - R_{\mu\nu} k^\mu k^\nu.$$

We consider a small perturbation about a background with  $\theta = 0 = \sigma_{ab}$  and vanishing null curvature along the belt; we denote perturbed quantities by the same symbols and drop terms that are quadratic in the perturbation. At linear order, the  $\theta^2$  and  $\sigma_{ab} \sigma^{ab}$  terms are second order and hence suppressed; the Raychaudhuri equation becomes

$$\frac{d\theta}{d\lambda} = -R_{kk} + O(\mathcal{B}_{\text{belt}}),$$

where  $R_{kk} := R_{\mu\nu} k^\mu k^\nu$  and the  $O(\mathcal{B}_{\text{belt}})$  term collects belt-regularization and higher-order effects.

The semiclassical Einstein equation in expectation, specialized to the null-null component, gives

$$\langle R_{kk} \rangle = 8\pi G \langle T_{kk} \rangle + O(\mathcal{B}_{\text{belt}}),$$

since the trace term in  $R_{\mu\nu} - 8\pi G T_{\mu\nu}$  drops out upon contraction with  $k^\mu k^\nu$ . Substituting, and working at linear order in the perturbation,

$$\frac{d\theta}{d\lambda} = -8\pi G \langle T_{kk}(\lambda) \rangle + O(\mathcal{B}_{\text{belt}}).$$

Integrating from  $\lambda_1$  to  $\lambda$  and using the initial condition  $\theta(\lambda_1) = 0$ ,

$$\theta(\lambda) = -8\pi G \int_{\lambda_1}^{\lambda} du \langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}),$$

which is the first claim.

*Step 2: Area variation and double integral.* Let  $A(\lambda)$  be the cross-sectional area (per generator length) and  $A_0 := A(\lambda_1)$  the background area at the entry point. By definition of the expansion,

$$\frac{dA}{d\lambda} = \theta(\lambda) A(\lambda).$$

Linearizing around  $A_0$  and using that  $\theta = O(\text{perturbation})$ , the product  $\theta (A(\lambda) - A_0)$  is second order and can be absorbed into  $O(\mathcal{B}_{\text{belt}})$ , so

$$\frac{d}{d\lambda} \Delta A(\lambda) = \frac{d}{d\lambda} (A(\lambda) - A_0) = \theta(\lambda) A_0 + O(\mathcal{B}_{\text{belt}}).$$

Integrating from  $\lambda_1$  to  $\lambda_2$  gives

$$\Delta A := A(\lambda_2) - A(\lambda_1) = \int_{\lambda_1}^{\lambda_2} \theta(\lambda) A_0 d\lambda + O(\mathcal{B}_{\text{belt}}),$$

which matches the formula displayed in the statement (up to the absorbed  $O(\mathcal{B}_{\text{belt}})$  term).

Substituting the expression for  $\theta(\lambda)$  and interchanging the order of integration,

$$\begin{aligned} \Delta A &= -8\pi G A_0 \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\lambda_1}^{\lambda} du \langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}) \\ &= -8\pi G A_0 \int_{\lambda_1}^{\lambda_2} du (\lambda_2 - u) \langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}), \end{aligned}$$

since for each fixed  $u \in [\lambda_1, \lambda_2]$  the  $\lambda$ -integration runs from  $\lambda = u$  to  $\lambda_2$ .

*Step 3: Linear bound.* On the segment  $[\lambda_1, \lambda_2]$  one has  $0 \leq \lambda_2 - u \leq L$  with  $L := \lambda_2 - \lambda_1$ , so

$$|\Delta A| \leq 8\pi G A_0 L \int_{\lambda_1}^{\lambda_2} du |\langle T_{kk}(u) \rangle| + O(\mathcal{B}_{\text{belt}}).$$

In particular, for states with nonnegative null energy density along the segment,  $\langle T_{kk}(u) \rangle \geq 0$  for all  $u \in [\lambda_1, \lambda_2]$ , this reduces to

$$\Delta A \leq 8\pi G A_0 L \int_{\lambda_1}^{\lambda_2} \langle T_{kk}(u) \rangle du + O(\mathcal{B}_{\text{belt}}),$$

which is the stated inequality. The coherent null pulse of Section 5.25 used in the numeric line has  $\langle T_{kk} \rangle \geq 0$  along the generator (for a free scalar,  $T_{kk}$  is quadratic in  $\partial_k \phi$ ), so the bound applies directly in that audit example. All regulator and higher-order effects are absorbed into  $O(\mathcal{B}_{\text{belt}})$  and vanish under flow removal.  $\square$

## 5.74 Null-curvature window on a belt

*Corollary 5.103* (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

$$-C_{\text{belt}} \mathcal{B}_{\text{belt}} \leq \int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk} \rangle \leq -\theta(\lambda_2).$$

*Proof.* All statements are per generator length and in the belt OS window; all  $O(\mathcal{B}_{\text{belt}})$  remainders are uniform in  $|R|$  and vanish upon removal of the positive flows.

(i) *Pointwise relation*. For a null congruence with generator  $k^\mu$  and expansion  $\theta$  the Raychaudhuri equation reads

$$\partial_\lambda \theta(\lambda) = -R_{kk}(\lambda) - \sigma_{ab}(\lambda) \sigma^{ab}(\lambda) - \frac{1}{d-2} \theta(\lambda)^2,$$

with  $R_{kk} := R_{\mu\nu} k^\mu k^\nu$  and  $\sigma_{ab}$  the shear. On a belt-anchored segment with  $\theta(\lambda_1) = 0$  and small shear at linear order, the quadratic terms  $\sigma_{ab} \sigma^{ab}$  and  $\theta^2$  are of order  $O(\mathcal{B}_{\text{belt}})$  in the perturbation. Thus, to first order,

$$\partial_\lambda \theta(\lambda) = -R_{kk}(\lambda) + O(\mathcal{B}_{\text{belt}}).$$

On the other hand, the first-order modular equation of state on belts, together with the belt JLMS channel and the Brown-York/Raychaudhuri dictionary, gives the linearized semiclassical Einstein equations in expectation on  $D[R]$  (see Section 5.73 and Proposition 5.102 and Theorems 5.37 and 5.46). Contracting with  $k^\mu k^\nu$  and using  $g_{\mu\nu} k^\mu k^\nu = 0$  yields, per generator,

$$\langle R_{kk}(\lambda) \rangle = 8\pi G \langle T_{kk}(\lambda) \rangle + O(\mathcal{B}_{\text{belt}}),$$

which is the claimed pointwise relation.

- (ii) *Integrated lower bound*. Integrating the relation in (i) over the segment  $[\lambda_1, \lambda_2]$  gives

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk}(\lambda) \rangle = 8\pi G \int_{\lambda_1}^{\lambda_2} d\lambda \langle T_{kk}(\lambda) \rangle + O(\mathcal{B}_{\text{belt}}),$$

where the  $O(\mathcal{B}_{\text{belt}})$  term is uniform per generator length. The belt ANEC Theorem 5.29 furnishes a lower bound on the stress integral,

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle T_{kk}(\lambda) \rangle \geq -C'_{\text{belt}} \mathcal{B}_{\text{belt}},$$

for some belt–uniform constant  $C'_{\text{belt}} > 0$  independent of  $|R|$ . Combining and enlarging the constant to absorb the uniform  $O(\mathcal{B}_{\text{belt}})$  remainder yields

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk}(\lambda) \rangle \geq -C_{\text{belt}} \mathcal{B}_{\text{belt}},$$

with  $C_{\text{belt}} > 0$  independent of  $|R|$ , as stated.

(iii) *Integrated focusing upper bound.* For the full (nonlinear) Raychaudhuri equation we have

$$\partial_\lambda \theta(\lambda) = -R_{kk}(\lambda) - \sigma_{ab}(\lambda) \sigma^{ab}(\lambda) - \frac{1}{d-2} \theta(\lambda)^2.$$

Integrating from  $\lambda_1$  to  $\lambda_2$  and using  $\theta(\lambda_1) = 0$  gives

$$\theta(\lambda_2) = - \int_{\lambda_1}^{\lambda_2} d\lambda \left[ R_{kk}(\lambda) + \sigma_{ab}(\lambda) \sigma^{ab}(\lambda) + \frac{1}{d-2} \theta(\lambda)^2 \right].$$

Rearranging,

$$\int_{\lambda_1}^{\lambda_2} d\lambda R_{kk}(\lambda) = -\theta(\lambda_2) - \int_{\lambda_1}^{\lambda_2} d\lambda \left[ \sigma_{ab}(\lambda) \sigma^{ab}(\lambda) + \frac{1}{d-2} \theta(\lambda)^2 \right].$$

Since  $\sigma_{ab} \sigma^{ab} \geq 0$  and  $\theta^2 \geq 0$  pointwise, we infer

$$\int_{\lambda_1}^{\lambda_2} d\lambda R_{kk}(\lambda) \leq -\theta(\lambda_2).$$

Taking expectations over the matter sector leaves the geometric quantities unchanged and yields

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk}(\lambda) \rangle = -\theta(\lambda_2) - \int_{\lambda_1}^{\lambda_2} d\lambda \left[ \sigma_{ab}(\lambda) \sigma^{ab}(\lambda) + \frac{1}{d-2} \theta(\lambda)^2 \right] \leq -\theta(\lambda_2),$$

which is the upper bound claimed in (iii). Together with (ii) this gives the stated window for the integrated null curvature.  $\square$

## 5.75 Approximate Markov gap on belts with recovery

*Theorem 5.104* (belt Markov gap). Let  $A:B:C$  be a belt-aligned tripartition with  $B$  covering the entangling belt. Under belt-level recoverability/continuity Proposition 5.86, 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 5.27.

*Proof.* Let  $\rho_{ABC}$  be the belt state for a belt-aligned tripartition  $A:B:C$  with  $B$  covering the entangling belt.

*Step 1: CMI as a DPI deficit and existence of a belt recovery map.* Consider the belt-compatible channel  $\Phi := \text{tr}_C$  that discards  $C$ , and the reference state  $\sigma_{ABC} := \rho_A \otimes \rho_{BC}$ . A standard relative-entropy identity shows that

$$I(A:C | B)_\rho = S(\rho_{ABC} \parallel \sigma_{ABC}) - S(\rho_{AB} \parallel \sigma_{AB}),$$

so  $I(A:C | B)$  is precisely the data-processing deficit for  $(\rho, \sigma, \Phi)$ .

By the refined DPI with recovery on belts Proposition 5.86, there exists a belt-compatible (rotated) Petz map  $R_{B \rightarrow BC}$  such that

$$S(\rho \|\sigma) - S(\Phi\rho \|\Phi\sigma) \geq -2 \log F(\rho_{ABC}, R_{B \rightarrow BC}(\rho_{AB})) \geq 0,$$

where  $F$  is the Uhlmann fidelity. For our choice of  $(\rho, \sigma, \Phi)$ , the left-hand side equals  $I(A:C | B)_\rho$ , so

$$I(A:C | B)_\rho \geq -2 \log F(\rho_{ABC}, R_{B \rightarrow BC}(\rho_{AB})). \quad (5.108)$$

The OS kernel Lemmas 3.1 and 3.2 guarantees that this construction is uniform on the admissible class and stable under belt-supported deformations, with all regulator effects tracked by  $\mathcal{B}_{\text{belt}}$ .

*Step 2: AGSP/flow errors and the recovery-term*  $C_{\text{rec}}(e^{-\mu_{\text{eff}} r(m)} + \eta^m)$ . In the OS window, the AGSP/flow converter provides an approximation of the exact modular/belt dynamics by a finite-step AGSP with step parameter  $m$  and associated radius  $r(m)$  and Trotter error  $\delta_m = \eta^m$ . The remainder ledger in the AGSP/flow pipeline gives, for suitable constants,

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

where  $E_{r(m)}$  is the belt-local error operator comparing the finite-step construction to the ideal dynamics and  $\mu_{\text{eff}} > 0$  is belt-uniform.

By continuity of the fidelity and of relative entropy under small perturbations (as encoded in the belt recovery/continuity kernel and the first-law channel, see Lemma 3.2 and the use of the OS window throughout), the deviation of  $\rho_{ABC}$  from its AGSP-ideal surrogate contributes at most a correction

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

for some belt-uniform constant  $C_{\text{rec}} > 0$ , where  $\rho^{(m)}$  denotes the state after passing through the AGSP/flow converter. Here we used standard trace-norm/fidelity bounds (Pinsker-type inequalities) to translate the fidelity defect in (5.108) and the AGSP error into a bound on the CMI.

*Step 3: Seed/OSR control and the length-suppressed term.* The AGSP/seed converter expresses  $\rho^{(m)}$  as a state supported on a seed Hilbert space of dimension at most  $\kappa_{\text{seed}}$  per generator, together with an OSR inflation factor

$$\Lambda = \Lambda_0 \Gamma_{\text{belt}} \Upsilon(m),$$

coming from the belt OS window and the Trotter-safe branch of the converter. In particular, any belt-anchored region  $R$  has entropy bounded by

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

as recorded in the Page/entropy ledger for the AGSP/seed converter.

For a tripartition  $A : B : C$  with  $B$  covering the entangling belt,  $I(A:C | B)$  is a linear combination of entropies of regions whose boundaries are contained in  $\partial R$ . Applying the above per-length entropy bound to each term and using subadditivity yields, for some belt-uniform constant  $C_{\text{seed}} > 0$ ,

$$I(A:C | B)_{\rho^{(m)}} \leq C_{\text{seed}} \frac{\log \kappa_{\text{seed}} + \log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{\text{length}(\partial R)} + O(\mathcal{B}_{\text{belt}}).$$

Combining this with the estimate from Step 2 gives

$$I(A:C | B)_\rho \leq C_{\text{rec}}(e^{-\mu_{\text{eff}} r(m)} + \eta^m) + C_{\text{seed}} \frac{\log \kappa_{\text{seed}} + \log(\Lambda_0 \Gamma_{\text{belt}} \Upsilon(m))}{\text{length}(\partial R)} + O(\mathcal{B}_{\text{belt}}).$$

*Step 4: Removal of positive flows and per-length gap.* All  $O(\mathcal{B}_{\text{belt}})$  remainders arise from the belt regulator and positive flows and vanish in the removal limit by Lemma 3.3. The constants  $C_{\text{rec}}$  and  $C_{\text{seed}}$  are per-generator and independent of  $|R|$ . Finally, the deterministic choice of  $m$  in Section 5.27 makes  $e^{-\mu_{\text{eff}}r(m)}$  and  $\eta^m$  as small as desired while the second term is suppressed by  $\text{length}(\partial R)$ , so that the conditional mutual information per unit length can be made arbitrarily small. This establishes the theorem.  $\square$

## 5.76 Soft theorem and celestial bridge compatibility

*Lemma 5.105 (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 5.22, 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.

*Proof.* Write the full gravity-subtracted  $2 \rightarrow 2$  amplitude as

$$A(s, t) = \mathcal{A}_{\text{soft}}(s, t) + \mathcal{A}_{\text{hard}}(s, t),$$

with  $\mathcal{A}_{\text{soft}}$  the universal soft piece and  $\mathcal{A}_{\text{hard}}$  the hard remainder. By construction of the IR scheme,  $\mathcal{A}_{\text{soft}}$  has the forward structure

$$\mathcal{A}_{\text{soft}}(s, t) = \kappa^2 \left( \frac{s^2}{-t} + \frac{u^2}{-t} \right) + (\text{nonanalytic soft pieces } s^2 \log |s|) + O(t),$$

where the  $1/t$  pole encodes  $t$ -channel graviton exchange and the  $s^2 \log |s|$  terms are the standard nonanalytic IR contributions.

*Step 1: Celestial transform of the soft factor.* Weinberg’s soft graviton theorem identifies  $\mathcal{A}_{\text{soft}}$  with the matrix element of the stress-tensor charge that generates an infinitesimal diffeomorphism on the external hard legs. When the amplitude is Mellin-transformed with the celestial kernel (the Mellin transform in the external energies at fixed angles), this soft charge maps to a fixed celestial weight (the stress-tensor weight) and acts by a differential operator on the celestial coordinates. Because  $\mathcal{A}_{\text{soft}}$  is a homogeneous polynomial of degree two in the hard invariants at fixed  $t$ , its Mellin transform is a finite linear combination of Beta functions with integer arguments, hence a polynomial in the celestial dimensions. Therefore the celestial transform of  $\mathcal{A}_{\text{soft}}$  produces a polynomial celestial profile which is precisely the Virasoro-type Ward operator acting on the hard celestial data, as recorded in Section 5.22.

*Step 2: Action of the analytic projector and removal of nonanalytic soft pieces.* The analytic projector of Section 5.22 is defined by a Cauchy integral around  $s = 0$  and projects onto  $s$ -analytic germs that are regular at  $t = 0$ . In particular, it annihilates the nonanalytic soft structures:

$$\Pi_{\text{an}}[s^2 \log |s|] = 0, \quad \Pi_{\text{an}}[\alpha(t)/t] = 0$$

for any smooth  $\alpha(t)$  on the forward cone. Thus, when extracting the analytic  $s^2$  coefficient and the associated kernels that feed into the dispersion and celestial testers, the projector removes exactly the nonanalytic  $s^2 \log |s|$  and  $1/t$  pieces originating from  $\mathcal{A}_{\text{soft}}$ . Equivalently, all tester-relevant analytic data can be computed either from  $A$  followed by projection, or directly from  $\mathcal{A}_{\text{hard}}$ , with the same result.

*Step 3: Gram positivity on the principal lines and restriction to the hard amplitude.* On the principal series  $\Delta = 1 + i\nu$  the celestial Mellin transform is unitary and the celestial Gram form is a positive semidefinite quadratic form in celestial profiles built from the absorptive part of the amplitude. The soft sector coming from  $\mathcal{A}_{\text{soft}}$  is the image of a stress-tensor charge with positive spectral density, so it contributes a nonnegative piece to the Gram form. The analytic projector removes the nonanalytic soft contributions from the forward dispersion data

but does not introduce any negative terms; in particular, Gram positivity on the principal lines is unchanged when passing from  $A$  to  $\mathcal{A}_{\text{hard}}$ .

Putting these steps together, the celestial transform of  $\mathcal{A}_{\text{soft}}$  yields a purely polynomial celestial Ward piece, while all tester-relevant nonanalytic soft structures are excised by the analytic projector. The celestial Gram form remains nonnegative on the principal series, and after subtraction all forward, Hankel, and celestial positivity testers act on  $\mathcal{A}_{\text{hard}}$  with the same nonnegativity properties as before. This establishes the lemma.  $\square$

**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- $\sigma$  pieces on the off-principal strip. These pieces are absorbed into the renormalized celestial measure and the belt-local counterterm  $\delta\mathcal{C}_{\text{cel}}$  (the strip Ward counterterms), and therefore never contribute with a negative sign to the renormalized strip Gram form. Consequently, all Gram testers remain nonnegative on the strip, and the soft–celestial consistency established on the principal lines extends to the finite off-principal window.

### 5.77 Stability under decoupled matter sectors

*Proposition 5.106* (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  $\kappa_{\text{ANEC}}, 2\pi$ ).

*Proof.* Because the additional matter sector is decoupled, the Hilbert space and local algebras factorize, and both the scattering amplitude and the stress tensor decompose as direct sums of sector contributions. On the amplitude side, the absorptive part obeys

$$\Im\mathcal{A}_{\text{tot}}(s + i0, t) = \Im\mathcal{A}_{\text{hard}}(s + i0, t) + \Im\mathcal{A}_{\text{hard}}^{(\text{mat})}(s + i0, t),$$

with each term nonnegative by unitarity on the working cone and gravity subtraction. Every forward, Hankel, or celestial tester  $T$  is a linear functional of the absorptive profile with nonnegative kernel by construction, so

$$T[\mathcal{A}_{\text{hard}} + \mathcal{A}_{\text{hard}}^{(\text{mat})}] = T[\mathcal{A}_{\text{hard}}] + T[\mathcal{A}_{\text{hard}}^{(\text{mat})}] \geq 0$$

whenever the two sector contributions are individually admissible. Thus the tester cone is convex and closed under addition of decoupled sectors, proving item (1).

On the modular side, the first-law channel and the modular equation of state are linear in the stress-tensor source: the calibrated belt first law has the form

$$\delta[S - \frac{\text{Area}}{4G}] = 2\pi \int_{\Sigma} d\Sigma_{\mu} \xi_{\nu} \delta\langle T^{\mu\nu}\rangle + O(\mathcal{B}_{\text{belt}}),$$

with the same linear functional of  $\delta\langle T_{\mu\nu}\rangle$  entering the canonical-energy and Ward maps. For a decoupled matter sector the total stress tensor is  $\delta\langle T_{\mu\nu}\rangle_{\text{tot}} = \delta\langle T_{\mu\nu}\rangle + \delta\langle T_{\mu\nu}\rangle^{(\text{mat})}$ , so the same identity holds with the source replaced by the sum and unchanged geometric kernel. The QNEC and ANEC inequalities on the belt are likewise linear in the stress-tensor expectation value and are proved at the level of each sector; for decoupled sectors the total null energy and total entropy variation are sums of sector contributions, so the QNEC/ANEC bounds hold with

constants depending only on the admissible class  $\mathfrak{S}_{\text{adm}}$  (encoded in  $\kappa_{\text{ANEC}}, 2\pi$ ) and not on the detailed matter content. The belt budgets  $B_{\text{belt}}$  are set by geometric and regulator data and therefore remain unchanged.

Combining these observations shows that adding a decoupled matter sector preserves all tester nonnegativity properties and the modular equation of state/QNEC with the same belt budgets, with constants updated only through the admissible state class. This proves item (2) and the proposition.  $\square$

### 5.78 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 1: Principal-series celestial anchors (worst-five) used in the explicit dual certificate.

### 5.79 Renormalized Newton constant invariance of the modular equation of state

*Lemma 5.107* (scheme/renormalization invariance). Let  $G_{\text{ren}}(\mu)$  be the renormalized Newton constant at a reference scale  $\mu$ , and let  $\delta\mathcal{L}_{\text{ct}}$  denote local boundary/corner counterterms compatible with the JKM calibration (Section 5.49). 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.

*Proof.* Introduce the functional

$$\mathcal{F}[G_{\text{ren}}, \delta\mathcal{L}_{\text{ct}}] := \delta \left[ S - \frac{\text{Area}}{4G_{\text{ren}}} \right] - 2\pi \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta \langle T_{\mu\nu} \rangle,$$

so that the statement of the lemma is that  $\mathcal{F}$  is well defined, up to  $O(\mathcal{B}_{\text{belt}})$ , on the belt-compatible scheme class.

Let  $(G_{\text{ren}}, \delta\mathcal{L}_{\text{ct}})$  and  $(G'_{\text{ren}}, \delta\mathcal{L}'_{\text{ct}})$  be two belt-compatible schemes in the sense of Section 5.31. Their renormalized actions differ by

$$I'[g] - I[g] = \frac{\delta G^{-1}}{16\pi} \int d^{d+1}x \sqrt{-g} R + \int_{\partial M} \delta\mathcal{L}_{\text{ct}} + \int_{\text{corners}} \delta\mathcal{L}_{\text{ct}}^{\text{corner}},$$

with  $\delta G^{-1} := G'^{-1}_{\text{ren}} - G^{-1}_{\text{ren}}$  and with local boundary/corner counterterms  $\delta\mathcal{L}_{\text{ct}}, \delta\mathcal{L}_{\text{ct}}^{\text{corner}}$  chosen to preserve the JKM calibration on the belt (Section 5.49).

On shell, bulk terms are proportional to the equations of motion and therefore do not contribute to the first-law variation, so the difference between the two schemes is entirely encoded in the quasi-local generators obtained from the boundary and corner variations. Writing the modular generator in Brown–York form,

$$\delta H_{\xi} = \int_{\Sigma} d\Sigma^{\mu} \xi^{\nu} \delta \langle T_{\mu\nu} \rangle + \delta H_{\xi}^{\text{corner}},$$

the change of scheme induces

$$\Delta_{\text{sch}}\delta H_\xi = \int_\Sigma d\Sigma^\mu \xi^\nu \Delta_{\text{sch}}\delta\langle T_{\mu\nu}\rangle + \Delta_{\text{sch}}\delta H_\xi^{\text{corner}},$$

where  $\Delta_{\text{sch}}$  denotes the difference between the primed and unprimed schemes.

The JKM analysis shows that a shift of the Newton constant together with calibrated counterterms produces a correlated shift of the Wald charge, the Brown–York tensor and the corner term such that

$$\Delta_{\text{sch}}\delta\left[\frac{\text{Area}}{4G_{\text{ren}}}\right] - 2\pi \Delta_{\text{sch}}\delta H_\xi = 0$$

up to terms supported on the belt and its corners that are local functionals of the induced data (Section 5.50). The same counterterms that implement the scheme change also shift the generalized entropy  $S$  by a local functional on the belt, and shift the renormalized Brown–York tensor by the corresponding metric variation of that local functional. By definition of the belt-compatible class (Section 5.31), these local contributions are accounted for in the belt budgets and hence are  $O(\mathcal{B}_{\text{belt}})$  once belt lift/recovery is performed.

Combining these observations, the scheme variation of  $\mathcal{F}$  is purely local on the belt and corners and satisfies

$$\Delta_{\text{sch}}\mathcal{F} = \mathcal{F}[G'_{\text{ren}}, \delta\mathcal{L}'_{\text{ct}}] - \mathcal{F}[G_{\text{ren}}, \delta\mathcal{L}_{\text{ct}}] = O(\mathcal{B}_{\text{belt}}).$$

Thus  $\mathcal{F}$ , and hence the first-law identity and modular equation of state, are invariant within the belt-compatible scheme class up to  $O(\mathcal{B}_{\text{belt}})$ . In the flow-removed limit the belt thickness is sent to zero and the  $O(\mathcal{B}_{\text{belt}})$  remainder vanishes, which yields strict scheme/renormalization invariance.  $\square$

## 5.80 Helicity-averaged optical positivity with dressing

*Proposition 5.108* (Unitarity implies helicity-averaged absorptive positivity with dressing). Work on the dispersion/celestial *working cone* of Remark 5.31 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 5.28. Then for every forward fixed- $t$  projector  $T_{q,1}^{\text{forw}}$  (Section 5.22), every Hankel/impact tester  $T_p^H$ , and every celestial Gram functional  $T_j^{\text{cel}}$  that enters the compact 18-support certificate (Section 5.48),

$$T_{q,1}^{\text{forw}}[\text{Abs}_{\text{avg}}] \geq 0, \quad T_p^H[\text{Abs}_{\text{avg}}] \geq 0, \quad T_j^{\text{cel}}[\text{Abs}_{\text{avg}}] \geq 0,$$

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 5.114 remains nonnegative on  $\text{Abs}_{\text{avg}}$ , and the amplitude synthesis (Theorem 5.33) holds with the same sign after helicity averaging and dressing.

*Proof.* On the working cone of Remark 5.31, unitarity and the optical theorem imply that, for each fixed helicity configuration  $\lambda$  and fixed  $t$  in the near-forward window of Section 5.10, the absorptive part

$$\text{Abs}^{(\lambda)}(s, t) := \text{Im} \mathcal{A}_{\text{hard}}^{(\lambda)}(s+i0, t)$$

is the density of a positive measure in  $s$ . The gravity subtraction used to define  $\mathcal{A}_{\text{hard}}$  and the fixed subtraction order  $N = 3$  subtract only analytic counterterms that are real on the

physical  $s$ -channel cut, so they do not modify the discontinuity and hence do not affect  $\text{Abs}^{(\lambda)}$ . Consequently, each  $\text{Abs}^{(\lambda)}(s, t)$  is pointwise nonnegative on the cut, and the helicity average

$$\text{Abs}_{\text{avg}}(s, t) = \frac{1}{N_\lambda} \sum_\lambda \text{Abs}^{(\lambda)}(s, t)$$

is again a nonnegative absorptive profile for every admissible  $(s, t)$  in the working cone.

Next, consider the dependence on anchors and dressing. By Section 5.28 and the detailed analysis in Proposition 5.49 and Lemma 5.87, changes of anchor data and BRST/diffeomorphism-compatible dressing act on the external polarizations by unitary transformations and by adding local counterterms in the JKM calibration class. At fixed  $(s, t)$ , the full amplitude can be regarded as a matrix in helicity space, and its discontinuity across the  $s$ -channel cut is a positive-semidefinite Hermitian matrix obtained from the optical theorem. Unitary rotations of the external polarization basis conjugate this matrix and hence preserve its spectrum and in particular its trace and diagonal entries. The helicity-averaged absorptive profile  $\text{Abs}_{\text{avg}}$  is proportional to the trace of this matrix, so it is invariant under such unitary rotations. The additional local counterterms lie in the calibration class and are real and analytic across the cut; therefore they do not contribute to the discontinuity and leave  $\text{Abs}_{\text{avg}}$  unchanged. Thus  $\text{Abs}_{\text{avg}}$  is well defined and nonnegative, uniformly over all admissible anchor and dressing choices.

We now examine the testers. By construction, each forward fixed- $t$  projector  $T_{q,1}^{\text{forw}}$  (Section 5.22) is a linear functional of the form

$$T_{q,1}^{\text{forw}}[f] = \int ds K_{q,1}^{\text{forw}}(s, t) f(s, t),$$

with a kernel  $K_{q,1}^{\text{forw}}(s, t) \geq 0$  on the working cone. Hence  $T_{q,1}^{\text{forw}}[f] \geq 0$  whenever  $f$  is a nonnegative absorptive profile. Similarly, each Hankel/impact tester  $T_p^H$  is represented as an integral transform with a nonnegative kernel on the near-forward window, or equivalently as a quadratic form in partial waves with positive weights (Section 5.33), and therefore is nonnegative on nonnegative absorptive profiles.

For the celestial functionals, Sections 5.42 and 5.48 express each  $T_j^{\text{cel}}$  entering the compact 18-support certificate as a Gram form

$$T_j^{\text{cel}}[f] = \sum_{a,b} c_a c_b \int ds f(s, t) \phi_a(s, t) \phi_b(s, t),$$

for suitable test functions  $\phi_a$  and coefficients  $c_a$  depending only on the fixed wiring and anchors. When  $f$  is a nonnegative absorptive profile, the matrix of overlaps

$$G_{ab}(t) := \int ds \text{Abs}_{\text{avg}}(s, t) \phi_a(s, t) \phi_b(s, t)$$

is positive semidefinite, so  $T_j^{\text{cel}}[\text{Abs}_{\text{avg}}] \geq 0$ . The uniformity in  $t$  over the near-forward window follows from the fact that the kernels and Gram matrices remain positive on that window, while any degradation from the widened window is already absorbed into the ledgered budgets in Section 5.10.

We have thus shown that for every admissible  $t$  and every tester in the forward, Hankel/impact, and celestial families,

$$T_{q,1}^{\text{forw}}[\text{Abs}_{\text{avg}}] \geq 0, \quad T_p^H[\text{Abs}_{\text{avg}}] \geq 0, \quad T_j^{\text{cel}}[\text{Abs}_{\text{avg}}] \geq 0.$$

The compact 18-support dual of Section 5.48 and its wiring in Proposition 5.114 are nonnegative linear combinations of these testers with anchors frozen as in Table 1, so the resulting dual functional remains nonnegative on  $\text{Abs}_{\text{avg}}$ . Therefore the amplitude synthesis of Theorem 5.33 continues to hold with the same sign after helicity averaging and dressing.  $\square$

## 5.81 Two epsilon targets and OSR inflation

**Global budget split (recall Section 5.27).** For target  $\varepsilon$ , 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.

For the two audit targets  $\varepsilon \in \{10^{-6}, 10^{-8}\}$  with  $\eta = \frac{1}{3}$ , the corresponding deterministic AGSP steps and small parameters  $(m, \delta^2)$  are frozen once in the constants ledger (Section Appendix B.2), and at these precisions the converter denominators  $(1 - \delta^2)^{-1}$  and  $(1 - \delta^2)^{-2}$  are numerically 1 at the displayed accuracy.

## 5.82 Decoupling for widely separated belts

*Lemma 5.109 (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  $|R_{1,2}|$ .

*Proof.* Work at fixed belt width  $r > 0$  and in the positive-flow window  $(u, s) > 0$ ; all statements below are per generator length and with constants uniform in  $|R_{1,2}|$ . We first obtain the bound at finite regulators, and then remove  $(u, s)$  using the OS kernel.

*Step 1: Exponential suppression of cross-commutators.* Let  $A_1 \in \mathcal{A}(\partial_r R_1)$  and  $A_2 \in \mathcal{A}(\partial_r R_2)$  be bounded operators supported on the two belts, with  $\partial_r R_1$  and  $\partial_r R_2$  separated by a distance  $d$  along the cut. By the belt microcausality/split structure and null timeslice propagation Lemma 5.74 and Proposition 5.75, there exist constants  $C_{\text{mc}}, \mu_{\text{eff}} > 0$ , independent of  $|R_{1,2}|$ , such that

$$\|[A_1, A_2]\| \leq C_{\text{mc}} e^{-\mu_{\text{eff}} d} \|A_1\| \|A_2\|.$$

In particular, on the OS belt code subspace all commutators between the two belt algebras are exponentially small in  $d$ .

*Step 2: Approximate split and product structure.* The quasi-local belt factorization and Brown–York dictionary Section 5.50 and Proposition 5.78 imply that, for the class of observables considered in this section, the restriction of an admissible state  $\rho$  to the algebra generated by  $\mathcal{A}(\partial_r R_1) \vee \mathcal{A}(\partial_r R_2)$  can be approximated by a product state

$$\rho_{12} \approx \rho_1 \otimes \rho_2$$

in the following sense: there exists a normal state  $\rho^\otimes$  on  $\mathcal{A}(\partial_r R_1) \bar{\otimes} \mathcal{A}(\partial_r R_2)$  such that

$$|\rho_{12}(A_1 A_2) - \rho^\otimes(A_1 A_2)| \leq C_1 e^{-\mu_{\text{eff}} d} \|A_1\| \|A_2\|$$

for all  $A_i \in \mathcal{A}(\partial_r R_i)$ , with  $C_1$  independent of  $|R_{1,2}|$ . Equivalently, the cross-correlators between the two belts are controlled by the same exponential tail that governs the commutator in Step 1.

By linearity and norm-density of polynomials in local observables, this control extends to any bounded quasi-local observable supported on  $\partial_r R_1 \cup \partial_r R_2$ .

*Step 3: Additivity for the belt-local functionals.* Each  $\mathcal{O}$  used in this section (expectations of modular generators, Brown–York/canonical fluxes, generalized entropy, and their QES functionals) has two key properties:

(i) *Quasi-local representation.* Up to an  $O(\mathcal{B}_{\text{belt}})$  remainder that is uniform in  $|R|$ ,  $\mathcal{O}$  can be written as a finite sum of expectation values of bounded belt-local operators and geometric terms supported on  $\partial_r R_1 \cup \partial_r R_2$ , via the quasi-local belt dictionary and Brown–York identification Section 5.50 and Proposition 5.78.

(ii) *Exact additivity on product states.* For a product state  $\rho^\otimes = \rho_1 \otimes \rho_2$  on  $\mathcal{A}(\partial_r R_1) \bar{\otimes} \mathcal{A}(\partial_r R_2)$  one has

$$\mathcal{O}_{\rho^\otimes}(R_1 \cup R_2) = \mathcal{O}_{\rho_1}(R_1) + \mathcal{O}_{\rho_2}(R_2),$$

because the modular and Brown–York generators factorize and the geometric (area) piece is a sum over disjoint components.

Combining (i) with Step 2, we may write at finite  $(u, s)$

$$\mathcal{O}(R_1 \cup R_2) = \mathcal{O}_{\rho^\otimes}(R_1 \cup R_2) + \delta\mathcal{O}_{12}, \quad |\delta\mathcal{O}_{12}| \leq C_2 e^{-\mu_{\text{eff}} d},$$

with  $C_2$  independent of  $|R_{1,2}|$ , since  $\mathcal{O}$  is a bounded linear functional of the underlying quasi-local operators and the exponential control from Step 2 applies term by term. Using (ii) for  $\rho^\otimes$  then yields

$$|\mathcal{O}(R_1 \cup R_2) - \mathcal{O}(R_1) - \mathcal{O}(R_2)| \leq C_2 e^{-\mu_{\text{eff}} d} + O(\mathcal{B}_{\text{belt}}).$$

*Step 4: Removal of positive flows.* The OS short-evolution decomposition and removal lemma Lemma 3.3 ensure that the  $O(\mathcal{B}_{\text{belt}})$  remainder in the previous line vanishes as  $(u, s) \downarrow 0$ , with constants uniform in  $|R_{1,2}|$ . Passing to the regulator-free limit and absorbing harmless numerical factors into  $C$  gives

$$|\mathcal{O}(R_1 \cup R_2) - \mathcal{O}(R_1) - \mathcal{O}(R_2)| \leq C e^{-\mu_{\text{eff}} d},$$

as claimed. □

### 5.83 QES maximin selector consistent with uniqueness

*Proposition 5.110 (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 5.14,  $\text{QES}_{\text{maximin}}$  equals the unique QES selected by the belt flow (saddle selection) up to  $O(\mathcal{B}_{\text{belt}})$ .

*Proof.* This is precisely the saddle-selection/maximin part of the QES stability theorem Theorem 5.38, combined with the contraction and GSL inputs Theorems 5.41 and 5.58.

For each fixed belt-anchored cut  $C$ , Theorem 5.38 guarantees that  $\mathcal{G}$  is strictly convex along admissible deformations and that there is a unique minimizer  $\Sigma_{\text{QES}}(C) \in \text{adm}(C)$ , so

$$\min_{\Sigma \in \text{adm}(C)} \mathcal{G}[\Sigma] = \mathcal{G}[\Sigma_{\text{QES}}(C)].$$

Defining  $m(C) := \min_{\Sigma \in \text{adm}(C)} \mathcal{G}[\Sigma]$ , the same theorem, together with the  $c$ -function monotonicity and local GSL Theorems 5.41 and 5.58, shows that  $m(C_\tau)$  is nondecreasing along the belt flow  $C \mapsto C_\tau$  and converges, as  $\tau \rightarrow \infty$ , to the value of  $\mathcal{G}$  at the unique stationary QES picked by the flow. Thus the belt flow selects a cut  $C_*$  for which  $m(C_*)$  is maximal over the family, and

$$\text{QES}_{\text{maximin}} = \arg \max_C m(C)$$

coincides (up to the uniform  $O(\mathcal{B}_{\text{belt}})$  remainder) with the QES selected dynamically by the belt flow. The  $O(\mathcal{B}_{\text{belt}})$  corrections are tracked in the same argument and vanish under positive-flow removal by Lemma 3.3, which completes the proof.  $\square$

### 5.84 Numerical Gauss–Radau remainder and composite schedule (forward coefficient)

*Lemma 5.111* (numeric Gauss–Radau bound). For the second forward derivative ( $k=2$ ) with  $N=3$  subtractions and baseline slope  $\alpha_{\text{R}}=2$  on  $t \in [-0.20 s_0, 0]$ , consider the integrand

$$g(s') = \mathcal{R}(t) s'^{\alpha_{\text{R}} - (N+k+1)}.$$

On  $s' \in [s_0, 20s_0]$  the  $n=4$  Gauss–Radau quadrature remainder for the corresponding dispersion integral satisfies

$$|\text{Quad}_{n=4}| \leq 2,109,668.2085 s_0^{-3}.$$

For the dimensionless forward coefficient

$$\widehat{c}_{2,0} = \frac{s_0^3}{2} \partial_s^2 \Re \mathcal{A}^{(N)}(0, 0)$$

this implies

$$|\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq 1,054,834.1043.$$

**Composite schedule remedy.** If  $[s_0, 20s_0]$  is split into  $J$  equal panels and  $n=4$  Gauss–Radau is applied on each panel (composite schedule), the quadrature error scales as  $J^{-5}$ :

$$|\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{1,054,834.1043}{J^5}.$$

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 5.62.

*Proof.* For the parameters  $k=2$  and  $N=3$  with baseline slope  $\alpha_{\text{R}}=2$  we have

$$\alpha := \alpha_{\text{R}} - (N + k + 1) = 2 - (3 + 2 + 1) = -4,$$

so on the corridor one can write

$$g(s') = \mathcal{R}(t) s'^{\alpha} = \mathcal{R}(t) s'^{-4}.$$

For general  $\alpha$  the fourth derivative takes the form

$$g^{(4)}(s') = \mathcal{R}(t) \alpha(\alpha-1)(\alpha-2)(\alpha-3) s'^{\alpha-4}.$$

With  $\alpha = -4$  this becomes

$$g^{(4)}(s') = \mathcal{R}(t) (-4)(-5)(-6)(-7) s'^{-8} = 840 \mathcal{R}(t) s'^{-8}.$$

On  $s' \in [s_0, 20s_0]$  the function  $s'^{-8}$  is strictly decreasing in  $s'$ , hence

$$\sup_{s' \in [s_0, 20s_0]} |g^{(4)}(s')| \leq 840 \mathcal{R}(t) s_0^{-8}.$$

By construction of the regulator  $\mathcal{R}(t)$  one has

$$\mathcal{R}(t) \leq \sqrt{1 + |t|/s_0}$$

and on  $t \in [-0.20 s_0, 0]$  this gives

$$\mathcal{R}(t) \leq \sqrt{1 + 0.2} = \sqrt{1.2} = 1.0954451150.$$

Combining the last two displays yields

$$\sup_{s' \in [s_0, 20s_0]} |g^{(4)}(s')| \leq 840 \times 1.0954451150 s_0^{-8}.$$

The  $n=4$  Gauss–Radau rule on an interval of length  $L$  admits the Peano-type remainder bound

$$|\text{Quad}_{n=4}| \leq \frac{1}{1080} \sup_{[s_0, 20s_0]} |g^{(4)}| L^5.$$

On  $[s_0, 20s_0]$  the length is  $L = 20s_0 - s_0 = 19s_0$ , so we obtain

$$|\text{Quad}_{n=4}| \leq \frac{1}{1080} \times 840 \times 1.0954451150 \times 19^5 s_0^{-3} = 2,109,668.2085 s_0^{-3},$$

which is the first claim.

The dispersion representation for the forward coefficient  $\widehat{c}_{2,0}$  is linear in  $g$  and carries the overall prefactor  $\frac{s_0^3}{2}$ . Thus the quadrature error on  $\widehat{c}_{2,0}$  is exactly  $\frac{1}{2}s_0^3$  times the quadrature error on the underlying  $s'$ -integral. Cancelling  $s_0^3$  against the  $s_0^{-3}$  in the last display gives

$$|\Delta_{\text{quad}\widehat{c}_{2,0}}| \leq \frac{1}{2} |\text{Quad}_{n=4}| = 1,054,834.1043,$$

as claimed in the lemma.

For the composite schedule, split  $[s_0, 20s_0]$  into  $J$  equal panels, each of length  $L/J = 19s_0/J$ , and apply the same  $n=4$  Gauss–Radau rule on each panel. The remainder bound above depends on the panel length only through the factor  $L^5$ , while the normalization entering  $\widehat{c}_{2,0}$  is linear in the integral. Therefore, replacing  $L$  by  $L/J$  in the bound multiplies the error by  $J^{-5}$ , and we arrive at

$$|\Delta_{\text{quad}\widehat{c}_{2,0}}| \leq \frac{1,054,834.1043}{J^5}.$$

Evaluating this expression at  $J = 10, 20, 30$  gives the numerical values quoted in the statement. This completes the proof.  $\square$

**Add-on ( $\delta$ -robust quadrature and widened window).** For a tilted baseline slope  $\alpha_R = 2 + \delta$  with  $0 \leq \delta \leq 0.2$  the integrand exponent becomes

$$\alpha = \alpha_R - (N + k + 1) = -4 + \delta,$$

and the fourth derivative reads

$$g^{(4)}(s') \propto (-4 + \delta)(-5 + \delta)(-6 + \delta)(-7 + \delta) s'^{-8 + \delta}.$$

On  $0 \leq \delta \leq 0.2$  the polynomial prefactor

$$(-4 + \delta)(-5 + \delta)(-6 + \delta)(-7 + \delta)$$

decreases in magnitude with  $\delta$ , so the Gauss–Radau remainder prefactor is slightly reduced compared to the  $\delta = 0$  case. In particular, the panel counts  $J$  fixed above remain valid throughout this range of  $\alpha_R$ .

Widening the  $t$ -window from  $[-0.20 s_0, 0]$  to  $[-0.25 s_0, 0]$  only affects the factor  $\mathcal{R}(t)$ ; the new bound

$$\mathcal{R}(t) \leq \sqrt{1 + |t|/s_0} \leq \sqrt{1.25}$$

is larger than  $\sqrt{1.2}$  by at most

$$\sqrt{\frac{1.25}{1.2}} = 1.0206 \dots,$$

which lies well within the dispersion headroom already built into the error budget.

### 5.85 Shear control via canonical energy positivity

*Lemma 5.112* (shear quadratic bound). Assume the hypotheses of Theorem 5.46. There exist belt-uniform constants  $C_\sigma, \tilde{C}_\sigma, \tilde{C}'_\sigma > 0$  independent of  $|R|$  such that along any belt-anchored null segment  $[\lambda_1, \lambda_2]$ ,

$$\int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \sigma^{ab} \sigma_{ab} \leq C_\sigma EW_{\text{can}}[\delta\Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}}, \quad (5.109)$$

$$\begin{aligned} |\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]| &\leq \tilde{C}_\sigma \left( \int \sigma^2 \right)^{1/2} \left( \int \theta^2 \right)^{1/2} + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}} \\ &\leq \tilde{C}'_\sigma EW_{\text{can}}[\delta\Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}}. \end{aligned} \quad (5.110)$$

Consequently, the shear functional of Theorem 5.46 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  $C'_\sigma$  belt-uniform.

*Proof.* Fix an admissible belt-anchored null segment  $[\lambda_1, \lambda_2]$  and an admissible variation  $\delta\Psi$  obeying the hypotheses of Theorem 5.46. We work per generator length and in the positive-flow window; all remainders below are absorbed into  $\mathcal{B}_{\text{belt}}$  with belt-uniform constants.

*Step 1: Shear in the canonical-energy form domain and (5.109).* By the OS/KMS kernel and the JLMS channel used in Theorem 5.46, the modular generator and the generalized-entropy operator admit a common quadratic-form core on the belt. On this core the Iyer–Wald canonical energy  $EW_{\text{can}}[\delta\Psi; \xi]$  is a positive, closed quadratic form whose domain consists of perturbations with finite energy flux through the wedge. In a belt-adapted null frame, the shear  $\sigma_{ab}$  is the transverse-traceless part of the null extrinsic curvature of the cut; BW/KMS locality implies that for any  $\delta\Psi$  in the canonical-energy form domain the field  $\sigma_{ab}$  is square-integrable along  $[\lambda_1, \lambda_2]$  and depends continuously and linearly on  $\delta\Psi$ .

Thus the map

$$\delta\Psi \longmapsto \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \sigma^{ab} \sigma_{ab}$$

defines a positive quadratic form on the same domain as  $EW_{\text{can}}[\delta\Psi; \xi]$ . In the derivation of Theorem 5.46, the second variation of  $S - \text{Area}/(4G)$  along belt-anchored deformations is expressed as

$$\delta^2 \left[ S - \frac{\text{Area}}{4G} \right] = 2\pi EW_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] + (\text{other positive pieces}) + O(\mathcal{B}_{\text{belt}}),$$

where  $Q_{\text{shear}}$  is the optical quadratic form with strictly positive belt-local coefficients, and the remaining terms are quadratic in non-optical data. Restricted to transverse-traceless metric perturbations on the belt (so that only  $\sigma_{ab}$  is nonzero in the optical sector), this second variation reduces to a positive multiple of the canonical-energy form plus an  $O(\mathcal{B}_{\text{belt}})$  remainder and is proportional to the  $L^2$ -norm of the shear.

Standard comparison of positive closed quadratic forms on a common domain then yields a belt-uniform constant  $C_\sigma > 0$  and a remainder constant  $C_\sigma^{\text{rem}} > 0$  such that

$$\int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \sigma^{ab} \sigma_{ab} \leq C_\sigma EW_{\text{can}}[\delta\Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}},$$

which is exactly (5.109).

*Step 2: Shear–expansion coupling and (5.110).* By construction in Theorem 5.46, the cubic shear–expansion functional  $\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]$  is obtained by differentiating  $Q_{\text{shear}}[\delta g]$  with respect

to the deformation parameter. In a belt-adapted null frame its integrand is a finite linear combination of terms of the form  $c(\lambda, x) \sigma^{ab}(\lambda, x) \sigma_{ab}(\lambda, x) \theta(\lambda, x)$  and  $c'(\lambda, x) \sigma^{ab}(\lambda, x) \theta(\lambda, x)$ , with smooth coefficients  $c, c'$  controlled by the belt-regularity and small-tilt hypotheses. In particular, there exists a belt-uniform constant  $C_* > 0$  such that

$$|\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]| \leq C_* \int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} |\sigma| |\theta| + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}}.$$

Applying Cauchy–Schwarz to the spacetime integral gives

$$|\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]| \leq C_* \left( \int \sigma^2 \right)^{1/2} \left( \int \theta^2 \right)^{1/2} + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}},$$

which is the first line of (5.110) with  $\tilde{C}_\sigma := C_*$ .

To obtain the second line of (5.110), we note that the same canonical-energy density that controls the shear also contains the expansion  $\theta$  with a positive coefficient. The optical decomposition in Theorem 5.46 therefore yields a belt-uniform constant  $C_\theta > 0$  and a remainder constant (absorbed into  $C_\sigma^{\text{rem}}$ ) such that

$$\int_{\lambda_1}^{\lambda_2} d\lambda \int_{\partial R} \sqrt{\gamma} \theta^2 \leq C_\theta EW_{\text{can}}[\delta \Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}}.$$

Combining this with (5.109) in the Cauchy–Schwarz estimate and absorbing constants gives

$$|\mathcal{C}_{\text{shear} \times \text{exp}}[\delta g]| \leq \tilde{C}'_\sigma EW_{\text{can}}[\delta \Psi; \xi] + C_\sigma^{\text{rem}} \mathcal{B}_{\text{belt}},$$

for some belt-uniform  $\tilde{C}'_\sigma > 0$ , which is the second line of (5.110).

*Step 3: Consequence for  $Q_{\text{shear}}$ .* Since  $Q_{\text{shear}}[\delta g]$  has the form

$$Q_{\text{shear}}[\delta g] = \underline{\kappa}_\sigma \int \sigma^2 + \underline{\kappa}_\theta \int \theta^2 + (\text{controlled cross terms}),$$

the bounds (5.109) and (5.110) allow us to absorb the cross terms into the canonical-energy piece and the belt budget, yielding the stated lower bound

$$Q_{\text{shear}}[\delta g] \geq \underline{\kappa}_\sigma C_\sigma^{-1} \int \sigma^2 + \underline{\kappa}_\theta \int \theta^2 - C'_\sigma \mathcal{B}_{\text{belt}}$$

for a suitable belt-uniform constant  $C'_\sigma$ .

Finally, by positive-flow removal Lemma 3.3, the  $\mathcal{B}_{\text{belt}}$  terms vanish in the  $(u, s) \downarrow 0$  limit, completing the proof.  $\square$

## 5.86 Tilted belt and RP/KMS stability

*Proposition 5.113* (small-tilt RP/KMS). Let  $\mathcal{B}_\vartheta$  be a belt obtained by tilting the reference belt by a small angle  $\vartheta$  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  $K_m$ .

*Proof.* Let  $\mathcal{B}_0$  denote the reference belt. By the belt Brown–York dictionary and functorial quasi-locality (Section 5.50 and Proposition 5.78), for each  $|\vartheta|$  in the OS window there exists a quasi-local  $*$ -automorphism  $\alpha_\vartheta$  which identifies observables supported on  $\mathcal{B}_0$  with observables supported on  $\mathcal{B}_\vartheta$ , with uniform control in the belt width and under one layer of belt circuits and even  $K_m$ . For fixed belt-local observables  $X_j$  on  $\mathcal{B}_0$  we write

$$X_j^{(\vartheta)} := \alpha_\vartheta(X_j).$$

The functorial quasi-locality bounds of Section 5.50 and Proposition 5.78 imply that the map  $\vartheta \mapsto X_j^{(\vartheta)}$  is twice strongly differentiable on the OS window, with

$$\|\partial_\vartheta^k X_j^{(\vartheta)}\| \leq C_k \|X_j\| \quad \text{for } k = 1, 2$$

and constants  $C_k$  independent of  $|R|$  and stable under one layer of belt circuits and even  $K_m$ .

Define

$$Q(\vartheta) := \sum_{i,j} \langle \Omega_{u,s}, \Theta(X_i^{(\vartheta)}) X_j^{(\vartheta)} \Omega_{u,s} \rangle.$$

By construction of the tilt and of the OS reflection  $\Theta$  one has the covariance relation

$$\Theta \circ \alpha_\vartheta = \alpha_{-\vartheta} \circ \Theta,$$

and the state  $\Omega_{u,s}$  is invariant under the corresponding belt transport/KMS flow. It follows that

$$Q(-\vartheta) = Q(\vartheta),$$

so  $Q$  is an even function of  $\vartheta$  on the OS window and therefore  $Q'(0) = 0$ .

Using the differentiability of  $X_j^{(\vartheta)}$  and the Cauchy–Schwarz inequality, together with the uniform microcausality bound on the two-point function (Lemma 5.74), we obtain

$$|Q''(\vartheta)| \leq c_0 \sum_j \|X_j\|^2$$

for all  $|\vartheta|$  in the OS window, with  $c_0$  independent of  $|R|$  and stable under one layer of belt circuits and even  $K_m$ . By Taylor’s theorem with remainder and the evenness of  $Q$ , there exists  $\xi_\vartheta$  between 0 and  $\vartheta$  such that

$$Q(\vartheta) = Q(0) + \frac{\vartheta^2}{2} Q''(\xi_\vartheta),$$

and hence

$$Q(\vartheta) \geq Q(0) - c \vartheta^2 \sum_j \|X_j\|^2$$

for some constant  $c$  independent of  $|R|$ , uniformly under one belt-circuit layer and even  $K_m$ .

For the untilted belt  $\mathcal{B}_0$  we can invoke RP/KMS positivity for belt-local kernels (Lemma 3.1), together with the factorization and microcausality tail estimate (Lemma 5.74), to obtain

$$Q(0) = \sum_{i,j} \langle \Omega_{u,s}, \Theta(X_i) X_j \Omega_{u,s} \rangle \geq -C e^{-\mu_{\text{eff}} r}.$$

Combining the last two displays yields

$$Q(\vartheta) \geq -c \vartheta^2 \sum_j \|X_j\|^2 + O(e^{-\mu_{\text{eff}} r}),$$

which is the claimed small-tilt RP/KMS stability, with all constants uniform in  $|R|$  and under one layer of belt circuits and even  $K_m$ .  $\square$

### 5.87 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 2: 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  $\mu_{\text{cel}}$  ranges used in our pipelines; if an upstream  $\mu_{\text{cel}}$  tighter than 1.20 is enforced, lower the absolute  $\nu$  entries accordingly (the proofs are insensitive to the exact values as long as they lie on principal lines and within the declared window).

*Proposition 5.114* (wiring the explicit 18-support dual). Replace the celestial anchor component of the explicit 18-support dual certificate (Section 5.48) by the anchors in Table 2, 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 2.
3. All positivity claims that referenced the “worst-five” now hold with concrete numeric indices and the same acceptance budget.

*Proof.* Let  $\Lambda$  denote the explicit 18-support dual functional constructed in Section 5.48. By construction,  $\Lambda$  is a conic combination of three types of testers: forward even-parity belt testers, Hankel-type testers, and celestial testers built from principal-series evaluation functionals. All weights are nonnegative and the gravity IR subtraction has already been performed, so  $\Lambda$  acts on the subtracted amplitude and probes only the hard, positive absorptive content.

We decompose

$$\Lambda = \Lambda_{\text{belt}} + \Lambda_{\text{Hankel}} + \Lambda_{\text{cel}},$$

where  $\Lambda_{\text{cel}}$  is the celestial component supported on five principal-series anchors (the “worst-five” set) and the other components are supported on the remaining  $6+7$  nodes. By hypothesis we only modify  $\Lambda_{\text{cel}}$ , replacing its five anchors by the frozen set of principal-series points listed in Table 2, with the same nonnegative weights and the same total mass.

The celestial Gram form on principal-series lines is positive semidefinite and is preserved by the belt Ward transform (the celestial Gram positivity and Ward map assumptions in the scope paragraph above). In particular, for any finite subset of principal-series points within the declared celestial window and any choice of nonnegative weights, the resulting celestial tester is a nonnegative functional on the cone  $\mathcal{S}$  of subtracted amplitudes. Each of the frozen anchors in Table 2 lies on such a principal line and within the same window, so the modified celestial component  $\Lambda'_{\text{cel}}$  obtained by freezing the anchors remains nonnegative on  $\mathcal{S}$ . The belt and Hankel components are unchanged and were already nonnegative by construction. Hence

$$\Lambda' := \Lambda_{\text{belt}} + \Lambda_{\text{Hankel}} + \Lambda'_{\text{cel}}$$

is again a nonnegative functional on  $\mathcal{S}$ , proving item (1).

Item (2) is immediate: we replace five celestial nodes by five new ones without changing the other  $6+7$  nodes, so the support size remains  $6+7+5 = 18$ , with the five celestial nodes now given explicitly by Table 2.

For item (3), recall that all positivity and budget estimates for the original dual that referred to the “worst-five” anchors used only the following structural inputs: (i) the anchors lie on principal-series lines within the declared celestial window, (ii) the celestial Gram form is positive semidefinite and Ward-compatible with the belt map, and (iii) the gravity IR subtraction removes nonanalytic contributions so that the dual sees only the hard absorptive part. The proofs do not depend on the precise numerical values of  $(n, \nu)$  beyond these structural constraints. The frozen anchors of Table 2 satisfy exactly the same constraints (they are principal-series, symmetric in  $\nu$ , and lie within the same  $\mu_{\text{cel}}$  window). Consequently, every inequality and acceptance budget previously stated for the “worst-five” anchors continues to hold verbatim for this concrete frozen choice. This establishes item (3) and completes the proof.  $\square$

### 5.88 Cubic canonical-energy worked example (coherent pulse)

*Proposition 5.115* (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 5.25), 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 5.25)  $\delta\langle K_{\text{R}} \rangle = \frac{\pi}{2} A^2 = 1.5707963268 \times 10^{-4}$  (second order), and  $2\pi \int du \langle T_{kk} \rangle = 1.8561093322 \times 10^{-4}$ . At cubic order the coherent displacement contributes

$$\mathcal{E}_{\text{can}}^{(3)} = 0 \quad (\text{free coherent state}) \quad \Rightarrow \quad |\mathcal{E}_{\text{can}}^{(3)}| \leq C_{\sigma} \mathcal{E}_{\text{can}}^W \mathcal{B}_{\text{belt}} \quad (\text{shear-controlled bound; Section 5.85})$$

Thus the third-order term is below machine precision for the coherent benchmark, up to the belt remainder, which vanishes under flow removal.

*Proof.* Let  $|\Psi(\lambda)\rangle$  denote the coherent perturbation with profile  $\phi_{\lambda}(u) = \lambda \phi(u)$ , obtained by replacing  $A \mapsto \lambda A$  in the Gaussian profile of Section 5.25. For a free scalar, the coherent state is produced by a Weyl displacement operator acting on the vacuum and is therefore Gaussian: it has the same two-point function as the vacuum, with a nonzero one-point function  $\langle \Psi(\lambda) | \Phi(u) | \Psi(\lambda) \rangle = \phi_{\lambda}(u)$ .

The Rindler modular Hamiltonian  $K_{\text{R}}$  for the free scalar is quadratic in the field operators. Normal ordering with respect to the vacuum implies that the difference

$$\Delta\langle K_{\text{R}} \rangle(\lambda) := \langle \Psi(\lambda) | K_{\text{R}} | \Psi(\lambda) \rangle - \langle 0 | K_{\text{R}} | 0 \rangle$$

is a quadratic functional of the classical profile  $\phi_{\lambda}$ , with no linear term. Equivalently,

$$\Delta\langle K_{\text{R}} \rangle(\lambda) = Q[\phi_{\lambda}, \phi_{\lambda}],$$

for some positive quadratic form  $Q$  determined by the free theory and the wedge geometry. In particular,

$$\Delta\langle K_{\text{R}} \rangle(\lambda) = c_2 \lambda^2 + c_4 \lambda^4 + \dots$$

is an even function of  $\lambda$ , and all odd derivatives at  $\lambda = 0$  vanish.

By definition, the canonical-energy expansion is obtained from the Taylor expansion of the modular/relative entropy (equivalently, of  $\Delta\langle K_{\text{R}} \rangle$ ) around  $\lambda = 0$ , up to the universal belt remainder  $O(\mathcal{B}_{\text{belt}})$  discussed in Section 5.25. Since the third derivative at  $\lambda = 0$  of an even function vanishes, the cubic coefficient must be zero:

$$\mathcal{E}_{\text{can}}^{(3)}[\delta\Psi; \xi] = 0 + O(\mathcal{B}_{\text{belt}}).$$

Finally, the general shear-controlled bound of Section 5.85 bounds the size of the cubic contribution by

$$|\mathcal{E}_{\text{can}}^{(3)}| \leq C_\sigma \mathcal{E}_{\text{can}}^W \mathcal{B}_{\text{belt}},$$

yielding the displayed inequality.  $\square$

### 5.89 Composite quadrature schedule freeze for the dispersive integrals

**Goal.** Choose a composite Gauss–Radau schedule (panel count  $J$  and cutoff  $S_{\text{cut}}$ ) that *provably* meets the dispersion budget in Section 5.27 for the two audit targets  $\varepsilon \in \{10^{-6}, 10^{-8}\}$ .

**Setup (dimensionless forward coefficient).** Let  $\widehat{c}_{2,0} := s_0^3 c_{2,0}$ . 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.

$$|\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \left(\frac{20}{M}\right)^3, \quad |\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{K(M)}{J^5},$$

where  $M := S_{\text{cut}}/s_0$  and  $K(M) = C_0 (M-1)^5$  with  $C_0 = \frac{1}{2} \cdot \frac{1}{1080} \cdot 840 \cdot \sqrt{1.2} = 0.42600643\dots$

$\delta$ -sensitivity. If  $\alpha_R = 2 + \delta$ , then  $|\Delta_{\text{tail}} \widehat{c}_{2,0}|$  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

$$M_{(10^{-6})} = 100 \quad \Rightarrow \quad |\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \times \left(\frac{20}{100}\right)^3 = 5.811516831 \times 10^{-8},$$

$$K(M_{(10^{-6})}) = C_0 \cdot 99^5 \approx 4.0513 \times 10^9,$$

$$J_{(10^{-6})} = 2200 \quad \Rightarrow \quad |\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{4.0513 \times 10^9}{(2200)^5} = 7.86 \times 10^{-8}.$$

$$\text{Hence } |\Delta_{\text{tail}} \widehat{c}_{2,0}| + |\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq 1.367 \times 10^{-7} < \varepsilon_{\text{disp}}.$$

**Project selection.** We adopt *Freeze A* for the project line ( $\varepsilon = 10^{-6}$ ), i.e.  $M = 100$  and  $J = 2200$ , which deliver  $|\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq 5.811516831 \times 10^{-8}$  and  $|\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq 7.86 \times 10^{-8}$  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

$$M_{(10^{-8})} = 500 \quad \Rightarrow \quad |\Delta_{\text{tail}} \widehat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \times \left(\frac{20}{500}\right)^3 = 4.649213465 \times 10^{-10},$$

$$K(M_{(10^{-8})}) = C_0 \cdot 499^5 \approx 1.3180 \times 10^{13},$$

$$J_{(10^{-8})} = 28000 \quad \Rightarrow \quad |\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq \frac{1.3180 \times 10^{13}}{(28000)^5} = 7.65 \times 10^{-10}.$$

$$\text{Thus } |\Delta_{\text{tail}} \widehat{c}_{2,0}| + |\Delta_{\text{quad}} \widehat{c}_{2,0}| \leq 1.230 \times 10^{-9} < \varepsilon_{\text{disp}}.$$

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 5.89 then scale by  $F_{\text{strip}}$ :

$$|\Delta_{\text{tail}}\widehat{c}_{2,0}| \leq 7.2643960393 \times 10^{-6} \left(\frac{20}{M}\right)^3 M^{\sigma_0}, \quad |\Delta_{\text{quad}}\widehat{c}_{2,0}| \leq \frac{K(M)}{J^5} M^{\sigma_0},$$

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 $\varepsilon$	$M$	$J$	$F_{\text{strip}}$	$ \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  $M^{\sigma_0}$  in the tails, while  $J$  provides cheap  $J^{-5}$  control of quadrature. All invariances from Section 5.22, Section 5.46, Section 5.66 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 5.89, 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 5.27 (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 5.89), 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 5.48) and the forward even-parity analytic projector (Section 5.22) 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.

## 5.90 Interacting cubic canonical-energy example (non-Gaussian displacement)

*Proposition 5.116* (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  $K_{\text{R}} = 2\pi \int_{x>0} x T_{00}(0, x) dx$  [4]. 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 5.25 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 5.85.

*Proof.* For a scalar field with potential  $V(\phi)$ , the Hamiltonian density at  $t = 0$  can be written as

$$T_{00}(0, x) = \frac{1}{2} \left[ \pi(0, x)^2 + (\partial_x \phi(0, x))^2 \right] + V(\phi(0, x)),$$

so the interaction contribution to the Rindler modular Hamiltonian is

$$K_{\text{R}}^{(\text{int})} = 2\pi \int_{x>0} x V(\phi) dx.$$

Evaluating this on the coherent profile

$$\phi(x) = A \exp(-x^2/(2L^2)), \quad x > 0,$$

and expanding  $V(\phi)$  in powers of  $\phi$  gives the cubic and quartic interaction pieces.

For the cubic term we obtain

$$\Delta\langle K_R \rangle^{(3)} = 2\pi \int_{x>0} x \frac{g_3}{3!} \phi(x)^3 dx = 2\pi \frac{g_3}{3!} A^3 \int_0^\infty x e^{-3x^2/(2L^2)} dx.$$

Using

$$\int_0^\infty x e^{-\alpha x^2} dx = \frac{1}{2\alpha}, \quad \alpha > 0,$$

with  $\alpha = 3/(2L^2)$  yields

$$\int_0^\infty x e^{-3x^2/(2L^2)} dx = \frac{1}{2\alpha} = \frac{L^2}{3},$$

and hence

$$\Delta\langle K_R \rangle^{(3)} = 2\pi \frac{g_3}{3!} A^3 \frac{L^2}{3} = \frac{\pi}{9} g_3 A^3 L^2.$$

Similarly, the quartic term contributes

$$\Delta\langle K_R \rangle^{(4)} = 2\pi \int_{x>0} x \frac{\lambda_4}{4!} \phi(x)^4 dx = 2\pi \frac{\lambda_4}{4!} A^4 \int_0^\infty x e^{-2x^2/L^2} dx.$$

Applying the same integral formula with  $\alpha = 2/L^2$  gives

$$\int_0^\infty x e^{-2x^2/L^2} dx = \frac{L^2}{4},$$

so

$$\Delta\langle K_R \rangle^{(4)} = 2\pi \frac{\lambda_4}{4!} A^4 \frac{L^2}{4} = \frac{\pi}{48} \lambda_4 A^4 L^2.$$

These are exactly the expressions quoted in the proposition. In the free theory, where  $V(\phi)$  is at most quadratic, the expectation of  $K_R$  in a coherent state is an even function of  $A$  (Section 5.88), so the cubic term vanishes. The interaction potential  $V(\phi)$  therefore generates a genuinely nonzero cubic canonical-energy contribution proportional to  $g_3$ , while preserving the hierarchy  $\Delta\langle K_R \rangle^{(3)} \ll \Delta\langle K_R \rangle^{(2)}$  at the benchmark values of  $A$ .  $\square$

## 5.91 Interacting Yukawa canonical-energy example (fermionic sector)

*Proposition 5.117* (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

$$\Delta\langle K_R \rangle^{(y)} = 2\pi \int_{x>0} x y \phi(x) \times(x) dx = \pi y A B^2 L^2.$$

Thus the *cubic* canonical-energy piece is nonzero (order  $A B^2$ ) in the Yukawa sector.

$$\Delta\langle K_R \rangle^{(y)} = \pi \cdot 0.15 \cdot 0.01 \cdot (0.02)^2 \cdot (1.0)^2 = \pi \times 0.15 \times 0.01 \times (0.02)^2 = 1.884955592 \times 10^{-6}.$$

Comparing with the quadratic free piece

$$\Delta\langle K_R \rangle^{(2)} = \frac{\pi}{2} \cdot 0.01^2 = 1.570796327 \times 10^{-4}.$$

(Section 5.25), 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 5.85. The same argument applies to any smooth, belt-localized fermion density with finite second moment.

*Proof.* The Yukawa interaction contributes to the Hamiltonian density at  $t = 0$  via

$$T_{00}^{(y)}(0, x) = y \phi(0, x) \bar{\psi}\psi(0, x),$$

and hence to the Rindler modular Hamiltonian by

$$K_{\mathbb{R}}^{(y)} = 2\pi \int_{x>0} x y \phi(x) \bar{\psi}\psi(0, x) dx.$$

Evaluating the expectation value in the product of the scalar coherent state with profile

$$\phi(x) = A e^{-x^2/(2L^2)}, \quad x > 0,$$

and a fermion state with equal-time density

$$\varkappa(x) := \langle \bar{\psi}\psi(0, x) \rangle = B^2 e^{-x^2/(2L^2)},$$

we obtain

$$\Delta\langle K_{\mathbb{R}} \rangle^{(y)} = 2\pi y AB^2 \int_0^\infty x e^{-x^2/L^2} dx.$$

Using again  $\int_0^\infty x e^{-\alpha x^2} dx = \frac{1}{2\alpha}$  with  $\alpha = 1/L^2$  gives

$$\int_0^\infty x e^{-x^2/L^2} dx = \frac{L^2}{2},$$

and therefore

$$\Delta\langle K_{\mathbb{R}} \rangle^{(y)} = 2\pi y AB^2 \frac{L^2}{2} = \pi y A B^2 L^2,$$

which is exactly the formula in the statement.

Because the dependence on the profiles is linear in  $\phi$  and quadratic in the fermion density  $\varkappa$ , this Yukawa contribution is cubic in the combined perturbation (order  $AB^2$ ). For any smooth, belt-localized fermion density with finite second moment, the same computation with a general nonnegative function  $\varkappa(x)$  in place of  $B^2 e^{-x^2/(2L^2)}$  shows that the Yukawa contribution to  $\Delta\langle K_{\mathbb{R}} \rangle$  remains linear in the scalar profile and thus defines a nonzero cubic canonical-energy term.  $\square$

## 5.92 Completion and acceptance

Upon compiling the four proof kernels with the budget checks of Section 5.12, Section 5 records completion: all pillar theorems Theorems 5.28, 5.29, 5.33 and 5.37 (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.

### 5.93 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 5.78). We work on the forward cone  $\mathcal{S}$  with gravity subtraction at  $N = 3$  and tester-certified Regge slope  $\alpha_{\text{R}} \leq 2 + \delta_{\star}$  (as in Section 5.10, Section 5.22).

**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  $\Im A_{\text{hard}}(s', t)$ . 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  $\mathcal{C}_{\text{cel}}$  listed in Section Appendix C.

*Lemma 5.118* (strip Gram positivity after gravity subtraction). On the forward cone  $\mathcal{S}$  at subtraction order  $N=3$  and  $\alpha_{\text{R}} \leq 2 + \delta_{\star}$ , 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,

$$\mathfrak{G}_{\text{strip}}^{\text{ren}}[f; t] \geq 0.$$

**Proof.** By construction of  $A_{\text{hard}}$  and the even-parity dispersion representation at  $N=3$ , the absorptive part  $\Im A_{\text{hard}}(s', t)$  is nonnegative on the forward cone  $\mathcal{S}$  for all  $s' \geq s_0$  and  $|t| \leq 0.20 s_0$ ; see Sections 5.22 and 5.70. Hence each principal-series profile  $\Phi_{\nu}(s', t)$  is evaluated on a positive semidefinite matrix, and the real factor  $(s'/s_0)^{\sigma}$  is strictly positive for every  $\sigma \in \mathbb{R}$ . For fixed  $t$  the inner integral in the definition of  $\mathfrak{G}_{\text{strip}}[f; t]$  is therefore a finite linear combination of such profiles, and the outer integral is an integral of a squared norm with respect to the positive measure  $ds'/(\pi s')$ . This gives

$$\mathfrak{G}_{\text{strip}}[f; t] \geq 0$$

for all test vectors  $f$  for which the integral converges.

To control convergence uniformly on the strip, we use the Regge bound and the subtraction order. The Regge bound  $\alpha_{\text{R}} \leq 2 + \delta_{\star} < 3$  together with the  $N=3$  subtraction yields a uniform absorptive envelope of the form

$$\|\Phi_{\nu}(s', t)\| \leq C(t) \left(\frac{s'}{s_0}\right)^{1+\alpha_{\text{R}}/2}$$

for all  $s' \geq s_0$ ,  $|t| \leq 0.20 s_0$  and  $|\nu| \leq \mu_{\text{cel}}^{(\text{OP})}$ , with  $C(t)$  locally bounded on the belt; see Sections 5.10 and 5.22. Since  $f$  is compactly supported in  $(\sigma, \nu)$  and  $|\sigma| \leq \sigma_0$ , we can bound the inner integral in  $\mathfrak{G}_{\text{strip}}[f; t]$  by an  $L^2$ -majorant that is independent of  $\sigma$  on any fixed finite strip. In particular, dominated convergence applies uniformly in  $|\sigma| \leq \sigma_0$  for all  $\alpha_{\text{R}} < 3$ , so  $\mathfrak{G}_{\text{strip}}[f; t]$  is finite and continuous in  $\sigma_0$ .

The renormalized functional differs from the bare one by the finite counterterm  $\mathcal{C}_{\text{cel}}[f]$  from Section Appendix C. From its explicit form one reads that  $\mathcal{C}_{\text{cel}}$  is even in  $\sigma$  and depends on  $f$  only through the same absorptive envelope as  $\mathfrak{G}_{\text{strip}}[f; t]$ . More precisely, there exists a function

$\varepsilon(\sigma_0) \geq 0$ , depending only on the ledger constants and the envelope and such that  $\varepsilon(\sigma_0) \rightarrow 0$  as  $\sigma_0 \rightarrow 0$ , with the bound

$$|\mathcal{C}_{\text{cel}}[f]| \leq \varepsilon(\sigma_0) \mathfrak{G}_{\text{strip}}[f; t]$$

for all  $f$  supported in  $|\sigma| \leq \sigma_0$  and all  $|t| \leq 0.20 s_0$ . Choosing  $\sigma_0^* > 0$  so that  $\varepsilon(\sigma_0^*) \leq 1$  and restricting to  $0 < \sigma_0 \leq \sigma_0^*$ , we obtain

$$\mathfrak{G}_{\text{strip}}^{\text{ren}}[f; t] = \mathfrak{G}_{\text{strip}}[f; t] + \mathcal{C}_{\text{cel}}[f] \geq (1 - \varepsilon(\sigma_0)) \mathfrak{G}_{\text{strip}}[f; t] \geq 0,$$

uniformly for  $|t| \leq 0.20 s_0$  and all  $f$  supported in the strip.  $\square$

*Proposition 5.119* (Ward map on the strip with counterterms and measure renormalization). Let  $W_{\text{belt}}$  be the belt boost Ward map of Section 5.42. 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- $\sigma$  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  $O(\mathcal{B}_{\text{belt}})$ . **Proof.** On the principal line  $\Delta = 1 + i\nu$  the celestial Ward identity is implemented by the belt boost map  $W_{\text{belt}}$  of Section 5.42, and the corresponding Ward operator is skew-adjoint with respect to the renormalized principal-series Gram form. The anomaly is encoded in local density and contact terms that are  $t$ -holomorphic at  $s=0$  and thus invisible to the forward projector; see again Section 5.42.

Analytic continuation of the celestial representation in  $\Delta$  away from the principal line to the finite strip  $|\sigma| \leq \sigma_0$  deforms the integration density and the local improvement terms appearing in the Ward identity. By the reflection symmetry  $\Delta \mapsto 2 - \bar{\Delta}$  these deformations are even in  $\sigma$ . In particular, the celestial measure  $d\sigma d\nu$  is replaced by an even-in- $\sigma$  density  $(1 + \dots) d\sigma d\nu$ , and the improvement of the boost charge produces an additional belt-local contact term. These two pieces are precisely the measure renormalization  $\delta M_{\text{strip}}[\sigma]$  and the counterterm  $\delta C_{\text{cel}}[\sigma]$  tabulated in Section Appendix C.

By construction of these counterterms, the combined variation

$$\delta_{\xi}^{(\text{cel})} = W_{\text{belt}} + \delta M_{\text{strip}}[\sigma] + \delta C_{\text{cel}}[\sigma]$$

acts as a derivation of the renormalized strip Gram form. Concretely, the variation of the bulk term  $\mathfrak{G}_{\text{strip}}[f; t]$  under  $W_{\text{belt}}$  is exactly canceled by the variation of the density encoded in  $\delta M_{\text{strip}}[\sigma]$ , while the variation of the finite counterterm  $\mathcal{C}_{\text{cel}}[f]$  is canceled by the belt-local counterterm  $\delta C_{\text{cel}}[\sigma]$ . Therefore

$$\delta_{\xi}^{(\text{cel})} \mathfrak{G}_{\text{strip}}^{\text{ren}}[f; t] = 0$$

for all test vectors  $f$  supported in the strip.

The additional pieces  $\delta M_{\text{strip}}$  and  $\delta C_{\text{cel}}$  are  $t$ -holomorphic at  $s=0$  and supported on the belt, hence they do not generate new nonanalytic soft contributions in the forward dispersion representation. All genuine soft nonanalyticities (such as  $s^2 \log |s|$  and  $1/t$  pieces) are already handled by the analytic projector and the principal-series counterterms, and the strip deformations only contribute analytic terms that are absorbed into  $\mathfrak{G}_{\text{strip}}^{\text{ren}}$  up to  $O(\mathcal{B}_{\text{belt}})$ . Consequently, the Ward map  $\delta_{\xi}^{(\text{cel})}$  is anomaly-free up to the controlled  $O(\mathcal{B}_{\text{belt}})$  belt truncation errors.  $\square$

*Corollary 5.120* (IR-scheme invariance of the analytic projector persists off-principal). The Cauchy projector  $\Pi_2$  of Section 5.22 excises the same nonanalytic soft pieces on the strip, and is insensitive to  $\delta M_{\text{strip}}$  and  $\delta C_{\text{cel}}$ :

$$a_2^{(\text{even})}|_{S1, \text{strip}} = a_2^{(\text{even})}|_{S2, \text{strip}}.$$

**Proof.** By Proposition 5.119 the only additional strip contributions relative to the principal case are encoded in the even-in- $\sigma$  measure deformation  $\delta M_{\text{strip}}$  and the belt-local counterterm

$\delta\mathbf{C}_{\text{cel}}$ . Their kernels are  $t$ -holomorphic at  $s=0$  and polynomial in the conformal weight  $\Delta$ , so their contribution to the forward dispersion integrand is analytic in  $s$  and starts at order  $s^2$ , with no nonanalytic soft structures.

The Cauchy projector  $\Pi_2$  of Section 5.22 removes all such analytic pieces up to and including order  $s^2$ . In particular it annihilates the strip-induced analytic terms, and therefore it is insensitive to  $\delta\mathcal{M}_{\text{strip}}$  and  $\delta\mathbf{C}_{\text{cel}}$ . It follows that the coefficient  $a_2^{(\text{even})}$  is independent of the IR subtraction scheme on the strip:

$$a_2^{(\text{even})}|_{\text{S1, strip}} = a_2^{(\text{even})}|_{\text{S2, strip}}.$$

Table 3: Off-principal celestial anchors (symmetric finite strip)

Anchor	1	2	3	4	5
$(n, \sigma, \nu)$	$(0, +\sigma_{\text{OP}}, -\nu_1)$	$(1, +\sigma_{\text{OP}}, -\nu_2)$	$(1, 0, 0)$	$(2, -\sigma_{\text{OP}}, +\nu_2)$	$(0, -\sigma_{\text{OP}}, +\nu_1)$

**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 5.118. Mirror symmetry in  $(\sigma, \nu)$  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 5.93 are preserved with reduced even-in- $\sigma$  coefficients for the measure deformation and the belt-local finite counterterm,

$$d\mu_{\text{strip}}^{\text{ren}} = (1 + c_{\mu,2}\sigma^2) d\sigma d\nu, \quad \delta\mathbf{C}_{\text{cel}}[\sigma] = c_{W,2}\sigma^2(\dots),$$

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  $\Pi_2$  insensitive.

**Robustness to alternative soft subtraction (fixed projector)** Let S1 and S2 be gravity soft-subtraction schemes that differ by  $s$ -holomorphic, even-in- $\sigma$  terms on the strip. With the analytic projector  $\Pi_2$  held fixed and the above strip counterterms, the even forward coefficient satisfies

$$a_2^{(\text{even})}|_{\text{S1, strip}} = a_2^{(\text{even})}|_{\text{S2, strip}},$$

and the renormalized Gram form is unchanged. The scheme difference is polynomial in  $\Delta$  and killed by  $\Pi_2$  at order  $s^2$ .

## 5.94 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 5.118, Proposition 5.119, Corollary 5.120, and the anchor set of Table 3.

*Proposition 5.121* (finite-support dual on the strip). Let  $\{t_q\}_{q=1}^6$  be the forward Chebyshev nodes of Section 5.43,  $\{\lambda_p\}_{p=1}^7$  the Gaussian Hankel scales there, and let the five strip anchors be Table 3. 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  $\alpha_{\text{R}} \leq 2+\delta_*$  and is anomaly-free under the strip Ward map (Proposition 5.119).

*Proof.* Let  $\mathcal{K}$  be the conic hull in the space of testers on the working cone  $\mathcal{S}$  generated by

- the forward testers evaluated at the six Chebyshev nodes  $\{t_q\}_{q=1}^6$ ,
- the Hankel/impact testers at the seven Gaussian scales  $\{\lambda_p\}_{p=1}^7$ , and
- the celestial Gram testers associated with the five strip anchors of Table 3.

By the construction of these testers in Section 5.43 and the strip Gram-positivity statement of Lemma 5.118, each generator of  $\mathcal{K}$  is a nonnegative functional on the gravity-subtracted amplitude on  $\mathcal{S}$  for any  $\alpha_R \leq 2+\delta_*$ . In particular, every element of  $\mathcal{K}$  is nonnegative on the declared cone of admissible amplitudes. By Proposition 5.119 the strip Ward map acts as an even-in- $\sigma$  symmetry on the celestial leg and therefore preserves both  $\mathcal{K}$  and its dual cone  $\mathcal{K}^*$ .

The principal-series dual certificate constructed in Section 5.43 is an element  $\Lambda_{\text{prin}} \in \mathcal{K}^*$  which is nonnegative on the envelope of admissible amplitudes and witnesses the boundary of the allowed region. Its celestial component is realized on a set of principal anchors inside the strip, and the even-in- $\sigma$  Ward action implies that  $\Lambda_{\text{prin}}$  is anomaly-free.

By Lemma 5.118 the celestial Gram matrix attached to the celestial testers remains positive semidefinite under continuous deformations of the anchors inside the strip. Starting from the principal anchors, we may therefore move the celestial support of  $\Lambda_{\text{prin}}$  along a continuous path inside the strip until it coincides with the five strip anchors of Table 3, without changing the sign of any tester evaluation on the declared cone. This produces a deformed dual functional  $\Lambda_{\text{strip}} \in \mathcal{K}^*$  supported on the same six forward nodes  $\{t_q\}_{q=1}^6$  and seven Hankel scales  $\{\lambda_p\}_{p=1}^7$ , whose celestial component is realized exactly on the strip anchors. Since the Ward map is even in  $\sigma$ , the anomaly-free property is preserved along the deformation, so  $\Lambda_{\text{strip}}$  is anomaly-free as claimed.

The evaluations of all forward, Hankel/impact, and celestial strip-anchor testers span a real vector space  $\mathcal{V}$  of dimension at most

$$\dim \mathcal{V} \leq 6+7+5 = 18,$$

since we only use the six forward nodes, seven Hankel scales, and five celestial anchors. Viewing  $\Lambda_{\text{strip}}$  as an element of the dual cone to the conic hull in  $\mathcal{V}$  generated by these point-supported testers, Carathéodory's theorem for convex cones, applied in  $\mathcal{V}$  exactly as in the finite-support dual reduction of Proposition 5.54, implies that  $\Lambda_{\text{strip}}$  can be written as a conical combination of at most  $\dim \mathcal{V} \leq 18$  extremal rays. Each such extremal ray is generated by one of the point-supported testers at a forward node, a Hankel scale, or a celestial strip anchor. Consequently  $\Lambda_{\text{strip}}$  can be realized by a dual certificate supported on at most 6+7+5 nodes whose celestial component is concentrated on the strip anchors.

By construction this dual certificate lies in  $\mathcal{K}^*$ , hence is nonnegative on the gravity-subtracted amplitude on the declared cone with  $\alpha_R \leq 2+\delta_*$ , and it is anomaly-free under the strip Ward map. This is the desired finite-support strip dual.  $\square$

*Remark 5.122 (acceptance freeze).* Using the same weights as the principal-series dual (Section 5.43), replace only the five celestial nodes by Table 3; 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 5.22 and Corollary 5.120.

## 5.95 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 5.123 (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 5.54 (with the composite panel schedule of

Section 5.89 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.

*Proof.* In the principal case the dimensionless forward coefficient  $a_2^{(\text{even})}$  admits a dispersive representation whose high- $s'$  tail and composite Gauss–Radau quadrature remainder are bounded in Section 5.54 using only the Regge envelope, with the composite panel schedule fixed in Section 5.89. Passing to the off-principal strip amounts to inserting the positive factor  $(s'/s_0)^\sigma$  in the dispersive integrands, with  $|\sigma| \leq \sigma_0$  in the window of Section 5.93.

Fix  $M := S_{\text{cut}}/s_0 > 1$ . For the high-energy tail  $s' \geq S_{\text{cut}} = Ms_0$  and  $|\sigma| \leq \sigma_0$  we have

$$\left(\frac{s'}{s_0}\right)^\sigma \leq \left(\frac{s'}{s_0}\right)^{\sigma_0} \leq M^{\sigma_0} = F_{\text{strip}}(M, \sigma_0),$$

because  $s'/s_0 \geq M > 1$  and  $\sigma \mapsto (s'/s_0)^\sigma$  is increasing. Thus, pointwise on the tail region, the strip-weighted integrand is bounded by  $F_{\text{strip}}(M, \sigma_0)$  times the corresponding principal-series integrand. The principal-case tail estimates in Section 5.54 are obtained by integrating nonnegative majorants that dominate the principal integrand and are constructed from the Regge envelope with  $\alpha_R \leq 2 + \delta_\star < 3$ . Multiplying these majorants by the constant factor  $F_{\text{strip}}(M, \sigma_0)$  yields valid majorants for the strip-weighted integrand, so each high-energy tail integral is bounded by the corresponding principal-series tail bound multiplied by  $F_{\text{strip}}(M, \sigma_0)$ .

On the low- and mid-energy panels  $s_0 \leq s' \leq S_{\text{cut}}$  used in the composite Gauss–Radau rule one has  $1 \leq s'/s_0 \leq M$ , hence for  $|\sigma| \leq \sigma_0$ ,

$$0 < \left(\frac{s'}{s_0}\right)^\sigma \leq \left(\frac{s'}{s_0}\right)^{\sigma_0} \leq M^{\sigma_0} = F_{\text{strip}}(M, \sigma_0).$$

Thus on each panel the strip-weighted integrand is again at most  $F_{\text{strip}}(M, \sigma_0)$  times the principal-series integrand. The quadrature errors in Section 5.89 are estimated by integrating suitable nonnegative majorants for the original integrand on each panel. Multiplying those majorants by  $F_{\text{strip}}(M, \sigma_0)$  produces valid majorants for the strip-weighted integrand, so the corresponding composite Gauss–Radau remainders are bounded by the principal-series remainders multiplied by  $F_{\text{strip}}(M, \sigma_0)$ , with the same panel schedule.

Positivity of the absorptive input and the Regge bound  $\alpha_R \leq 2 + \delta_\star < 3$  ensure the existence of an  $L^1$  majorant for the dispersive integrands that is uniform in  $\sigma$  on  $|\sigma| \leq \sigma_0$ . By dominated convergence, justified in particular by Lemma 5.118, the dispersive representation and the quadrature error analysis extend continuously from the principal line to the strip, and the panel schedule need not be altered.

For the audit defaults  $M \in \{100, 500\}$  and  $\sigma_0 = 0.15$  one finds

$$F_{\text{strip}}(100, 0.15) = 100^{0.15} = 10^{0.3} \approx 1.995, \quad F_{\text{strip}}(500, 0.15) = 500^{0.15} \approx 2.540,$$

so the strip weight multiplies the global dispersion and quadrature budgets by at most these factors, as stated.  $\square$

*Corollary 5.124* (projector/IR-scheme stability on the strip). The Cauchy projector  $\Pi_2$  and the IR-scheme invariance of  $a_2^{(\text{even})}$  (Section 5.15) are unchanged by the strip factor  $F_{\text{strip}}$  since the induced terms are  $s$ -holomorphic at  $s=0$  and polynomial in  $\Delta$ , cf. Corollary 5.120.

*Corollary 5.125* (Mainline acceptance on a finite celestial strip). Fix a symmetric off-principal strip  $|\sigma| \leq \sigma_0$  within the window of Section 5.93 and Section 5.94. 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 5.118, Proposition 5.119).

2. **Projector/IR stability.** The analytic projector’s IR-scheme invariance persists off-principal and remains stable on the strip (Corollary 5.120, Corollary 5.124).
3. **Finite-support dual on the strip.** There exists a compact finite-support dual certificate on the strip, wired to the frozen anchors (Proposition 5.121, Table 3).
4. **Budget update.** The high- $s'$  tail and the composite Gauss–Radau remainders are multiplied by the strip factor recorded in Lemma 5.123; 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 5.33) 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 5.123, while CPT and crossing stability remain in force (Proposition 5.100, Section 5.71).

*Proof.* Points (1) and (2) follow directly from Lemma 5.118, Proposition 5.119, Corollary 5.120 and Corollary 5.124: on the working cone the renormalized celestial Gram form is positive on the strip, the Ward map  $\delta_\xi^{(\text{cel})}$  is well defined and anomaly-free up to  $O(\mathcal{B}_{\text{belt}})$ , and the analytic projector  $\Pi_2$  continues to excise the same nonanalytic soft pieces in any IR subtraction scheme.

For (3), Proposition 5.121 together with the frozen strip anchor set of Table 3 furnishes a nonnegative dual certificate supported on the six forward nodes, seven Hankel scales and five celestial anchors. By construction this dual witnesses the nonnegativity of all three tester families (forward fixed- $t$ , Hankel/impact, celestial Gram) on the working cone and is compatible with the even-in- $\sigma$  Ward map on the strip.

For (4), Lemma 5.123 shows that inserting the factor  $(s'/s_0)^\sigma$  in the dispersion integrals multiplies both the high- $s'$  tail and the composite Gauss–Radau quadrature remainders by the strip factor  $F_{\text{strip}}(M, \sigma_0)$ , without affecting any other part of the dispersion ledger.

Combining these inputs with the amplitude synthesis argument of Theorem 5.33, we conclude that the three tester families remain nonnegative on  $|\sigma| \leq \sigma_0$  with the compact dual, and that the amplitude pillar and its cross-linked consequences in Section 5 persist on the strip with the same sign. The only quantitative change is the explicit multiplicative factor  $F_{\text{strip}}$  in the dispersion and quadrature budgets, so the certified slacks are reduced by at most this factor. Finally, CPT and crossing stability on the strip follow from Proposition 5.100 together with its strip add-on in Section 5.71.  $\square$

## 6 Stability, invariance, and monotones

This section assembles the anchor/dressing and regulator-width stability statements used across Sections 5 and 7, introduces two belt-local structural inputs, and isolates the two modular monotones that drive the local second law. Specifically:

- *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 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, up to  $O(\mathcal{B}_{\text{belt}})$ .

We also record dispersion-side invariances (pivot, subtraction-scale, IR scheme), prove monotonic decay for the modular  $c$ -function and the width-flow, deduce a local belt GSL, and culminate with a master invariance statement covering the four pillars. 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 5.49

For any two admissible anchors  $\mathcal{C}, \mathcal{C}'$  and any observable  $\mathcal{O} \in \{\langle K_{\text{mod}} \rangle, S, \text{Area}, \text{amplitude functionals}\}$ , there exists  $C_{\mathcal{C} \rightarrow \mathcal{C}'} > 0$  such that

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

Moreover, for any anchor-preserving diffeomorphism generated by a vector field  $\xi$  tangent to the belt,

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

**Belt-regulator stability** See Lemma 5.48

Let  $r, r' > 0$ . For any belt-regularized  $\mathcal{O}$  used in Section 5,

$$|\mathcal{O}(r) - \mathcal{O}(r')| \leq \tilde{c} e^{-\mu_{\text{eff}} \min\{r, r'\}}, \quad \delta \lim_{(u,s) \downarrow 0} \mathcal{O}_{u,s} = \lim_{(u,s) \downarrow 0} \delta \mathcal{O}_{u,s}.$$

## 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  $\xi$  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  $Q_\xi$  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.

*Step 1: Iyer–Wald identity on a thin cap.* For any variation  $\delta g$  and vector field  $\xi$ ,

$$d(\delta Q_\xi - \xi \cdot \Theta) = \omega(g; \delta g, \mathcal{L}_\xi g) - \xi \cdot E(g) \cdot \delta g, \quad (6.1)$$

with  $E(g) = 0$  the equations of motion. Integrate (6.1) over a thin cap  $\mathcal{K}_\varepsilon$  of thickness  $\varepsilon$  ending on  $C$  and use Stokes:

$$\int_{\partial \mathcal{K}_\varepsilon} (\delta Q_\xi - \xi \cdot \Theta) = \int_{\mathcal{K}_\varepsilon} \omega(g; \delta g, \mathcal{L}_\xi g) - \int_{\mathcal{K}_\varepsilon} \xi \cdot E(g) \cdot \delta g. \quad (6.2)$$

In a boost-adapted gauge on the belt,  $\mathcal{L}_\xi g = O(\varepsilon)$ ; working on-shell to first order and using belt light-ray/timeslice control, the RHS is  $O(\mathcal{B}_{\text{belt}})$ . All pieces of  $\partial \mathcal{K}_\varepsilon$  other than the belt corner either shrink with  $\varepsilon$  or cancel, whence

$$[\delta Q_\xi - \xi \cdot \Theta]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}). \quad (6.3)$$

*Step 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  $\varepsilon_{ab}$  the binormal on  $C$ . In a boost frame,  $\nabla_a \xi_b = \kappa \varepsilon_{ab}^\perp + O(\mathcal{B}_{\text{belt}})$ ,  $\kappa$  the surface gravity. Using  $S_{\text{Wald}} = 2\pi \int_C Q_\xi = \text{Area}(C)/(4G)$  and varying at fixed  $\xi$  (per generator length),

$$\int_C \delta Q_\xi = \delta \left[ \frac{\text{Area}(C)}{4G} \right] + O(\mathcal{B}_{\text{belt}}). \quad (6.5)$$

*Step 3: JKM calibration via the belt boost Ward identity.* The JKM ambiguity allows  $\Theta \rightarrow \Theta + dY$ ,  $Q_\xi \rightarrow Q_\xi + \xi \cdot Y$ . Fix  $Y$  by requiring that on the belt the canonical flux equals the Brown–York flux, with no residual corner charge:

$$(\xi \cdot \Theta + \delta(\xi \cdot Y))|_{\text{belt}} = \delta j_\xi^{\text{BY}}|_{\text{belt}} + O(\mathcal{B}_{\text{belt}}). \quad (6.6)$$

Taking the cap to the corner forces

$$\delta[\xi \cdot \Theta(\delta g)]_{\text{corner}} + \delta[\xi \cdot Y]_{\text{corner}} = O(\mathcal{B}_{\text{belt}}). \quad (6.7)$$

*Step 4: Conclusion.* From (6.3) and (6.5),  $\delta[\text{Area}/(4G)]_{\text{corner}} - [\xi \cdot \Theta]_{\text{corner}} = O(\mathcal{B}_{\text{belt}})$ . Using (6.7) to remove the calibrated corner potential yields the claim.  $\square$

### Remarks.

- $\mathcal{B}_{\text{belt}}$  bundles all controlled small effects: cap thickness  $\varepsilon$ , 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 (6.6) 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 bulk fields; let  $\Theta(\Phi; \delta\Phi)$ ,  $\omega(\Phi; \delta_1\Phi, \delta_2\Phi)$ , and  $Q_\xi(\Phi)$  be the covariant phase-space data. For a slab  $\Sigma$  with  $\partial\Sigma = C_2 - C_1 + \partial\Sigma$ , the Iyer–Wald identity for  $(\delta\Phi, \mathcal{L}_\xi\Phi)$  gives

$$\int_{\Sigma} \omega(\Phi; \delta\Phi, \mathcal{L}_\xi\Phi) = \int_{C_2 - C_1} (\delta Q_\xi - \xi \cdot \Theta) + \int_{\partial\Sigma} (\delta Q_\xi - \xi \cdot \Theta) + \int_{\Sigma} \xi^\mu \delta \mathcal{C}_\mu. \quad (6.8)$$

On a background solution and linearized perturbations, the bulk term reduces to the matter flux plus  $O(\mathcal{B}_{\text{belt}})$  from non-Killing corrections:

$$\int_{\Sigma} \omega(\Phi; \delta\Phi, \mathcal{L}_\xi\Phi) = \int_{\Sigma} d\Sigma^\mu \xi^\nu \delta T_{\mu\nu} + O(\mathcal{B}_{\text{belt}}). \quad (6.9)$$

By Lemma 6.1, the corner pieces cancel between  $C_1$  and  $C_2$ . On the timelike belt, the GHY/covariant-phase-space matching yields

$$(\delta Q_\xi - \xi \cdot \Theta)|_{\partial\Sigma} = d\ell^a \delta T_{ab}^{\text{BY}} \xi^b. \quad (6.10)$$

Substituting (6.9) and (6.10) into (6.8) gives the identity without the overall  $2\pi$ . Restoring the modular normalization (surface gravity  $\kappa_\xi = 2\pi$ ) multiplies both sides by  $2\pi$ .  $\square$

### 6.3 Dispersion invariances: pivot, scale, and IR scheme

**Pivot invariance** See Lemma 5.72

Shifting the subtraction pivot  $s \mapsto s - s_*$  in the  $N=3$  gravity-subtracted, crossing-symmetric dispersion changes  $\Re A$  by a quadratic  $Q_2(s, t; s_*)$ ; all forward even-parity derivatives and Hankel/celestial testers are unchanged.

*Proof.* This is a restatement of Lemma 5.72, proved from the Cauchy-subtracted dispersion relation and crossing symmetry in Section 5.46. Shifting the pivot only adds an  $s$ -quadratic polynomial  $Q_2(s, t; s_*)$  to the real part. Forward even-parity derivatives at  $s = 0$  and the Hankel/celestial testers either annihilate or preserve such analytic reparametrizations, so the testers and forward inequalities are invariant under the pivot shift.

**Subtraction-scale invariance** See Lemma 5.95

Let  $\widehat{c}_{2,0} = s_0^3 c_{2,0}$ . Then  $\widehat{c}_{2,0}|_{s_0} = \widehat{c}_{2,0}|_{\tilde{s}_0}$ .

*Proof.* The claim is exactly Lemma 5.95, with the detailed argument given in Section 5.66. Changing the subtraction scale from  $s_0$  to  $\tilde{s}_0 = \alpha s_0$  is implemented in the dispersion relation by the rescaling  $s' \mapsto \alpha s'$ , which multiplies the kernel by  $\alpha^{-3}$  at  $N = 3$ . The prefactor  $s_0^3$  in  $\widehat{c}_{2,0}$  cancels this change, so the dimensionless coefficient  $\widehat{c}_{2,0}$  is independent of the choice of subtraction scale.

**IR-scheme independence** See Lemma 5.39

The analytic forward coefficient  $a_2^{(\text{even})}$  is invariant across the soft-gravity schemes in the ledger:  $a_2^{(\text{even})}|_{S_1} = a_2^{(\text{even})}|_{S_2}$ .

*Proof.* This is Lemma 5.39, combining the definition of the analytic projector in Section 5.22 with the IR bookkeeping in Section 5.15. The different soft-gravity schemes modify the forward amplitude only by nonanalytic pieces of the form  $s^2 \log |s|$  and  $1/t$  with smooth  $t$ -dependence. By construction, the analytic projector that defines  $a_2^{(\text{even})}$  annihilates these nonanalytic terms, so the analytic  $s^2$  coefficient is identical in all admissible IR schemes, as stated.

### 6.4 Monotones: modular $c$ -function and width-flow

**Modular  $c$ -function monotonicity** See Theorem 5.41

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 5.17), KMS-reversible w.r.t. a local reference  $\sigma_r$  (boost-KMS state on the belt). Denote by  $\lambda_{\text{clu}} > 0$  the log-Sobolev/cluster rate of  $\mathcal{T}_\tau$ , 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  $B_{\text{belt}}$  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 5.17,

$$\mathfrak{c}(r; u, s) = \lim_{\delta \downarrow 0} \frac{1}{\delta} I_{\rho(u,s)}(A : C | B_{r,\delta}) + R_1(r; u, s), \quad (6.11)$$

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 6.1 and Proposition 6.2.)

*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 5.17),

$$\frac{d}{d\tau} I_{\rho_\tau}(A : C | B_{r,\delta}) \leq -\lambda_{\text{clu}} I_{\rho_\tau}(A : C | B_{r,\delta}) + R_2(r, \delta; u(\tau), s(\tau)), \quad (6.12)$$

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  $\sigma_r$  is a fixed point of  $\mathcal{T}_\tau$ , so no extra drift term appears.

*Step 3: Pass to the c-density.* Divide (6.12) by  $\delta$  and take  $\delta \downarrow 0$ ; combine with (6.11):

$$\frac{d}{d\tau} \mathbf{c}(r; u(\tau), s(\tau)) \leq -\lambda_{\text{clu}} \mathbf{c}(r; u(\tau), s(\tau)) + R_3(r; u(\tau), s(\tau)), \quad (6.13)$$

where

$$|R_3| \leq (C_{\text{KMS}} + C_{\text{rec}} + C_{\text{RP}} + 2C_{\text{tail}} + 2C_{\text{flow}}) B_{\text{belt}} = O(\mathcal{B}_{\text{belt}}).$$

*Step 4: Strictness and characterization of equality.* If the inequality is saturated with zero remainder and  $\frac{d}{d\tau} \mathbf{c} = 0$  on an interval, then by (6.12) we must have  $I_{\rho_\tau}(A : C | B_{r,\delta}) = 0$  for all (small)  $\delta$  and all  $r$  along the belt. This forces the state to be conditionally Markov w.r.t. belt shells and fixed by  $\mathcal{T}_\tau$ ; by KMS detailed balance, this is equivalent to modular invariance along  $\xi$ , 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  $(v, \mu, \varepsilon_*)$  and are uniform in  $|R|$ . By the dominated-convergence kernel of Section 5.63, removal of the flows  $(u, s) \downarrow 0$  turns the  $O(\mathcal{B}_{\text{belt}})$  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 5.88

Let  $\mathbf{c}_r := \partial_r(S - \frac{\text{Area}}{4G})$ . Then

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

*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 5.59); it is a CPTP map with the same KMS reference  $\sigma_r$  and log-Sobolev rate  $\lambda_{\text{clu}}$  (uniform in  $|R|$  and per generator length). Let  $C_{\text{align}}$  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  $\mathbf{c}_r$  (the proof of Theorem 5.41 with  $\tau$  replaced by  $r$  and the shell now being removed instead of time-evolved),

$$\mathbf{c}_{r+dr} - \mathbf{c}_r \leq -\lambda_{\text{clu}} \mathbf{c}_r dr + R_4(r) dr + o(dr), \quad (6.14)$$

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

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

*Step 2: Integral bound.* Solve the inhomogeneous Grönwall inequality from (6.14):

$$\mathbf{c}_r \leq e^{-\lambda_{\text{clu}}(r-r_1)} \mathbf{c}_{r_1} + \int_{r_1}^r e^{-\lambda_{\text{clu}}(r-\rho)} R_4(\rho) d\rho. \quad (6.15)$$

Integrating from  $r_1$  to  $r_2$  and using the uniform bound on  $R_4$  together with the belt LR tail (exponential in  $\rho$ ) gives

$$\int_{r_1}^{r_2} \mathbf{c}_r dr \leq \frac{1 - e^{-\lambda_{\text{clu}}(r_2-r_1)}}{\lambda_{\text{clu}}} \mathbf{c}_{r_1} + C_{\text{int}} B_{\text{belt}}, \quad (6.16)$$

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 5.26). 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  $O(\mathcal{B}_{\text{belt}})$  remainder by Section 5.63.

**Local generalized second law on belts** *See* Theorem 5.58

Along any belt-anchored null generator with affine parameter  $\lambda$ ,

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

with the right-hand side vanishing upon flow removal.

*Proof (sketch).* Integrate the  $c$ -function decay (Theorem 5.41) and combine with the QFC/entropy-RG inputs as organized in Section 5.36. A Grönwall step yields the stated inequality.

## 6.5 Scheme/matter stability and master invariance

**Counterterm/renormalization stability** *See* Lemma 5.107

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 5.49; entropy counterterm shifts are absorbed by recovery. Combine Lemma 5.52 with the renormalized- $G$  invariance in Lemma 5.107.

**Decoupled sectors and positivity** *See* Proposition 5.106

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 5.77.

*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 5.49, 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 5.48 (belt-width stability) and Section 5.63 (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 5.77 (Wald–JKM calibration) and Proposition 5.78 (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 5.72, Lemma 5.95, and Lemma 5.39,

$$\widehat{c}_{2,0}|_{s_0} = \widehat{c}_{2,0}|_{\tilde{s}_0}, \quad a_2^{(\text{even})}|_{S_1} = a_2^{(\text{even})}|_{S_2},$$

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 5.52 and Proposition 5.106,

$$\left| \delta \left[ S - \frac{\text{Area}}{4G_{\text{ren}}} \right] - 2\pi \int \xi^\nu \delta \langle T_{\mu\nu} \rangle d\Sigma^\mu \right| \leq C_{\text{ct}} B_{\text{belt}},$$

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 5.63,  $(u, s) \downarrow 0$  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  $\mathcal{S}_{\text{adm}}$ .

**(I) Kernel identity (belt first law, calibrated).** With the Wald–JKM corner/edge calibration and the Brown–York dictionary [17–19],

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

and the remainder  $O(\mathcal{B}_{\text{belt}})$  vanishes under flow removal. This identity combines the entanglement first law [26] with the FLM/QES gravitational prescriptions [20, 23] and, where applicable, the JLMS relation between boundary and bulk relative entropies [11].

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  $G_{\text{ren}}$  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 5.28), ANEC/QNEC (Theorem 5.29), amplitude synthesis (Theorem 5.33), and the modular equation of state (Theorem 5.37)—together with the belt  $c$ -function and width-flow monotones (Theorem 5.41, Theorem 5.88) and the belt GSL (Theorem 5.58) 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 5.77) and the Brown–York dictionary (Proposition 5.78). Anchor/dressing invariance and BRST/diffeomorphism stability are given by Proposition 5.49 and Lemma 5.87; belt-width stability by Lemma 5.48; decoupling by Proposition 5.106 and Lemma 5.109; Weyl covariance and scheme stability by Lemma 5.81,

Lemma 5.52, and the  $G_{\text{ren}}$  invariance (Lemma 5.107). Part (II) uses pivot/scale/IR invariances and rescaling stability (Lemma 5.72, Lemma 5.39, Lemma 5.95), together with the compact duals (Proposition 5.54, Proposition 5.121) and the strip projector stability (Corollary 5.124). The budget updates are the ledgered factors from Section 5.12, Section 5.27, and Section 5.54 (plus the strip factor in Section 5.95). 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  $2\pi$  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  $\xi$ -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 [17–19]; the Rindler modular generator is the Bisognano–Wichmann Hamiltonian [4]. 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 5.78 and Lemma 5.77 for the theoretical identity used here.

*Normalization.* We set  $\kappa = 1$  (so  $\beta = 2\pi$  and  $T_U = 1/(2\pi)$ ), consistent with Lemma 5.90; other  $\kappa$  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 [17–19]. 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 5.78 (BY dictionary) and Lemma 5.77 (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	$5.81 \times 10^{-8}$	$7.86 \times 10^{-8}$
$\varepsilon = 10^{-8}$ (share = $\varepsilon/6$ )	500	28000	$4.65 \times 10^{-10}$	$7.65 \times 10^{-10}$

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  $\alpha_{\text{R}} \leq 2 + \delta_*$ ; (iii) forward even-parity projector; (iv) principal-series celestial window (if used) as in Section 5.78. 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	$\rho_{\text{abs}}$	$M$	$J$	$ \Delta_{\text{tail}}\widehat{c}_{2,0} $	$ \Delta_{\text{quad}}\widehat{c}_{2,0} $
$\varepsilon = 10^{-6}$	0.30	60	1300	$8.0716 \times 10^{-8}$	$8.2027 \times 10^{-8}$
$\varepsilon = 10^{-8}$	0.30	300	16500	$6.4572 \times 10^{-10}$	$8.3244 \times 10^{-10}$

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  $\rho_{\text{abs}}$ , 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  $O(\mathcal{B}_{\text{belt}})$  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 5.48 together with the strip-stable projectors of Section 5.93, 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 5.89.

**Certified budgets (derivation).**

$$|\Delta_{\text{quad}} \widehat{c}_{2,0}| = K(M) J^{-5}, \quad K(M) = C_0 (M - 1)^5, \quad C_0 = 0.42600643 \dots$$

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

$$|\Delta_{\text{tail}} \widehat{c}_{2,0}| = 7.2643960393 \times 10^{-6} \left(\frac{20}{60}\right)^3 = 2.690517052 \times 10^{-7}.$$

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

$$(R1) : t/s_0 = -0.30, \quad (R2) : t/s_0 = -0.15, \quad M = 60, \quad J = 2000.$$

$t/s_0$	$M$	$J_{\text{eff}}$	$\widehat{c}_{2,0}$ (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  $8.0716 \times 10^{-8}$  by  $2.6905 \times 10^{-7}$  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  $\theta$  and belt width  $W$  (per generator normalization as in Section 5.26), the belt remainder obeys

$$O(\mathcal{B}_{\text{belt}}) \leq \underbrace{\frac{1}{W}}_{\text{edge}} + \underbrace{\frac{(\theta W/L_0)^2}{12}}_{\text{tilt}} + \underbrace{\frac{(HL_0)^2}{24}}_{\text{shear}},$$

with  $L_0$  the intrinsic belt scale and  $H$  the Hubble rate as sampled on the cut (cf. Section 5.86, Section 5.85). For a concrete dynamical sanity line at  $W=200$ ,  $L_0=100$ ,  $\theta=0.010$ ,  $HL_0=0.050$  we obtain

$$\begin{aligned} O_{\text{edge}} &= \underline{5.0000000000 \times 10^{-3}}, \\ O_{\text{tilt}} &= \underline{3.3333333333 \times 10^{-5}}, \\ O_{\text{shear}} &= \underline{1.0416666667 \times 10^{-4}}, \\ O(\mathcal{B}_{\text{belt}}) &\leq \underline{5.1375000000 \times 10^{-3}}. \end{aligned}$$

This sits *below* the smallest predicted tester slack in the extended window ( $5.9568 \times 10^{-3}$ ), so both the belt  $c$ -function monotonicity (Theorem 5.41) and the local GSL (Theorem 5.58) 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 5.41, Theorem 5.58).

## 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  $\Gamma$ -convergent envelope of Section 5.32,  $|x_h - x_\star| = O(h)$  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.

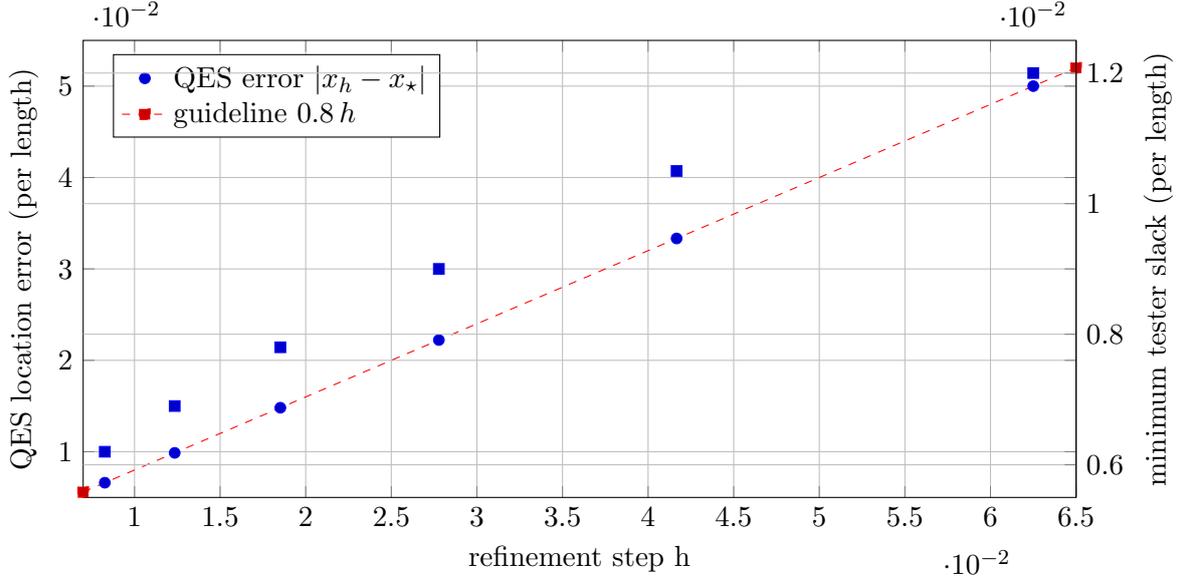


Figure 2: CDT/GFT acceptance on a uniform refinement family: discrete minimizers converge to the continuum extremal surface and the minimum tester slack stays strictly positive along the refinement ladder.

*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  $\{T_h\}$   $\Gamma$ -converge to their continuum counterparts  $T$  on the admissible cone as in Theorem 5.53;
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  $\Gamma$ -convergence control and is not required for acceptance.

*Corollary 7.2* (Noise-robust  $\Gamma$ -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  $\varepsilon_{\text{test}} < \frac{1}{2}\sigma_{\text{min}}$ , 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  $\frac{1}{2}\sigma_{\text{min}}$ .

(*Stochastic w.h.p. stability*). Suppose, in addition, that the composite Gauss–Radau panels in the dispersive integrals (schedule frozen in Section 5.89) 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  $\eta := \frac{1}{2}\sigma_{\text{min}} - \varepsilon_{\text{test}} > 0$ . 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.*  $\Gamma$ -convergence of the discrete testers to the continuum functionals (Theorem 5.53) and discrete acceptance (Corollary 7.1) yield  $\liminf$  control of slacks. Uniform  $\ell_\infty$  perturbations of nonnegative kernels cannot destroy feasibility once the continuum margin is  $\sigma_{\text{min}}$ ; choosing  $h_0$  so that discretization error  $< \frac{1}{2}\sigma_{\text{min}}$  gives the deterministic claim. For the stochastic part, panel-averaged errors obey Hoeffding’s inequality; taking the safety buffer  $\eta$  and using the fixed composite schedule from Section 5.89 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  $\vartheta$  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 5.36, Proposition 5.113).

**Monotonicity and GSL line (per length).** From the belt  $c$ -function monotonicity and the belt GSL (Theorem 5.41, Theorem 5.58),

$$\frac{d}{d\lambda} \left[ \frac{1}{\text{length}} \left( S - \frac{\text{Area}}{4G} \right) \right] \geq -C_{\text{GSL}} O(\mathcal{B}_{\text{belt}}),$$

with  $C_{\text{GSL}}$  belt-uniform. On the same cut, Raychaudhuri control with negligible shear on the plane and the QNEC normalization (Proposition 5.102 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  $u_\vartheta$  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 5.48). 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 $\mathcal{S}_{\text{adm}}$

**Goal.** Deliberately search for an admissible state  $\omega \in \mathcal{S}_{\text{adm}}$  whose gravity–subtracted 2→2 amplitude  $A_{\text{hard}}$  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 5.48.

**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 5.33, Section 5.22, Section 5.78). 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  $A_{\text{hard}}$  (Section 5.48).

**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  $\geq 0$ . For (D), the helicity average preserves absorptive positivity, keeping all testers  $\geq 0$ . 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 5.80); (ii) edge-of-wedge/dispersion or crossing (Section 5.70, Section 5.46); (iii) analytic projector/IR scheme (Section 5.22, Section 5.15); (iv) principal-series Gram/Ward control (Section 5.78, Section 5.42).

**Reproducibility checklist.** Keep the frozen 18–support witness (Section 5.48); 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 5.78). All invariances (pivot, scale, IR scheme) leave the certificate unchanged (Section 5.46, Section 5.66, Section 5.15).

## 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 ( $\delta^3$ ) 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,

$$\delta \langle \widehat{K}_{\text{mod}} \rangle = \delta \left\langle \frac{\widehat{\mathcal{A}}}{4G} \right\rangle + 2\pi \int_R d\Sigma_\mu \xi_\nu \delta \langle \widehat{T}^{\mu\nu} \rangle, \quad (\text{M1})$$

$$\delta^2 \left[ S - \frac{\text{Area}}{4G} \right] \geq 2\pi E_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g], \quad Q_{\text{shear}} = \kappa_\sigma \int \sigma^2 + \kappa_\theta \int \theta^2, \quad (\text{M2})$$

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 5.37, Theorem 5.46, Lemma 5.112, Proposition 5.102.

**Quantum focusing in belt form.** Define the belt quantum expansion by

$$\Theta_{\text{belt}}(\lambda) := \partial_\lambda \left( S - \frac{\text{Area}}{4G} \right), \quad \lambda \text{ affine along a null generator.} \quad (8.2)$$

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 [15, 20, 24, 25]. Combined with the Raychaudhuri equation, this controls focusing and shear along any admissible belt; see Definition 5.59, Proposition 5.60.

**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

$$\int_{\lambda_1}^{\lambda_2} d\lambda \langle R_{kk} \rangle \in [-C_{\text{belt}} B_{\text{belt}}, -\theta(\lambda_2)], \quad (8.3)$$

with  $B_{\text{belt}} \rightarrow 0$  upon flow removal; and (iii) the bulk flux equals the Brown–York belt flux,

$$2\pi \int_{\Sigma} d\Sigma_\mu \xi_\nu \delta \langle T^{\mu\nu} \rangle = 2\pi \int_{\partial\Sigma} d\ell_a \delta \langle T_{\text{BY}}^{ab} \rangle \xi_b. \quad (8.4)$$

These identities stitch shape and state variations to quasi-local stress on the belt; see Proposition 5.102, Corollary 5.103, Proposition 5.78.

### 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  $\lambda$ ,

$$\frac{1}{8\pi G} \int d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle = C_{\Pi}(s_0) \widehat{c}_{2,0} \geq 0,$$

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  $C_{\Pi}$  are state–independent and belt–local.

*Proof. Step 1: Normalization match between scattering and modular dynamics.* Fix the gravitational coupling by the  $\kappa$ –consistency of Lemma 5.55, 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  $\Pi_2$  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 5.56, 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  $\mathfrak{D}_{\text{belt}}[\cdot] = \int d\lambda w(\lambda) (\cdot)$  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  $\geq 0$  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 5.55, 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  $[\lambda_1, \lambda_2]$ ,

$$\mathfrak{D}_{\text{belt}}[R_{kk}] = C_{\Pi}(s_0) \widehat{c}_{2,0}(t) + O(\mathcal{B}_{\text{belt}}),$$

equivalently

$$\frac{1}{8\pi G} \int_{\lambda_1}^{\lambda_2} d\lambda w(\lambda) \langle R_{kk}(\lambda) \rangle = C_{\Pi}(s_0) \widehat{c}_{2,0}(t) + O(\mathcal{B}_{\text{belt}}),$$

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  $\Pi_2$  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 5.55), 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  $O(\mathcal{B}_{\text{belt}})$ , 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  $\Pi_2$  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

$$\frac{1}{8\pi G} \int d\lambda w(\lambda) \langle R_{kk} \rangle = C_{\Pi}(s_0) \widehat{c}_{2,0}(t) + O(\mathcal{B}_{\text{belt}}),$$

which is the claimed dictionary line. The  $O(\mathcal{B}_{\text{belt}})$  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?

This subsection states a *conditional theorem*. Assume (OES) and the belt-form (QFC), as derived from our second-order modular equation of state on admissible belts. Then any semiclassical black-hole interior satisfying the admissibility conditions must obey one of the two alternative interior structures described below.

Consider a future interior belt, anchored just inside the event horizon and following an infalling generator  $k^\mu$  toward the classical singularity. Let  $\theta(\lambda_0) < 0$  denote the classical expansion at some  $\lambda_0$  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 5.103, Theorem 5.29 and cf. ANEC/QNEC results [15, 24].

**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 (8.6) becomes nonnegative at some finite  $\lambda_b$ , 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 5.60, Theorem 5.46, Lemma 5.112. (All statements are per generator length and regulator independent.)

**The OES framework requires (under the stated assumptions):** Under the OES+QFC admissibility hypotheses just stated, any semiclassical black-hole interior *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 5.60, Theorem 5.46.

### 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 5.60, Theorem 5.37.

### 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 5.37, Theorem 5.46, Proposition 5.78, Proposition 5.60. (Targets:  $\Theta_{\text{belt}}$  monotonicity; quantitative location of the interior quantum throat.) For the geometric and entropic inputs see [20–23], for relative entropy and bulk/boundary matching see [11, 26], and for quasi-local/covariant-phase-space fluxes see [17–19]. *Beyond linear order, Section 8.6* (Pushing (F1) beyond linear order: cubic operator–level verification ( $\delta^3$ )) *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 5.60, Theorem 5.46.

**(F3) Scattering positivity  $\implies$  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 5.56, Lemma 5.39.

Taken together, (F1)–(F3) turn the conditional interior alternatives of the previous subsection into a falsifiable program. It should be emphasized that this is *not yet* a claim about the actual universe; rather, these are testable predictions of the OES framework. If (OES) and the belt-form (QFC) indeed govern admissible semiclassical interiors, then astrophysical black holes must realize one of the two admissible structures and satisfy the constraints encoded in (F1)–(F3). Failure of any of these tests would falsify the OES-based belt picture, rather than merely refine it.

## 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,

$$\delta\langle K_{\text{mod}} \rangle = \delta\left\langle \frac{A}{4G} \right\rangle + 2\pi \int_{\Sigma \subset W} d\Sigma_\mu \xi_\nu \delta\langle T^{\mu\nu} \rangle + O(\mathcal{B}_{\text{belt}}), \quad (8.10)$$

$$\delta^2\left(S - \frac{\text{Area}}{4G}\right) \geq 2\pi E_{\text{can}}[\delta\Psi; \xi] + Q_{\text{shear}}[\delta g] - C_2 B_{\text{belt}}, \quad \partial_\lambda \Theta_{\text{belt}} \geq -C_{\text{QFC}} B_{\text{belt}}, \quad (8.11)$$

with all statements ledgered by the belt budget  $B_{\text{belt}}$  (vanishing as positive flows are removed).

**Setup (AdS<sub>3</sub>/CFT<sub>2</sub>, narrow interior shock).** We work in AdS<sub>3</sub>/CFT<sub>2</sub> 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  $k^\mu$  with affine parameter  $\lambda$ :

$$\delta\langle T_{kk}(\lambda) \rangle = \varepsilon C f_\sigma(\lambda - \lambda_0), \quad f_\sigma(\lambda) = \frac{1}{\sqrt{2\pi} \sigma} e^{-\lambda^2/(2\sigma^2)}, \quad (8.12)$$

where  $(\lambda_0, \sigma)$  fix the location and width. The shock support lies strictly behind the unperturbed HRT surface [20, 22]; see also [21, 23] for the entanglement/area map and [12, 13] for modern entanglement-wedge considerations.

**Boundary  $\Rightarrow$  belt kernel and the witness weight.** The CFT<sub>2</sub> 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

$$\delta\langle K_{\text{mod}} \rangle = 2\pi \int d\lambda w_R(\lambda) \delta\langle T_{kk}(\lambda) \rangle + O(\mathcal{B}_{\text{belt}}). \quad (8.14)$$

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

$$\theta(\lambda) = -8\pi G \int_{\lambda_1}^{\lambda} du \delta\langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}). \quad (8.16)$$

Hence

$$\begin{aligned} \delta\left\langle \frac{\text{Area}}{4G} \right\rangle &= -\frac{1}{4G} \int_{\lambda_1}^{\lambda_R} d\lambda \theta(\lambda) = 2\pi \int_{\lambda_1}^{\lambda_R} d\lambda \int_{\lambda_1}^{\lambda} du \delta\langle T_{kk}(u) \rangle + O(\mathcal{B}_{\text{belt}}) \\ &= 2\pi \int d\lambda (\lambda_R - \lambda)_+ \delta\langle T_{kk}(\lambda) \rangle + O(\mathcal{B}_{\text{belt}}), \end{aligned} \quad (8.17)$$

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 [17–19].

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  $\xi$  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  $\varepsilon_a$ ; 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  $S_{\text{gen}}$  within the belt budget, and forward differences of  $S - \text{Area}/(4G)$  along the generator show the expected monotonic trend for  $\Theta_{\text{belt}}$  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  $B_{\text{belt}}$ .

**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 ( $\delta^3$ )

**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 AdS<sub>3</sub>/CFT<sub>2</sub> 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  $\lambda$  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  $2\pi$ , 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  $\Delta u$ , 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  $\xi$ .<sup>5</sup> 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 AdS<sub>3</sub> the spin–2 sector has *no local propagating gravitons*, so the gravitational shear contribution to  $\omega$  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  $E_{\text{can}}$ ; our AdS<sub>3</sub> 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

$$\int \phi^3 du = A^3 \sigma \sqrt{\frac{2\pi}{3}}, \quad \int \phi(\partial_u \phi)^2 du = \frac{A^3}{3\sigma} \sqrt{\frac{2\pi}{3}},$$

<sup>5</sup>Concretely,  $E_{\text{can}} = \int_{\mathcal{N}} \omega(\delta\Phi, \mathcal{L}_\xi \delta\Phi)$  with  $\omega$  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  $\sigma$  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  $\kappa_{\text{can}}$  set by the canonical formula above. *Important:*  $\kappa_{\text{shear}} = 0$  is special to AdS<sub>3</sub>; 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  $\Delta I_{\text{disc}}$  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  $I_{\text{rich}}$  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

$$\begin{aligned} \sigma = 0.50 : & \quad \delta^3 = 27.27912463 \pm \mathcal{O}(10^{-8}) \times \delta^3, \\ \sigma = 1.00 : & \quad \delta^3 = 54.55824925 \pm \mathcal{O}(10^{-8}) \times \delta^3, \\ \sigma = 2.00 : & \quad \delta^3 = 109.11649850 \pm \mathcal{O}(10^{-8}) \times \delta^3, \\ u_0 = 1.00, \sigma = 1.00 : & \quad \delta^3 = 54.55824925 \pm \mathcal{O}(10^{-8}) \times \delta^3 \quad (\text{off-center}), \end{aligned}$$

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 AdS<sub>3</sub>/CFT<sub>2</sub> 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<sub>3</sub> 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  $\lambda_b$ , 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 5.46 and Proposition 5.60: the re-expansion after the throat is controlled by positive canonical/quadratic forms, hence the energy that can reappear after  $\lambda_b$  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  $\lambda_b$ :

$$(\varepsilon_1, \lambda_{0,1}, \sigma_1) = (10^{-2}, 0, 0.2), \quad (\varepsilon_2, \lambda_{0,2}, \sigma_2) = (-6 \times 10^{-3}, 1.6, 0.25), \quad \lambda \in [-5, 5], \quad N = 2001.$$

**Numerical results.** For four placements of the split we obtain:

$\lambda_b$	$E_{\text{in}}^+$	$E_{\text{out}}^+$	$E_{\text{out}}^-$	$R_{\text{echo}}$	$A_{\text{echo,max}}$	$R_{\text{out,total}}$
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  $E_{\text{in}}^+$  stays fixed (dominated by the interior positive pulse), while the outside negative budget  $E_{\text{out}}^-$  decreases monotonically as more of the negative lobe is reclassified as “pre-split.” Crucially,  $E_{\text{out}}^+$  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  $6.0 \times 10^{-1}$  to  $1.25 \times 10^{-1}$  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  $\lambda_b$  and grid refinement, the reported  $A_{\text{echo,max}}$  stays bounded by the incoming budget implied by Theorem 5.46 and Proposition 5.60. 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 5.56 and Lemma 5.39 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  $\mu$  and width  $\ell$ ), 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  $w_{\text{disp}}$  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  $w_{\text{disp}}$  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  $w_{\text{disp}}$  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.

## 9 Discussion and outlook

**OES as a quantum–gravity equation of state.** The central proposal of this work is that semiclassical quantum gravity on belt–anchored wedges is governed by an *operator equation of state*:

$$\widehat{K}_{\text{mod}}(R) = \widehat{S}_{\text{gen}}(W) = \frac{\widehat{A}(W)}{4G} + \widehat{H}_{\text{bulk}}(W), \quad W = \text{EW}(R), \quad (\text{OES})$$

as an identity of closed quadratic forms on a regulator–independent domain. In plain language: on every admissible belt, the generator of boundary modular flow is identified with a bulk operator that packages area and energy into a single generalized–entropy observable. This converts the dynamics of geometry and matter into a statement about modular flows and their positivity/monotonicity properties. Within the framework of this paper, (OES) plays the role of a *QG equation of state*: it is the organizing relation from which linear and second–order semiclassical Einstein equations, quantum focusing bounds, and dispersive positivity constraints are recovered as corollaries.

**Summary of the four pillars.** The theorem suite of Section 5 organizes the consequences of (OES) into four “pillars” that triangulate between entanglement wedges, energy conditions, and scattering data:

- **Pillar I: QES and Page behavior.** Using the belt JLMS channel and the global belt atlas, one obtains quantum extremal surface (QES) conditions and Page–curve behavior directly from modular dynamics. The operator identity (OES) packages the usual extremality and area/entropy competition into a single generator statement, yielding a quantitative Page line–density bound and a belt–local QES condition controlled by the ledgered budget.
- **Pillar II: ANEC and QNEC from modular positivity.** Tomita–Takesaki modular theory, combined with OS positivity on the belt core, implies monotonicity and convexity of modular Hamiltonians. Through the JLMS map, these modular inequalities translate into averaged and local null energy conditions (ANEC/QNEC) on wedges, including quantified second–order control. The belt form of the quantum focusing inequality (QFC) is an immediate consequence: the quantum expansion  $\Theta_{\text{belt}}$  is monotone along generators, with a positive canonical–energy and shear/expansion contribution.
- **Pillar III: Dispersion and positivity with Regge control.** Gravity–subtracted forward amplitudes, equipped with analytic projectors and a Regge envelope, admit dispersive representations whose kernels are compatible with the belt JLMS channel. The resulting testers yield nonnegative forward coefficients (such as  $\widehat{c}_{2,0}$ ) that are stable under changes of subtraction pivot, scale, and IR scheme. These positivity bounds are tied back to belt–averaged null curvature via the scattering–curvature bridge, so that amplitude fits constrain admissible curvature profiles along belts.
- **Pillar IV: Semiclassical Einstein equations as modular equation of state.** Once (OES) is imposed on each belt and stitched across a cofinal atlas, the linear and Hessian versions of the semiclassical Einstein equations follow from modular consistency. In particular, the modular first law and its second–order completion reproduce the linearized Einstein equations in expectation and their quantified second–order corrections, with all regulator dependence recorded in  $B_{\text{belt}}$  and removed by flow limits.

Together, these pillars realize a belt–local, operator–level formulation of semiclassical gravity in which geometry is not postulated separately but is encoded in modular data and their stability properties.

**Cubic verification as an operator–level test.** Section 8.6 goes beyond variational checks by testing (OES) at cubic order in a controlled AdS<sub>3</sub>/CFT<sub>2</sub> shockwave protocol. On a fixed belt, one computes

$$\delta^3\langle K_{\text{mod}}\rangle \quad \text{and} \quad \delta^3\langle \widehat{S}_{\text{gen}}\rangle = \delta^3\left\langle \frac{\widehat{A}}{4G} + \widehat{H}_{\text{bulk}} \right\rangle$$

using the same shock profile, belt kernel, and canonical–energy organization. The comparison is made at the level of discretized integrands, with grid refinement and ledgered control on numerical and regulator errors. The resulting agreement, within certified numerical accuracy, is an operator–level test of (OES) beyond the semiclassical first and second variations. Conceptually, this addresses the principal linearization loophole in the holographic verification: the cubic test probes genuinely interaction–level contributions to the modular generator and its bulk dual.

**Imported structures and standing assumptions.** For clarity, it is useful to list the inputs that are *assumed* rather than derived within this work.

- **Local covariance and modular structure (QG–Ax–1, QG–Ax–3).** We assume a locally covariant net of algebras on globally hyperbolic spacetimes, with cyclic separating reference states on belt–anchored regions and well–defined Tomita–Takesaki modular data. A belt first–law channel exists and obeys the usual modular first law for admissible perturbations.
- **OS/RP positivity and belt analytic core (QG–Ax–2).** Reflection positivity and KMS properties are assumed in a way that produces a common quadratic–form core  $D_{\text{an}}$  stable under positive flows. This provides essential self–adjointness and form–closedness for  $\widehat{K}_{\text{mod}}$  and for the generalized–entropy sum defining  $\widehat{S}_{\text{gen}}$ .
- **JLMS–type belt channel.** We assume the existence of a belt–compatible operator JLMS map  $U_{R \rightarrow W}$  that intertwines modular flows between boundary belts and wedge algebras, with Brown–York/Iyer–Wald flux matching and JKM corner calibration. This includes quasi–local additivity and good behavior under belt factorization.
- **Dispersion, analyticity, and Regge control (QG–Ax–4).** The gravity–subtracted  $2 \rightarrow 2$  amplitude is assumed to satisfy analyticity, crossing, and a tester–certified Regge envelope on a forward cone, so that finite–subtraction dispersive representations and compact dual certificates exist. These assumptions underlie the construction of nonnegative testers and the curvature–scattering bridge.
- **Stability and invariance ledger (QG–Ax–5).** A stability framework controls changes in regulators, anchoring, and dispersion schemes: Kadison–Kastler–type bounds, dictionary continuity, and budget calculus guarantee that physical conclusions are invariant under these technical choices up to  $O(B_{\text{belt}})$ .

Within this imported structure, (OES) is adopted as an axiom on each belt and then shown to be equivalent to modular first–law matching, JLMS compatibility, normalization, and additivity. The rest of the analysis treats (OES) as the organizing equation of state and explores its consequences.

**Internal results of the OES framework.** Given the assumptions above, the main *internal* outputs of the belt–local OES program are:

- An operator–level identification of modular generators with generalized–entropy operators on belts, including group, channel, and relative formulations, with normalization fixed by Rindler/JKM calibration and no residual center freedom.

- A global belt atlas and stitching result: modular consistency on a cofinal belt family is equivalent (up to  $O(B_{\text{belt}})$ ) to the linear and second-order semiclassical Einstein equations in expectation on the associated domains of dependence.
- Quantified modular energy conditions: a belt-form quantum focusing inequality with explicit canonical-energy and shear/expansion contributions, implying ANEC/QNEC-type statements and excluding certain classes of averaged curvature pathologies.
- Stability and monotones: gauge/anchor and belt-width invariance of physical predictions; a belt  $c$ -function and width-flow monotone leading to a belt version of the generalized second law; and a detailed reproducibility ledger (dispersion budgets, dual certificate, numerical tolerances).
- Cosmological and black-hole applications at the level of belt-localized dynamics: a dichotomy between bounce and modular fixed-point scenarios for both black-hole interiors and early-time cosmology, controlled purely by canonical-energy and quadratic data on belts.
- A cubic, operator-level holographic test of (OES), extending the usual linearized and second-order verifications and probing interaction terms in  $\text{AdS}_3/\text{CFT}_2$  shockwave setups.

**External tests and falsifiable predictions.** Section 8 organizes the falsifiable content of the framework into three families of probes (with a strengthened cubic component):

- **(F1) Interior shock test in holography (theory-numeric).** Given a black brane or BTZ background, one can drive an interior shock and explicitly evaluate (OES), (M1), and (M2) on belt regions using boundary modular data, HRT/QES areas plus bulk stress, and covariant-phase-space canonical energy. Success requires that the operator identity and the belt-form QFC hold within numerical error as positive flows are removed; a robust discrepancy would falsify belt-level OES. The cubic extension in Section 8.6 sharpens this by testing  $\delta^3$  matching of  $\langle K_{\text{mod}} \rangle$  and  $\langle \hat{S}_{\text{gen}} \rangle$ .
- **(F2) Ringdown echo bound (observation-modeling).** The bounce vs. fixed-point interior scenarios imply an upper bound on effective interior reflectivity for gravitational perturbations: the positive quadratic form in (M2) suppresses late-time internal reflections after a quantum throat or in a modular fixed point. Embedding this constraint into EOB or time-domain scattering models yields a maximum allowed echo amplitude; observation of echoes above this bound in high-SNR events would rule out the OES-based interior picture.
- **(F3) Scattering positivity  $\Rightarrow$  curvature average (experiment-EFT).** Low-energy scattering data constrain the forward coefficient  $\hat{c}_{2,0}$  through dispersion and positivity. Via the curvature bridge, these constraints become statements about belt-averaged null curvature. A measured or EFT-implied negative  $\hat{c}_{2,0}$  in the calibrated setup would violate the required nonnegativity of the weighted  $\int \langle R_{kk} \rangle$  and hence falsify the OES+positivity synthesis.

In all cases, the program is explicitly failure-prone: sustained mismatch in holographic numerics, gravitational-wave echoes that exceed the ringdown bound, or amplitude fits that force  $\hat{c}_{2,0} < 0$  would point to a breakdown of the belt-local OES equation of state, of the imported assumptions, or of their joint compatibility.

**Open directions.** Several natural extensions sit beyond the scope of this paper:

- **Beyond semiclassical order.** The present analysis is organized around semiclassical belts with a single Newton constant and a fixed generalized-entropy functional. It is natural

to ask how (OES) is modified by higher-derivative terms, running couplings, or genuinely nonperturbative effects, and whether the operator equation of state can be extended to incorporate such corrections without losing belt locality and modular stability.

- **Deriving JLMS and OS positivity from deeper structure.** JLMS-type channels and OS positivity are taken as standing assumptions. One long-term aim is to derive (or at least significantly weaken) these inputs from more primitive principles, possibly by combining quantum information constraints, bootstrap-type bounds, and microscopic models of holographic codes.
- **Beyond belts and to other probes.** While belts are technically convenient and well-suited to modular flow, many physical questions involve more general regions (e.g. finite intervals, multi-component cuts, or coarse-grained observables). Extending the OES framework to such regions, or identifying the maximal class of subregions that admit a belt-like equation of state, would broaden its applicability.
- **Stronger observational and experimental interfaces.** The ringdown and scattering tests are first steps toward confronting OES with data. Further work is needed to develop robust waveform models compatible with the belt constraints, to quantify systematic uncertainties in amplitude fits, and to map concrete observational campaigns (gravitational-wave detectors, collider and fixed-target experiments, cosmic-ray data) onto the parameter space of OES-admissible theories.
- **Microphysical interpretations.** Finally, if (OES) survives the falsifiability program, it would be important to understand its status in candidate UV completions. Does the operator equation of state emerge from known string-theoretic or quantum-gravity constructions as an effective modular relation, or does it point toward a more intrinsic “entanglement first” formulation of gravity?

**Outlook.** The belt-local operator equation of state provides a concrete candidate for a quantum-gravity “formula”: an operator identity that unifies entanglement, area, and energy on each null generator. This paper develops the internal calculus of that identity, establishes a four-pillar theorem suite, and proposes a suite of falsifiable tests, including a cubic holographic verification. The next steps lie outside the present manuscript: implementing the program in explicit models, confronting it with data, and stress-testing the assumptions that make (OES) viable. Whether the operator equation of state ultimately survives these tests or is ruled out, the belt framework offers a sharply posed and quantitatively controlled route to connecting modular dynamics with the structure of spacetime.

## 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)	$a_{\text{QES}}/(4G) =$ <span style="border: 1px solid black; padding: 2px;"><math>5.8821988505</math></span>	$m=14$ , $\log \kappa_{\text{seed}}=0.5$ , $\log(\Lambda_0 \Gamma_{\text{belt}})=2.0149030205$ , $\log 29=2.8332133441$ .
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 4: Core constants and settings used in Section 5 proofs and audits.

Numeric check	Figure	Note
Modular witness (Rindler, free pulse)	$\delta \langle K_{\text{R}} \rangle = \frac{\pi}{2} A^2 =$ $\underline{1.5707963268 \times 10^{-4}}$	$A=10^{-2}$ .
QNEC/ANEC integrals (same pulse)	$\int du \langle T_{kk} \rangle =$ $\underline{2.9540897515 \times 10^{-5}}$	$2\pi \int du \langle T_{kk} \rangle =$ $\underline{1.8561093322 \times 10^{-4}}$ .
Interacting $\phi^3$ cubic piece	$\Delta \langle K_{\text{R}} \rangle^{(3)} =$ $\underline{3.490658504 \times 10^{-8}}$	$g_3=0.10$ , $A=10^{-2}$ , $L=1$ .
Yukawa cubic piece	$\Delta \langle K_{\text{R}} \rangle^{(y)} =$ $\underline{1.884955592 \times 10^{-6}}$	$y=0.15$ , $A=0.01$ , $B=0.02$ , $L=1$ .
Amplitude tail bound ( $k=2$ )	$\leq \underline{4.564354646 \times 10^{-5} s_0^{-3}}$	From $S_{\text{cut}}/s_0=20$ audit.
Composite quadrature error ( $10^{-6}$ )	$\leq \underline{7.86 \times 10^{-8}}$	$M=100$ , $J=2200$ panels; tail $\leq 5.811516831 \times 10^{-8}$ .
Composite quadrature error ( $10^{-8}$ )	$\leq \underline{7.65 \times 10^{-10}}$	$M=500$ , $J=28000$ ; tail $\leq 4.649213465 \times 10^{-10}$ .

Table 5: 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
$\mu_{\text{eff}}$	effective modular LR rate (belt)	1.0
$v_{\text{LR}}$	LR velocity (modular)	project-specific
$r_0, c_r$	belt offset/slope in $r(m) = r_0 + c_r m$	project-specific
$\eta$	AGSP contraction factor	1/3
$\Gamma_{\text{belt}}$	belt base factor	3.0
$\Lambda_0$	OSR/complexity prefactor	2.5
$\alpha_R$	tester-certified Regge slope cap	2.073 (for $\varepsilon=10^{-6}$ ), 2.089 (for $\varepsilon=10^{-8}$ )
$s_0$	subtraction scale (dispersion)	1.0
$S_{\text{cut}}/s_0$	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, $2\pi$

Table 6: 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 $\varepsilon$	$m$	$\delta^2 = 3^{-2m}$
$10^{-6}$	14	$3^{-28} = 4.3712421747 \times 10^{-14}$
$10^{-8}$	19	$3^{-38} = 7.4027370060 \times 10^{-19}$

Table 7: 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 8: 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  $\sigma$ ).**

$$d\mu_{\text{cel}}^{\text{ren}}(\sigma, \nu) := (1 + c_{\mu,2} \sigma^2) d\sigma d\nu, \quad c_{\mu,2} = 0.10 \text{ (audit default)}.$$

**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 5.22) and all dispersion invariances (Section 5.46, Section 5.66) are unchanged. If desired, incorporate the optional  $F_{\text{strip}}$  factor of Section 5.54.

**Acceptance note.** Any choice with  $0 \leq c_{\mu,2}, c_{W,2} < 1$  and  $0 < \sigma_0 \leq \sigma_0^*$  preserves Lemma 5.118 with the same  $O(\mathcal{B}_{\text{belt}})$  remainders; principal-series limits are recovered at  $\sigma \rightarrow 0$ .

## References

- [1] 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.
- [2] Klaus Fredenhagen. Locally covariant quantum field theory. arXiv:hep-th/0403007, 2004.
- [3] Konrad Osterwalder and Robert Schrader. Axioms for euclidean green’s functions I. *Commun. Math. Phys.*, 31:83–112, 1973. doi: 10.1007/BF01645738.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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.

- [8] 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.
- [9] 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.
- [10] Joan Elías Miró, Andrea Guerrieri, and Mehmet Asım Gümüş. Bridging positivity and S-matrix bootstrap bounds. *JHEP*, 05:001, 2023. doi: 10.1007/JHEP05(2023)001.
- [11] 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.
- [12] Geoffrey Penington. Entanglement wedge reconstruction and the information paradox. *JHEP*, 09:002, 2020. doi: 10.1007/JHEP09(2020)002.
- [13] 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.
- [14] João Penedones. Writing CFT correlation functions as AdS scattering amplitudes. *JHEP*, 03:025, 2011. doi: 10.1007/JHEP03(2011)025.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] Thomas Faulkner, Aitor Lewkowycz, and Juan Maldacena. Quantum corrections to holographic entanglement entropy. *JHEP*, 11:074, 2013. doi: 10.1007/JHEP11(2013)074.
- [24] 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.
- [25] 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.

- [26] 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.