Prepared for submission to JHEP

Information clustering and geometric pressure in AdS_3

Feiyu Deng

Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing, China

E-mail: dengfy@itp.ac.cn

Abstract: Joint quantum storage saves memory through correlations, while cooperation requires access to several laboratories. We formulate this competition at a fixed finite factorizing regulator, with specified component algebras, independently recoverable block memories and a nonnegative merger price. The optimal net saving is the established entropy-partition objective evaluated on actual component marginals. In the leading Einstein-AdS_3 vacuum RT entropy model, this optimum has a geometric representation. For fixed finitely many separated intervals, the minimum of free perimeter plus positive homology area, relative to separately optimized single-interval caps, satisfies $F^{\epsilon}_{\lambda} - F^{\epsilon}_{\lambda,\mathrm{disc}} = -qG^{\mathrm{RT}}_{\tau} + O(c\epsilon/\delta)$, where $q = \sqrt{1-\lambda^2}$ and $\tau = c\pi\lambda/(3q)$. The ultraviolet estimate is uniform on compact subcritical source intervals, and the minimization is over positive finite-perimeter regions with exact trace and compact essential support. Entropy-preserving component refinement and a global calibration establish the representation. Replacing RT entropies by actual-state entropies requires separate error bounds. Collective, nested and free-Dirac examples distinguish the partition effects. First-derivative estimates determine the geometric area jump, without extending the result to finite-c derivatives. The representation extends to the noncompact thermal line with a charged horizon. A compact-BTZ annulus instead gives a finite mismatch at the unchanged tariff.

Keywords: AdS-CFT Correspondence, Quantum Entanglement, Conformal Field Theory

Contents

1 Introduction **1**

2 Component algebras and the storage task **6**

3 From entropy partitions to positive homology components **7**

4 A pressure representation of optimal storage **9**

5 Collective phases, nested groups and area response **14**

6 Thermal states and the compact annulus obstruction **21**

7 Discussion **25**

A Coding and entropy-partition properties **27**

B Access networks and harmless refinement **30**

C Entropy-preserving refinement **33**

D Global arc minimum and its evaluation **35**

E Exact trace and cutoff-contact elimination **37**

F Uniform ultraviolet expansion **39**

G Oblique-cutoff BV reduction and planar horizon confinement **41**

H Global minimization on the compact cylinder **46**

I Numerical methods and checks **49**

1 Introduction

Correlations reduce the memory needed to store a quantum source when separated laboratories can encode it jointly. If joint access carries a cost, the storage problem also requires a choice of which laboratories cooperate. The relevant quantity is the memory saving after that cost is paid. Its definition requires the component algebras, the allowed encoding and decoding operations, and the charged resources; the spectrum of the joint density matrix alone does not determine it.

Consider a quantum field theory at a fixed finite factorizing regulator, with specified component algebras on $A = \bigcup_{i=1}^{n} A_i$. The separated maximal connected components are the laboratories. Their algebras are supplied as part of the task, without assuming a canonical continuum tensor factorization. A partition groups the laboratories into cooperating blocks, each with its own encoder, memory and decoder. Every block must be recoverable from its own memory. For a set of labels B, write ρ_B for the reduced state on $A_B = \bigcup_{i\in B} A_i$, with $\rho_i = \rho_{\{i\}}$. The memory saved by jointly storing this block, relative to separate storage of its components, is its total correlation,

$$T(B) = \sum_{i\in B} S(\rho_i) - S(\rho_B) = D\left(\rho_B \,\middle\|\, \bigotimes_{i\in B} \rho_i\right). \tag{1.1}$$

Here S and D denote the von Neumann entropy and quantum relative entropy, measured in nats. Joint access costs $\tau(|B| - 1)$ per source use, with $\tau \geq 0$. Thus every additional laboratory in a block has the same per-use price, and the charge is extensive in the number of source preparations. For this access model, the minimum combined cost of storage and access, C_τ, and the net saving relative to separate storage, G_τ, satisfy

$$C_\tau = \sum_i S(\rho_i) - G_\tau = \min_{\mathcal{P}\vdash[n]} \left[\sum_{B\in\mathcal{P}} S(\rho_B) + \tau(n - |\mathcal{P}|)\right]. \tag{1.2}$$

The label set is $[n] = \{1, \ldots, n\}$; $\mathcal{P} \vdash [n]$ denotes a set partition and $|\mathcal{P}|$ its number of blocks. The saving can therefore be written as

$$\boxed{G_\tau(\rho_A; \{A_i\}) = \max_{\mathcal{P}\vdash[n]} \left\{\sum_{B\in\mathcal{P}} T(B) - \tau(n - |\mathcal{P}|)\right\}.} \tag{1.3}$$

The maximization includes every set partition, including blocks with nonconsecutive labels. In (1.1), the actual marginals fix the product state; there is no auxiliary reference state to choose. The state, component algebras, block-storage rules and access price then determine the optimization.

This objective is the residual-entropy partition construction of info-clustering, evaluated with von Neumann entropy in place of Shannon entropy [1–4]. In particular, $C_\tau - n\tau$ is the Dilworth truncation of the residual-entropy function $B \mapsto S(\rho_B) - \tau$, evaluated on $[n]$ [1, arXiv v3, section III-C, eqs. (3.10)–(3.13)]. The established partition optimization and hierarchy describe how correlation savings compete with access costs. Strong subadditivity gives the submodularity needed for the quantum entropy input [5].

The storage rates refer to independent repetitions of a fixed source. Each decoder recovers its block and its correlations with a purifying reference from its own memory, without decoder side information or shared auxiliary entanglement. Schumacher compression and its converse give the entropy sum in (1.2) [6, 7]. Section 2 states the resource assumptions and realizes the merger count through an all-pairs access network. It also separates regulator refinements that preserve the laboratories from changes of the access task itself. The resulting saving is defined for general regulated states.

In the leading Einstein-AdS$_3$ vacuum RT entropy model, the same partition optimization can be computed geometrically. Keeping the regulator and component algebras fixed, we use RT block entropies as its input and represent the optimum by a perimeter-plus-area minimum. On a time-symmetric Poincaré slice with metric

$$\mathrm{d}s^2 = R^2 \frac{\mathrm{d}x^2 + \mathrm{d}z^2}{z^2}, \qquad z \geq \epsilon, \qquad c = \frac{3R}{2G}, \tag{1.4}$$

we prescribe the interior cutoff trace A and define

$$F^{\epsilon}_{\lambda}(A) = \inf_{E \in \mathcal{A}_{\epsilon}(A)} \frac{L_{\mathrm{free}}(E) + (\lambda/R)V(E)}{4G}, \qquad 0 \leq \lambda < 1. \tag{1.5}$$

The class $\mathcal{A}_{\epsilon}(A)$ comprises positive finite-perimeter regions with exact interior trace A and compact essential support. Each region lies almost everywhere in some compact subset of the closed cutoff half-plane $z \geq \epsilon$; the class has no common containing box. Free perimeter excludes the prescribed cutoff intervals, and the area is the positive homology area. The source λ/R changes both the boundary curves and their preferred connectivity: increasing pressure shrinks caps and can eliminate bridges between components. We compare this minimum with $F^{\epsilon}_{\lambda,\mathrm{disc}}$, the sum of the independently optimized single-interval cap actions at the same source.

For every fixed finite union of separated vacuum intervals, we prove the relation

$$\begin{aligned} F^{\epsilon}_{\lambda}(A) - F^{\epsilon}_{\lambda,\mathrm{disc}}(A) &= -q\, G^{\mathrm{RT}}_{\tau}(A_1, \ldots, A_n) + O(c\epsilon/\delta), \\ q = \sqrt{1-\lambda^2}, &\qquad \tau = \frac{c\pi\lambda}{3q}. \end{aligned} \tag{1.6}$$

Consecutive endpoints are separated by at least $\delta > 0$. The superscript RT specifies that every block entropy is taken from the leading semiclassical Einstein-AdS$_3$ vacuum RT model [8, 9]. With endpoints and n fixed, the error is uniform as λ ranges over a compact subset of $[0, 1)$. Equation (1.6) therefore gives the optimal net memory saving as the relative geometric saving divided by q. The relation between τ and λ matches the specified access price to the geometric source. Both are inputs; G_{τ} is the information quantity represented by the minimum. Replacing actual-state entropies by their RT values requires separate entropy-error bounds.

The proof first relates field blocks to homology components. Graph min-cut results ensure that disjoint subsystems admit disjoint min-cuts and that disconnected minimizing components remain min-cuts for their traces [10, section 2.4, Lemmas 3–4]. We give the continuous finite-anchor separation argument and its ultraviolet limit. Each block entropy already minimizes over all RT channels. When a minimizing homology region splits, refinement of the block preserves its entropy and cannot increase the cooperation charge. Unrestricted field partitions can thus be compared through positive homology components.

The geometric minimum is then obtained from calibrated directed pressure arcs. Minimizing their total cost over endpoint assignments produces disjoint arcs with the required orientations. A finite-perimeter lower bound rules out a smaller action from rough boundaries, extra cutoff contacts or loops. This establishes the global comparison, including tied channels, beyond the stationary-circle calculation.

The optimized arc action fixes the source-to-price relation. Its ultraviolet expansion consists of the RT logarithm multiplied by q, a common endpoint contribution and a signed orientation contribution. The orientation count determines the number of positive disk components. Subtracting the independently optimized caps cancels the common terms and leaves the entropy-partition objective with the tariff in (1.6). The subtraction compares optimized cooperating and no-cooperation configurations at the same source.

Collective and nested examples show why the full partition domain matters. Three components can save memory collectively although every individual pair merge loses, and a block with nonconsecutive labels can be strictly optimal. A free-Dirac benchmark separates this pooling effect from the stronger leading-RT possibility of positive triple saving with vanishing pair mutual informations. The pressure representation associates changes in the cooperating groups with finite area jumps and determines their cutoff corrections. It extends to the noncompact thermal line when the oblique cutoff and charged horizon are included. Compact BTZ permits an annular region in the same perimeter-plus-area family: its Euler characteristic no longer counts laboratory groups. The resulting finite mismatch at the unchanged tariff limits the geometric representation while leaving the storage optimum well defined.

Subregion volume already has information-compression interpretations. Abt et al. relate subregion complexity to the tensor-network complexity of an optimal compression map [11, arXiv v2, sections I and III.A]. Their construction separates the Hilbert-space size at the RT cut from the number of tensors implementing the embedding; for a general network, this tensor count only bounds circuit complexity from above. Their subsequent kinematic-space construction expresses vacuum subregion volumes in terms of boundary entanglement data [12]. These are direct antecedents of the information/volume interpretation. Compression-map complexity, asymptotic memory rate and recurring access rental measure different resources. Constrained state merging also assigns an operational optimal cost to curve lengths in a different task [13].

The variational argument draws on several established results. Info-clustering gives the partition objective and hierarchy [1], and graph min-cuts give the component structure [10]. Planar least-gradient theory distinguishes BV rotation from the conversion of an optimal flow into a transport plan [14, arXiv v1, section 1 and Proposition 2.3]. The anisotropic least-gradient/Beckmann equivalence requires a bounded contractible Lipschitz domain and uniformly elliptic norm integrands, with zero boundary mass needed for the reverse minimizer construction [15, arXiv v1, section 3, Theorem 3.4]. Under the respective strict-convexity hypotheses, exact-trace existence is known for BV data, including finite jumps with atomic derivatives [14, arXiv v1, Proposition 5.1][15, arXiv v1, Theorem 7.1]. Least-gradient functionals with a volume source have been studied on bounded domains subject to barrier, regularity and smallness assumptions [16]. These hypotheses do not yield the nonreversible exact-trace result on the unbounded hyperbolic domain considered here. Constant-geodesic-curvature magnetic arcs and the distinct Randers constructions with constant flag curvature are also known [17, 18].

Our result evaluates the specified block-storage and recurring-access optimum as a global geometric minimum. It uses the established compression principle and partition

objective. The additional statements are the tariff-matched nonzero-pressure minimum with exact binary trace, its explicit coefficients, the global endpoint assignment and the cutoff-contact classification of all minimizers, together with first-derivative control and the compact-topology obstruction. The directed cost depends on position and is nonreversible. Although the cutoff domain is unbounded, every admissible competitor has compact essential support. Attainment and the classification of all minimizers hold in this class; extension to its finite-action closure remains unproved. The following comparison is limited to the relevant distinctions from the cited antecedents.

Established result	Use and further result in this paper
Info-clustering [1, arXiv v3, section III-C]	The residual-entropy partitions and hierarchy are used as established. We derive the pressure representation of the optimum and its source-to-price coefficients.
Min-cut refinement [10, section 2.4, Lemmas 3–4]	Graph models give disjoint cuts and component additivity. Appendix C proves the continuous finite-anchor statement and takes its ordered ultraviolet limit.
BV/Beckmann and transport [14, 15]; source functional [16, equation (1.3)]	The anisotropic norm-cost equivalence assumes a bounded contractible domain; its reverse minimizer construction requires zero boundary mass. We use calibration and current decomposition for the nonreversible hyperbolic reduction, positive-region reconstruction and classification of all contacts. Binary jump data are already covered by earlier results.
Integral-current representation [19, Proposition 2.4, Theorem 5.3]	The decomposition theorem is used unchanged. Appendix E derives the endpoint, weighted-mass and one-form consequences needed here; appendices G and H treat the remaining ambient spaces.
Magnetic and Randers curves [17, 18]	The constant-curvature mechanisms are established. Appendix D computes the cutoff-directed action and equality conditions for this problem, without adding a general Randers theorem.
Compression and RT-volume topology [11, 12, 20]	Earlier work supplies compression-map interpretations, entropy-based volume formulas, Euler-characteristic jumps and annular sectors. For the distinct memory-rate/access task, we derive the global pressure representation and tariff dictionary, control its response and exhibit the compact mismatch at the unchanged tariff.

The representation depends on the specified access-resource model; it does not determine a unique microscopic access price in quantum field theory.

Sections 2–4 develop the storage task, its RT partition reduction and the pressure minimum. Section 5 gives collective and nested phases and their area response. Section 6 extends the calculation to the thermal domain and identifies the compact obstruction. The coding, variational and topological proofs appear in appendices A–H. The companion note *Field constructions and their composition limits* treats spectral, interval-family, supplied-probe, density-geometry and replica constructions as separate tasks. They provide no additional proof of the actual-state saving.

2 Component algebras and the storage task

At a fixed finite factorizing regulator, consider the actual state on $\mathcal{H}_A = \bigotimes_i \mathcal{H}_{A_i}$. Its laboratories A_i are the maximal connected components of the prescribed region. All information rates use natural logarithms. For m independent preparations of the source, choose a partition $\mathcal{P}$ and assign an encoder and decoder to each block. A decoder acts only on its block's stored memory. It must recover the block, including correlations with a purifying reference, with entanglement fidelity tending to one. Decoder side information, shared auxiliary entanglement, uncounted memory and joint decoding between blocks are excluded.

For these resources, the noiseless source-coding theorem fixes the minimum sum of memory rates to $\sum_{B\in\mathcal{P}} S(\rho_B)$. Schumacher compression achieves the rate of each block, and the memory-dimension converse gives the matching lower bound. Contractivity and a telescoping estimate ensure vanishing joint recovery error when the number of blocks is fixed. Separate storage costs $\sum_i S(\rho_i)$, leaving a saving of $\sum_B T(B)$. The product of the source marginals is also the unique product state minimizing relative entropy [21]. Appendix A gives the coding and relative-entropy arguments. The coding limit is taken at fixed regulator and fixed theory, before any ultraviolet or large-c comparison. Independence is assumed between repeated preparations, with no spatial independence requirement.

We charge $m\tau(|B|-1)$ memory-equivalent nats for joint access to a block. The same per-use price applies to each additional laboratory, so the full partition incurs $m\tau(n-|\mathcal{P}|)$. Minimizing the combined memory and access rates gives (1.2), and comparison with separate storage gives (1.3). The recurring nature of the charge matters: a fixed installation fee would vanish upon division by m, changing the asymptotic task.

A homogeneous access network realizes the merger charge. Give each laboratory free internal coherent control, and allow any pair to rent an ideal quantum link for the batch at cost $m\tau$. The active links remain fixed throughout the batch. Within a connected group, registers may be routed to one laboratory for joint encoding or decoding and then returned. Traffic, capacity, latency, temporary workspace and driving energy are not charged separately. Write H for the graph of active links on the laboratory labels, $E(H)$ for its edge set, and $\mathcal{P}(H)$ for the partition into connected components. The number of rented links is $|E(H)|$. Connecting k laboratories requires at least $k-1$ links, with equality for a spanning tree. Therefore

$$\min_{H:\,\mathcal{P}(H)=\mathcal{P}} |E(H)| = \sum_{B\in\mathcal{P}} (|B|-1) = n - |\mathcal{P}|. \tag{2.1}$$

Connectivity thus fixes the number of charged links, while their homogeneous per-use price is an assumption of the model. State merging and quantum communication obey different resource laws that depend on the state [22, 23].

The choice of laboratories precedes the choice of partition. Free coherent control within each connected component of A, together with the absence of a free relay through the complement, makes the maximal connected intervals the chargeable laboratories. Subdivision of regulator cells leaves them unchanged if it preserves these free-control domains. Treating the subdivisions as independent laboratories changes the access task. Appendix B states

the conditions for state-and-algebra-preserving refinement and gives the optimization for a weighted access graph.

For example, let laboratories 1 and 3 share a uniform classical bit, with 2 pure. At price $\tau = 0.4$, all-pairs access favors $\{1,3\} \mid \{2\}$ and yields $\log 2 - 0.4 > 0$. Restricting the network to nearest neighbors, or assigning complete-graph weights $(w_{12}, w_{23}, w_{13}) = (1,1,2)$, makes separate storage optimal with zero saving. Only the available access resource has changed; the source state is identical. The competing costs are given in appendix B.

With two laboratories, the choice is between joint and separate storage:

$$G_\tau = (I(A_1 : A_2) - \tau)_+. \tag{2.2}$$

For three or more laboratories, total correlation can cover a block's collective charge even when every pair fails this test. This requires optimization over all set partitions. Restricting to pairwise-greedy merges or consecutive labels would impose an additional constraint on the task.

The submodular partition structure fixes the response to access price. The saving is nonnegative, bounded by the total correlation of the union, monotone under product local quantum channels, and additive between independent groups of components. As τ increases, G_τ remains continuous, convex and nonincreasing, while the finest optimal partitions refine. These statements are proved in appendix A. Entropy counterterms that are additive over the specified components cancel in the difference between each partition cost and separate storage. Consequently, if the finitely many block relative entropies have finite continuum limits, so does G_τ, even when the individual storage rates diverge.

3 From entropy partitions to positive homology components

To evaluate the storage optimum in a holographic vacuum, each nonempty $B \subseteq [n]$ must be assigned its marginal entropy $S(\rho_B)$, minimized over admissible Ryu–Takayanagi geodesic channels. Order the interval endpoints ∂A_B along the boundary. We denote by $\mathrm{NC}(\partial A_B)$ their noncrossing perfect matchings compatible with positive homology to A_B. Each endpoint belongs to one pair, and pairs do not interlace. For a pair (a,b), the ordering $a < b$ corresponds to $x_a < x_b$. In a common endpoint convention, the leading semiclassical Einstein-AdS_3 vacuum entropy is

$$S_{\mathrm{RT}}(A_B) = \frac{c}{6} \min_{M_B \in \mathrm{NC}(\partial A_B)} \sum_{(a,b)\in M_B} 2 \log \frac{x_b - x_a}{\epsilon}. \tag{3.1}$$

The holographic-vacuum assumptions are essential to this entropy model; large central charge by itself is insufficient [8, 9]. For all $2n$ ordered endpoints, the corresponding matching set is abbreviated as $\mathrm{NC}(2n)$. We use ϵ for both the entropy convention and the radial anchors only in the ultraviolet comparison; no microscopic finite-wall duality is assumed.

The field partition imposes no geometric connectivity constraint. A block's minimizing RT region may split into positive components. Refining the block along their traces

preserves the entropy sum and cannot increase a nonnegative tariff. This is the component-additivity mechanism of min-cuts [10, section 2.4, Lemmas 3–4]. The graph result supplies the structural step; the finite-anchor argument below establishes disjointness and exact component traces in the continuous geometry, including tied minima. Refinement acts on marginal entropies already minimized over all their channels. Minimizing the unrestricted field cost can thus select a geometric connectivity pattern.

For a noncrossing matching of all endpoints, let $L_{\mathrm{RT}}(M)$ be its dimensionless logarithmic length and $\chi(M)$ its number of positive homology components. A geodesic minimum has no compact interior holes, so its Euler characteristic equals this count. Separate positive caps inside cutoff-connected negative gaps remain components.

Theorem 3.1 (RT partition reduction)**.** *For every $\tau \geq 0$ and every fixed finite set of distinct interval endpoints,*

$$G_\tau^{\mathrm{RT}} = \sum_i S_{\mathrm{RT}}(A_i) - \min_{M \in \mathrm{NC}(2n)} \left[\frac{c}{6} L_{\mathrm{RT}}(M) + \tau\big(n - \chi(M)\big) \right]. \tag{3.2}$$

A minimizing field partition has a noncrossing representative after entropy-preserving component refinement. The statement concerns values and does not assume unique optimizers.

Begin at finite anchors, with geodesic entropies S_{geo}. Fix a field partition $\mathcal{P}$ and select an entropy-minimizing region for each block. These regions have disjoint interiors: cutting and reassembling free boundaries at a crossing preserves their total length, whereas rounding the resulting corners strictly reduces it. Geodesic uniqueness and distinct anchors also exclude tangencies. Each positive component minimizes entropy for its own trace. Otherwise, uniting a shorter replacement with the remaining components would lower the block perimeter by subadditivity, contradicting minimality. Component refinement therefore leaves the entropy sum unchanged and gives a matching with $\chi(M) \geq |\mathcal{P}|$. For $\tau \geq 0$, this implies

$$\frac{c}{6} L_{\mathrm{geo}}(M) + \tau\big(n - \chi(M)\big) \leq \sum_{B \in \mathcal{P}} S_{\mathrm{geo}}(A_B) + \tau\big(n - |\mathcal{P}|\big). \tag{3.3}$$

Taking minima bounds the geometric cost from above by the field cost.

For the reverse inequality, take any noncrossing matching. Its positive components define a field partition $\mathcal{P}_M$ with $|\mathcal{P}_M| = \chi(M)$. Their boundaries provide admissible entropy channels for the corresponding blocks. Minimizing each marginal entropy over all admissible channels can only lower the component cost, giving

$$\begin{aligned} \min_{\mathcal{P}} \left[\sum_{B \in \mathcal{P}} S_{\mathrm{geo}}(A_B) + \tau(n - |\mathcal{P}|) \right] &\leq \sum_{B \in \mathcal{P}_M} S_{\mathrm{geo}}(A_B) + \tau(n - \chi(M)) \\ &\leq \frac{c}{6} L_{\mathrm{geo}}(M) + \tau(n - \chi(M)). \end{aligned} \tag{3.4}$$

The minimum over M gives the opposite inequality and hence equality at finite anchors.

To pass to the ultraviolet, keep the endpoint set fixed. There are finitely many candidates, each with n edges, and $2\,\mathrm{arsinh}(\Delta/2\epsilon) = 2\log(\Delta/\epsilon) + O(\epsilon^2/\delta^2)$ holds uniformly.

Subtract the common divergence before minimizing to obtain (3.2). Appendix C proves the separation argument in full and gives a counterexample for negative tariff.

Noncrossing does not require consecutive blocks. For $\{1,3\} \mid \{2\}$, an outer positive region joins the first and third intervals, while a separate middle cap lies in its cutoff-connected negative gap. Such a nested partition is the strict optimum in the geometry of section 5. The remaining task is to generate the component price from a single perimeter-plus-area minimum at the corresponding source.

4 A pressure representation of optimal storage

The component count in theorem 3.1 still enters G_τ^{RT} through an explicit cooperation price. A local area source reproduces that price once the free boundaries are optimized. Nonzero pressure displaces those boundaries from their geodesic positions, so perimeter and positive area must be varied together. We keep the storage task fixed and derive a geometric representation of its optimum.

Work on $H_\epsilon = \{(x,z) : z > \epsilon\}$ with cutoff wall $\Gamma_\epsilon = \{z = \epsilon\}$. Competitors are measurable positive regions $E \subset H_\epsilon$, identified modulo two-dimensional null sets. Compact essential support requires, for each competitor, a compact $K \subset \overline{H}_\epsilon$ with $E \subset K$ almost everywhere. This K may vary with E; no common containing box is fixed. Impose $\mathbf{1}_E \in BV(H_\epsilon)$ and

$$\mathrm{tr}_{\Gamma_\epsilon} \mathbf{1}_E = \mathbf{1}_A \quad \mathcal{H}^1\text{-almost everywhere on } \Gamma_\epsilon. \tag{4.1}$$

Here the interior BV trace is obtained almost everywhere from averages over shrinking upper half-balls. These conditions define $\mathcal{A}_\epsilon(A)$ without restrictions on topology, boundary-component number or regularity. Compact essential support is imposed on competitors; it is not inferred for arbitrary finite-action regions reaching infinity. Additional cutoff contacts are allowed and will be excluded for minimizers.

The variational setting is weighted, anisotropic least-gradient theory [15, 24]; functionals with a volume source have also been studied [16]. For a positive 0/1 region, the action depends only on the region and its reduced boundary. Null-set changes such as zero-width cracks or appended filaments therefore create no extra boundary components. In unit-radius variables, define

$$\begin{aligned} P_H(E;H_\epsilon) &= \int_{\partial^* E \cap H_\epsilon} \frac{\mathrm{d}\mathcal{H}^1}{z}, \qquad \mathrm{Area}_H(E) = \int_E \frac{\mathrm{d}x\,\mathrm{d}z}{z^2}, \\ J_\lambda^\epsilon(A) &= \inf_{E \in \mathcal{A}_\epsilon(A)} \{P_H(E;H_\epsilon) + \lambda\,\mathrm{Area}_H(E)\}. \end{aligned} \tag{4.2}$$

The dimensional quantities satisfy $L_{\mathrm{free}} = RP_H$, $V = R^2\,\mathrm{Area}_H$ and $J_\lambda^\epsilon = 6F_\lambda^\epsilon/c$. The perimeter charge includes only the free boundary. The cutoff data are fixed by the interior trace in (4.1); changing a representative's literal intersection with the wall neither changes that trace nor introduces an independently variable zero-cost boundary.

4.1 The global pressure minimum

Write the line and area forms as

$$\alpha = \frac{\sqrt{\mathrm{d}x^2 + \mathrm{d}z^2}}{z}, \qquad \beta = \frac{\mathrm{d}x}{z}, \qquad \mathrm{d}\beta = \frac{\mathrm{d}x \wedge \mathrm{d}z}{z^2}. \tag{4.3}$$

In planar least-gradient theory, rotation of the BV derivative produces a Beckmann flow with boundary source and sink measures [14, 15]. Here the derivative of the interval trace gives unit atoms at the endpoints. Stokes' theorem expresses the area through an oriented boundary integral, yielding the directed cost $\alpha - \lambda\beta$. This cost is nonreversible, unlike the symmetric norm or Riemannian costs in those least-gradient results. Nonreversible costs and constant-curvature magnetic arcs are established Randers/isoholonomic ingredients [17, 18]. The cited least-gradient existence results already cover BV boundary data with atomic derivatives under their respective hypotheses. The present reduction combines a position-dependent Randers cost on the unbounded hyperbolic half-plane with exact trace and no wall mass, positive binary-region realizability, and classification of every cutoff contact. We prove the assignment formula in this domain:

$$J_\lambda^\epsilon = \frac{\lambda W}{\epsilon} + \min_\pi \sum_{i=1}^{n} d_\lambda^\epsilon(u_i, v_{\pi(i)}), \qquad W = \sum_i (v_i - u_i), \tag{4.4}$$

where each directed edge has cost

$$d_\lambda^\epsilon(u, v) = \inf_{C:(u,\epsilon)\to(v,\epsilon),\, z\geq\epsilon} \int_C (\alpha - \lambda\beta). \tag{4.5}$$

Because every terminal endpoint occurs once, the width term is common to all assignments. We construct a region attaining the minimum assignment cost and prove that this cost bounds every finite-perimeter competitor from below.

For an edge, set $\Delta = |v - u|$, $d = \Delta/2$ and $\sigma = \operatorname{sgn}(v - u)$. The directed distance is

$$H_\lambda(d) = q \operatorname{arsinh} \frac{qd}{\epsilon} + \lambda \arcsin \frac{\lambda d}{\sqrt{d^2 + \epsilon^2}}, \tag{4.6}$$

$$d_\lambda^\epsilon(u, v) = 2H_\lambda(d) - 2\sigma\lambda \arctan \frac{d}{\epsilon}. \tag{4.7}$$

Adding the oriented cutoff width gives its action

$$f_{\lambda,\epsilon}(u, v) = 2H_\lambda(d) + 2\sigma\lambda \left(\frac{d}{\epsilon} - \arctan \frac{d}{\epsilon} \right). \tag{4.8}$$

The minimizing curve has constant geodesic curvature. Absorb its orientation into $\kappa = \sigma\lambda$; the Euclidean circle parameters are

$$r_\kappa = \frac{\sqrt{\epsilon^2 + q^2 d^2} + \kappa\epsilon}{q^2}, \qquad z_c = -\kappa r_\kappa, \qquad \kappa = \sigma\lambda. \tag{4.9}$$

Use the entire circle segment above the cutoff. Reverse arcs may overhang horizontally and would be lost in a graph-restricted variation. Figure 1 shows the competition between connected and disconnected pressure minima for two intervals.

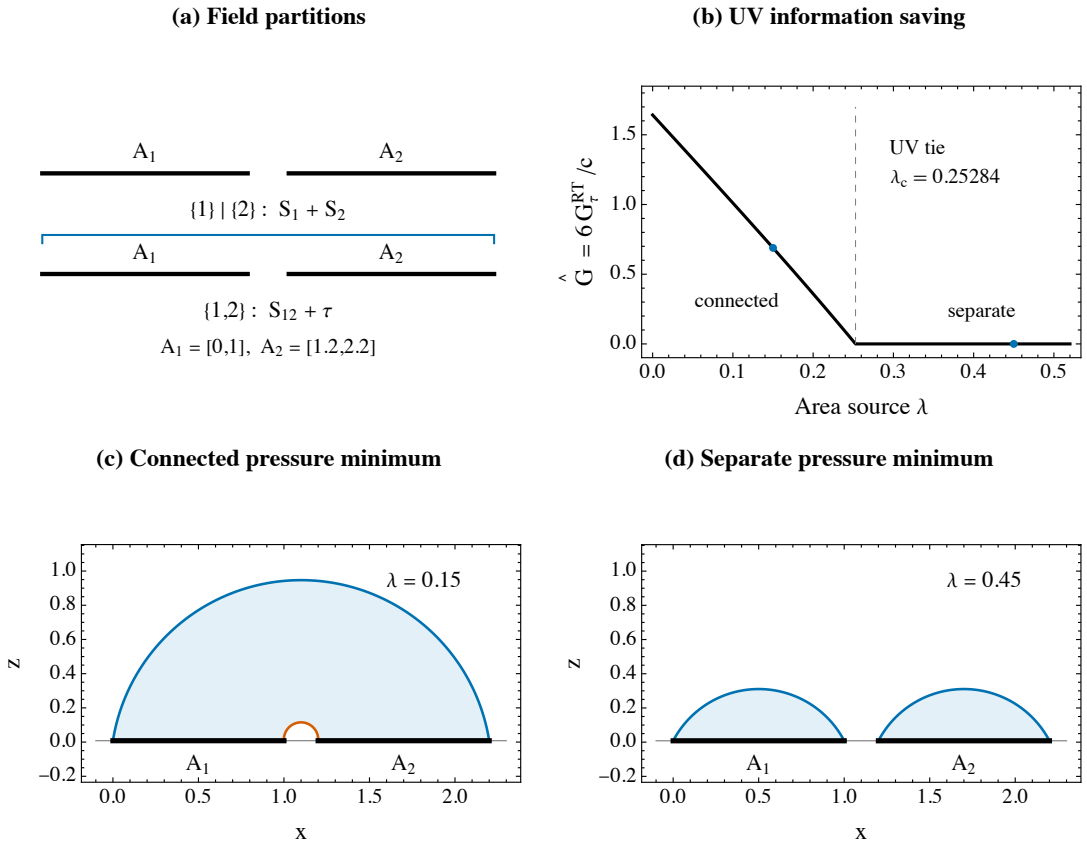


Figure 1. Storage choices and pressure minima for $A_1 = [0,1]$ and $A_2 = [1.2, 2.2]$. (a) The boundary grouping problem compares separate storage, of cost $S_1 + S_2$, with joint storage, of cost $S_{12}+\tau$. (b) In the leading-RT ultraviolet model the saving is $\widehat{G} = 6G_\tau^{\mathrm{RT}}/c = 2[\log(25/11) - \pi\lambda/q]_+$, where $q = \sqrt{1-\lambda^2}$ and $\tau = c\pi\lambda/(3q)$. The two costs tie at $\lambda_c = \log(25/11)/\sqrt{\pi^2 + \log^2(25/11)} \simeq 0.25284$. Dots indicate the sources shown below. (c,d) Exact finite-cutoff global pressure minima for $\lambda = 0.15$ and 0.45 at $\epsilon = 0.008$. Both endpoint assignments are compared; the selected regions are shaded by their positive homology area. The orange reverse arc bounds the excluded cutoff-connected gap, and thick segments denote the fixed trace. The optimized circular arcs are drawn in equal-aspect Poincaré coordinates. The transition in (b) belongs to the ultraviolet information limit; its λ_c need not coincide with the finite-cutoff transition.

A global calibration controls arbitrary paths, including overhanging reverse arcs. In Fermi coordinates (ρ, t) about a suitable geodesic, choose a globally defined endpoint function ψ obeying

$$\alpha^2 = \mathrm{d}\rho^2 + \cosh^2\rho\,\mathrm{d}t^2, \qquad \beta + \sinh\rho\,\mathrm{d}t = \mathrm{d}\psi \tag{4.10}$$

These identities give the calibration inequality

$$\alpha - \kappa\beta \geq q\,\mathrm{d}t - \kappa\,\mathrm{d}\psi. \tag{4.11}$$

Saturation requires $\rho = -\operatorname{artanh}\kappa$ and motion forward in t. Integration gives a sharp endpoint lower bound, with equality only on the stated arc up to increasing parametrization

and pauses. Reflecting the construction covers the reverse direction. Appendix D supplies the coordinates, equality analysis and definite integrals, establishing (4.7) for arbitrary paths.

Take an endpoint assignment of minimum total cost and the calibrated arc for every edge. If two arcs intersect, exchange their outgoing tails. This preserves both the total directed action and the endpoint bijection. Each exchanged path must saturate its endpoint bound: strict inequality would allow replacement by a calibrated arc and lower the assignment cost. Distinct oriented tangents produce a corner or reversing cusp after exchange, violating the calibration equality conditions. Equal tangents force coincidence by uniqueness for the constant-curvature initial-value problem. The connected segment of that circle above the cutoff then has the same two anchors, contradicting distinct endpoints. Hence every minimizing assignment consists of disjoint, simple arcs whose interiors lie strictly above the wall.

To reconstruct the positive region, label the ordered cutoff endpoints $x_1, \dots, x_{2n}$. Odd labels are starts and even labels are ends. Closing each arc with its cutoff chord produces caps that are disjoint or nested. Outermost arcs point left to right. Directions alternate with depth, since preceding sibling subtrees contain an even number of endpoints. Taking the points covered by an odd number of caps gives a region with trace A and

$$\partial\Omega = A_{\rightarrow} - \sum_i C_i, \qquad \mathrm{Area}_H(\Omega) = \frac{W}{\epsilon} - \sum_i \int_{C_i} \beta, \qquad P_H(\Omega) = \sum_i \int_{C_i} \alpha. \tag{4.12}$$

Its action equals the assignment value. Reverse inner arcs subtract enclosed area while contributing positive perimeter, as the oriented boundary of a positive region requires.

Lemma 4.1 (Finite-perimeter reduction and cutoff-contact elimination)**.** *For $0 \le \lambda < 1$, the infimum in (4.2) is attained and equals (4.4). Every minimizer, up to null sets, is a regular positive region bounded by the fixed cutoff intervals and exactly n disjoint calibrated free arcs. Its geometric free boundary meets Γ_ϵ precisely at ∂A, and each meeting is transverse. Consequently, any admissible region with a genuine additional cutoff contact admits a strictly lower-action replacement with the same BV trace. This includes tangential contacts; no local corner-rounding hypothesis is required.*

For the lower bound, extend $\mathbf{1}_E$ by zero below the cutoff. The reversed free-boundary current has the prescribed endpoint boundary and carries no wall mass. Its mass-additive decomposition consists of n anchor-to-anchor paths and closed curves. Each path obeys the calibration bound; each closed curve costs at least $(1-\lambda)$ times its hyperbolic length. The constructed assignment region attains this bound. Equality excludes cycles and forces every path onto a calibrated arc, eliminating additional wall contacts. The trace and current argument in appendix E also handles tangencies, accumulating contacts and transversality at the endpoints.

The order of minimization matters: first minimize over endpoint assignments, then reconstruct the positive region. Deforming a nonwinning noncrossing geodesic matching edge by edge can produce intersecting arcs. Such deformations need not be admissible embedded pressure channels; the global assignment minimum supplies the admissible representative.

4.2 Uniform limit and the cooperation tariff

At fixed edge length and subcritical source, the minimized edge action has the expansion

$$f_{\lambda,\epsilon}(u,v) = \frac{\lambda(v-u)}{\epsilon} + 2q\log\frac{\Delta}{\epsilon} + 2(q\log q + \lambda\arcsin\lambda) - \pi\lambda\sigma + O(\epsilon/\Delta). \tag{4.13}$$

Summing signed widths gives $\lambda W/\epsilon$. Every edge has the same endpoint contribution $2(q\log q + \lambda\arcsin\lambda)$, leaving the orientation term to distinguish channels. A positive component meeting k interval traces has one forward outer arc and $k-1$ reverse gap arcs. Thus

$$\sum_{e\in M}\sigma_e = 2\chi(M) - n. \tag{4.14}$$

Positive caps nested in cutoff-connected gaps contribute to this component count.

With fixed finite n, all possible edge lengths are at least δ, so the remainder is uniform over the finite assignment set. Every unrestricted assignment winner is noncrossing at each cutoff; restricting the minimum to that subset therefore leaves its value unchanged. After subtraction of the common terms, the uniform limit determines the leading minimum even at channel ties. Embeddedness of nonwinning deformations is unnecessary. Combining (4.13) and (4.14) yields

$$\begin{aligned} F_\lambda^\epsilon = \frac{c}{6}\left[\frac{\lambda W}{\epsilon} + 2n(q\log q + \lambda\arcsin\lambda) + \pi\lambda n\right] \\ + \min_{M\in\mathrm{NC}(2n)}\left[\frac{cq}{6}L_{\mathrm{RT}}(M) - \frac{c\pi\lambda}{3}\chi(M)\right] + O(c\epsilon/\delta). \end{aligned} \tag{4.15}$$

Partition dependence survives only through the RT length and positive component count. Matching these terms to theorem 3.1 determines the same price for every channel:

$$q\tau = \frac{c\pi\lambda}{3}, \qquad \tau = \frac{c\pi\lambda}{3\sqrt{1-\lambda^2}}. \tag{4.16}$$

Theorem 4.2 (Positive-area correspondence)**.** *Let A have fixed finite n and distinct nondegenerate interval endpoints, and let F_λ^ϵ be defined by (1.5) on the compact-essential-support class $\mathcal{A}_\epsilon(A)$. In the leading vacuum RT entropy model, for $0 \le \lambda < 1$,*

$$F_\lambda^\epsilon - F_{\lambda,\mathrm{disc}}^\epsilon = -q\,G_\tau^{\mathrm{RT}} + O(c\epsilon/\delta), \qquad q = \sqrt{1-\lambda^2}, \tag{4.17}$$

with the price (4.16). The remainder is uniform on compact subintervals of $[0,1)$ at fixed endpoints and fixed n. The disconnected baseline is the sum of the separately pressure-optimized single-interval caps.

The separately optimized caps contribute $q\sum_i S_{\mathrm{RT}}(A_i)$ in geodesic entropy and have n positive components. Subtracting their baseline from (4.15) cancels the fixed-trace and common endpoint terms. Theorem 3.1 then yields (4.17). Appendix F gives the raw normalization and remainder bounds. For a single interval, $G_\tau = 0$: the optimized cap and its reference share the entire pressure response. The source-free limit $\lambda = 0$ recovers the RT region.

Boundary displacement accounts for both the factor q and the common endpoint function. The signed orientation term counts positive components, converting their competition with entropy into the tariff. Holding RT geodesics fixed while evaluating area would discard this displacement and define a different variational problem.

To separate geometric cutoff errors from errors in the entropy input, suppose $S(\rho_B) = S_{\mathrm{RT}}(A_B) + e_B$ for every nonempty subset and set $E = \max_{B\neq\varnothing} |e_B|$. Comparing the finite partition minima then gives

$$|G_\tau - G_\tau^{\mathrm{RT}}| \leq 2nE. \tag{4.18}$$

The actual-state saving can therefore replace G_τ^{RT} in (4.17) at an additional value error of at most $2qnE$. This bound, derived in appendix F, gives no control of source derivatives. The area response requires its own derivative estimate.

5 Collective phases, nested groups and area response

Three components already admit collective mergers and spatially nonconsecutive groups. Their complete price envelopes exhibit both effects and supply the stable branches for the area-response calculation. We use the normalization

$$\widehat{S} = \frac{6}{c}S, \qquad \widehat{G} = \frac{6}{c}G, \qquad \widehat{\tau} = \frac{6}{c}\tau = \frac{2\pi\lambda}{q}. \tag{5.1}$$

The hats express entropies and savings in units of $c/6$, leaving the physical central charge unchanged. The last equality follows from the holographic source-to-price dictionary. For the free-field example below, $\widehat{\tau}$ is an independently varied access price.

For two intervals, let ℓ_1, ℓ_2 be their widths and g their separation, and define

$$\eta = \frac{\ell_1\ell_2}{(\ell_1+g)(\ell_2+g)}. \tag{5.2}$$

The leading-RT mutual information is $(c/3)[\log(\eta/(1-\eta))]_+$ [9]. Combining the field threshold (2.2) with the pressure correspondence gives

$$\begin{aligned} G_\tau^{\mathrm{RT}} &= \frac{c}{3}\left[\log\frac{\eta}{1-\eta} - \frac{3\tau}{c}\right]_+, \\ F_\lambda^\epsilon - F_{\lambda,\mathrm{disc}}^\epsilon &= -\frac{cq}{3}\left[\log\frac{\eta}{1-\eta} - \frac{\pi\lambda}{q}\right]_+ + O(c\epsilon/\delta). \end{aligned} \tag{5.3}$$

At leading ultraviolet order, the connected branch minimizes the pressure action when

$$\eta > \eta_{\mathrm{crit}}(\lambda) = \frac{1}{1+\exp(-\pi\lambda/q)}. \tag{5.4}$$

The zero-source limit recovers the transition at $\eta = 1/2$. Positive pressure increases the correlation needed to support a bridge. This two-interval result fixes the normalization for the additional partition choices available to three components.

5.1 Collective and nested partitions

For three components, denote the strongest pair mutual information by $p = \max_{i<j} I(A_i : A_j)$ and the total correlation by $t = T(123)$. Maximizing over the five partitions yields the full envelope

$$G_\tau = \max\{0, p - \tau, t - 2\tau\}, \qquad t = \sum_{i<j} I(A_i : A_j) - I_3, \tag{5.5}$$

Here $I_3 = S_1 + S_2 + S_3 - S_{12} - S_{13} - S_{23} + S_{123}$. A pair attaining p has mutual information $t - p$ with the remaining component, hence $t \geq p$. If $t > 2p$, the triple passes directly to singletons at $\tau = t/2$ and is optimal below that price. It remains profitable throughout $p < \tau < t/2$, where every pair merger loses. If $p \leq t < 2p$, a strongest pair is optimal between $\tau = t - p$ and $\tau = p$. At $t = 2p$, all three branch types meet. Thus $t > 2p$ characterizes the collective regime, without requiring $I_3 \neq 0$.

An explicit leading-RT example consists of three unit intervals separated by gaps 0.2:

$$(u_1, v_1, u_2, v_2, u_3, v_3) = (0, 1, 1.2, 2.2, 2.4, 3.4). \tag{5.6}$$

The adjacent pairs have $\widehat{I}_{12} = \widehat{I}_{23} = 1.641961104$, and the distant pair has $\widehat{I}_{13} = 0$. The union has total correlation

$$\widehat{T}(123) = -2\log(3.4 \times 0.2^2) = 3.990200786. \tag{5.7}$$

At $\widehat{\tau} = 1.8$, every pair merger has negative payoff, whereas joint storage of the triple saves

$$\widehat{G}_{1.8} = 3.990200786 - 2(1.8) = 0.390200786 > 0. \tag{5.8}$$

At this price, each adjacent-pair partition yields -0.158038896, the distant pair yields -1.8, and singletons yield zero. These exhaust the competing partitions. The triple remains the unique optimum until its direct transition at $\widehat{\tau} = 1.995100393$. A greedy search restricted to profitable pair mergers therefore misses this phase.

The same optimization admits spatially nonconsecutive blocks. For example, take

$$(u_1, v_1, u_2, v_2, u_3, v_3) = (0, 1, 1.001, 1.002, 1.003, 2.003). \tag{5.9}$$

Joining the first and third components gives

$$\widehat{T}(13) = -2\log(2.003 \times 0.003) = 10.228993867. \tag{5.10}$$

At $\widehat{\tau} = 3$, the partition $\{1, 3\} \mid \{2\}$ has saving 7.228993867, above the triple's 6.426218445; both adjacent-pair payoffs are negative and the singleton payoff is zero. The minimizing region has an outer positive component joining A_1 and A_3 around a cutoff-connected negative gap, with a separate A_2 cap inside that gap. This nested configuration obeys the same matching conditions as the other channels. Figure 2 shows the pressure-minimizing region, and figure 3 gives the complete information envelopes for both examples.

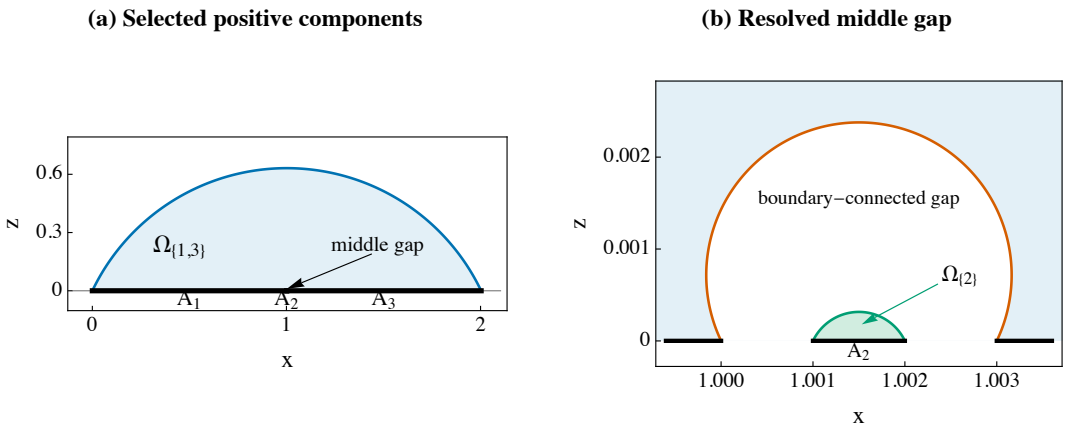


Figure 2. Pressure-minimizing positive homology region for the nested example (5.9). The outer component connects A_1 to A_3 around a cutoff-connected negative gap containing a separate positive A_2 cap. The right panel enlarges the gap. The minimum is taken over all six directed assignments at source $\lambda = 3/\sqrt{4\pi^2 + 9}$, corresponding to $\widehat{\tau} = 3$, with $\epsilon = 10^{-6}$. Both panels use equal-aspect coordinates and show the computed pressure arcs. Theorem 4.2 relates this finite-cutoff geometry to the information task in the ultraviolet limit.

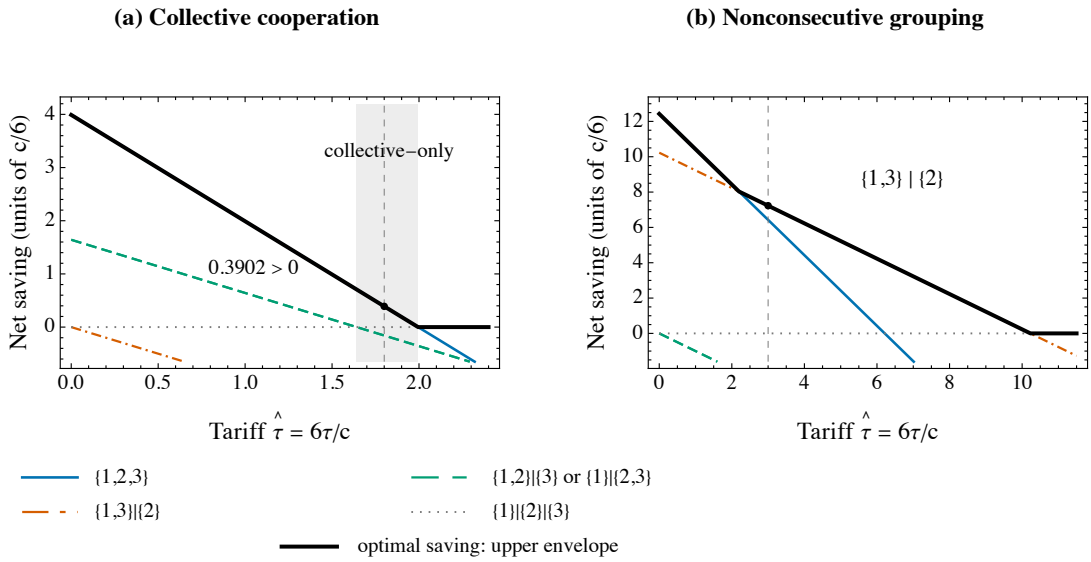


Figure 3. Complete price envelopes for the three-component examples, in units of $c/6$. Each branch gives a fixed partition's net saving relative to separate storage; the thick black curve is the maximum over all five partitions. Symmetry makes the two adjacent-pair branches coincide. (a) For the geometry (5.6), a collective interval has negative pair payoffs and a profitable triple. (b) For the geometry (5.9), the nonconsecutive partition $\{1, 3\} \mid \{2\}$ is strictly optimal over a price interval. Vertical markers locate $\widehat{\tau} = 1.8$ and 3, respectively. Leading vacuum RT entropies determine the branches after cancellation of the common ultraviolet terms.

5.2 Pairwise-additive and holographic correlation data

Collective saving can also arise from pairwise-additive entropy data. The ungauged massless Dirac vacuum on the line provides a comparison with the holographic case when each union is reduced to its fermionic mode algebra. Its replica prescription fixes Neveu–Schwarz conditions on the standard basis cycles, without the spin-structure sum used in the parity-gauged theory [25]. Disconnected spin reductions obtained through a Jordan–Wigner transformation refer to different algebras and need not share these entropies [26].[1] One full Dirac field has $c = 1$, and N independent copies give $c = N$. Subtracting the continuum union entropy from the sum of the single-interval entropies gives [28, 29]

$$T_D(B) = \sum_{\substack{i<j \\ i,j\in B}} I_D(i:j), \qquad I_D(i:j) = \frac{c}{3}\log\frac{(u_j-u_i)(v_j-v_i)}{(u_j-v_i)(v_j-u_i)}. \tag{5.11}$$

Each block therefore gains the sum of its internal pair mutual-information weights. Including the merger charge gives a graphical partition optimization [1].

For three equal-width intervals whose adjacent gap-to-width ratio is $r > 0$, these pair weights take the form

$$a_D = \widehat{I}_{12} = \widehat{I}_{23} = 2\log\frac{(1+r)^2}{r(2+r)}, \qquad b_D = \widehat{I}_{13} = 2\log\frac{(2+2r)^2}{(1+2r)(3+2r)}. \tag{5.12}$$

Since $a_D > b_D > 0$, the normalized envelope (5.5), with $p = a_D$ and $t = 2a_D + b_D$, reduces to

$$\widehat{G}_{\tau,D} = [2a_D + b_D - 2\widehat{\tau}]_+, \qquad \widehat{\tau}_{\mathrm{off},D} = a_D + b_D/2. \tag{5.13}$$

Every pair payoff is negative on $a_D < \widehat{\tau} < a_D + b_D/2$, while the triple remains profitable. The saving is fully determined by the pair mutual informations, with $I_3 = 0$. Pairwise-additive entropy data can thus support a collective optimum that a pairwise-greedy rule misses. Neither the collective transition nor $I_3 = 0$ is a criterion for genuine multipartite entanglement.

Leading RT in the same equal-interval geometry supports a stronger effect: a triple can save even when every pair mutual information vanishes. The five endpoint-channel

[1] At a finite fermionic regulator, the component algebras satisfy the canonical anticommutation relations and combine through the graded tensor product. For physical parity-superselected states and channels, the fermionic source-coding theorem fixes the minimum block rate at $S(\rho_B)$ [27, Theorem III.1]. We express its base-two rates in nats. The entanglement-fidelity recovery criterion controls every block dilation, including correlations with external systems [27, Propositions III.1–III.2]. The source is even, hence commutes with total parity. Admissible coding maps are fermionic completely positive trace-preserving maps whose Kraus operators have definite parity, as specified in Proposition II.1 of that source. These operators may have either parity; they need not all be even. Extending the maps by the identity in the graded sense gives physical parallel compositions. For a fixed finite number of blocks, let δ_B bound the trace-norm error of the extended recovery map on the original joint source and its purifying reference. A telescoping expansion of the composition minus the identity, followed by trace-norm contractivity for each term, bounds the joint error by $\sum_B \delta_B \to 0$. This joint recovery justifies the block-rate optimization and its declared merger charge for the fermionic task. The construction uses graded fermionic products throughout and makes no identification with ordinary disjoint spin tensor factors.

products are 1, $r(2+r)$, $r(2+r)$, $r^2(3+2r)$ and $(1+2r)(3+2r)$. The last exceeds one. The fourth is minimal for $r < 1/2$, and the singleton product is minimal for $r \geq 1/2$. These channel comparisons give

$$\widehat{p}_{\mathrm{RT}} = [-2\log(r(2+r))]_+, \quad \widehat{t}_{\mathrm{RT}} = [-2\log(r^2(3+2r))]_+, \quad \widehat{I}_{13,\mathrm{RT}} = 0. \tag{5.14}$$

For $\widehat{p}_{\mathrm{RT}} > 0$, $\widehat{t}_{\mathrm{RT}} - 2\widehat{p}_{\mathrm{RT}} = 2\log[(2+r)^2/(3+2r)] > 0$. The same strict inequality holds when the pair value is zero and the triple value is positive. The full price envelope is therefore $\widehat{G}_\tau^{\mathrm{RT}} = [\widehat{t}_{\mathrm{RT}} - 2\widehat{\tau}]_+$. Throughout the strict window

$$\sqrt{2} - 1 < r < \frac{1}{2}, \qquad 0 < \widehat{\tau} < -\log[r^2(3+2r)] \tag{5.15}$$

all pair mutual informations vanish at leading RT order and the triple is the unique optimum. Its total correlation obeys $\widehat{t}_{\mathrm{RT}} = -\widehat{I}_3$, so the saving depends on the known nonextensivity of holographic mutual information [30]. The strict channel and payoff inequalities ensure stability under sufficiently small endpoint perturbations.

At $r = 0.45$, the switching price is $\widehat{\tau}_{\mathrm{off,RT}} = 0.2360388393\ldots$, or $\lambda_c = 0.03754026783\ldots$ in source units by (5.1). Relative to the same-source optimized caps, the ultraviolet pressure bridge persists for $0 < \lambda < \lambda_c$. This contrasts pairwise-additive cooperation with a leading-RT saving that cannot be reconstructed from the three pair values. The endpoint ratio is shared, while the pair data differ between the theories. Exact factorization of the finite-c marginals does not follow from vanishing leading-RT pair values, and the comparison provides no entanglement witness.

Figure 4 collects the full price envelopes at $r = 0.45$. Both models retain a profitable triple after every pair payoff becomes negative. Vanishing pair mutual informations occur only in the leading-RT panel at these endpoints.

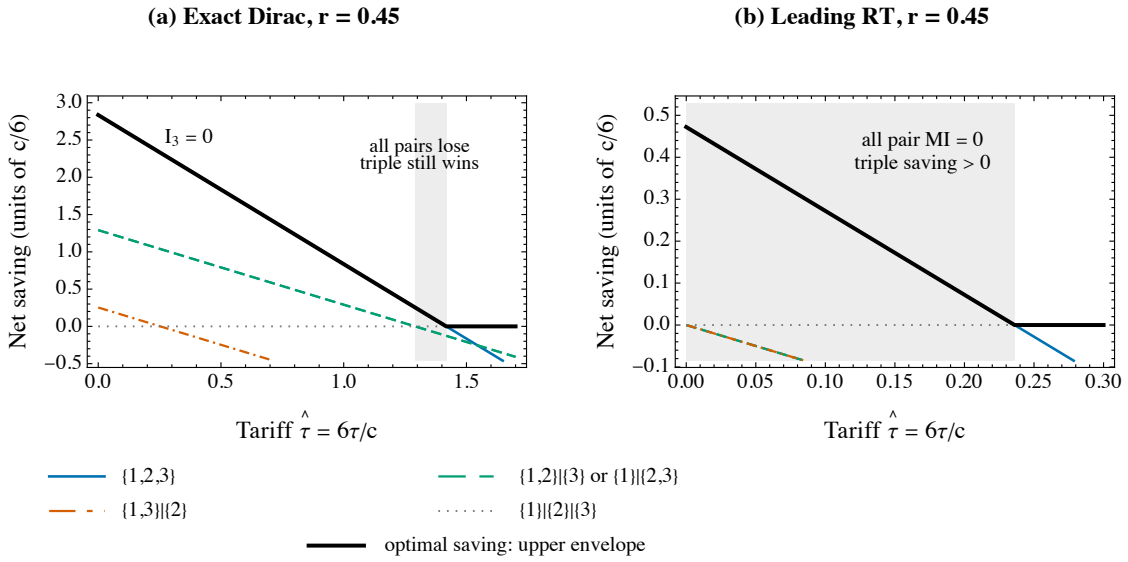


Figure 4. Full continuum price envelopes for three equal intervals with gap-to-width ratio $r = 0.45$. (a) For the ungauged massless Dirac vacuum on the fermionic mode algebras, the total correlation is pairwise additive and $I_3 = 0$. The triple remains optimal on $a_D < \widehat{\tau} < a_D + b_D/2$, where all pair payoffs are negative. (b) Leading Einstein-AdS$_3$ vacuum RT gives zero pair mutual informations and a profitable triple on $0 < \widehat{\tau} < 0.2360388393\ldots$. The thin branches are the five partition payoffs, and the thick curve is their upper envelope in units of $c/6$. Adjacent-pair branches coincide in (a), and all three pair branches coincide in (b). Gray shading identifies the collective intervals. The panels hold the endpoint ratio fixed and use different tariff ranges; their pair correlation data are not matched. The comparison concerns continuum entropy models and carries no entanglement-witness interpretation. The access price in the Dirac panel is varied independently.

5.3 Area response and cutoff corrections

As the tariff increases, the principal-partition hierarchy refines the finest optimal partition. Monotonicity of (5.1) carries this refinement to the source-selected positive components in the leading ultraviolet theory. Within a generic phase with a unique finest partition, information groups that have separated remain separate as pressure increases. The optimized value stays continuous through a transition, although its slope can change. Finite-cutoff corrections can shift the crossing.

On a stable branch with optimal partition $\mathcal{P}$, the unpriced total correlation and merger count are

$$T_{\mathcal{P}} = \sum_{B\in\mathcal{P}} T(B), \qquad m_{\mathcal{P}} = n - |\mathcal{P}|, \qquad G_\tau = T_{\mathcal{P}} - \tau m_{\mathcal{P}}. \tag{5.16}$$

The response uses leading-RT entropies, for which the relative branch action tends to $-qT_{\mathcal{P}} + (c\pi\lambda/3)m_{\mathcal{P}}$. To differentiate this limit, convergence of the value must be supplemented by a derivative estimate. Appendix F derives a uniform first-derivative expansion from the exact arc action on compact subcritical source intervals. For an open interval

with a stable winning matching, the envelope theorem then yields

$$\lim_{\epsilon\to 0}\frac{V^{\epsilon}_{\mathcal{P}}-V^{\epsilon}_{\mathrm{disc}}}{4GR}=\frac{\lambda}{q}T_{\mathcal{P}}+\frac{c\pi}{3}m_{\mathcal{P}}. \tag{5.17}$$

The reference area comes from caps optimized separately at the same pressure source.

Consider a transition at $0<\lambda_c<1$ between stable phases with merger counts m_- and m_+ below and above the crossing. Equality of the branch actions fixes $T_- - T_+ = c\pi\lambda_c(m_- - m_+)/(3q_c)$, with $q_c=\sqrt{1-\lambda_c^2}$. First taking the ultraviolet limit within each stable phase, then approaching the transition from either side, gives

$$\boxed{V_- - V_+ = \frac{2\pi R^2}{1-\lambda_c^2}(m_- - m_+).} \tag{5.18}$$

The common cutoff divergence cancels in this limiting one-sided area difference. In the pressure representation, each lost cooperation link contributes the stated area jump. At a multi-branch crossing, $m_\pm$ denote the merger counts of the phases active immediately on either side. The jump vanishes when those counts coincide. At zero source, the theorem applies only to the right-sided response.

Equation (5.18) gives the nonzero-pressure extension of the Gauss–Bonnet area jump [11], with boundary deformation encoded in $1/(1-\lambda_c^2)$. For the two-interval geometry in figure 1, $\lambda_c = 0.2528355461\ldots$ predicts an area jump $6.7122728059\ldots R^2$. Independent integration of the circles at the actual finite-cutoff crossing yields $6.7122703665\ldots R^2$ at $\epsilon = 10^{-7}$, consistent with the derivative calculation for this geometry. Control of the area derivative comes from the arc-action expansion; the RT entropy-value bound and the value error in theorem 4.2 alone are insufficient.

Expanding the exact edge action gives the first cutoff corrections:

$$\begin{aligned} f_{\lambda,\epsilon}(u,v) = {}& \frac{\lambda(v-u)}{\epsilon}+2q\log\frac{\Delta}{\epsilon}+2(q\log q+\lambda\arcsin\lambda)-\pi\lambda\sigma \\ & +\frac{4\sigma\lambda\epsilon}{\Delta}+\frac{2(1-2\lambda^2)}{q}\frac{\epsilon^2}{\Delta^2}+O(\epsilon^3/\Delta^3). \end{aligned} \tag{5.19}$$

The coefficients are derived in appendix F. Subtracting the optimized caps from a stable winning matching M gives

$$F^{\epsilon}_{\lambda}-F^{\epsilon}_{\lambda,\mathrm{disc}}+qG^{\mathrm{RT}}_{\tau}=\frac{c}{6}\,4\lambda\epsilon\left[\sum_{e\in M}\frac{\sigma_e}{\Delta_e}-\sum_i\frac{1}{v_i-u_i}\right]+O(c\epsilon^2/\delta^2). \tag{5.20}$$

The linear coefficient is generically nonzero: reverse gap arcs and forward cap arcs contribute with opposite signs, precluding its absorption into a common endpoint term. Figure 5 tests this expansion against the exact finite-cutoff minimum.

Replacing the boundary logarithms with finite-anchor geodesic entropies (C.3) changes the quadratic edge coefficient to $-2\lambda^2/q$. The linear residual persists, so the information correspondence remains inexact at a finite wall. At a leading-order phase degeneracy, the cutoff corrections must be minimized over the tied branches; an arbitrary tied branch need not continue to the finite-cutoff minimum.

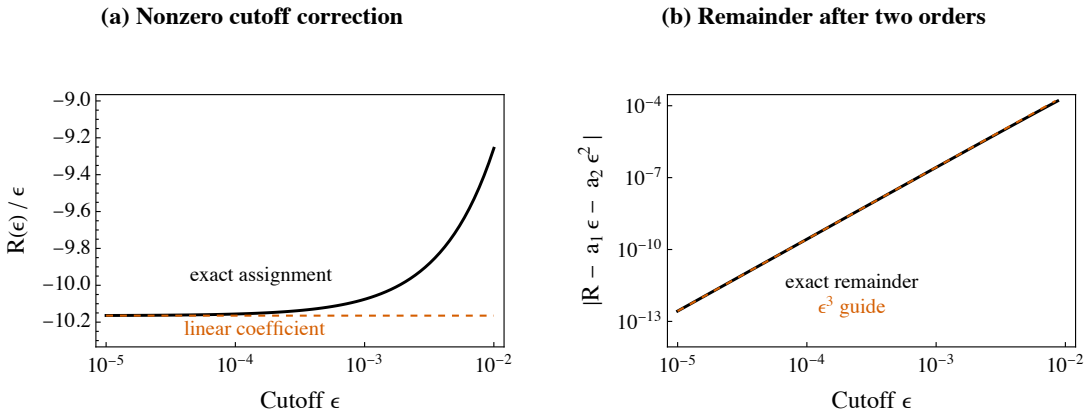


Figure 5. Normalized finite-cutoff residual for (5.6) at $\lambda = 0.2$, with unit interval width. We denote the residual by $R(\epsilon) = 6[F^{\epsilon}_{\lambda} - F^{\epsilon}_{\lambda,\mathrm{disc}} + qG^{\mathrm{RT}}_{\tau}]/c$. The exact pressure minimum is taken over all six directed assignments and compared with an independent optimization of the RT partition saving. (a) $R(\epsilon)/\epsilon$ approaches the nonzero linear coefficient a_1 in (5.20). (b) Removing $a_1\epsilon + a_2\epsilon^2$ leaves a cubic remainder; a_2 follows from (5.19) by summing the winning edges and subtracting the singleton caps. The dashed ϵ^3 line is a guide normalized at the smallest cutoff, with no fitted coefficient. High-precision arithmetic is used at 80 cutoffs in $[10^{-5}, 10^{-2}]$. The curves test one fixed geometry; the all-component statement rests on the analytic derivation.

Adding a local entropy-scheme term a_i to each original component shifts every partition cost by $\sum_i a_i$, which cancels from G_τ. The geometric subtraction similarly removes the fixed-trace area divergence and common endpoint terms. A finite term depending on the number of interior components would alter the phase competition; the positive-area functional used here contains no such term.

6 Thermal states and the compact annulus obstruction

Temperature changes the RT endpoint cost, while the static planar-BTZ exterior remains locally hyperbolic. To extend the pressure representation, we must also transform its variational domain: the cutoff becomes oblique, and the minimizing region must remain outside the horizon.

Let $z_h = \beta_T/(2\pi)$, where β_T is the inverse temperature. The metric of the right planar-BTZ exterior is

$$\frac{\mathrm{d}s^2}{R^2} = \frac{\mathrm{d}x^2 + \mathrm{d}z^2/(1 - z^2/z_h^2)}{z^2}, \qquad x \in \mathbb{R}, \quad \epsilon \le z \le z_h. \tag{6.1}$$

The prescribed boundary region is a fixed finite union of separated finite intervals on the infinite line. There is no spatial identification or winding sector. Admissible positive regions have compact essential support, finite perimeter and the prescribed interior BV cutoff trace. Horizon segments in the reduced boundary contribute to the free perimeter. The same minimum is obtained by imposing zero horizon trace; appendix G proves this equivalence by showing that the minimizing curves are confined to the exterior.

For a single interval, the leading planar-BTZ RT entropy is

$$S_{\mathrm{RT},\beta_T}([u,v]) = \frac{c}{3}\log\left[\frac{\beta_T}{\pi\epsilon}\sinh\frac{\pi(v-u)}{\beta_T}\right]. \tag{6.2}$$

The entropy of a union minimizes the sum of these endpoint costs over admissible non-crossing matchings, including disconnected channels [31, 32]. Denote by $G_\tau^{\mathrm{RT},\beta_T}$ the saving (1.3) obtained from these thermal entropies, with all set partitions allowed.

Theorem 6.1 (Planar thermal correspondence)**.** *Let $0 \le \lambda < 1$, $q = \sqrt{1-\lambda^2}$, and let the interval endpoints and $z_h > 0$ be fixed, and suppose $e = \epsilon/z_h < q$. In the compact-essential-support charged-horizon BV domain of appendix G, the pressure minimum is attained. Every minimizer has exactly n disjoint calibrated free arcs, no additional cutoff contacts, and no horizon contact. In the ultraviolet limit,*

$$F^{\epsilon}_{\lambda,\beta_T}(A) - F^{\epsilon}_{\lambda,\beta_T,\mathrm{disc}}(A) = -qG_\tau^{\mathrm{RT},\beta_T}(A_1,\dots,A_n) + O\left(c\frac{\epsilon}{z_h}\right), \qquad \tau = \frac{c\pi\lambda}{3q}. \tag{6.3}$$

The implicit constant depends on the fixed dimensionless endpoints x_a/z_h and is uniform for λ in compact subsets of $[0,1)$. The disconnected baseline is the sum of separately pressure-minimized single-interval actions in the same domain.

The map

$$U = e^{x/z_h}\sqrt{1-z^2/z_h^2}, \qquad Z = e^{x/z_h}z/z_h \tag{6.4}$$

is an isometry from the exterior to $U \ge 0$ in the upper half-plane. It sends the horizon to the geodesic $U = 0$ and the cutoff to an oblique ray. Applying calibration and the finite-perimeter reduction with trace on this ray gives a minimizing assignment of alternating caps. The outermost caps are forward and lie strictly at $U > 0$. A reverse arc enclosed by such a cap could reach the horizon only by crossing the cap, which a minimizing assignment excludes. The minimizing region is therefore confined to the exterior. Its component count fixes the same tariff as in the vacuum, giving (6.3). Temperature-independent RT-volume topology terms were obtained in [11]; the argument here establishes the pressure-optimized minimum on the stated domain.

For two intervals of widths ℓ_1, ℓ_2, gap g and total span $T = \ell_1 + g + \ell_2$, define

$$K_{\beta_T} = \frac{\sinh(\pi\ell_1/\beta_T)\sinh(\pi\ell_2/\beta_T)}{\sinh(\pi g/\beta_T)\sinh(\pi T/\beta_T)}. \tag{6.5}$$

The corresponding saving and phase condition are

$$G_\tau^{\mathrm{RT},\beta_T} = \frac{c}{3}\left[\log K_{\beta_T} - \frac{\pi\lambda}{q}\right]_+, \tag{6.6}$$

$$\text{connected phase} \iff \log K_{\beta_T} > \frac{\pi\lambda}{q}. \tag{6.7}$$

If $\log K_{\beta_T} > 0$, the ultraviolet critical source is

$$\lambda_c = \frac{\log K_{\beta_T}}{\sqrt{\pi^2 + (\log K_{\beta_T})^2}}. \tag{6.8}$$

If this logarithm is nonpositive, no connected phase survives at positive source; equality denotes a tie. Temperature thus enters through the conformal cross-ratio without changing the tariff. The exact finite-cutoff action and its nonzero residual are derived in appendix G. The theorem covers finite intervals on the infinite thermal line with charged horizon boundary. Infinite intervals, compact quotients, free horizon attachments and general excited states require separate arguments.

6.1 The compact annulus obstruction

Compact BTZ admits an annular homology sector bounded at zero source by complement geodesics and the horizon [11, 20]. An annulus is connected and has Euler characteristic zero. The orientation term therefore ceases to count laboratory groups, which was the step converting area into a merger charge. We test the unchanged tariff against the independently optimized compact pressure problem. Figure 6 displays the contrast with the planar geometry.

Use $J = 6F/c$ and $\widehat{G} = 6G/c$, set $z_h = 1$, and identify $x \sim x + H$ with $H = 10$. Choose

$$A_1 = [0, 4.9], \qquad A_2 = [5, 9.9], \qquad \lambda = 0.1, \tag{6.9}$$

Each complementary gap then has length 0.1. Charging all non-cutoff boundary leaves the essential inner loop free to move away from the horizon. Its minimum contribution is

$$\min_{\rho \geq 0} H(\cosh\rho - \lambda \sinh\rho) = qH, \qquad \rho_* = \operatorname{artanh}\lambda. \tag{6.10}$$

The bound $\cosh\rho - \lambda\sinh\rho \geq q$ establishes this minimum. Its location lies inside the cutoff when $e < q$. The loop must therefore be optimized along with the anchored arcs; fixing it at its zero-source position changes the variational problem.

Write $d_\pm(t)$ for the exact thermal directed edge action at separation t and orientation $\pm$. After subtracting the common fixed-trace term, the candidate actions are

$$\begin{aligned} J_{\text{two}} &= 2d_+(4.9), \\ J_{\text{bridge}} &= d_+(9.9) + d_-(0.1), \\ J_{\text{annulus}} &= qH + 2d_-(0.1). \end{aligned} \tag{6.11}$$

The annulus attains the global minimum for sufficiently small cutoff. The proof in appendix H includes every integer winding sector. Current decomposition first separates the winding contribution; $\partial_t d_\pm > q$ then reduces the infinite lift search to eight lower bounds. The annular bound is attained because the reverse gap arcs remain separated from each other and from the essential loop. Applying the same argument to each singleton gives its direct disk cap as the independently minimized reference.

The resulting leading-RT mutual information and storage saving are

$$\widehat{I} = 4\log\frac{\sinh 2.45}{\sinh 0.05} - 10 = 8.9787760550\ldots, \qquad \widehat{G}_\tau = \widehat{I} - \frac{2\pi\lambda}{q} > 0. \tag{6.12}$$

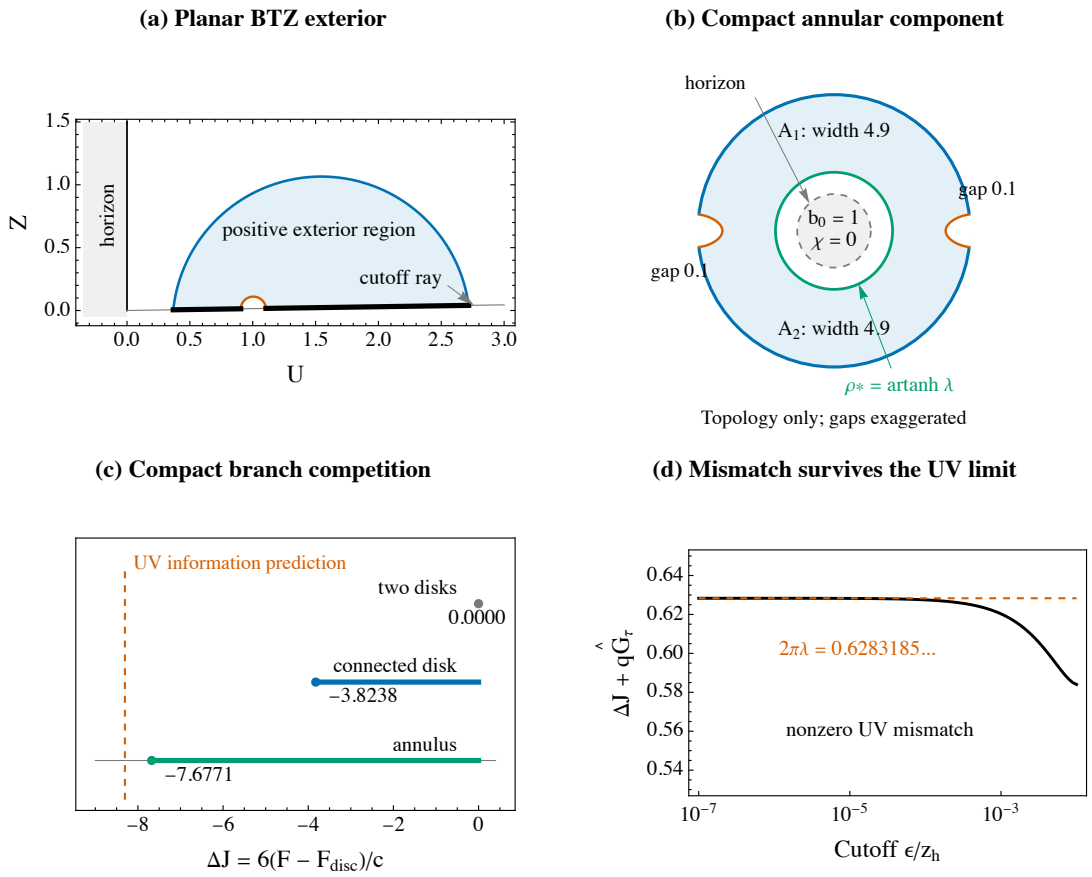


Figure 6. Horizon confinement on the line and its compact-topology obstruction. (a) Pressure-minimizing curves for $A_1/z_h = [-1, -0.1]$, $A_2/z_h = [0.1, 1]$, $\lambda = 0.1$ and $\epsilon/z_h = 0.015$, shown in upper-half-plane coordinates $U = e^{x/z_h}\sqrt{1 - z^2/z_h^2}$, $Z = e^{x/z_h} z/z_h$. The positive region lies strictly at $U > 0$, bounded away from the geodesic horizon $U = 0$; the cutoff is an oblique ray. Equal-aspect coordinates display the exact curves, with their lengths measured in the hyperbolic metric. (b) Topology schematic for $x \sim x + 10z_h$ and $A/z_h = [0, 4.9] \cup [5, 9.9]$. Gaps are exaggerated and the radial direction has no metric scale. The annulus has $b_0 = 1$ and $\chi = 0$. Its essential inner loop minimizes the action at $\rho_* = \operatorname{artanh} \lambda$, outside the horizon $\rho = 0$. (c) For $\lambda = 0.1$ and $\epsilon/z_h = 10^{-6}$, the exact branch differences $\Delta J = 6(F - F_{\lambda,\mathrm{disc}})/c$ are 0, -3.8237963, and -7.6771407 for two disks, either symmetry-related connected disk, and the annulus. Leading-RT ultraviolet entropies with the unchanged tariff instead give $-q\widehat{G}_\tau = -8.3054508$. (d) The cutoff dependence approaches $\Delta J + q\widehat{G}_\tau \to 2\pi\lambda$, a finite mismatch $c\pi\lambda/3$ in F. All free-boundary length, including any horizon boundary, is charged. The annular minimum therefore obstructs extension of the planar representation with the same tariff.

The pressure and storage expressions consequently approach different limits:

$$\lim_{e\to 0}(J_{\text{annulus}} - J_{\text{two}}) = -q\widehat{I} + 4\pi\lambda = -7.6771323139\ldots,$$
$$-q\widehat{G}_\tau = -q\widehat{I} + 2\pi\lambda = -8.3054508446\ldots. \tag{6.13}$$

The difference is $2\pi\lambda$ in J, or $c\pi\lambda/3$ in F, and remains finite in the ultraviolet limit. Two disks have $\chi = 2$ and the annulus has $\chi = 0$: the geometric change carries two Euler units, while the storage task charges one merger. Common endpoint and fixed-trace counterterms cancel from the comparison and cannot remove the discrepancy. Both sides already use their global optima and independently minimized singleton reference. For this leading-RT ultraviolet example, the unchanged storage formula overestimates the relative saving and hence underestimates the pressure-minimized action. The storage optimum remains valid for its specified resources; disk topology is the assumption that fails in its compact geometric representation.

7 Discussion

For fixed component algebras at a finite factorizing regulator, the net saving from independently recoverable block memories and recurring joint access is the entropy-partition optimum G_τ. Its geometric representation is the main result of this paper. In the leading-RT vacuum and noncompact thermal models, G_τ^{RT} equals the difference between the separately optimized cap action and the global pressure minimum, divided by q, in the ultraviolet limit. Entropy-preserving refinement identifies the competing positive components; calibration and current decomposition establish the minimum over the specified compact-essential-support domains. Replacing G_τ^{RT} by the actual-state saving requires the separate error control in (4.18).

Three components already distinguish this optimization from pairwise merging. The ungauged free Dirac field supports collective cooperation with pairwise-additive entropy data. In the leading-RT window (5.15), the triple has positive saving even though all pair mutual informations vanish; those three pair values cannot determine the saving. A nonconsecutive block can also be uniquely optimal. Both possibilities require the full partition domain. Since classical and quantum correlations both contribute, G_τ is not an entanglement witness.

The result depends on which laboratories and resources are charged. Under ideal internal control, an all-pairs rental network realizes the linear merger count. Weighted graphs and state-dependent communication costs define different optimizations. Even the choice of laboratories changes the answer: a Bell pair within one laboratory has $G_\tau = 0$, whereas separately charging its two halves gives $(2\log 2 - \tau)_+$. The source-to-price dictionary represents the specified access model; it neither fixes a unique microscopic price nor identifies memory-equivalent rental with physical work. Compression-map complexity [11] and constrained-merging entanglement cost [13] likewise charge resources distinct from the memory rate and recurring access used here.

The annular counterexample isolates the topological requirement. Nested positive caps retain disk topology; an annulus separates Euler characteristic from the number of

laboratory groups. Its finite ultraviolet mismatch persists after independently minimizing the singleton reference. Replacing vacuum endpoint distances by thermal ones therefore cannot extend the representation to compact BTZ. Adding a topology-dependent charge would change the storage task and require a new operational definition.

The supplementary note treats complementary spectra, interval-entropy composition, supplied conformal probes, density-matrix geometry and replica channel weights. Each uses different inputs and composition rules. Its finite-cutoff supplied-probe identity assumes prepared probes and a tracked phase reference, and does not establish the actual-state storage result. Here coding is performed first at fixed finite factorizing regulator and theory. Continuum relative-entropy control and the leading-RT approximation are separate requirements.

Further questions. Finite-c control is needed before the geometric response can be compared with the actual storage task near a phase boundary. Equation (4.18) bounds errors in G_τ from uniform block-entropy errors, but controls neither source derivatives nor changes of optimizer. An order-one mutual information in a leading-RT disconnected two-interval phase, for example, keeps the actual saving positive until τ reaches that value. With $\tau = c\pi\lambda/3 + O(c\lambda^3)$, this produces a $\lambda = O(1/c)$ boundary layer. Its crossover must be determined from CFT entropies before extracting an area response. The large-c limit must precede identification of the zero-source classical response; away from that limit, a phase gap or suitable one-sided derivative control is required.

Local access constraints pose a different question. With finite link capacity, distance-dependent rental or state-dependent communication costs, which savings are achievable, and what converse fixes their optimum? The weighted-graph construction in appendix B specifies one starting point. A geometric comparison would then have to represent this independently fixed task, with its own access rule and price.

For compact space, one can instead ask whether another independently defined information task admits a pressure representation. Joint prediction relative to a specified same-theory thermal reference on a connected laboratory hull is one candidate. It needs full-hull access and a regional reference prescription beyond the disconnected state. A model-regret penalty would have to follow from the allowed predictors, actions and prices and yield its own operational optimum. Choosing that penalty to cancel the annular mismatch would leave the central question unanswered. Any compact correspondence for this prediction task would be a separate result.

The geometric domain also leaves concrete open problems. A no-escape or truncation theorem would be needed to replace compact essential support by arbitrary finite-action exact-trace competitors. Uniform control as gaps close, the number of components grows, or λ approaches one is likewise open: the present estimates fix finite n and separated endpoints, restrict λ to a compact subcritical interval, and require $e < q$ thermally. Pressure minimization selects the exposed perimeter–area tradeoff; it does not recover every fixed-area branch or stationary saddle. Lorentzian, higher-dimensional and end-of-the-world-brane extensions each require a specified variational principle and a separate comparison with the information task.

Acknowledgments

Feiyu Deng is supported by the National Natural Science Foundation of China (NSFC) Project No. 12547135. The author used OpenAI models and Hermes Agent (Nous Research) for manuscript drafting and revision, literature searches, consistency checks and computational assistance, and thanks OpenAI and the Hermes Agent team for developing these tools. The author is responsible for all scientific judgments, calculations and conclusions.

A Coding and entropy-partition properties

The storage cost in (1.2) follows from independently recoverable compression of the block marginals. We derive this rate and its consequences for general finite-dimensional states. At a finite regulator, C_τ includes the stored memory and the access charge. The reference cost is the optimal storage rate with cooperation forbidden.

Fix a finite-dimensional factorizing regulator and a theory, and hold both fixed throughout the coding limit. The subsystem Hilbert spaces are

$$\mathcal{H}_A = \bigotimes_{i=1}^{n} \mathcal{H}_{A_i}, \qquad A_i = [u_i, v_i], \qquad u_1 < v_1 < \cdots < u_n < v_n. \tag{A.1}$$

The atoms are the maximal connected components of the prescribed region. For $B \subseteq [n]$, set $A_B = \bigcup_{i\in B} A_i$ and $\rho_B = \mathrm{Tr}_{A\setminus A_B}\, \rho_A$. We use natural logarithms and measure information rates in nats. The state may have arbitrary correlations between components.

For a partition $\mathcal{P}$, assign an encoder and a decoder to each block. After m independent source repetitions, the encoder for B compresses $\rho_B^{\otimes m}$ into a quantum memory $M_{B,m}$. The decoder has access only to that memory. Its entanglement fidelity must tend to one, so recovery preserves correlations with a purifying reference as well as the marginal state. Decoder side information, auxiliary shared entanglement, uncounted classical memory and cross-block decoding are excluded.

The optimal memory rate for these independently recoverable blocks is

$$Q(\mathcal{P}) = \inf \limsup_{m\to\infty} \frac{1}{m} \sum_{B\in\mathcal{P}} \log \dim M_{B,m} = \sum_{B\in\mathcal{P}} S(\rho_B). \tag{A.2}$$

Schumacher compression attains the entropy rate for each marginal [6, 33]. To recover the full source, apply the block recovery channels to a purification of that source. A telescoping expansion and trace-distance contractivity bound the resulting error by the sum of the local recovery errors. The number of blocks is fixed, so this sum tends to zero with the local entanglement-fidelity errors. The product recovery therefore preserves interblock correlations as well as local states.

For the converse, let R_B purify the block source. Recovery with asymptotically unit entanglement fidelity requires

$$2mS(\rho_B) - o_{\mathrm{reg},c}(m) \le I(R_B : M_{B,m}) \le 2 \log \dim M_{B,m}. \tag{A.3}$$

Data processing and entropy continuity give the first inequality at the fixed finite factorizing regulator and fixed theory, including fixed c. The second follows from the memory dimension. Thus $o_{\mathrm{reg},c}(m)/m \to 0$ as $m \to \infty$ with these data held fixed. Here reg labels the coding regulator; ϵ is the radial cutoff in the geometric entropy convention. To exchange the coding limit with ultraviolet regulator removal or a large-c limit, one would need uniform dimension control or an energy-constrained continuity bound, together with control of the recovery-error scaling. We make no such exchange. The continuum limit of G_τ instead uses finite limits of the finitely many block relative entropies. Evaluation with leading RT entropies additionally uses the stated semiclassical assumptions and the entropy-value estimate (4.18). Summing the memory bounds over blocks proves the converse to (A.2), in agreement with the general noiseless-coding converse [7]. The fidelity condition refers to the specified source and its purification; it imposes no worst-case recovery requirement on every state in the full Hilbert space.

Joint operation of a block with $|B|$ constituents incurs the charge $m\tau(|B|-1)$, extensive in the number of source uses. The factor $|B|-1$ counts the links of a minimal abstract cooperation tree. This count by itself specifies neither a physical communication protocol nor its cost. A fixed one-time setup fee would vanish in the rate limit and define a different storage problem. The cost (1.2) follows by adding this extensive charge to (A.2).

Subtracting separate storage expresses the saving in terms of the block total correlation:

$$T(B) := D\left(\rho_B \middle\| \bigotimes_{i\in B} \rho_i\right) = \sum_{i\in B} S(\rho_i) - S(\rho_B). \tag{A.4}$$

To obtain the entropy expression, expand the logarithm of the product into the sum of the factor logarithms and evaluate each term in its marginal. For an arbitrary product state $\bigotimes_{i\in B} \sigma_i$, the same calculation gives

$$D\left(\rho_B \middle\| \bigotimes_{i\in B} \sigma_i\right) = T(B) + \sum_{i\in B} D(\rho_i\|\sigma_i).$$

With the usual support convention, positivity of the marginal relative entropies makes the product of the actual marginals the unique minimizing product state [21]. No auxiliary reference state enters. The optimal saving is therefore

$$G_\tau = \max_{\mathcal{P}\vdash[n]} \sum_{B\in\mathcal{P}} \left[T(B) - \tau(|B|-1)\right]. \tag{A.5}$$

Each partition retains the total correlations within its blocks and pays for the access needed to use them. Classical correlations contribute to this saving, which consequently does not define an entanglement measure.

The relative-entropy form gives two properties of the saving. At fixed $\tau \geq 0$, it is monotone under product local quantum channels:

$$G_\tau\left((\bigotimes_i \Lambda_i)(\rho_A)\right) \leq G_\tau(\rho_A). \tag{A.6}$$

On each block, the channel maps the product of the input marginals to the product of the output marginals. Data processing can therefore only decrease each payoff in (A.5), and the same holds for their maximum. For two independent groups of atoms in a state $\rho_U \otimes \rho_V$, the saving is additive:

$$G_\tau(\rho_U \otimes \rho_V) = G_\tau(\rho_U) + G_\tau(\rho_V). \tag{A.7}$$

If a block meets both groups, its total correlation is the sum of the contributions from U and V. Splitting the block between the groups preserves this entropy saving and removes a charge τ. Thus an optimal partition can be chosen with every block contained in one group, and the two maximizations separate. Both properties use the stated channels and independent states without additional shared resources.

Subadditivity and comparison with the singleton and one-block partitions give the bounds

$$S(\rho_A) \le C_\tau \le \sum_i S(\rho_i), \qquad 0 \le G_\tau \le D\left(\rho_A \middle\| \bigotimes_i \rho_i\right). \tag{A.8}$$

For two components the two possible partitions give

$$C_\tau = \min\{S_1 + S_2, S_{12} + \tau\}, \qquad G_\tau = (I(A_1 : A_2) - \tau)_+. \tag{A.9}$$

Joint encoding is strictly preferred precisely when the mutual-information saving exceeds the access charge. With more components, block total correlations enter the optimization, and pairwise tests alone need not determine the optimum.

For a fixed state, the finitely many partition costs are affine functions whose minimum is C_τ and whose slopes are $n - |\mathcal{P}|$. Their lower envelope is continuous, nondecreasing and concave; the saving G_τ is nonincreasing and convex. At zero charge, $C_0 = S(\rho_A)$, whereas the singleton partition minimizes the cost at sufficiently large charge. Wherever the optimal partition is unique over an interval,

$$\frac{\mathrm{d}C_\tau}{\mathrm{d}\tau} = n - |\mathcal{P}_{\mathrm{opt}}|, \qquad \frac{\mathrm{d}G_\tau}{\mathrm{d}\tau} = -(n - |\mathcal{P}_{\mathrm{opt}}|). \tag{A.10}$$

At a transition, the slopes of the minimizing partitions give the one-sided derivatives.

Removing the common term $n\tau$ puts the cost in the form

$$D_\tau := C_\tau - n\tau = \min_{\mathcal{P} \vdash [n]} \sum_{B \in \mathcal{P}} [S(\rho_B) - \tau]. \tag{A.11}$$

For Shannon entropy, this expression is the residual-entropy Dilworth truncation of [1]. Its principal partitions arise from submodular optimization [2–4]. Quantum strong subadditivity [5] makes $B \mapsto S(\rho_B)$ submodular, so the same partition construction yields a refinement hierarchy. Setting the empty-set value to zero makes $S(B) - \tau$ intersecting-submodular, which suffices without full submodularity of the shifted function. The finest minimizing partitions refine as τ increases. This mathematical construction does not transfer classical secret-key or communication-for-omniscience interpretations to quantum states: those interpretations require their own source models.

Equation (A.5) depends on the state together with its specified component algebras. Product unitaries preserve the saving, while a unitary that mixes components can change the block marginals. The comparison with geometry must retain this distinction. For separated regions in a continuum theory, one can use a common split or factorizing regulator. If all the finitely many block relative entropies have finite limits, G_τ also has a finite limit: the ultraviolet entropy terms cancel within each block payoff, and the maximum is over a finite set. Neither the individual memory rates nor C_τ need remain finite.

B Access networks and harmless refinement

We now specify an access resource that gives the charge in G_τ. At a finite factorizing regulator, a graph records the free coherent-control links between cells. Contract each connected free-control domain to a laboratory. When local control within A is free, time and gate depth are uncharged, and the complement provides no free relay, the laboratories are the maximal connected intervals A_i. Only the final stored memory is charged within a laboratory; coherent control across a gap requires rented access. The source is the actual reduced field state, repeated independently m times. A cell refinement that preserves the free-control domains leaves these laboratories unchanged. Charging the subdivisions as independent laboratories would define a different task.

Allow any pair of laboratories to rent an ideal bidirectional quantum link. Fix the active graph H for a coding batch. Each active link costs $m\tau$ in natural-log memory units, regardless of traffic, separation or processing time. A connected component of the graph permits arbitrary joint encoding and decoding. Distinct components have no cooperation, shared auxiliary resources, decoder side information or joint decoding. Connectivity provides the required joint access: finite-dimensional systems can be routed along a tree to one laboratory, processed jointly and returned. Capacity, gate depth, energy and latency are unconstrained, and the rental covers both encoding and recovery.

For a fixed H, Schumacher compression with independently recoverable block memories gives

$$\inf \limsup_{m\to\infty} \frac{1}{m} \left[\sum_{B\in\mathcal{P}(H)} \log\dim M_{B,m} + m\tau|E(H)| \right] = \sum_{B\in\mathcal{P}(H)} S(\rho_B) + \tau|E(H)|. \tag{B.1}$$

The fidelity condition and memory converse are those of (A.2). The noiseless coding theorem fixes the entropy term; the link charge is an independent resource assumption [6, 7, 33]. A connected graph with k vertices needs at least $k-1$ edges, with equality for a spanning tree. Removing cycles preserves joint access. Hence, for $\tau \geq 0$ and a prescribed partition,

$$\min_{H:\,\mathcal{P}(H)=\mathcal{P}} |E(H)| = \sum_{B\in\mathcal{P}} (|B|-1) = n - |\mathcal{P}|. \tag{B.2}$$

All pairwise links are available, so every partition can be realized. Minimizing (B.1) over the active graph gives C_τ; its difference from separate storage is G_τ. Connectivity determines the dependence on block cardinality, while the homogeneous price and its extensivity are

assumptions of the access model. The charged resource is the number of independent connections. It does not count communicated qubits or give a universal energy or work cost for joining field regions.

Changing the allowed links or their prices changes the optimum. If the available graph is Γ and its dimensionless edge weights w_e are nonnegative, the same spanning-tree argument gives

$$C_{\tau,\Gamma,w} = \min_{\substack{\mathcal{P}:\,\Gamma[B] \text{ connected} \\ \text{for every } B\in\mathcal{P}}} \sum_{B\in\mathcal{P}} \left[S(\rho_B) + \tau\,\mathrm{MST}_w(\Gamma[B])\right], \tag{B.3}$$

The minimum spanning-tree weight is zero for a singleton. To see the effect of the access graph, consider

$$\rho_{123} = \tfrac{1}{2}(|000\rangle\langle 000| + |101\rangle\langle 101|)$$

at $\tau = 0.4$ nats. Laboratories 1 and 3 share a uniform bit, while laboratory 2 is pure. With uniform complete-graph access, the optimum is $\{1,3\} \mid \{2\}$ at cost $\log 2 + 0.4 < 2\log 2$. The triple costs $\log 2 + 0.8 > 2\log 2$, and either adjacent-pair partition costs $2\log 2 + 0.4$. A nearest-neighbor path excludes the nonconsecutive block and makes singletons optimal. Complete-graph metric weights $w_{12} = w_{23} = 1$, $w_{13} = 2$ also favor singletons, since either the correlated pair or the triple costs $\log 2 + 0.8$. A free relay through laboratory 2, or a charge for one virtual link independent of its physical route, would alter the resource. Thus uniform all-pairs access is a substantive idealization for spatial field regions.

A one-time installation fee has a different asymptotic effect. A fixed cost $\kappa|E(H)|$ contributes $\kappa|E(H)|/m$ to the rate and vanishes in the coding limit. If the available network is connected, the resulting optimum is $C = S(\rho_A)$, with saving $G = \sum_i S(\rho_i) - S(\rho_A)$. A nonzero asymptotic tariff therefore requires an order-m operating or access cost, or a separately priced maintained network resource whose charge is not amortized as a one-time fee.

Communication tasks count other resources. Under free classical communication, quantum state merging has net entanglement cost $S(A|B) = S(AB) - S(B)$, which may be negative [22]. Fully quantum Slepian–Wolf transfer has quantum communication rate $I(A:R)/2$, where R is a purifying reference, and produces the associated entanglement output [23]. Distributed encoding with a common decoder has its own nontrivial rate regions [23, 34]. These results do not determine the state-independent access price τ. The classical residual-entropy Dilworth truncation already has the entropy-partition form: $C_\tau - n\tau$ is the minimum of $\sum_B [H(B) - \tau]$ [1, arXiv v3, eqs. (3.10)–(3.13)]. This antecedent establishes the partition structure; the present access operations and their prices must be specified separately.

The storage task takes τ as an independent input. Its geometric representation follows by matching the ultraviolet action, whose entropy and component coefficients are q and $-c\pi\lambda/3$, to $q[\sum_B S(\rho_B) - \tau|\mathcal{P}|]$. The coefficients agree when

$$q = \sqrt{1-\lambda^2}, \qquad q\tau = \frac{c\pi\lambda}{3}. \tag{B.4}$$

This dictionary applies in the leading-RT ultraviolet regime and allows a geometric evaluation of the independently priced storage task. It neither derives a microscopic link price nor singles out this access model.

Proposition B.1 (Free-control domains and harmless refinement)**.** *Let V be the cells of a finite factorizing regulator and let $\Gamma_0 = (V, E_0)$ specify free bidirectional coherent access. Assume that arbitrary local operations, local ancillary systems, and operations of arbitrary finite depth along E_0 are available, with processing time uncharged. Let $V_1, \dots, V_n$ be the connected components of Γ_0. The operational atoms are*

$$\mathcal{H}_i = \bigotimes_{v \in V_i} \mathcal{H}_v.$$

Charging additional access links between these quotient atoms gives $\tau(n - |\mathcal{P}|)$ for a prescribed partition when all quotient pairs are available at the same nonnegative per-use price τ.

Suppose a finer cell graph has connected fibers over each original cell, represents every original free edge by a free path, and introduces no free path between previously different components. Its free-control components correspond bijectively to the V_i; thus the tariff-bearing atoms are unchanged. If, in addition, its state and atom algebras are identified with the original ones by local isometries and local pure ancillas, then C_τ and G_τ are unchanged. Independent mixed ancillas local to each atom also leave G_τ unchanged.

Proof. Choose a spanning tree in each connected component. Coherent registers can be moved along this tree to one vertex, processed jointly and returned, since time, temporary workspace, traffic and finite circuit depth are unrestricted. Local ancillas and partial traces realize the corresponding channels. Conversely, distinct components have no free edge between them, so the allowed free operations cannot provide coherent access across their cut. The connected components are thus the maximal domains of free joint control. Each added edge reduces the component count by at most one. Joining k domains consequently requires at least $k - 1$ links, and a tree achieves this number. Summing the minimum link counts over the target blocks gives the tariff.

In the finer graph, the connected fibers and the free paths representing original edges connect every refined cell over each V_i. Since no new free path joins different V_i, these are exactly the new components, proving the bijection. Local isometries preserve the nonzero spectra of all block marginals, and pure local ancillas add no entropy. Both optimized values are therefore unchanged. If the added local factors σ_i are independent mixed states, a block entropy gains $\sum_{i \in B} S(\sigma_i)$. Every partition then gains the same amount $\sum_i S(\sigma_i)$, which also appears in separate storage and cancels from G_τ. □

For a regulated one-dimensional region A with nearest-neighbor control and no relay through the complement, the graph components are the maximal connected intervals of A. This identification uses the assumption of free local control; it cannot be inferred from spatial topology alone. Propagation-time constraints, gate costs, restricted controls or access through the complement can alter either the free-control domains or the resource needed to join them.

The proposition holds at a finite regulator. A genuine ultraviolet refinement introduces physical degrees of freedom and need not relate the block states by local isometries. Exact invariance of C_τ or G_τ therefore need not follow, although the graph hypotheses still preserve the component labels. For fixed finitely many separated components, a finite continuum limit of G_τ requires the block relative-entropy limits of section 2. This limit neither constructs finite-time continuum control nor resolves continuum factorization.

Arbitrarily subdividing a laboratory into independently charged atoms does not preserve the positive-price saving. A pure Bell pair within one atom has $G_\tau = 0$; assigning its halves to two separately charged atoms gives $G_\tau = (2\log 2 - \tau)_+$. Such a subdivision violates the state- and algebra-preserving refinement hypotheses of the proposition.

C Entropy-preserving refinement

The reduction rests on component additivity of minimum cuts [10, section 2.4, Lemmas 3–4]. Here we prove it for continuous planar geodesics, allowing a block's entropy-minimizing region to be disconnected. We use the RT entropy (3.1) and the shifted cost $D_\tau = C_\tau - n\tau$ in (A.11); the partition-independent term $n\tau$ is thus removed throughout.

For a block $B = \{i_1 < \cdots < i_k\}$, the connected homology channel consists of the outer arc (u_{i_1}, v_{i_k}) and the gap arcs $(v_{i_r}, u_{i_{r+1}})$. Their logarithmic cost is

$$K(B) = \frac{c}{3}\left[\log\frac{v_{i_k} - u_{i_1}}{\epsilon} + \sum_{r=1}^{k-1}\log\frac{u_{i_{r+1}} - v_{i_r}}{\epsilon}\right]. \tag{C.1}$$

The marginal entropy satisfies $S_{\mathrm{RT}}(A_B) \le K(B)$, with strict inequality whenever a disconnected channel is cheaper. The two channels can also tie at a transition. Replacing every field block by (C.1) would therefore fail to preserve the entropy minimum. The passage to bulk connected components must allow the block itself to split.

For a matching M of the $2n$ endpoints, write $L_{\mathrm{RT}}(M)$ for the logarithmic sum in (3.1) with the factor $c/6$ omitted, and $\chi(M)$ for the number of positive homology components. A geodesic minimum has no compact interior holes, so this count equals its Euler characteristic. Nesting remains possible: a negative gap that reaches the cutoff may contain a separate positive cap.

Theorem C.1 (RT partition reduction)**.** *For every $\tau \ge 0$ and every fixed finite set of distinct interval endpoints,*

$$D_\tau[S_{\mathrm{RT}}] = \min_{M\in\mathrm{NC}(2n)}\left[\frac{c}{6}L_{\mathrm{RT}}(M) - \tau\chi(M)\right]. \tag{C.2}$$

The field minimum over all set partitions has a noncrossing minimizing representative after entropy-preserving component refinement. Equality is a statement about values; optimal partitions need not be unique.

Proof. Begin with finite anchors and replace the logarithmic edge costs by the geodesic lengths

$$\ell_\epsilon(a,b) = 2\,\mathrm{arsinh}\,\frac{|x_b - x_a|}{2\epsilon}. \tag{C.3}$$

First consider entropy-minimizing regions for disjoint label blocks. If two such regions E and F have crossing free geodesics, replace them by $E \setminus F$ and $F \setminus E$. Their disjoint cutoff traces are preserved. Cutting and reassembling the free boundaries preserves the total length but produces corners at the crossings. Rounding these corners strictly reduces that length, contradicting minimality of both original regions for their prescribed traces.

An interior tangency is also excluded. Hyperbolic geodesics with the same tangent at a point coincide, and a complete geodesic meets the horocycle in at most two points. Coincident segments would therefore share anchors, which disjoint label blocks cannot do. The regions have no compact free loops and have distinct cutoff traces; planar separation then rules out overlapping interiors without a boundary crossing. Their interiors are thus disjoint. The argument applies to any choice among tied RT minima and requires no channel uniqueness.

For an arbitrary field partition $\mathcal{P}$, choose a minimizing region for each block. These regions have disjoint interiors, so their union has trace A and free length equal to the sum of the block minima. Decompose the union into positive connected components. Each component minimizes the perimeter for its own cutoff trace: if a shorter competitor existed, its union with the other components of the original block would retain the block trace, while perimeter subadditivity would give a perimeter strictly below the block minimum. It follows that $S_{\mathrm{geo}}(A_B)$ is the sum of the component entropies. This decomposition refines the partition while preserving the entropy sum and gives $\chi \geq |\mathcal{P}|$. For $\tau \geq 0$, it therefore obeys

$$\frac{c}{6} L_{\mathrm{geo}}(M) - \tau\chi(M) \leq \sum_{B \in \mathcal{P}} S_{\mathrm{geo}}(A_B) - \tau|\mathcal{P}|. \tag{C.4}$$

Thus every field partition yields a geometric matching with no greater cost. Minimizing over the two domains gives the first inequality between their optimal values.

For the reverse inequality, take any noncrossing matching. Its $\chi(M)$ positive components define a partition $\mathcal{P}_M$ of the interval labels. Each component boundary is an admissible entropy surface for its block, so minimizing the block entropies can only lower the sum of their lengths:

$$D_\tau[S_{\mathrm{geo}}] \leq \sum_{B \in \mathcal{P}_M} S_{\mathrm{geo}}(A_B) - \tau|\mathcal{P}_M| \leq \frac{c}{6} L_{\mathrm{geo}}(M) - \tau\chi(M). \tag{C.5}$$

Minimizing over M gives the opposite inequality and proves equality at finite anchors. To pass to the logarithmic RT model, use the expansion of (C.3), $2\log(|x_b - x_a|/\epsilon) + O(\epsilon^2/\delta^2)$. For fixed distinct endpoints, this error is uniform over all edges. The partition and matching sets are finite, and each candidate uses n edges, so the error is also uniform under minimization. After subtracting the common logarithmic divergence, the limit $\epsilon \to 0$ gives (C.2). □

The construction obtains the noncrossing representative by refining blocks whose minimizing regions are disconnected. It does not require a connected wedge for each original block. A nontrivial entropy-preserving split strictly decreases the cost when $\tau > 0$ and can tie when $\tau = 0$. The tariff sign determines whether such refinement is favorable.

Noncrossing partitions can contain nonconsecutive labels. For example, $\{1,3\} \mid \{2\}$ joins the first and third intervals while placing a separate cap on the second. That cap lies in a negative gap of the larger region which is connected to the cutoff. Restricting the optimization to consecutive blocks would exclude this configuration, even though it is the strict optimum for the geometry in section 5.

A negative tariff can reverse (C.4). Consider $A_1 = [0,1]$ and $A_2 = [10,11]$ in the logarithmic RT model, with entropies and tariffs measured in units of $c/6$. Subtracting the common singleton entropy makes both the one-block and two-block entropy costs vanish. At $\widehat{\tau} = 6\tau/c = -1$, the field minimum $\widehat{D}_\tau$ is 1, whereas the geometric minimum of $L_{\mathrm{RT}} - \widehat{\tau}\chi$ is 2, since the geometry counts the two disconnected components. The mismatch survives the ultraviolet limit. Thus analytic continuation of an individual pressure arc to negative source does not extend the information correspondence beyond nonnegative source.

D Global arc minimum and its evaluation

Constant-geodesic-curvature arcs solve the hyperbolic magnetic or isoholonomic problem [17]; the Randers geodesic and global-minimizer results of [18] concern the distinct constant-flag-curvature navigation setting. For the present cutoff problem, we derive the directed action and its all-path bound directly, including reverse overhangs. Reflection puts an edge in left-to-right form from $(m-d,\epsilon)$ to $(m+d,\epsilon)$ with $\kappa = \sigma\lambda$. Translate m to zero and set $q = \sqrt{1-\kappa^2}$. For a graph, the action and its translation first integral are

$$\frac{\sqrt{1+z'^2}-\kappa}{z}\,\mathrm{d}x, \qquad \frac{1}{z}\left(\frac{1}{\sqrt{1+z'^2}}-\kappa\right) = \text{constant}. \tag{D.1}$$

On the forward graph segment $|\theta| < \pi/2$ of $x = r\sin\theta$, $z = r(\cos\theta - \kappa)$, the first integral equals $1/r$, so the circle is stationary. Global minimality, including reverse overhangs beyond a single graph, requires the calibration below.

The endpoint condition yields

$$r = r_\kappa = \frac{\sqrt{\epsilon^2 + q^2 d^2} + \kappa\epsilon}{q^2}, \qquad \theta_0 = \arctan\frac{d}{\epsilon} - \arcsin\frac{\kappa d}{\sqrt{d^2+\epsilon^2}}, \tag{D.2}$$

and the entire arc is

$$x = r\sin\theta, \qquad z = r(\cos\theta - \kappa), \qquad -\theta_0 \le \theta \le \theta_0. \tag{D.3}$$

The angle is on the continuous real branch with $0 < \theta_0 < \pi$ and $\cos\theta_0 = \kappa + \epsilon/r$, giving $z > \epsilon$ in the interior. For $\lambda > 0$, a reverse edge has $\kappa = -\lambda$ and overhangs precisely when $\theta_0 > \pi/2$, or equivalently $\lambda d > \epsilon$. Formula (D.3) covers both sides of this threshold.

To construct the global calibration, set $a_0 = qr$ and define on the whole upper half-plane

$$\rho = \operatorname{arsinh}\frac{x^2+z^2-a_0^2}{2a_0 z}, \qquad t = \frac{1}{2}\log\frac{(x+a_0)^2+z^2}{(x-a_0)^2+z^2},$$
$$\psi = \arctan\frac{x+a_0}{z} + \arctan\frac{x-a_0}{z}. \tag{D.4}$$

The denominators are positive for $z > 0$, and the sum of arctangents is a smooth real function. For the pullback, write

$$D_\pm = (x \pm a_0)^2 + z^2, \quad Q = D_+ D_-, \quad U = a_0^2 + z^2 - x^2, \quad V = 2xz. \tag{D.5}$$

Since $Q = U^2 + V^2$, differentiation gives

$$\begin{aligned} \mathrm{d}t &= \frac{2a_0}{Q}(U\,\mathrm{d}x - V\,\mathrm{d}z), \\ \mathrm{d}\psi &= \frac{2z(x^2+z^2+a_0^2)}{Q}\mathrm{d}x - \frac{2x(x^2+z^2-a_0^2)}{Q}\mathrm{d}z, \\ \mathrm{d}\rho &= \frac{V\,\mathrm{d}x + U\,\mathrm{d}z}{z\sqrt{Q}}, \qquad \cosh\rho = \frac{\sqrt{Q}}{2a_0 z}. \end{aligned} \tag{D.6}$$

The $\mathrm{d}z$ coefficient of $\mathrm{d}\psi - \sinh\rho\,\mathrm{d}t$ vanishes, while its $\mathrm{d}x$ coefficient is

$$\frac{2z^2(x^2+z^2+a_0^2) - (x^2+z^2-a_0^2)U}{zQ} = \frac{1}{z}. \tag{D.7}$$

The orthogonal numerators $V\,\mathrm{d}x + U\,\mathrm{d}z$ and $U\,\mathrm{d}x - V\,\mathrm{d}z$, with $U^2 + V^2 = Q$, give the metric identity. Thus

$$\alpha^2 = \mathrm{d}\rho^2 + \cosh^2\rho\,\mathrm{d}t^2, \qquad \beta = \mathrm{d}\psi - \sinh\rho\,\mathrm{d}t. \tag{D.8}$$

The pointwise identity

$$\cosh^2\rho - (q - \kappa\sinh\rho)^2 = (q\sinh\rho + \kappa)^2 \tag{D.9}$$

implies $\alpha \geq (q - \kappa\sinh\rho)\mathrm{d}t$ for either sign of $\mathrm{d}t$. Substitution into (D.8) gives (4.11). Equality on a nonzero tangent requires $\mathrm{d}\rho = 0$, $q\sinh\rho + \kappa = 0$, and $\mathrm{d}t > 0$. The equality locus is the circle (D.3); its cutoff-to-cutoff segment lies in the constrained domain. The full-half-plane bound is therefore sharp for $z \geq \epsilon$, and the equality conditions give uniqueness up to increasing parametrization and pauses.

The action can now be evaluated along this arc. Its length is

$$L = \int_{-\theta_0}^{\theta_0} \frac{\mathrm{d}\theta}{\cos\theta - \kappa} = \frac{4}{q}\,\mathrm{artanh}\left[\sqrt{\frac{1+\kappa}{1-\kappa}}\tan\frac{\theta_0}{2}\right]. \tag{D.10}$$

The artanh argument is below one because the endpoints have positive z. Denote it by y. The endpoint identity

$$\frac{2y}{1-y^2} = \frac{q\sin\theta_0}{\cos\theta_0 - \kappa} = \frac{qd}{\epsilon} \tag{D.11}$$

gives

$$L = \frac{2}{q}\,\mathrm{arsinh}\,\frac{qd}{\epsilon}, \qquad \int_C \beta = \int_{-\theta_0}^{\theta_0} \frac{\cos\theta\,\mathrm{d}\theta}{\cos\theta - \kappa} = 2\theta_0 + \kappa L. \tag{D.12}$$

Hence $\int(\alpha - \kappa\beta) = q^2 L - 2\kappa\theta_0$. Substituting (D.2) gives (4.7); restoring the signed cutoff width gives (4.8). The endpoint data fix the inverse-trigonometric branches, including for reverse overhangs. The result retains the source deformation in both the optimizing arc and its action.

E Exact trace and cutoff-contact elimination

The assignment construction supplies a regular admissible region and hence an upper bound. Its relation to boundary transport follows the planar least-gradient/Beckmann structure [14, 15]. For the present nonreversible cost and exact binary trace, we prove lemma 4.1 by deriving the reverse inequality on the full admissible class, including rough boundaries and additional wall contacts. The path estimate is the calibration of appendix D.

Proof. Extend $\mathbf{1}_E$ by zero below the cutoff and denote its integration two-current by $[E]$, with orientation $\mathrm{d}x \wedge \mathrm{d}z$. Let $A_{\rightarrow}$ be the sum of cutoff interval currents directed from u_i to v_i. The BV trace and Gauss–Green formulas [35, Chapter 3] give, for the prescribed trace,

$$\begin{aligned} T_E &:= A_{\rightarrow} - \partial[E], \qquad \|T_E\|(\Gamma_\epsilon) = 0, \\ \partial T_E &= \sum_{i=1}^{n} (\delta_{(v_i,\epsilon)} - \delta_{(u_i,\epsilon)}). \end{aligned} \tag{E.1}$$

Here T_E is an integral one-current and $\|T_E\|$ is its Euclidean mass measure. The wall part of $\partial[E]$ equals $A_{\rightarrow}$, leaving T_E as the reversed free boundary with zero wall mass. With weighted mass $\mathbf{M}_H(T) = \int z^{-1}\mathrm{d}\|T\|$, Stokes' formula yields

$$\mathbf{M}_H(T_E) = P_H(E; H_\epsilon), \qquad \mathrm{Area}_H(E) = \frac{W}{\epsilon} - T_E(\beta). \tag{E.2}$$

The compact-essential-support condition defining $\mathcal{A}_\epsilon(A)$ makes each competitor's integration current $[E]$, its boundary $\partial[E]$ and T_E compactly supported, independently of null-set representatives. Since $z \geq \epsilon$ on the support, the weighted integrals are finite. Stokes' formula applies with a smooth test cutoff equal to one there; no exhaustion of an unbounded competitor or boundary-at-infinity term enters the argument.

Application of the current decomposition. Apply the decomposition in the complete ambient Euclidean plane. Compact essential support and finite weighted perimeter imply finite Euclidean perimeter, since z is bounded above on the support and bounded below by ϵ. The zero-extension trace formula then makes $[E]$ an integral two-current of finite mass with integral boundary. Subtracting its boundary from the finite interval current gives an integral one-current T_E with finite normal mass $\mathbf{N}(T_E) = \mathbf{M}(T_E) + \mathbf{M}(\partial T_E)$ and boundary mass $2n$ by (E.1).

Theorem 5.3 of [19], valid for integral one-currents in complete metric spaces, decomposes T_E into a countable sum of currents of injective Lipschitz curves or injective Lipschitz loops, with equality of normal masses. Proposition 2.4 of the same source gives equality of the mass measures of both the currents and their boundaries. Thus $\|T_E\| = \sum_j \|T_j\|$ and $\|\partial T_E\| = \sum_j \|\partial T_j\|$. Every summand is supported in $\operatorname{spt} T_E \subset \overline{H}_\epsilon$: a point of a nonconstant curve outside this support would contribute positive mass to a ball of zero total mass.

An injective nonclosed curve has boundary $\delta_{\mathrm{end}} - \delta_{\mathrm{start}}$ and boundary mass two. The total boundary consists of the $2n$ distinct prescribed unit atoms, so measure additivity excludes endpoints elsewhere and multiple curve endpoints at an anchor. Current additivity

fixes the signs: the negative atoms u_i are starts and the positive atoms v_i are ends. There are therefore exactly n open curves C_i, defining a permutation π. Pairing two starts or two ends, or cancelling endpoints at an intermediate point, is excluded. The remaining summands Z_k are loops.

Integrating the bounded nonnegative weight z^{-1} against the mass-measure equality gives $\mathbf{M}_H(T_E) = \sum_j \mathbf{M}_H(T_j)$. To evaluate the one-form, insert a smooth compact test cutoff equal to one near $\operatorname{spt} T_E$ into $\beta = z^{-1}\mathrm{d}x$. Current additivity gives $T_E(\beta) = \sum_j T_j(\beta)$, with absolute convergence following from $|T_j(\beta)| \le \mathbf{M}_H(T_j)$. These consequences for the compactly supported current make no statement about ideal-boundary behavior. The decomposition therefore yields

$$\begin{aligned} P_H(E; H_\epsilon) + \lambda \operatorname{Area}_H(E) &= \frac{\lambda W}{\epsilon} + \sum_{i=1}^{n} \int_{C_i} (\alpha - \lambda\beta) + \sum_k \int_{Z_k} (\alpha - \lambda\beta) \\ &\ge \frac{\lambda W}{\epsilon} + \sum_{i=1}^{n} d_\lambda^\epsilon(u_i, v_{\pi(i)}) + (1-\lambda) \sum_k \ell_H(Z_k) \\ &\ge \frac{\lambda W}{\epsilon} + \min_\pi \sum_{i=1}^{n} d_\lambda^\epsilon(u_i, v_{\pi(i)}). \end{aligned} \tag{E.3}$$

Every decomposed curve lies in $\overline{H}_\epsilon$, with intermediate cutoff contacts allowed. The first inequality uses the all-path calibration for open curves and $|\beta| \le \alpha$ for loops; all series converge absolutely. The bound thus applies to rough exact-trace competitors without a trace-preserving smooth approximation or an assumed regularity theorem for minimizers.

The explicit assignment competitor attains (E.3). Conversely, for any minimizer, equality excludes every nonzero closed curve since $1-\lambda > 0$. Each of the n open curves saturates its endpoint calibration and is therefore the unique circular arc of (D.3), up to increasing parametrization and pauses. The assignment also minimizes the assignment action, so the preceding uncrossing argument makes its arcs disjoint.

The geometric free boundary of a BV region is $\operatorname{spt} T_E$; unlike the topological boundary, it is independent of null-set modifications. On a calibrated arc, $0 < \theta_0 < \pi$ and $\cos\theta > \cos\theta_0$ for $-\theta_0 < \theta < \theta_0$, hence $z > \epsilon$ throughout its interior. At either endpoint the absolute vertical component of its Euclidean unit tangent is $\sin\theta_0 = d/r_\kappa > 0$. Therefore

$$\operatorname{spt} T_E \cap \Gamma_\epsilon = \partial A, \tag{E.4}$$

with transverse endpoint contact, including at $\lambda = 0$ and for reverse arcs with overhangs. The alternating-cap region and E have the same boundary current. Their compactly supported characteristic functions therefore have identical distributional gradients and agree almost everywhere. The regular closed representative is the closure of its open interior, without hanging one-dimensional pieces or zero-width cracks. This classifies all minimizers and their contacts. A region with a genuine additional contact cannot minimize by (E.4); replacement by an assignment minimizer preserves (4.1) and strictly lowers the action. □

Cutoff intervals in A belong to the fixed wall boundary. Appending a wall segment adds only a null decoration and cannot connect geometric components. Genuine accumulation of free boundary on a complementary gap is admissible for competitors and excluded

for minimizers by (E.4). Because the conclusion concerns free-perimeter support, it also excludes genuine tangential or accumulating contacts carrying no separate point-mass length. Tied assignments may give several regular minimizers, all with the same contact property. The replacement is global and trace-preserving; the argument does not supply a local corner-rounding surgery at every contact.

F Uniform ultraviolet expansion

Uniformity over the finite endpoint set is needed before taking the ultraviolet limit of the minimum. Set $x = \epsilon/\Delta$ and restrict λ to a compact subinterval of $[0, 1)$. The functions entering the exact edge action (4.8) have the expansions

$$\begin{aligned} \operatorname{arsinh} \frac{q}{2x} &= \log \frac{q}{x} + \frac{x^2}{q^2} + O(x^4), \\ \arcsin \frac{\lambda}{\sqrt{1+4x^2}} &= \arcsin \lambda - \frac{2\lambda}{q} x^2 + O(x^4), \\ \arctan \frac{1}{2x} &= \frac{\pi}{2} - 2x + \frac{8}{3} x^3 + O(x^5). \end{aligned} \tag{F.1}$$

Inserting the coefficients of (4.8) gives (5.19). The quadratic contributions from H_λ sum to $2(1 - 2\lambda^2)x^2/q$, while the orientation-dependent linear term comes entirely from the arctangent. On a fixed compact subcritical source interval, the Taylor remainder bounds are uniform over the finitely many endpoint assignments.

For the minimum, if $J_M^\epsilon = J_M^{(0),\epsilon} + r_M^\epsilon$ and $|r_M^\epsilon| \le K\epsilon/\delta$ for every assignment M, then

$$\left| \min_M J_M^\epsilon - \min_M J_M^{(0),\epsilon} \right| \le K\epsilon/\delta. \tag{F.2}$$

The constant absorbs fixed n. This gives uniform control of the minimum value, including at ties between leading assignments. A strict leading gap fixes the winning assignment for sufficiently small cutoff, allowing the higher-order expansion term by term. At a degeneracy, the correction must itself be minimized over the tied set.

The area response requires a separate derivative estimate. The two chain-rule terms in H_λ cancel, leaving the exact identity

$$\partial_\lambda f_{\lambda,\epsilon}(u, v) = -\frac{2\lambda}{q} \operatorname{arsinh} \frac{qd}{\epsilon} + 2 \arcsin \frac{\lambda d}{\sqrt{d^2 + \epsilon^2}} + 2\sigma \left(\frac{d}{\epsilon} - \arctan \frac{d}{\epsilon} \right), \qquad d = \frac{\Delta}{2}. \tag{F.3}$$

Applying the same elementary expansions to this identity gives

$$\partial_\lambda f_{\lambda,\epsilon}(u, v) = \frac{v-u}{\epsilon} - \frac{2\lambda}{q} \log \frac{q\Delta}{\epsilon} + 2 \arcsin \lambda - \pi\sigma + \frac{4\sigma\epsilon}{\Delta} + O(\epsilon^2/\Delta^2). \tag{F.4}$$

The remainder is uniform over the finite edge set and a compact subinterval of $[0, 1)$. This establishes C^1 control on each fixed branch independently of the value estimate (F.2).

An optimal matching with a strict leading gap on an open source interval remains optimal on every compact subinterval for sufficiently small cutoff. For this regular pressure

minimum, differentiating the optimized action gives $\partial_\lambda F^\epsilon = V^\epsilon/(4GR)$. Subtract the separately optimized singleton caps in (F.4) and use (4.14). The stable refined field partition is then identified by theorem 3.1, yielding (5.17).

At a transition, first apply the argument in the two adjacent stable phases and then take the one-sided limits. The common baseline is smooth and cancels. Equality of the limiting actions gives the tie relation preceding (5.18), and hence the jump formula in the stated order of limits. These estimates give no derivative control for a subleading actual-state entropy error, a closing phase gap, or the joint limit $\lambda \to 1$ and $\epsilon \to 0$.

The exact finite-anchor geodesic length instead expands as

$$2\,\mathrm{arsinh}\,\frac{\Delta}{2\epsilon} = 2\log\frac{\Delta}{\epsilon} + 2\frac{\epsilon^2}{\Delta^2} + O(\epsilon^4/\Delta^4). \tag{F.5}$$

Subtracting q times this length from the pressure edge action leaves the quadratic coefficient

$$\frac{2(1-2\lambda^2)}{q} - 2q = -\frac{2\lambda^2}{q}. \tag{F.6}$$

The linear coefficient is unchanged, so the finite-anchor entropy convention cannot make the pressure dictionary exact. At zero source the linear term vanishes, as required by the ordinary geodesic limit. The factors of $1/q$ also show why the joint critical-source and ultraviolet limit requires separate control.

Apply the partition theorem with $q\tau = c\pi\lambda/3$ and recall $D_\tau = C_\tau - n\tau$. The absolute normalization is then fixed as follows.

Proposition F.1 (Absolute normalization)**.** *Let A have fixed finite n and distinct nondegenerate interval endpoints, and let F^ϵ_λ be defined by (1.5) on the compact-essential-support class $\mathcal{A}_\epsilon(A)$. In the leading vacuum RT entropy model, for $0 \le \lambda < 1$,*

$$F^\epsilon_\lambda(A) = q\,C^{\mathrm{RT}}_\tau(A) + B^\epsilon_\lambda(A) + O(c\epsilon/\delta), \qquad \tau = \frac{c\pi\lambda}{3q}, \tag{F.7}$$

where the common source and endpoint term is

$$B^\epsilon_\lambda(A) = \frac{c}{6}\left[\frac{\lambda W}{\epsilon} + 2n(q\log q + \lambda\arcsin\lambda) - \pi\lambda n\right]. \tag{F.8}$$

The remainder is uniform on compact subintervals of $[0,1)$ at fixed endpoints and fixed n. If $F^\epsilon_{\lambda,\mathrm{disc}}$ is the sum of the separately pressure-optimized single-interval cap actions, then

$$F^\epsilon_\lambda - F^\epsilon_{\lambda,\mathrm{disc}} = -q\,G^{\mathrm{RT}}_\tau + O(c\epsilon/\delta). \tag{F.9}$$

Combining (4.15) with theorem 3.1 gives the absolute formula. The disconnected caps share the term (F.8) and contribute $q\sum_i S_{\mathrm{RT}}(A_i)$ in entropy, proving (4.17). For $n = 1$, the saving vanishes and (F.7) becomes the single-cap expansion. At $\lambda = 0$, it gives the usual RT minimum.

The passage from the RT entropy model to actual vacuum entropies requires a separate assumption. In the chosen endpoint convention, write

$$S(\rho_B) = S_{\mathrm{RT}}(A_B) + e_B, \qquad E = \max_{\varnothing \neq B \subseteq [n]} |e_B|. \tag{F.10}$$

Each partition has total error at most nE. Comparing the minima therefore gives

$$|C_\tau - C_\tau^{\mathrm{RT}}| \le nE, \qquad |G_\tau - G_\tau^{\mathrm{RT}}| \le 2nE. \tag{F.11}$$

The error in (F.7) thus increases by at most qnE, and that in (4.17) by at most $2qnE$. For the finitely many subsets, $E = o(c)$ suffices for the leading large-c result. Controlled $O(1)$ entropy corrections produce an $O(n)$ value error, separate from the geometric ultraviolet error.

G Oblique-cutoff BV reduction and planar horizon confinement

The thermal cutoff maps to an oblique ray in the hyperbolic plane. To prove theorem 6.1, we minimize on the inward side of this ray with the prescribed BV trace. A global arc calibration reduces the finite-perimeter problem to endpoint assignments. The minimizing regions then lie strictly outside the horizon, so the enlarged domain and the physical exterior have the same minimum.

G.1 The oblique cutoff and charged horizon

In the coordinates $X = x/z_h$ and $z/z_h = \operatorname{sech}\rho$, the thermal metric becomes

$$\frac{\mathrm{d}s^2}{R^2} = \mathrm{d}\rho^2 + \cosh^2\rho\,\mathrm{d}X^2, \qquad 0 \le \rho \le \rho_c = \operatorname{arcosh}(1/e), \qquad e = \epsilon/z_h. \tag{G.1}$$

Writing $s_e = \sqrt{1-e^2}$, the isometry $(U, Z) = (e^X \tanh\rho, e^X \operatorname{sech}\rho)$ sends the cutoff to

$$\Gamma_e = \{a(s_e, e) : a > 0\}, \qquad a = e^{x/z_h}, \tag{G.2}$$

and the horizon to $H_0 = \{U = 0, Z > 0\}$. We first minimize over the hyperbolic domain on the inward side of the cutoff, allowing extension through the horizon:

$$D_e = \{(U,Z) : Z > 0,\ N(U,Z) := -eU + s_e Z > 0\}, \qquad \mathrm{d}s_H^2 = (\mathrm{d}U^2 + \mathrm{d}Z^2)/Z^2. \tag{G.3}$$

The boundary of this domain in $\mathbb{H}^2$ is Γ_e. Let $A_e \subset \Gamma_e$ be the image of the prescribed intervals. The class $\widetilde{\mathcal{A}}_e(A_e)$ consists of positive sets, identified modulo null sets, with $\mathbf{1}_E \in BV(D_e)$, exact interior trace $\mathbf{1}_{A_e}$ on Γ_e, and compact essential support in $\overline{D}_e$ viewed as a subset of $\mathbb{H}^2$. Thus each competitor is contained almost everywhere in some compact $K \subset \mathbb{H}^2$; K may depend on the competitor. Its support stays away from $Z = 0$; Euclidean boundedness would not ensure this near the ray's ideal endpoint.

Restricting to the exterior gives the charged-horizon problem

$$\begin{aligned} \mathcal{A}_e^+(A_e) &= \{E \in \widetilde{\mathcal{A}}_e(A_e) : E \subset \{U \ge 0\} \text{ a.e.}\}, \\ J_{\lambda,e}^+ &= \inf_{E \in \mathcal{A}_e^+(A_e)} \{P_H(E; D_e) + \lambda \operatorname{Area}_H(E)\}. \end{aligned} \tag{G.4}$$

Here $P_H(E; D_e) = \int_{\partial^* E \cap D_e} Z^{-1}\mathrm{d}\mathcal{H}^1$ and $\operatorname{Area}_H(E) = \int_E Z^{-2}\mathrm{d}U\,\mathrm{d}Z$. Taking the perimeter relative to D_e rather than $D_e \cap \{U > 0\}$ charges every reduced-boundary segment produced by zero extension across H_0. No endpoint atom or attachment segment on H_0 is free.

In physical coordinates, compact essential support requires bounded x, while the radial coordinate already satisfies $\epsilon \le z \le z_h$.

Imposing zero exterior BV trace on H_0 defines the smaller class $\mathcal{A}_e^{+,0}$. The minimizer constructed below stays a positive distance from H_0, so the minima over $\widetilde{\mathcal{A}}_e$, $\mathcal{A}_e^{+}$ and $\mathcal{A}_e^{+,0}$ coincide. This equality permits either a charged horizon or zero horizon trace. Allowing nonzero horizon trace without charging its perimeter changes the variational problem.

G.2 All-path minimum at the oblique wall

The hyperbolic length and area primitive are globally defined by

$$\alpha = \frac{\sqrt{\mathrm{d}U^2 + \mathrm{d}Z^2}}{Z}, \qquad b = \frac{\mathrm{d}U}{Z}, \qquad \mathrm{d}b = \frac{\mathrm{d}U \wedge \mathrm{d}Z}{Z^2}, \qquad |b| \le \alpha. \tag{G.5}$$

In the Fermi coordinates of (G.1), $b = \sinh\rho\,\mathrm{d}X + \mathrm{d}\psi$ with $\psi = \arcsin(\tanh\rho)$. The gauge term has equal endpoint values on the cutoff ray. Hence

$$\int_{A_{e,\to}} b = \frac{s_e W}{\epsilon}, \qquad W = \sum_i (v_i - u_i). \tag{G.6}$$

We orient the wall current $A_{e,\to}$ toward increasing a, placing D_e on its left.

For distinct physical endpoints in the order u, v, put $t = |v-u|/z_h$, $\sigma = \mathrm{sgn}(v-u)$, and $h = t/2$. The directed action minimized over all paths in $\overline{D}_e$ is

$$d_{\lambda,e}(u,v) = \inf_{C: p(u)\to p(v),\, C\subset\overline{D}_e} \int_C (\alpha - \lambda b), \qquad p(x) = e^{x/z_h}(s_e, e). \tag{G.7}$$

For $e < q = \sqrt{1-\lambda^2}$, we obtain

$$d_{\lambda,e}(u,v) = 2q\,\mathrm{arsinh}\,\frac{q\sinh h}{e} + 2\lambda\arcsin\frac{\lambda\sinh h}{\sqrt{\sinh^2 h + e^2}} - 2\sigma\lambda\arctan\frac{s_e\tanh h}{e}. \tag{G.8}$$

To derive this expression, set $a = e^{\min(u,v)/z_h}$, $b_0 = e^{\max(u,v)/z_h}$, $m = (a+b_0)/2$, $d = (b_0-a)/2$, and $\kappa = \sigma\lambda$. The constant-curvature circle through the anchors has radius r and center $(H, -\kappa r)$, with

$$r = \frac{\kappa e m + s_e\sqrt{q^2d^2 + e^2 a b_0}}{q^2 - e^2}, \qquad H = \frac{m + \kappa e r}{s_e}, \qquad H^2 - q^2 r^2 = a b_0. \tag{G.9}$$

When $e < q$, its radius is the unique positive root of $(q^2-e^2)r^2 - 2\kappa e m r - (d^2 + a b_0 e^2) = 0$. The center satisfies $H > qr > 0$: H starts positive at $e = 0$, remains continuous for $e < q$, and cannot vanish because of the last identity in (G.9).

After a horizontal translation, the Fermi calibration of appendix D applies with supporting geodesic radius qr. Its bound holds for all paths in $\mathbb{H}^2$, without any restriction on the anchor locations. Equality selects the constant-geodesic-curvature circle with the prescribed orientation, uniquely up to increasing parametrization and pauses. Reflection covers the reverse orientation by sending b to $-b$. It remains to locate the calibrated arc relative to the cutoff.

The circle meets the cutoff line at $a(s_e, e)$ and $b_0(s_e, e)$. Both ideal endpoints $H \pm qr$ are positive, giving $N < 0$ at either end of the full upper-half-plane arc. The segment between the cutoff intersections is therefore its unique connected part with $N \geq 0$, and satisfies $N > 0$ in the interior. At either endpoint, the absolute normal component of the Euclidean unit tangent is $d/r > 0$, since the cutoff chord has length $2d$. The calibrated arc is thus admissible, with exactly two transverse wall contacts. The argument requires no regularity assumption on competing paths.

The arc has Euclidean angular span and hyperbolic length

$$\begin{aligned} \Theta_\sigma &= 2\arctan\frac{s_e\tanh h}{e} - 2\sigma\arcsin\frac{\lambda\sinh h}{\sqrt{\sinh^2 h + e^2}} = 2\arccos\frac{\kappa + em/r}{s_e}, \\ \ell_{\lambda,e} &= \frac{2}{q}\operatorname{arsinh}\frac{q\sinh h}{e}. \end{aligned} \tag{G.10}$$

The integral of $\alpha - \lambda b$ is $q^2\ell_{\lambda,e} - \sigma\lambda\Theta_\sigma$, which yields (G.8). This is the optimized pressure action; evaluating area on an RT curve would give a different quantity. At $\lambda = 0$, the result is the finite-anchor geodesic length $2\operatorname{arsinh}(\sinh h/e)$.

G.3 Finite-perimeter reduction

For a competitor in $\widetilde{\mathcal{A}}_e$, extend its characteristic function by zero outside D_e in $\mathbb{H}^2$. Compact essential support defines a compactly supported integral two-current $[E]$ away from $Z = 0$, even when the free boundary is irregular or has additional wall contacts. The BV trace and Gauss–Green formula on the straight Lipschitz wall give [35, Chapter 3]

$$\begin{aligned} T_E &= A_{e,\to} - \partial[E], & \|T_E\|(\Gamma_e) &= 0, \\ \partial T_E &= \sum_{i=1}^{n}(\delta_{p(v_i)} - \delta_{p(u_i)}), & \mathbf{M}_H(T_E) &= P_H(E; D_e), \\ \operatorname{Area}_H(E) &= \frac{s_e W}{\epsilon} - T_E(b), & \mathbf{M}_H(T) &= \int Z^{-1}\mathrm{d}\|T\|. \end{aligned} \tag{G.11}$$

A smooth test cutoff equal to one on the compact support justifies Stokes' theorem. Any horizon segment of a physical competitor lies inside D_e and contributes to T_E, without adding boundary atoms.

Decompose this current in the complete Euclidean (U, Z) plane as in appendix E. On its compact support away from $Z = 0$, Euclidean and hyperbolic masses are comparable, and T_E has finite normal mass. Theorem 5.3 and Proposition 2.4 of [19] give equality of both current and boundary mass measures under decomposition. The distinct unit endpoint atoms yield exactly n source-to-sink curves, with no endpoint cancellation. Integration against the bounded weight Z^{-1} preserves mass additivity, while $|b| \leq \alpha$ ensures absolute convergence and additivity of the line integrals. Each curve lies in $\operatorname{spt} T_E \subset \overline{D}_e$.

Consequently,

$$\begin{aligned} P_H(E; D_e) + \lambda \operatorname{Area}_H(E) &= \frac{\lambda s_e W}{\epsilon} + \sum_{i=1}^{n} \int_{C_i} (\alpha - \lambda b) + \sum_k \int_{Z_k} (\alpha - \lambda b) \\ &\geq \frac{\lambda s_e W}{\epsilon} + \sum_{i=1}^{n} d_{\lambda,e}(u_i, v_{\pi(i)}) + (1-\lambda) \sum_k \ell_H(Z_k) \\ &\geq \frac{\lambda s_e W}{\epsilon} + \min_{\pi} \sum_i d_{\lambda,e}(u_i, v_{\pi(i)}). \end{aligned} \tag{G.12}$$

The open curves remain in $\overline{D}_e$, where the all-path calibration applies even if they touch the cutoff between their endpoints. Absolute convergence justifies both sums. No trace-preserving smooth approximation or contact regularity theorem enters this bound.

Take a minimizing assignment and its calibrated arcs. At an intersection of two arcs, exchanging their outgoing tails preserves the total directed action. Minimality then forces each exchanged path to saturate its endpoint bound. Distinct oriented tangents would produce a corner or reversing cusp, which violates the calibration equality conditions. Equal oriented tangents instead determine the same constant-curvature arc by uniqueness of the initial-value problem. Such a circle has only its two prescribed cutoff intersections, so this case would require duplicate anchors. The minimizing arcs are therefore disjoint.

Closing each arc with its cutoff-ray chord gives disjoint or nested caps. Along the ray, start and end anchors alternate. Their directions therefore alternate with nesting depth: outermost arcs are forward, the next are reverse gap arcs, and the pattern continues inward. Points covered by an odd number of caps form a compact positive region with exact trace A_e and boundary current $A_{e,\to} - \sum_i C_i$. This region attains the lower bound in (G.12), proving existence and the assignment formula

$$\widetilde{J}_{\lambda,e} = \frac{\lambda s_e W}{\epsilon} + \min_{\pi} \sum_i d_{\lambda,e}(u_i, v_{\pi(i)}). \tag{G.13}$$

Conversely, equality in (G.12) eliminates every nonzero closed curve and makes all n open curves the calibrated arcs of a minimizing assignment. Their free-boundary support meets Γ_e transversely and only at the prescribed anchors. The original minimizer and the alternating-cap region have the same integration-current boundary. Their compactly supported characteristic functions thus have equal distributional gradients and agree almost everywhere. This characterizes every minimizer modulo null sets. In particular, it excludes additional tangential contacts and accumulations of contacts, including those beyond the reach of local corner rounding. Null filaments and cracks do not affect the result.

G.4 Horizon confinement

An outermost arc has forward orientation and circle-center height $-\lambda r \leq 0$. Its circle is confined, for $Z > 0$, to

$$H - qr \leq U \leq H + qr, \qquad H - qr > 0. \tag{G.14}$$

Together with its positive-U cutoff chord, this arc bounds a closed cap contained in $U > 0$. All nested caps and free arcs lie inside the outermost caps. Thus every free arc of a minimizing region, including each reverse gap arc, stays outside H_0. Compactness and finite n imply a positive distance from each minimizing representative to H_0.

The minimizer on the enlarged domain therefore belongs to $\mathcal{A}_e^{+,0}$. Since $\mathcal{A}_e^{+,0} \subset \mathcal{A}_e^{+} \subset \widetilde{\mathcal{A}}_e$, the three infima equal (G.13), and every physical minimizer has the contact and confinement properties proved above. An independently minimized reverse arc can cross the horizon: in the UV, a gap g does so when $\lambda > \operatorname{sech}(g/(2z_h))$. Such an arc cannot occur in a globally minimizing positive assignment. Individual channel minima therefore need not be admissible exterior minima.

G.5 Ultraviolet information partition and finite-cutoff residual

At fixed nonzero endpoint separation, (G.8) expands as

$$d_{\lambda,e}(u,v) = 2q \log\left[\frac{2z_h}{\epsilon}\sinh\frac{|v-u|}{2z_h}\right] + 2(q\log q + \lambda \arcsin\lambda) - \sigma\pi\lambda \\ + 2\sigma\lambda\frac{\epsilon}{z_h}\coth\frac{|v-u|}{2z_h} + O((\epsilon/z_h)^2). \tag{G.15}$$

For fixed separated endpoints, there are finitely many assignments. The expansion is uniform over these assignments and over sources in compact subintervals of $[0,1)$. It therefore passes through minimization at the level of values, including at degeneracies of the leading term. Each exact minimizing assignment is noncrossing, with $\sum\sigma = 2\chi - n$ determined by its positive components.

To relate the result to the entropy partition, map the boundary endpoints monotonically to $a = e^{x/z_h}$. A thermal logarithmic edge is the vacuum edge between a, b_0 plus the local shift $-\log a - \log b_0$ and a common cutoff convention. Since every matching and every full information partition uses each endpoint once, these shifts leave theorem 3.1 unchanged. All set partitions remain allowed. Applying the theorem with $q\tau = c\pi\lambda/3$ to (G.13) gives

$$F^{\epsilon}_{\lambda,\beta_T}(A) = qC^{\mathrm{RT},\beta_T}_{\tau}(A) + \frac{c}{6}\left[\frac{\lambda s_e W}{\epsilon} + 2n(q\log q + \lambda\arcsin\lambda) - \pi\lambda n\right] + O(c\epsilon/z_h). \tag{G.16}$$

The separately minimized singleton caps have the same additive baseline. Subtracting their value proves (6.3). For an interval of length ℓ, the finite-cutoff minimum at $e < q$ is

$$F^{\epsilon}_{\lambda,\beta_T}([u,u+\ell]) = \frac{c}{6}\left[\frac{\lambda s_e \ell}{\epsilon} + d_{\lambda,e}(u,u+\ell)\right]. \tag{G.17}$$

For a stable minimizing matching, the first correction to (G.16) is

$$\frac{c}{6}\,2\lambda\frac{\epsilon}{z_h}\sum_{\text{edges}}\sigma\coth\frac{|v-u|}{2z_h}. \tag{G.18}$$

Subtracting the corresponding singleton sum gives a generally nonzero correction to the relative minimum. Near ties, these corrections must also be minimized over the tied branches;

the uniform estimate for the value is unchanged. As in the vacuum theorem, the entropy correspondence holds at leading UV order and does not give an exact finite-cutoff identity.

The arc and BV results require $e < q$ at finite cutoff. This holds for sufficiently small cutoff at every fixed $\lambda < 1$, without implying uniformity in a simultaneous $\lambda \to 1$, shrinking-gap or divergent-interval limit. Replacing leading RT entropies by the true thermal CFT entropies is governed separately by the finite-partition error bound (F.11); the geometric cutoff argument is unchanged.

H Global minimization on the compact cylinder

On the compact cylinder, a global minimum must also be compared across winding sectors. The RT region bounded by complement geodesics and the horizon is an annulus [20], whose Euler characteristic differs from its positive component count. We prove that an annular pressure region minimizes globally in the example below, and that its value disagrees with the unchanged information tariff.

Set $R = z_h = 1$ and normalize the observables as $J = 6F/c$, $\widehat{S} = 6S/c$, $\widehat{I} = 6I/c$, $\widehat{G} = 6G/c$ and $\widehat{\tau} = 6\tau/c$. This removes an overall scale without fixing the physical central charge. Impose $x \sim x + H$ with $H = 10$, and choose

$$A_1 = [0, 4.9], \qquad A_2 = [5, 9.9], \qquad \lambda = 0.1, \quad q = \sqrt{0.99}. \tag{H.1}$$

The two complementary gaps both have length $g = 0.1$. The exterior cylinder is $0 \le \rho \le \rho_c$, with metric $\mathrm{d}\rho^2 + \cosh^2\rho\,\mathrm{d}x^2$ and cutoff parameter $e = \operatorname{sech}\rho_c$. Positive competitors have exact BV trace A at the outer cutoff. All remaining boundary segments, including horizon segments, contribute to the perimeter. Equivalently, extend the characteristic function by zero through $\rho = 0$ before computing free perimeter. An essential inner boundary is allowed to move through the exterior. It has no free endpoint on the horizon and is not pinned there.

Denote the calibrated directed thermal edge action, excluding the fixed-trace term, by $d^e_\lambda(t, \sigma)$ for $t > 0$ and $\sigma = \pm 1$:

$$\begin{aligned} d^e_\lambda(t,\sigma) = {} & 2q \operatorname{arsinh}\frac{q\sinh(t/2)}{e} + 2\lambda \arcsin\frac{\lambda\sinh(t/2)}{\sqrt{\sinh^2(t/2)+e^2}} \\ & - 2\sigma\lambda \arctan\frac{\sqrt{1-e^2}\tanh(t/2)}{e}, \qquad e < q. \end{aligned} \tag{H.2}$$

We suppress the common fixed-trace contribution $\lambda \sinh\rho_c\,|A|$ in all channel comparisons below.

The global primitive $b_{\mathrm{cyl}} = \sinh\rho\,\mathrm{d}x$ satisfies $|b_{\mathrm{cyl}}| \le \alpha$, with $\mathrm{d}b_{\mathrm{cyl}}$ the area form. For any oriented closed curve of winding w,

$$\int (\alpha - \lambda b_{\mathrm{cyl}}) \ge qH|w|. \tag{H.3}$$

This follows from $\alpha \ge \cosh\rho\,|\mathrm{d}x|$ and $\cosh\rho - \lambda\sinh\rho \ge q$; reversing orientation can only increase the lower bound. Equality is attained by a positively winding circle at

$$\rho_\lambda = \operatorname{artanh}\lambda, \qquad H(\cosh\rho_\lambda - \lambda\sinh\rho_\lambda) = qH. \tag{H.4}$$

This circle lies below the cutoff if and only if $e < q$.

To apply current decomposition, work on the complete doubled cylinder $\rho \in \mathbb{R}$, $x \in \mathbb{R}/H\mathbb{Z}$, with metric $\mathrm{d}\rho^2 + \cosh^2\rho\,\mathrm{d}x^2$. Zero extension through both the horizon and outer cutoff makes $[E]$ a compactly supported integral two-current. Subtracting its boundary from the prescribed wall current gives an integral one-current of finite normal mass with precisely the $2n$ endpoint atoms. Theorem 5.3 and Proposition 2.4 of [19] apply in this complete metric space as in appendix E. The ambient mass is hyperbolic length, so equality of mass measures directly gives length additivity. Since the global primitive b_{cyl} has norm at most one, its line integrals also form an absolutely convergent series.

Choose a lift of each initial anchor. The corresponding open path then lifts uniquely to a path ending at $v_{\pi(i)} + k_i H$ for some integer k_i. Decomposition is performed on the finite-mass cylinder current, before lifting; no periodically repeated infinite-mass current is required. By Stokes' theorem the closed form $\mathrm{d}x$ integrates to zero on $\partial[E]$. Subtracting the prescribed wall contribution gives total cycle winding $-\sum_i k_i$. This sum converges because each loop has hyperbolic length at least $H|w|$ and the total length is finite. Applying the all-path calibration to the open paths and (H.3) to the cycles gives

$$\sum_i d^e_\lambda\Big(|v_{\pi(i)} + k_i H - u_i|, \operatorname{sgn}(v_{\pi(i)} + k_i H - u_i)\Big) + qH\left|\sum_i k_i\right|. \tag{H.5}$$

The bound is a relaxation of the region problem. It remains valid for sums of curves that intersect or cannot bound an admissible region; these sums are not claimed to be geometric saddles.

All integer lifts can be reduced to a finite comparison. Differentiating (H.2) with $h = t/2$ and $s_e = \sqrt{1-e^2}$ gives

$$\partial_t d^e_\lambda(t,\sigma) = \frac{\cosh h\sqrt{e^2 + q^2\sinh^2 h} - \sigma\lambda s_e e}{\sinh^2 h + e^2} > q, \qquad e < q. \tag{H.6}$$

To prove the strict inequality, subtract the square of $q(\sinh^2 h + e^2) + \sigma\lambda s_e e$ from the square of the first numerator term. The difference is $(\sinh^2 h + e^2)(qs_e - \sigma\lambda e)^2 > 0$. Whenever an edge has absolute lift displacement greater than H, removing one period lowers its action by more than qH and changes the cycle penalty by at most qH. Repeating this step leaves only the two lifts of each edge with absolute displacement below H.

There are two endpoint assignments for two intervals, each with four remaining lift combinations. Symmetry reduces their eight lower bounds to

$$D = 2d_+(4.9), \qquad C = d_+(9.9) + d_-(0.1), \qquad B = qH + 2d_-(0.1), \tag{H.7}$$

$$N = qH + d_+(4.9) + d_-(5.1), \qquad N_2 = 2qH + 2d_-(5.1), \qquad X = qH + 2d_+(9.9), \tag{H.8}$$

where C and N each occur twice, and $d_\pm(t) = d^e_\lambda(t, \pm 1)$. The first three expressions correspond to two-disk, connected-disk and annular candidates. The remaining expressions are lower bounds irrespective of mutual curve admissibility. Monotonicity together with $d_-(t) > d_+(t)$ gives $N > D$ and $N_2 > N$. Moreover, $d_+(9.9) - d_-(0.1) > 9.8q - 2\pi\lambda > 0$,

which implies $X > B$ for (H.1). The two remaining comparisons have positive UV limits:

$$\lim_{e\to 0}(D-B) = q\left[4\log\frac{\sinh 2.45}{\sinh 0.05} - 10\right] - 4\pi\lambda = 7.6771323139\ldots,$$
$$\lim_{e\to 0}(C-B) = 2q\log\frac{\sinh 4.95}{\sinh 0.05} - 10q - 2\pi\lambda = 3.8533402704\ldots. \tag{H.9}$$

Their signs can also be fixed without the decimal evaluations. Elementary exponential estimates give $\widehat{I}_{\mathrm{RT}} > 8$ and $2\log[\sinh 4.95/\sinh 0.05] - 10 > 4$. Together with $q > 0.99$, these bound the gaps below by 6.6 and 3.3, respectively. Continuity then preserves both inequalities for all sufficiently small e.

We next establish attainment of the lowest bound and minimize the singleton reference independently. The annular value B is realized by two reverse gap arcs and the optimized essential loop. These curves are disjoint: in the UV the gap arcs have minimum radial coordinate $3.5887524093\ldots$, while $\rho_\lambda = 0.1003353477\ldots$. Their separation follows directly from the circle geometry. Normalize the lifted endpoints of either gap to $\exp(\pm g/2)$ and write the center as $(h, -\kappa r)$, with $\kappa = -\lambda$ and $h^2 - q^2r^2 = 1$. The full circle obeys $|x - x_{\mathrm{centre}}| \le \operatorname{arsinh} r$ under $x = \log\sqrt{U^2+Z^2}$. In the UV, $\operatorname{arsinh} r = 0.0502516795\ldots$, whereas adjacent gap centers are separated by 5. The smallest ρ is determined by $h\sinh\rho - \kappa r = \cosh\rho$, yielding the radial separation quoted above. Both separations persist for sufficiently small cutoff. Removing the two gap caps and the inner collar leaves a positive annulus with the prescribed trace, saturating each lower bound contributing to B.

For a singleton, the one-path relaxation leaves $d_+(4.9)$ and $qH + d_-(5.1)$, with the latter strictly larger. The direct forward cap realizes the first value. In a centered lift it satisfies $\cosh x = h\tanh\rho - \lambda r\operatorname{sech}\rho$. The right-hand side increases strictly with $\rho \ge 0$, so the cap stays within its boundary interval in x. A width below H then ensures that the cap is embedded on the quotient. Thus D is the sum of independently minimized singleton values. At $\lambda = 0$, the same argument establishes the RT phases used below. In each case, attainment follows from an admissible region in addition to the sum of individual curve minima.

The UV comparisons and attainment are uniform for $0 \le \lambda \le 0.4$. Indeed, $q \ge \sqrt{0.84}$ and the coarse bounds above keep both gaps in (H.9) positive. The reverse-gap minimum is exactly $\operatorname{arsinh}[\operatorname{csch}(0.05)] - \operatorname{artanh}\lambda$, which remains above the loop radius $\operatorname{artanh}\lambda$, and its coordinate half-extent is at most $\operatorname{arsinh}[\sinh(0.05)/\sqrt{0.84}] < 0.055$. These bounds give a uniform UV source range around the example. Figure 6 shows the $\lambda = 0.1$ representative.

The global minimization also determines the entropy comparison. Each singleton uses its direct cap, whereas the union uses the two complementary gaps and the horizon. Hence

$$\widehat{I}_{\mathrm{RT}} = 4\log\frac{\sinh 2.45}{\sinh 0.05} - 10 = 8.9787760550\ldots > \widehat{\tau}, \qquad \widehat{\tau} = \frac{2\pi\lambda}{q}. \tag{H.10}$$

The intrinsic two-atom saving is still $\widehat{G}_\tau = (\widehat{I}_{\mathrm{RT}} - \widehat{\tau})_+$. Comparing it with the pressure minimum from (H.8), we find

$$\lim_{e\to 0}(J_\lambda - J_{\lambda,\mathrm{disc}}) = -q\widehat{I}_{\mathrm{RT}} + 4\pi\lambda = -7.6771323139\ldots,$$
$$-q\widehat{G}_\tau = -q\widehat{I}_{\mathrm{RT}} + 2\pi\lambda = -8.3054508446\ldots. \tag{H.11}$$

The UV discrepancy is $2\pi\lambda$, or $c\pi\lambda/3$ for general c in this geometry. Common local endpoint and fixed-trace counterterms cancel in the comparison and cannot remove it.

Gauss–Bonnet identifies the origin of this discrepancy. At zero source, for n fixed intervals with geodesic free boundaries,

$$\frac{V}{R^2} = \frac{|A|}{z_h}\sinh\rho_c + n\pi - 2\pi\chi(E) + o(1). \tag{H.12}$$

The two disks have $\chi = 2$ and the annulus has $\chi = 0$. Their area difference is $4\pi R^2$, while one information merger incurs a single tariff with normalized source-zero slope 2π. The first source derivatives therefore already disagree. Fixing the inner loop at the horizon changes the finite-source problem and leaves this first-order discrepancy intact. The planar theorem relies on equality of Euler characteristic and positive component count. This compact-BTZ regime violates that equality and hence the representation with the unchanged tariff; the example establishes no general compact-space representation.

I Numerical methods and checks

The numerical tests compare separate implementations of the partition problem, direct curve integrals and finite-lattice entropy calculations. They test the formulas on finite samples; the global correspondence follows from the analytic arguments above. Code, input parameters and instructions for reproducing the calculations and figures accompany the supplementary material.

Figures 3 and 4 show complete partition envelopes, including losing branches. Figure 6 combines pressure-minimizing curves, a topology schematic and the compact residual. The coordinate curves use the exact optimized arcs, while schematic topology panels carry no metric radial scale. This distinction permits separate checks of the geometric construction and of its numerical action.

The first partition calculation compares four envelopes on 3310 endpoint geometries at 37044 tariff probes. The sample contains 243 three-component and 2187 four-component separation-grid configurations, followed by 300, 300, 200 and 80 reproducible random configurations at $n = 3, 4, 5, 6$. Grid separations are chosen from $\{0.125, 1, 4\}$, and random separations are log-uniform in $[10^{-2}, 10^{2}]$. An independently implemented calculation tests 141 geometries at $n = 1, \ldots, 7$, with six tariffs and both boundary-logarithmic and finite-anchor geodesic costs, giving 1692 comparisons. The respective numerical assertion tolerances are 10^{-8} and 10^{-10}. These are implementation checks of theorem 3.1, not uniform error bounds over geometries or computing platforms.

The free-Dirac comparison starts from the known continuum union entropy

$$S_D(A_B) = \frac{c}{3}\left[\sum_{i,j\in B}\log|v_j - u_i| - \sum_{\substack{i<j\\ i,j\in B}}\log[(u_j - u_i)(v_j - v_i)] - |B|\log\epsilon\right], \tag{I.1}$$

up to component-additive regulator constants [28, 29]. Diagonal terms give the single-interval entropies; pairing the off-diagonal terms gives (5.11). We evaluate all 31 nonempty

subsets of an asymmetric five-interval geometry at 70-digit precision, compare the five RT channels at seven ratios, and compute finite hopping-chain union entropies in the CAR convention for widths $40, 80, 160, 320$ sites. The finite-lattice value of I_3 need not vanish. Its cutoff dependence tests approach to the continuum formula without establishing a convergence rate or a finite-c holographic correction.

For pressure, direct circle integrals are compared with the optimized edge actions. The collective and nested envelopes are evaluated in full, and both residual coefficients are tested on the 80-cutoff grid specified in figure 5. The area response is evaluated on the two sides of the actual finite-cutoff crossing. Thermal tests cover 200 random finite-component geometries and 96 near-wall cases. The compact calculation checks the winding-sector lower bounds and separation of the gap arcs from the optimized essential loop. These samples test distinct steps of the construction and carry no common maximum-error estimate; calibration and current decomposition provide the corresponding global statements.

References

[1] C. Chan, A. Al-Bashabsheh, Q. Zhou, T. Kaced and T. Liu, *Info-clustering: A mathematical theory for data clustering*, *IEEE Trans. Mol. Biol. Multi-Scale Commun.* **2** (2016) 64 [1605.01233].

[2] F.D.J. Dunstan, *Matroids and submodular functions*, *Q. J. Math.* **27** (1976) 339.

[3] S. Fujishige, *Polymatroidal dependence structure of a set of random variables*, *Information and Control* **39** (1978) 55.

[4] H. Narayanan, *The principal lattice of partitions of a submodular function*, *Linear Algebra and its Applications* **144** (1991) 179.

[5] E.H. Lieb and M.B. Ruskai, *Proof of the strong subadditivity of quantum-mechanical entropy*, *J. Math. Phys.* **14** (1973) 1938.

[6] B. Schumacher, *Quantum coding*, *Phys. Rev. A* **51** (1995) 2738.

[7] H. Barnum, C.A. Fuchs, R. Jozsa and B. Schumacher, *A general fidelity limit for quantum channels*, *Phys. Rev. A* **54** (1996) 4707 [quant-ph/9603014].

[8] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [hep-th/0603001].

[9] M. Headrick, *Entanglement Rényi entropies in holographic theories*, *Phys. Rev. D* **82** (2010) 126010 [1006.0047].

[10] S. Hernández-Cuenca, V.E. Hubeny and M. Rota, *The holographic entropy cone from marginal independence*, *JHEP* **09** (2022) 190 [2204.00075].

[11] R. Abt, J. Erdmenger, H. Hinrichsen, C.M. Melby-Thompson, R. Meyer, C. Northe et al., *Topological complexity in AdS_3/CFT_2*, *Fortsch. Phys.* **66** (2018) 1800034 [1710.01327].

[12] R. Abt, J. Erdmenger, M. Gerbershagen, C.M. Melby-Thompson and C. Northe, *Holographic subregion complexity from kinematic space*, *JHEP* **01** (2019) 012 [1805.10298].

[13] B. Czech, P. Hayden, N. Lashkari and B. Swingle, *The information theoretic interpretation of the length of a curve*, *JHEP* **06** (2015) 157 [1410.1540].

[14] S. Dweik and F. Santambrogio, *L^p bounds for boundary-to-boundary transport densities, and $W^{1,p}$ bounds for the BV least gradient problem in 2D*, *Calc. Var. Partial Differential Equations* **58** (2019) 31 [1805.00769].

[15] S. Dweik and W. Górny, *Optimal transport approach to sobolev regularity of solutions to the weighted least gradient problem*, *SIAM J. Math. Anal.* **55** (2023) 1916 [2112.13920].

[16] S. Dweik, "Least gradient problem with nonhomogeneous term." CVGMT preprint, manuscript dated 20 August 2026, retrieved 16 October 2026, CVGMT paper 7433.

[17] V.N. Berestovskii and I.A. Zubareva, *Sub-Riemannian distance on the Lie group $SO_0(2,1)$*, *St. Petersburg Math. J.* **28** (2017) 477 [1507.05554].

[18] C. Robles, *Geodesics in Randers spaces of constant curvature*, *Trans. Amer. Math. Soc.* **359** (2007) 1633 [math/0501358].

[19] P. Bonicatto, G. Del Nin and E. Pasqualetto, *Decomposition of integral metric currents*, *J. Funct. Anal.* **282** (2022) 109378 [2104.07593].

[20] V.E. Hubeny, H. Maxfield, M. Rangamani and E. Tonni, *Holographic entanglement plateaux*, *JHEP* **08** (2013) 092 [1306.4004].

[21] K. Modi, T. Paterek, W. Son, V. Vedral and M. Williamson, *Unified view of quantum and classical correlations*, *Phys. Rev. Lett.* **104** (2010) 080501 [0911.5417].

[22] M. Horodecki, J. Oppenheim and A. Winter, *Quantum state merging and negative information*, *Commun. Math. Phys.* **269** (2007) 107 [quant-ph/0512247].

[23] A. Abeyesinghe, I. Devetak, P. Hayden and A. Winter, *The mother of all protocols: Restructuring quantum information's family tree*, *Proc. R. Soc. A* **465** (2009) 2537 [quant-ph/0606225].

[24] R.L. Jerrard, A. Moradifam and A.I. Nachman, *Existence and uniqueness of minimizers of general least gradient problems*, *J. Reine Angew. Math.* **734** (2018) 71 [1305.0535].

[25] M. Headrick, A. Lawrence and M.M. Roberts, *Bose-fermi duality and entanglement entropies*, *J. Stat. Mech.* **2013** (2013) P02022 [1209.2428].

[26] A. Coser, E. Tonni and P. Calabrese, *Spin structures and entanglement of two disjoint intervals in conformal field theories*, *J. Stat. Mech.* **2016** (2016) 053109 [1511.08328].

[27] P. Perinotti, A. Tosini and L. Vaglini, *Shannon theory for quantum systems and beyond: Information compression for fermions*, in *The Quantum-Like Revolution*, A. Plotnitsky and E. Haven, eds., (Cham), Springer (2023), DOI [2106.04964].

[28] H. Casini and M. Huerta, *Remarks on the entanglement entropy for disconnected regions*, *JHEP* **03** (2009) 048 [0812.1773].

[29] H. Casini and M. Huerta, *Reduced density matrix and internal dynamics for multicomponent regions*, *Class. Quant. Grav.* **26** (2009) 185005 [0903.5284].

[30] P. Hayden, M. Headrick and A. Maloney, *Holographic mutual information is monogamous*, *Phys. Rev. D* **87** (2013) 046003 [1107.2940].

[31] S. Ryu and T. Takayanagi, *Aspects of holographic entanglement entropy*, *JHEP* **08** (2006) 045 [hep-th/0605073].

[32] P. Calabrese and J. Cardy, *Entanglement entropy and quantum field theory*, *J. Stat. Mech.* **2004** (2004) P06002 [hep-th/0405152].

[33] B. Schumacher, *Sending quantum entanglement through noisy channels*, quant-ph/9604023.

[34] C. Ahn, A.C. Doherty, P. Hayden and A.J. Winter, *On the distributed compression of quantum information*, *IEEE Trans. Inf. Theory* **52** (2006) 4349 [quant-ph/0403042].

[35] L. Ambrosio, N. Fusco and D. Pallara, *Functions of Bounded Variation and Free Discontinuity Problems*, Oxford Mathematical Monographs, Oxford University Press (2000), 10.1093/oso/9780198502456.001.0001.