Minimum Preparation-Energy Scaling for Spatial Separability in Scalar Quantum Fields

Michael Chamerski

Aathra Research LLC

For every fixed positive collar width, the separated regional pair in the scalar-field models considered here admits globally normal product states. We determine the sharp asymptotic order of the global mean-energy infima over PPT, separable, and product preparations as the collar width tends to zero. For the real scalar field in $d+1$ dimensions, $d \geq 3$, we prove

$$c_M \frac{\mathrm{Area}(S)}{\delta^d} \leq E^{(m)}_{\mathrm{PPT}}(\delta) \leq E^{(m)}_{\mathrm{Sep}}(\delta) \leq E^{(m)}_{\mathrm{prod}}(\delta) \leq C_M \frac{\mathrm{Area}(S)}{\delta^d}.$$

The Minkowski constants are uniform when $m\delta$ remains bounded. The same law holds on bounded-geometry ultrastatic backgrounds with a uniform spectral gap. The lower bound applies to every globally normal state that is positive under partial transpose in the algebraic sense, including every separable state. A matching upper bound is furnished by globally normal quasifree Hadamard product states.

For the sharp regional momentum cross block, we also determine the least positive diagonal covariance required to erase it. Its curved-interface cost is $\mathrm{Area}(S)\delta^{-d}c_d(m\delta)+O(\mathrm{Area}(S)\delta^{-(d-1)})$ uniformly for bounded $m\delta$, where c_d is an explicit strictly decreasing mass-collar crossover function. At fixed mass the leading coefficient is independent of interface curvature and equals $1/(12\pi^2)$ in $d=3$.

The result is stable under approximation. Let γ_M be the uniform local witness gap produced by the proof. Exact membership in the PPT, separable, or product class, and every fixed-tolerance approximation with $0 \leq \epsilon < \gamma_M$, have the same $\mathrm{Area}(S)/\delta^d$ energy order. The lower proof combines a finite local PPT witness with an interface frame estimate controlled by the Hamiltonian. In $3+1$ dimensions, exact membership and every such fixed-accuracy approximation cost order $\hbar c\ \mathrm{Area}(S)/\delta^3$. Microcausality remains exact; the law constrains preparation rather than transmission.

1 Introduction

Two spacelike separated regions carry commuting observable algebras. Their commutativity is the field-theoretic statement of causal independence. It is not statistical independence. A state is product only when every joint expectation factorizes, and it is separable only when it can be assembled as a convex mixture of such product expectations.

Michael Chamerski: mchamerski@aathra.org

At positive separation, the split property permits normal product states. Buchholz established product states for suitably separated massive free-field algebras, and Sanders constructed separable quasifree Hadamard states on broad curved backgrounds [1–3]. At zero separation, tangent local algebras can be maximally Bell correlated and normal separable states can cease to exist [4, 5]. The unresolved boundary is quantitative. How much global energy is required to keep two regional algebras separable, or even close to the PPT set, when a collar of width δ is still positive but tends to zero?

Mode counting fixes the candidate ultraviolet scale. A collar resolves wavelengths of order δ, each carrying energy of order $1/\delta$, while a smooth interface contains order $\mathrm{Area}(S)/\delta^2$ independent cells. Their product is $\mathrm{Area}(S)/\delta^3$ in three spatial dimensions. The theorem below turns this count into a regulator-free optimization law. Bény and coauthors anticipated the boundary-over-cutoff scale in a lattice productization argument, while Hollands and Sanders combined dilation with full-density interface packing to bound entanglement measures of a fixed vacuum state [5, 6]. Sanders' planar translation-invariant competitors have finite energy density but infinite total energy; finite-total-energy separable Hadamard states for bounded regions were left open [3]. None of these results minimizes the global Hamiltonian over a state that may vary with the collar, or gives a regulator-free PPT lower bound matched by a finite-total-energy exact-product upper bound. A varying state can reorganize all of its correlations, so a fixed-vacuum estimate does not answer that optimization problem. To our knowledge, the construction and fixed-density frame estimate below provide the two matching bounds for the first time.

The main result proves the dimensional prediction without a regulator. It is also stronger than a separability bound. We minimize over the algebraic PPT class introduced by Verch and Werner [7]. This class contains every separable state. Its use avoids any tensor-product assumption for type III local algebras.

The proof has a lower half and an upper half, and the two halves use different ideas. The lower half begins with one finite local witness. Moving its expectation from the vacuum side of the PPT boundary to the PPT side forces a fixed occupation of the finitely many one-particle modes seen by that witness. Translations of those mode spaces along the interface form a Bessel family, whose total number operator is bounded by δH. This gives the local energy count

$$\#\{\text{neutralized cells}\} \le C\delta\,\rho(H). \tag{1}$$

A smooth interface supplies a full-density family of order $\mathrm{Area}(S)/\delta^2$ cells. A state that is PPT on the full regional pair must neutralize the witness in every cell, which forces the cubic law.

The zero threshold in Eq. (1) is not essential. If a measurement protocol tolerates residual witness negativity ϵ, the same argument replaces the vacuum gap γ by $\gamma - \epsilon$. Averaging the cell witnesses then converts the count into an intrinsic state-norm distance from the full PPT set. Consequently, no fixed finite-energy state can be PPT, or remain within any fixed accuracy below the witness gap, at all sufficiently fine spatial resolutions.

The decisive phase-space fact is that the frame constant does not deteriorate as the number of interface cells grows. Dybalski's sharpened nuclearity estimates controlled discrete sums of translated local excitations when increasing mutual separation suppressed their off-diagonal terms [8, 9]. Here the cell separation remains $O(\delta)$. By resolving only the finite modes seen by the PPT witness, the one-particle estimate remains uniform at the full interface density $\mathrm{Area}(S)/\delta^{d-1}$.

The upper half is constructive. An exact heat-kernel factorization of the two vacuum cross-covariance blocks produces a positive covariance correction whose off-diagonal entries cancel the vacuum cross terms. We prove separately that this correction is trace class in

the vacuum one-particle topology, is the covariance of a Radon Gaussian distribution of coherent displacements, and has finite mean energy. Its Weyl characteristic functional factorizes on the two regional algebras, so the resulting globally normal Hadamard state is exactly product. Its energy is at most $C\,\mathrm{Area}(S)/\delta^3$. A flat-interface calculation gives an explicit massive upper function. The exponent is therefore optimal.

On ultrastatic backgrounds, the Hadamard tangent limit and curved frame replace exact dilation and translation, while the heat-semigroup completion supplies the upper bound. For the separate constrained problem of positive diagonal erasure of the sharp momentum block, a boundary-Hankel calculation determines the exact flat crossover coefficient and its curved-interface asymptotic.

Sections 4–6 prove the lower bound and robust exclusion law. Section 7 constructs the matching product state. The appendices supply the operator-algebraic witness approximation, the curved frame estimates, the trace-class completion integral, and the boundary-Hankel calculus.

2 Regional states and the energy problem

Let $O \mapsto \mathcal{A}(O)$ be a local net in a vacuum representation on $\mathcal{H}$. Spacelike separation gives

$$[\mathcal{A}(O_A), \mathcal{A}(O_B)] = 0. \tag{2}$$

The notation $\mathcal{A}_A \vee \mathcal{A}_B$ means the von Neumann algebra generated by the two commuting algebras.

A state η is product across $(\mathcal{A}_A, \mathcal{A}_B)$ when

$$\eta(AB) = \eta_A(A)\eta_B(B) \qquad (A \in \mathcal{A}_A,\ B \in \mathcal{A}_B). \tag{3}$$

It is separable when it belongs to the weak-star closure of the convex hull of product states. We use the weak-star closure because it is intrinsic to the algebras and includes the usual norm-closed normal separable class.

The algebraic PPT condition requires

$$\sum_{\alpha,\beta} \eta(A_\beta A_\alpha^* B_\alpha^* B_\beta) \geq 0 \tag{4}$$

for every finite choice $A_\alpha \in \mathcal{A}_A$ and $B_\alpha \in \mathcal{A}_B$ [7]. It reduces to positivity of the partial transpose for matrix algebras. Every separable state obeys Eq. (4).

Let $H \geq 0$ be the global Hamiltonian and let ρ be a globally normal density state. Its mean energy is

$$\rho(H) = \mathrm{Tr}(\rho H) \in [0, +\infty]. \tag{5}$$

For a regional pair depending on δ, define

$$E_{\mathrm{PPT}}(\delta) = \inf\{\rho(H) : \rho|_{\mathcal{A}_A \vee \mathcal{A}_B} \text{ is PPT}\}, \tag{6}$$

$$E_{\mathrm{Sep}}(\delta) = \inf\{\rho(H) : \rho|_{\mathcal{A}_A \vee \mathcal{A}_B} \text{ is separable}\}, \tag{7}$$

$$E_{\mathrm{prod}}(\delta) = \inf\{\rho(H) : \rho|_{\mathcal{A}_A \vee \mathcal{A}_B} \text{ is product}\}. \tag{8}$$

The infimum of an empty set is $+\infty$. The construction in Section 7 proves that the product set is nonempty at every positive collar in the models considered here.

Approximation requires a distance which remains meaningful for type III algebras. For a bounded linear functional φ on a von Neumann algebra $\mathfrak{M}$, its norm is

$$\|\varphi\|_{\mathfrak{M}^*} = \sup_{\|O\| \leq 1} |\varphi(O)|. \tag{9}$$

For density operators in finite dimensions this is the trace norm of their difference, without the conventional factor 1/2. Operationally, any single observable with $\|O\| \leq 1$ gives a lower bound on this distance. Let

$$\mathfrak{M}_\delta = \mathcal{A}(A_\delta) \vee \mathcal{A}(B_\delta). \tag{10}$$

For $\mathcal{C} \in \{\mathrm{PPT}, \mathrm{Sep}, \mathrm{prod}\}$, let $\mathcal{C}(\mathcal{A}(A_\delta), \mathcal{A}(B_\delta))$ denote the corresponding set of states relative to that ordered pair of commuting algebras, and define

$$D_{\mathcal{C}}(\rho; \delta) = \inf_{\eta \in \mathcal{C}(\mathcal{A}(A_\delta), \mathcal{A}(B_\delta))} \|\rho|_{\mathfrak{M}_\delta} - \eta\|_{\mathfrak{M}_\delta^*}, \tag{11}$$

and, for $\epsilon \geq 0$,

$$E_{\mathcal{C},\epsilon}(\delta) = \inf \left\{ \rho(H) : \begin{array}{l} \rho \text{ is a globally normal density state,} \\ D_{\mathcal{C}}(\rho; \delta) \leq \epsilon \end{array} \right\}. \tag{12}$$

The competing state η need not possess a density operator or a global normal extension. This makes Eq. (11) stronger than a comparison restricted to normal regional states. Because prod $\subset$ Sep $\subset$ PPT, the distances obey $D_{\mathrm{prod}} \geq D_{\mathrm{Sep}} \geq D_{\mathrm{PPT}}$. Each algebraic-PPT inequality is a weak-star closed linear inequality in the state. The PPT set is therefore convex and weak-star closed. The product set is closed by Eq. (3), and the separable set was defined with its weak-star closure. All three sets are consequently norm closed, so distance zero recovers exact membership.

We set $\hbar = c = 1$ until Section 9. In d spatial dimensions, let (Σ, h) be a complete Riemannian manifold. We write $\mathcal{H}^{d-1}$ for $(d-1)$-dimensional Hausdorff measure and set $\mathrm{Area}(S) = \mathcal{H}^{d-1}(S)$.

Regular collars. A regular collar is specified by a compact nonempty smooth two-sided hypersurface $S \subset \Sigma$ with positive tubular radius. In Fermi coordinates $(s, u) \in S \times (-u_0, u_0)$, the two spatial regions have the form

$$A_\delta \cap \mathcal{T} = \{(s, u) : -u_0 < u < 0\}, \qquad B_\delta \cap \mathcal{T} = \{(s, u) : \delta < u < u_0\}. \tag{13}$$

The regions are relatively compact. Outside the tubular neighborhood $\mathcal{T}$, their closures remain at a fixed positive distance. Uniform $C^{1,1}$ perturbations with positive reach and bounded Jacobians obey the same estimates. This is a geometric condition; no witness-cell collection is supplied as a hypothesis.

Uniform families. Family-uniform constants require common positive lower bounds on the injectivity and tubular radii, and uniform bounds on curvature, local interface graphs, Fermi Jacobians, and the derivatives of the potential used below. These bounds control both cell packing and overlap. For the upper construction, the fixed-distance parts of the cutoff supports must have volume at most $C\,\mathrm{Area}(S)$ and uniform cutoff seminorms. The same hypotheses give a common positive lower bound on every interface area. No upper area, intrinsic-diameter, or global annular-growth assumption is needed. For a single fixed interface the support conditions are automatic, and its fixed-distance $O(1)$ terms are absorbed into the displayed law after reducing δ_M.

The spacetime regions are the causal developments of A_δ and B_δ. In Minkowski spacetime, h is Euclidean and the scalar one-particle energy is

$$\varepsilon_m = (-\Delta + m^2)^{1/2}, \qquad H_m = \mathrm{d}\Gamma(\varepsilon_m). \tag{14}$$

Here $\mathrm{d}\Gamma$ denotes bosonic second quantization of a one-particle operator. On an ultrastatic spacetime $\mathbb{R} \times \Sigma$, let

$$L = -\Delta_h + V \geq \nu_0^2 > 0, \qquad \varepsilon = L^{1/2}, \qquad H = \mathrm{d}\Gamma(\varepsilon). \tag{15}$$

We assume bounded geometry of all orders, meaning a uniform positive injectivity radius and uniform bounds on curvature and its derivatives. We also assume a positive tubular radius and uniform bounds on all derivatives of V. Closed Σ is included. The spectral gap is used only to control the global infrared sector. The Minkowski theorem below also covers the massless case without a gap.

3 The sharp interface law

Theorem 1 (Sharp interface energy law)**.** *Let $d \geq 3$ and let A_δ, B_δ form a regular collar along S.*

In the real scalar field on $(d+1)$-dimensional Minkowski spacetime, fix $M < \infty$. There are constants $c_M, C_M, \delta_M > 0$ such that for every $m \geq 0$ and every $0 < \delta < \delta_M$ with $m\delta \leq M$,

$$\boxed{c_M \frac{\mathcal{H}^{d-1}(S)}{\delta^d} \leq E^{(m)}_{\mathrm{PPT}}(\delta) \leq E^{(m)}_{\mathrm{Sep}}(\delta) \leq E^{(m)}_{\mathrm{prod}}(\delta) \leq C_M \frac{\mathcal{H}^{d-1}(S)}{\delta^d}} \tag{16}$$

The same estimate holds on the ultrastatic scalar theory Eq. (15). Its constants are uniform on families with a common positive spectral gap, injectivity and tubular radii, uniform curvature, potential, local-interface, Fermi-Jacobian, and cutoff-seminorm bounds, and fixed-distance cutoff volume at most $C\,\mathrm{Area}(S)$.

More precisely, the proof produces local selfadjoint contractions $X_{\delta,k}$, with $\|X_{\delta,k}\| \leq 1$, and uniform constants $\gamma_M, C_{\mathrm{count},M}, c_{\mathrm{pack}} > 0$. Every PPT state has $\eta(X_{\delta,k}) \geq 0$, while the vacuum has $\omega(X_{\delta,k}) \leq -\gamma_M$. Every globally normal state of finite energy obeys

$$\#\{k : \rho(X_{\delta,k}) \geq 0\} \leq C_{\mathrm{count},M}\, \delta\, \rho(H). \tag{17}$$

The interface family contains at least $c_{\mathrm{pack}} \mathcal{H}^{d-1}(S) \delta^{-(d-1)}$ cells. Hence every fixed finite-energy state satisfies $\rho(X_{\delta,k}) < 0$ in every cell once δ is sufficiently small.

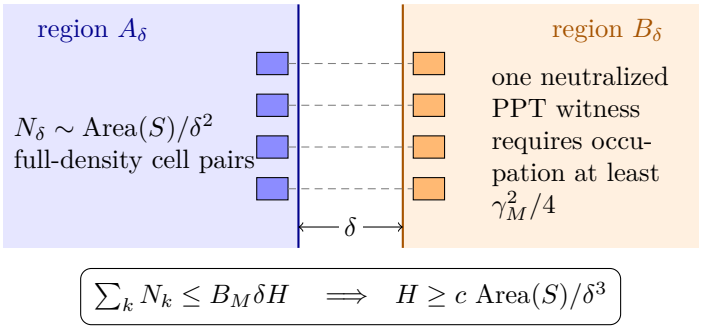


Figure 1: Shown for $d = 3$. The cubic law separates geometry from energy. A regular interface supplies order $\mathrm{Area}(S)/\delta^2$ cell pairs. Neutralizing the fixed local PPT witness costs a nonzero mode occupation in each cell, while the frame estimate bounds their sum by $B\delta H$.

Corollary 2 (Energy-constrained PPT exclusion)**.** *Under the hypotheses of Theorem 1, there are constants $\gamma_M, K_M > 0$ such that every globally normal finite-energy state obeys*

$$\boxed{D_{\mathrm{PPT}}(\rho;\delta) \geq \left[\gamma_M - K_M\sqrt{\frac{\delta^d \rho(H_m)}{\mathcal{H}^{d-1}(S)}}\right]_+.} \tag{18}$$

Here $[x]_+ = \max\{x, 0\}$. The same conclusion holds with the corresponding uniform constants on the ultrastatic families covered by the theorem.

Consequently, if $0 \leq \epsilon < \gamma_M$ and $D_{\mathrm{PPT}}(\rho;\delta) \leq \epsilon$, then

$$\rho(H_m) \geq \frac{(\gamma_M - \epsilon)^2}{K_M^2}\frac{\mathcal{H}^{d-1}(S)}{\delta^d}. \tag{19}$$

The distance to the separable or product set is no smaller. The same energy bound therefore holds when the regional state lies within distance ϵ of either of those smaller classes. Since an exact product competitor has distance zero from all three classes, the complete robust law is

$$\boxed{\frac{(\gamma_M - \epsilon)^2}{K_M^2}\frac{\mathcal{H}^{d-1}(S)}{\delta^d} \leq E^{(m)}_{\mathrm{PPT},\epsilon}(\delta) \leq E^{(m)}_{\mathrm{Sep},\epsilon}(\delta) \leq E^{(m)}_{\mathrm{prod},\epsilon}(\delta) \leq C_M\frac{\mathcal{H}^{d-1}(S)}{\delta^d}.} \tag{20}$$

The power δ^{-d} is sharp. In $d = 3$, the theorem reads

$$E_{\mathrm{Sep}}(\delta) \asymp \frac{\mathrm{Area}(S)}{\delta^3}. \tag{21}$$

The lower constant is a finite-mode witness constant. It is independent of Gaussianity. The upper constant comes from an explicit product covariance. They need not agree.

Corollary 2 shows that exact equality is not driving the divergence. Any fixed approximation tolerance $\epsilon < \gamma_M$ leaves the exponent unchanged. The dimensionless ratio $\delta^d\rho(H_m)/\mathcal{H}^{d-1}(S)$ controls how closely the regional state can approach the PPT set.

The square root in Eq. (18) has a concrete origin. A small vacuum–excitation superposition can change the expectation of a bounded witness by the square root of its occupation probability. A linear correction would therefore be false for unrestricted states. The proof keeps this dependence rather than hiding it inside an unspecified continuity constant.

The lower bound now reduces to two quantitative statements. Changing a bounded finite-mode witness by a fixed amount requires a fixed expected occupation of those modes. The translated occupations are bounded by δH. The next three sections prove these statements and count the interface cells. Section 7 then constructs the matching product state.

4 One finite witness and its occupation cost

4.1 Why the witness is chosen in a derivative subnet

An arbitrary compactly supported scalar smearing can carry a zero-momentum component. After multiplication by $\varepsilon^{-1/2}$, translated copies of such a mode need not form a Bessel family on a two-dimensional interface. Without removing this zero-momentum obstruction, a uniform full-density frame bound fails. We remove it at the source.

Write a real Cauchy datum as $F = (f, g)$, and denote the one-particle Hilbert space by $\mathcal{H}_1$. In Minkowski spacetime its one-particle vector is

$$K_m F = \frac{1}{\sqrt{2}} \left(\varepsilon_m^{-1/2} \widehat{f} + i \varepsilon_m^{1/2} \widehat{g} \right). \tag{22}$$

Choose an integer r large enough that $4r > d + 5$, and restrict the local Weyl generators to data

$$f = (-\Delta)^r f_0, \qquad g = (-\Delta)^r g_0, \qquad f_0, g_0 \in C_c^\infty(O \cap \Sigma). \tag{23}$$

Let $\mathcal{A}^{(r)}(O)$ be the subnet generated by these derivative Weyl operators.

Proposition 3 (Uniform finite derivative-mode witness)**.** *Fix two nonempty, relatively compact, spacelike separated reference cells O_A, O_B and a compact mass interval $0 \le \mu \le M$. There is a finite collection of selfadjoint derivative-Weyl polynomials such that, for each μ in that interval, one member X_μ satisfies*

$$\|X_\mu\| \le 1, \qquad \eta(X_\mu) \ge 0 \quad \textit{for every PPT state } \eta, \qquad \omega_\mu(X_\mu) \le -\gamma_M. \tag{24}$$

Here the displayed subscripts are to be read at the selected mass μ; after normalization the gap may be chosen uniformly,

$$\gamma_M > 0, \qquad G_M < \infty. \tag{25}$$

The second constant is a common bound for the finite Gram inverses. The frame constant B_M is derived from it in Section 5. On every bounded-geometry ultrastatic family in Theorem 1, the rescaled flat massless witness defines, in each sufficiently small Fermi cell, a finite Weyl polynomial $X_{\delta,k}$ which, after one family-uniform scalar normalization, is a selfadjoint contraction with the same PPT positivity and a uniform negative ground-state expectation.

The proof has four steps. First, the derivative subnet is still local and nonabelian. Second, its translated one-particle spaces are dense, so the usual analytic-continuation proof gives the Reeh–Schlieder property. Third, Theorem 6.2 of Verch and Werner then supplies completely positive maps from two qubit algebras into the two local algebras whose vacuum output is NPT [7, Thm. 6.2]. Finally, the maps can be approximated through their positive Choi matrices, not entry by entry. This last choice preserves complete positivity exactly while replacing all local operators by finite Weyl polynomials.

Appendix A checks these four points and the compact-mass uniformity. In the curved case, the same finite Choi sum-of-squares formulas remain completely positive for the local CCR. The Hadamard tangent limit is used only to transport the strict negative vacuum expectation [10, Sec. 5.3]. Thus the algebraic PPT inequality and the vacuum-gap estimate enter by logically separate routes.

Negativity of X_μ is stronger than nonseparability. It certifies failure of algebraic PPT through fixed local operations. The construction uses no massive dilation symmetry; time-zero Weyl relations identify the algebraic polynomial across masses, while vacuum expectations and finite Gram matrices vary continuously with μ.

4.2 A finite-mode witness has a finite occupation price

Let $\mathsf{E} \subset \mathcal{H}_1$ be the finite-dimensional complex one-particle span of the Weyl arguments appearing in X. Bosonic Fock space factorizes as

$$\mathcal{F}(\mathcal{H}_1) \simeq \mathcal{F}(\mathsf{E}) \otimes \mathcal{F}(\mathsf{E}^\perp). \tag{26}$$

The witness acts only on the first factor. Let P_{E} be the one-particle projection onto E and let

$$N_{\mathsf{E}} = \mathrm{d}\Gamma(P_{\mathsf{E}}) \tag{27}$$

be its number operator. Here $\mathrm{d}\Gamma$ denotes second quantization of a one-particle operator.

For an arbitrary globally normal state ρ, denote the reduced density operator on $\mathcal{F}(\mathsf{E})$ by ρ_{E}. Its vacuum vector is $|0_{\mathsf{E}}\rangle$. The probability of leaving the finite-mode vacuum is bounded by the expected number,

$$1 - \langle 0_{\mathsf{E}}|\rho_{\mathsf{E}}|0_{\mathsf{E}}\rangle \leq \rho(N_{\mathsf{E}}). \tag{28}$$

The trace-norm estimate against a pure state then gives

$$\|\rho_{\mathsf{E}} - |0_{\mathsf{E}}\rangle\langle 0_{\mathsf{E}}|\|_1 \leq 2\sqrt{\rho(N_{\mathsf{E}})}. \tag{29}$$

The subscript 1 denotes the trace norm. Since $\|X\| \leq 1$ and the vacuum expectation is at most $-\gamma$, the same estimate gives the state-independent inequality

$$\rho(X) \leq -\gamma + 2\sqrt{\rho(N_{\mathsf{E}})}. \tag{30}$$

This is the quantitative link between a measurable witness value and the occupation of its finite mode space. It holds whether or not the witness has crossed the PPT boundary. If $\rho(X) \geq 0$ while $\omega(X) \leq -\gamma$, the witness expectation changes by at least γ. Since $\|X\| \leq 1$,

$$\gamma \leq \|\rho_{\mathsf{E}} - |0_{\mathsf{E}}\rangle\langle 0_{\mathsf{E}}|\|_1 \leq 2\sqrt{\rho(N_{\mathsf{E}})}. \tag{31}$$

Thus every state that neutralizes the witness pays the uniform lower price

$$\boxed{\rho(N_{\mathsf{E}}) \geq \frac{\gamma^2}{4}.} \tag{32}$$

This step assumes neither Gaussianity nor a particle interpretation for the regional algebras. N_{E} counts only the finite global Fock modes used by the witness.

5 The energy-frame estimate

The occupation cost becomes an interface law only if the translated mode spaces can be added without overcounting. This is the role of the frame estimate.

For mass μ, let E_μ denote the finite one-particle span of the Weyl arguments in the chosen witness, and put $n_{\mathsf{E}} = \dim \mathsf{E}_\mu$. Let $e_{a,\mu}$, $1 \leq a \leq n_{\mathsf{E}}$, be an orthonormal basis of this space. Spatial translations on the one-particle space are denoted by $\tau(\mathbf{x})$. Define

$$q_{a,\mu} = \varepsilon_\mu^{-1/2} e_{a,\mu}. \tag{33}$$

Before orthonormalization, Eqs. (22) and (23) give the explicit form

$$q_{a,\mu}(\mathbf{p}) = \frac{1}{\sqrt{2}} \left[\frac{|\mathbf{p}|^{2r}}{\sqrt{|\mathbf{p}|^2 + \mu^2}} \widehat{f}_{0,a}(\mathbf{p}) + i|\mathbf{p}|^{2r} \widehat{g}_{0,a}(\mathbf{p}) \right]. \tag{34}$$

The finite Gram matrix is uniformly invertible on a compact mass interval after redundant arguments are removed. Its inverse square root therefore changes Eq. (34) only by uniformly bounded finite linear combinations.

For $|\alpha| \le d+1$, differentiate the products $\overline{q}_{a,\mu} q_{b,\mu}$ rather than either factor separately. Their worst low-momentum multiplier is $|\mathbf{p}|^{4r}/(|\mathbf{p}|^2+\mu^2)$. The choice $4r > d+5$ gives a mass-independent locally integrable majorant for all derivatives used below. At large momentum the test functions are Schwartz. It follows that

$$\sup_{0\le\mu\le M} \left\| \overline{q}_{a,\mu} q_{b,\mu} \right\|_{W^{d+1,1}(\mathbb{R}^d)} < \infty. \tag{35}$$

Repeated integration by parts gives

$$|\langle q_{a,\mu}, \tau(\mathbf{x}) q_{b,\mu}\rangle| \le C_M (1+|\mathbf{x}|)^{-(d+1)}. \tag{36}$$

Let Λ be a relatively separated subset of $\mathbb{R}^d$. Ambient packing bounds the number of its points in the nth unit shell by $O((1+n)^{d-1})$. Hence Eq. (36) is summable in every row of the block Gram matrix indexed by $(\mathbf{x}, a)$. Schur's test gives a constant B_M independent of the number of points such that

$$\sum_{\mathbf{x}\in\Lambda} \sum_{a=1}^{n_{\mathsf{E}}} |\tau(\mathbf{x}) q_{a,\mu}\rangle\langle \tau(\mathbf{x}) q_{a,\mu}| \le B_M \mathbf{1}. \tag{37}$$

Since $e_{a,\mu} = \varepsilon_\mu^{1/2} q_{a,\mu}$, this is equivalent, as a quadratic-form inequality, to

$$\boxed{\sum_{\mathbf{x}\in\Lambda} P_{\tau(\mathbf{x})\mathsf{E}_\mu} \le B_M \varepsilon_\mu.} \tag{38}$$

Second quantization preserves positive form inequalities, so

$$\sum_{\mathbf{x}\in\Lambda} \mathrm{d}\Gamma(P_{\tau(\mathbf{x})\mathsf{E}_\mu}) \le B_M H_\mu. \tag{39}$$

Equation (38) is the new phase-space input. It controls only the finite modes which cross the PPT boundary. It does not attempt to cover the unit ball of a local algebra, and its coefficient is therefore independent of the number of cells.

6 From a regular interface to the lower law

6.1 The interface supplies the cells

Fix a small reference pair of positively separated spatial cells in $\mathbb{R}^d$. A regular collar supplies scaled copies of this pair without an additional packing assumption. Cover the compact interface by finitely many graph patches on which the unit normals lie in a fixed narrow cone. Choose one oriented reference pair for each patch and shrink it once so that literal translates fit on the two sides throughout that patch. Partition the interface among the patches, then choose one globally maximal set of points $s_{\delta,k} \in S$ separated by $L\delta$ in the induced metric. Assign each selected center to its designated graph patch and orientation. Take L larger than the tangential diameter of every reference pair. Positive reach and uniform graph bounds place one translated pair at every chosen point and give uniformly bounded overlap across patch boundaries and distinct folds. Standard volume comparison on the compact hypersurface gives

$$N_\delta \ge c_{\mathrm{pack}} \mathcal{H}^{d-1}(S) \delta^{-(d-1)}. \tag{40}$$

Within each graph patch, the centers form a relatively separated Ahlfors family and the mode spaces are literal translates of one fixed oriented space. Equation (38) therefore applies patch by patch. Summing the finitely many frame bounds only changes the constant. This is exactly the geometric input used in the shell count below Eq. (36).

The assertion remains true for uniform $C^{1,1}$ collars. At scale δ, the graph error from the tangent plane is $O(\delta^2)$, while the free space inside a reference cell is $O(\delta)$. Reducing the fixed cell size once makes every copy fit for all sufficiently small collars.

6.2 Exact scaling in Minkowski spacetime

Let d_λ be the unitary dilation on the one-particle momentum space,

$$(d_\lambda \psi)(\mathbf{p}) = \lambda^{d/2} \psi(\lambda \mathbf{p}), \tag{41}$$

and let $D_\lambda = \Gamma(d_\lambda)$. It maps the mass-$\lambda m$ reference theory to scale λ in the mass-m theory. Direct substitution gives

$$d_\lambda \varepsilon_{\lambda m} d_\lambda^* = \lambda \varepsilon_m, \qquad D_\lambda H_{\lambda m} D_\lambda^* = \lambda H_m. \tag{42}$$

This is an isomorphism between different mass theories. It is not a massive dilation symmetry.

Choose $\lambda_\delta = \alpha_0 \delta$, with a fixed $\alpha_0 > 0$ small enough that the scaled reference cells fit inside Eq. (13). The reference mass is

$$\mu_\delta = m \lambda_\delta. \tag{43}$$

The condition $m\delta \le M$ keeps μ_δ in a compact interval. After conjugating Eq. (38) by d_{λ_δ}, the physical cell projections obey

$$\sum_{k=1}^{N_\delta} P_{\delta,k} \le B_M \lambda_\delta \varepsilon_m. \tag{44}$$

Define the corresponding physical cell number operators by

$$N_{\delta,k} = \mathrm{d}\Gamma(P_{\delta,k}). \tag{45}$$

Second quantization yields

$$\sum_{k=1}^{N_\delta} \rho(N_{\delta,k}) \le B_M \lambda_\delta \rho(H_m). \tag{46}$$

If the regional restriction of ρ is PPT, its restriction to every cell pair is PPT. It therefore neutralizes every local witness. Applying Eq. (32) in all cells gives

$$\frac{\gamma_M^2}{4} N_\delta \le B_M \lambda_\delta \rho(H_m). \tag{47}$$

Equations (40) and (47) prove

$$\rho(H_m) \ge \frac{c_{\mathrm{pack}} \gamma_M^2}{4 B_M \alpha_0} \frac{\mathcal{H}^{d-1}(S)}{\delta^d}. \tag{48}$$

This is the first inequality in Eq. (16).

For an arbitrary finite-energy state, fix a tolerance $0 \le \epsilon < \gamma_M$. Equation (30) shows that every cell satisfying $\rho(X_{\delta,k}) \ge -\epsilon$ must obey

$$\rho(N_{\delta,k}) \ge \frac{(\gamma_M - \epsilon)^2}{4}. \tag{49}$$

Apply Eq. (46) only to those cells. This gives the tolerant count

$$\boxed{\#\{k : \rho(X_{\delta,k}) \ge -\epsilon\} \le \frac{4B_M\alpha_0\delta}{(\gamma_M - \epsilon)^2}\rho(H_m).} \tag{50}$$

At $\epsilon = 0$, Eq. (50) is the exact neutralized-cell law claimed in Eq. (17). When its right side is smaller than one, every witness in the family remains below $-\epsilon$. The conclusion therefore survives a fixed measurement margin.

This count also controls approximation to PPT in each cell. Let $D_{\mathrm{PPT},k}(\rho;\delta)$ be the state-norm distance between the restriction of ρ to the kth cell pair and that cell's algebraic-PPT set. Since every cellwise PPT state has nonnegative expectation on $X_{\delta,k}$,

$$D_{\mathrm{PPT},k}(\rho;\delta) \ge [-\rho(X_{\delta,k})]_+. \tag{51}$$

Therefore

$$\#\{k : D_{\mathrm{PPT},k}(\rho;\delta) \le \epsilon\} \le \#\{k : \rho(X_{\delta,k}) \ge -\epsilon\}. \tag{52}$$

The tolerant witness count is thus also a bound on how many interface cells can lie within distance ϵ of PPT.

6.3 Robust distance from the PPT set

The tolerant cell count has an intrinsic regional form. Average the witnesses over the interface,

$$\overline{X}_\delta = \frac{1}{N_\delta}\sum_{k=1}^{N_\delta} X_{\delta,k}. \tag{53}$$

Each witness has norm at most one, so $\left\|\overline{X}_\delta\right\| \le 1$. If η is any algebraic-PPT state on $\mathfrak{M}_\delta$, its restriction to each cell is PPT. Therefore

$$\eta(\overline{X}_\delta) \ge 0. \tag{54}$$

This remains true even when η is singular or has no globally normal extension.

Write $n_k = \rho(N_{\delta,k})$. Averaging Eq. (30) and applying Cauchy–Schwarz gives

$$\begin{aligned}\rho(\overline{X}_\delta) &\le -\gamma_M + \frac{2}{N_\delta}\sum_{k=1}^{N_\delta}\sqrt{n_k} \\ &\le -\gamma_M + 2\sqrt{\frac{1}{N_\delta}\sum_{k=1}^{N_\delta} n_k}.\end{aligned} \tag{55}$$

The frame estimate and the geometric cell count now enter separately. Equations (46) and (40) yield

$$\rho(\overline{X}_\delta) \le -\gamma_M + 2\sqrt{\frac{B_M\alpha_0}{c_{\mathrm{pack}}}}\sqrt{\frac{\delta^d\rho(H_m)}{\mathcal{H}^{d-1}(S)}}. \tag{56}$$

Equation (56) becomes a distance statement by observable–state duality. For every PPT state η,

$$\begin{aligned}\|\rho|_{\mathfrak{M}_\delta} - \eta\|_{\mathfrak{M}_\delta^*} &\geq \left|\eta(\overline{X}_\delta) - \rho(\overline{X}_\delta)\right| \\ &\geq \left[-\rho(\overline{X}_\delta)\right]_+ \\ &\geq \left[\gamma_M - 2\sqrt{\frac{B_M\alpha_0}{c_{\mathrm{pack}}}}\sqrt{\frac{\delta^d\rho(H_m)}{\mathcal{H}^{d-1}(S)}}\right]_+ . \end{aligned} \tag{57}$$

Taking the infimum over η proves Corollary 2 with

$$K_M = 2\sqrt{\frac{B_M\alpha_0}{c_{\mathrm{pack}}}}. \tag{58}$$

Solving Eq. (18) for $\rho(H_m)$ gives Eq. (19). The exact lower bound is the special case $\epsilon = 0$. The derivation also explains the experiment represented by $\overline{X}_\delta$: choose one interface cell uniformly at random and perform its finite witness protocol. No collective type III measurement is required.

6.4 The curved one-particle frame

The flat proof used exact dilation to set the cell size and exact translation to compare different cells. Neither operation is a symmetry on a general ultrastatic slice. Fermi rescaling replaces the dilation, while a direct Sobolev synthesis estimate replaces translation invariance.

Work in Fermi coordinates centered at an interface point. Put $\Phi_{k,\delta}(y) = \exp_{s_{\delta,k}}(\delta y)$ on the fixed reference cell and define $J_{k,\delta}$ by $\Phi_{k,\delta}^*\mathrm{dvol}_h = \delta^d J_{k,\delta}(y)\mathrm{d}^d y$. Let $f_{0,j}, g_{0,j}$ be the undifferentiated flat primitives used in Eq. (23). To keep the cell notation readable, denote their density-corrected rescalings by $f_{k,j,\delta}, g_{k,j,\delta}$:

$$f_{k,j,\delta}(\Phi_{k,\delta}(y)) = \delta^{2r-(d+1)/2} J_{k,\delta}(y)^{-1/2} f_{0,j}(y), \tag{59}$$

$$g_{k,j,\delta}(\Phi_{k,\delta}(y)) = \delta^{2r-(d-1)/2} J_{k,\delta}(y)^{-1/2} g_{0,j}(y). \tag{60}$$

The powers have separate roles. After $(-\Delta_h)^r$ is applied once, the f-amplitude is $\delta^{-(d+1)/2}$ and the g-amplitude is $\delta^{-(d-1)/2}$. Thus the one-particle Gram matrix is scale invariant, whereas the energy-divided modes below have squared norm of order δ. The factor $J_{k,\delta}^{-1/2}$ removes the leading volume distortion.

The corresponding one-particle vectors are

$$s_{k,j,\delta} = \frac{1}{\sqrt{2}}\left(\varepsilon^{-1/2}(-\Delta_h)^r f_{k,j,\delta} + i\varepsilon^{1/2}(-\Delta_h)^r g_{k,j,\delta}\right). \tag{61}$$

Let $\mathsf{E}_{k,\delta}$ be their span, choose an orthonormal basis $e_{k,a,\delta}$, and put

$$q_{k,a,\delta} = \varepsilon^{-1/2} e_{k,a,\delta}. \tag{62}$$

Proposition 4 (Uniform curved frame)**.** *Under the bounded-geometry, potential, gap, collar, and chart bounds in Theorem 1, there are $\delta_0, B, g_-, C > 0$, independent of the interface point and of the member of the controlled family, such that for $0 < \delta < \delta_0$:*

1. *the Gram matrices of the vectors $s_{k,j,\delta}$ have all eigenvalues in a fixed interval $[g_-, C]$;*

2. on the form domain of $\varepsilon^{1/2}$,

$$\boxed{\sum_{k=1}^{N_\delta} P_{\mathsf{E}_{k,\delta}} \le B\delta\varepsilon.} \tag{63}$$

Appendix B proves both statements. The key estimate is local. For every finite coefficient array $(c_{k,j})$, the superposition of rescaled f profiles has H^{2r-1} norm squared at most $C\delta\sum|c_{k,j}|^2$. The analogous g superposition obeys the same bound in H^{2r}. Uniform ellipticity converts these estimates into a Bessel bound for the divided one-particle modes. Within-cell Gram conditioning then gives Eq. (63). The proof uses only bounded overlap of the witness supports and introduces no global shell or diameter assumption.

The same finite Choi sum-of-squares formulas built from the curved Weyl generators define CP maps in every cell. They preserve the algebraic PPT inequality exactly. The Hadamard tangent limit is needed only to transport the strict negative vacuum gap into every cell. Combining Eq. (63), the occupation cost, and Eq. (40) proves

$$E_{\mathrm{PPT}}(\delta) \ge c\mathcal{H}^{d-1}(S)\delta^{-d}. \tag{64}$$

All constants are uniform on the bounded geometric families stated in the theorem. The tolerant count and averaged-witness proof use only this frame inequality, the uniform gap, and the same geometric cell count. They therefore give Eq. (18) on the curved family with $K = 2\sqrt{B/c_{\mathrm{pack}}}$. No curved-space density matrix or regional tensor factor is introduced.

7 Matching product states

Sharpness requires exact product states at the same $\mathrm{Area}(S)/\delta^d$ order. We now construct them.

7.1 Positive completion of the cross covariance

For the ultrastatic scalar vacuum, the symmetric equal-time covariance on real observable test data $F = (f, g)$ is

$$C_0(F, F') = \frac{1}{2}\left(\langle f, \varepsilon^{-1} f'\rangle + \langle g, \varepsilon g'\rangle\right). \tag{65}$$

Choose smooth cutoffs equal to one on A_δ and B_δ. Their supports remain at distance at least $c_0\delta$, and their jth derivatives are $O(\delta^{-j})$. The cross covariance is the two-component smoothing operator

$$K_\delta = \frac{1}{2}\begin{pmatrix} \chi_A\varepsilon^{-1}\chi_B & 0 \\ 0 & \chi_A\varepsilon\chi_B \end{pmatrix}. \tag{66}$$

Both entries in Eq. (66) mean the unique extensions of the off-diagonal distribution kernels, initially evaluated on compactly supported smooth data. They are not unrestricted compositions with the generally unbounded operator ε. The separated cutoff supports make both extensions smoothing.

The completion is easiest to see for one rank-one cross term. Let $u_A = (u_\phi, u_\pi)$ and $v_B = (v_\phi, v_\pi)$ be smooth classical mean configurations supported on the two sides. We write

$$\ell_u(F) = \langle u_\phi, f\rangle + \langle u_\pi, g\rangle. \tag{67}$$

If a covariance contains the cross term $a\,\ell_{u_A}\ell_{v_B}$ and its transpose, then adding the positive square

$$a(\ell_{u_A} - \ell_{v_B})^2 \tag{68}$$

cancels those cross terms and adds only positive diagonal covariance. The mean energy carried by that square is $a(\|u_A\|_E^2 + \|v_B\|_E^2)/2$, where

$$\|u\|_E^2 = \|\varepsilon u_\phi\|_2^2 + \|u_\pi\|_2^2\,. \tag{69}$$

The heat-kernel construction below is an integral of this elementary operation.

Define the energy nuclear cost of the cross kernel by

$$\nu_E(K_\delta) = \inf \frac{1}{2}\sum_\nu a_\nu \left(\|u_\nu\|_E^2 + \|v_\nu\|_E^2\right), \tag{70}$$

where the infimum is over smooth support-preserving expansions

$$K_\delta(F_A, F_B) = \sum_\nu a_\nu \ell_{u_\nu}(F_A)\ell_{v_\nu}(F_B), \qquad a_\nu > 0. \tag{71}$$

Here u_ν and v_ν are supported on the A and B sides, respectively, and the energy sum in Eq. (70) is finite. The equality is an equality of bilinear forms on smooth regional test data, continuous in the energy projective tensor norm. The notation also includes Bochner integrals of rank-one forms; signs and finite component indices are absorbed into v_ν.

Extend u_ν, v_ν by zero to global mean configurations and put $w_\nu = u_\nu - v_\nu$. Every such expansion has the positive completion

$$Q_\delta(F, F') = \sum_\nu a_\nu \ell_{w_\nu}(F)\ell_{w_\nu}(F') \geq 0. \tag{72}$$

Adding Q_δ to the vacuum covariance preserves the CCR uncertainty inequality. Its off-diagonal block is $-K_\delta$, so the new covariance has zero cross block on $A_\delta \times B_\delta$. At this stage this is a positive quasifree covariance form. Global normality, exact factorization of the regional state, and identification of its added mean energy with the selected completion cost are separate operator statements. They are proved after the heat estimate and in Appendix C. For the selected expansion, write

$$\mathcal{E}(Q_\delta) = \frac{1}{2}\sum_\nu a_\nu \left(\|u_\nu\|_E^2 + \|v_\nu\|_E^2\right). \tag{73}$$

The infimum of this quantity over all admissible expansions is $\nu_E(K_\delta)$.

The relevant estimate is

$$\boxed{\nu_E(K_\delta) \leq C\mathcal{H}^{d-1}(S)\delta^{-d}.} \tag{74}$$

It follows from an exact heat-kernel factorization. Let $p_t(x, z)$ be the kernel of e^{-tL} and put

$$u_{t,z}^J(x) = \chi_J(x)p_{t/2}(x, z), \qquad J \in \{A, B\}. \tag{75}$$

The semigroup identity gives

$$\chi_A e^{-tL}\chi_B = \int_\Sigma u_{t,z}^A \otimes u_{t,z}^B \,\mathrm{dvol}_h(z). \tag{76}$$

Since $\chi_A\chi_B = 0$, functional calculus gives the exact off-diagonal representations

$$K_{\phi,\delta} = \frac{1}{2\sqrt{\pi}} \int_0^\infty t^{-1/2} \int_\Sigma u^A_{t,z} \otimes u^B_{t,z} \, \mathrm{dvol}_h(z) \mathrm{d}t, \tag{77}$$

$$K_{\pi,\delta} = -\frac{1}{4\sqrt{\pi}} \int_0^\infty t^{-3/2} \int_\Sigma u^A_{t,z} \otimes u^B_{t,z} \, \mathrm{dvol}_h(z) \mathrm{d}t. \tag{78}$$

The two heat vectors in a rank-one term can have very different norms: one is controlled by the distance to A and the other by the distance to B. Multiplying one vector by a positive number and the other by its reciprocal does not change their tensor product. We use this freedom to balance the two diagonal costs. Put

$$a(z) = \mathrm{dist}(z, \mathrm{supp}\, \chi_A), \qquad b(z) = \mathrm{dist}(z, \mathrm{supp}\, \chi_B). \tag{79}$$

Saloff-Coste's Gaussian estimate includes uniformly bounded lower-order terms, and the bounded-geometry derivative estimates cover the Laplacian [11, Thms. 6.1 and 6.5] [12, Thm. 4.1]. Uniform scaled interior parabolic estimates in normal charts extend the derivative bounds to the real potential V under the coefficient bounds stated above. The argument in Appendix C also shows that one Gaussian exponent $\eta > 0$ works for every spatial derivative order, while the prefactor depends on that order. For $0 < t \le 1$, define

$$r(t,z) = \exp\left[\frac{\eta}{4t}\left(a(z)^2 - b(z)^2\right)\right], \quad \widetilde{u}^A_{t,z} = r(t,z) u^A_{t,z}, \quad \widetilde{u}^B_{t,z} = r(t,z)^{-1} u^B_{t,z}. \tag{80}$$

For $t \ge 1$, set $r(t,z) = 1$. Reciprocal rescaling leaves every cross term unchanged, while the two Gaussian bounds acquire the same exponent, proportional to

$$\exp\left[-\frac{\eta}{2t}\left(a(z)^2 + b(z)^2\right)\right]. \tag{81}$$

Add the positive square built from $(\widetilde{u}_A, -\widetilde{u}_B)$ in the field sector and from $(\widetilde{u}_A, \widetilde{u}_B)$ in the momentum sector. Their cross terms cancel Eqs. (77) and (78) exactly.

Gaussian heat-kernel and gradient estimates then give

$$I_0(t) := \int_\Sigma \sum_{J=A,B} \|\widetilde{u}^J_{t,z}\|_2^2 \, \mathrm{dvol}_h(z) \le C\mathcal{H}^{d-1}(S) t^{-(d-1)/2} e^{-c\delta^2/t} + R_0(t), \tag{82}$$

$$\begin{aligned} I_1(t) &:= \int_\Sigma \sum_{J=A,B} \|L^{1/2}\widetilde{u}^J_{t,z}\|_2^2 \, \mathrm{dvol}_h(z) \\ &\le C(t^{-1} + \delta^{-2} + 1)\mathcal{H}^{d-1}(S) t^{-(d-1)/2} e^{-c\delta^2/t} + R_1(t) \end{aligned} \tag{83}$$

for $0 < t \le 1$. The remainders come from the fixed-distance part of the regions and have δ-independent weighted integrals. The δ^{-2} term is the cost of the cutoff gradients. For the Minkowski operator $L = -\Delta + m^2$, the last factor in Eq. (83) is $t^{-1} + \delta^{-2} + m^2$. The hypothesis $m\delta \le M$ absorbs the mass term into the same $C_M\delta^{-2}$ bound. The large-t integrals converge by the spectral gap on the curved background. In massless Minkowski space they converge by the explicit Euclidean heat kernel and relative compactness of the regions. Consequently,

$$\begin{aligned} \mathcal{E}(Q_\delta) &\le C \int_0^\infty \left[t^{-3/2} I_0(t) + t^{-1/2} I_1(t)\right] \mathrm{d}t \\ &\le C\mathcal{H}^{d-1}(S)\delta^{-d}. \end{aligned} \tag{84}$$

It follows that $\nu_E(K_\delta) \le \mathcal{E}(Q_\delta)$ also obeys Eq. (74). The two singular integrals are

$$\int_0^1 t^{-(d+2)/2} e^{-c\delta^2/t} \mathrm{d}t = O(\delta^{-d}), \qquad \delta^{-2} \int_0^1 t^{-d/2} e^{-c\delta^2/t} \mathrm{d}t = O(\delta^{-d}). \tag{85}$$

The second estimate is the point at which $d > 2$ enters the upper proof. Appendix C proves the overlap estimates. The part of the regions outside the tubular collar stays at fixed distance and contributes only $O(1)$.

For every fixed $\delta > 0$, differentiation in either kernel variable adds only powers of $t^{-1/2}$ and δ^{-1}. The scalar $r(t, z)$ is independent of those kernel variables. The common Gaussian exponent in Eq. (80) therefore survives every derivative. All differentiated small-t integrals converge, so the completed Cauchy covariance has a smooth compactly supported kernel. Its spacetime evolution is a smooth bisolution, whose wavefront set is empty. Adding it therefore leaves the Hadamard wavefront set unchanged [13].

Proposition 5 (Normal coherent realization of the completion)**.** *For every fixed $\delta > 0$, the positive heat-kernel completion Q_δ is the covariance of a centered Radon Gaussian distribution of coherent displacements in the vacuum Fock representation. Its coherent-state mixture is a globally normal density state ρ_δ with covariance $C_0 + Q_\delta$. The state is Hadamard, is exactly product across $\mathcal{A}(A_\delta)$ and $\mathcal{A}(B_\delta)$, and satisfies*

$$\rho_\delta(H) \le C\mathcal{H}^{d-1}(S)\delta^{-d}. \tag{86}$$

The constants have the uniformity asserted in Theorem 1.

Proof. A real Cauchy datum has the ultrastatic one-particle representative

$$KF = \frac{1}{\sqrt{2}} \left(\varepsilon^{-1/2} f + i\varepsilon^{1/2} g \right). \tag{87}$$

In Minkowski spacetime, spatial Fourier transformation turns this into Eq. (22). A classical mean configuration $u = (u_\phi, u_\pi)$ determines the real one-particle displacement

$$Z(u) = \frac{1}{\sqrt{2}} \left(\varepsilon^{-1/2} u_\pi - i\varepsilon^{1/2} u_\phi \right). \tag{88}$$

With the Weyl convention of Eq. (22), conjugation by the Fock displacement $D(Z(u))$ multiplies $W(F)$ by $e^{i\ell_u(F)}$, with ℓ_u as in Eq. (67). The real one-particle norm and the coherent mean energy are

$$\|Z(u)\|^2 = \frac{1}{2} \left(\langle u_\phi, \varepsilon u_\phi \rangle + \langle u_\pi, \varepsilon^{-1} u_\pi \rangle \right), \tag{89}$$

$$\langle Z(u), \varepsilon Z(u) \rangle = \frac{1}{2} \left(\|\varepsilon u_\phi\|_2^2 + \|u_\pi\|_2^2 \right) = \frac{1}{2} \|u\|_E^2 . \tag{90}$$

Regard the one-particle Hilbert space as a real Hilbert space with inner product $\mathrm{Re}\langle \cdot, \cdot \rangle$. For the selected rank-one completion in Eq. (72), define

$$T_\delta = \sum_\nu a_\nu |Z(w_\nu)\rangle\langle Z(w_\nu)|_{\mathbb{R}}, \tag{91}$$

with the sum replaced by the corresponding positive integral for the heat decomposition. This is the precise real one-particle covariance induced by Q_δ. It is trace class. In the gapped theory, Eq. (89) is bounded by $(2\nu_0)^{-1} \|u\|_E^2$. For Minkowski space, uniformly for

every $m \geq 0$, all completion vectors have support in one fixed compact set $\mathcal{K}_0$ and $\varepsilon_m^{-1} \leq |\mathbf{p}|^{-1}$. The Riesz-kernel representation and the Hardy–Littlewood–Sobolev inequality give

$$\begin{aligned}\langle u, |\mathbf{p}|^{-1} u\rangle &= a_d \iint_{\mathcal{K}_0\times\mathcal{K}_0} \frac{\overline{u(x)}u(y)}{|x-y|^{d-1}}\,\mathrm{d}x\,\mathrm{d}y \\ &\leq C_d \|u\|^2_{2d/(d+1)} \leq C_d|\mathcal{K}_0|^{1/d}\,\|u\|_2^2 \end{aligned} \tag{92}$$

by Hölder's inequality [14]. This controls the momentum-displacement contribution to the one-particle norm. For a field displacement, Cauchy–Schwarz and the ordinary compact-support Poincaré inequality give, uniformly in m, $\langle u, \varepsilon_m u\rangle \leq \|u\|_2\,\|\varepsilon_m u\|_2 \leq C_{\mathcal{K}_0}\,\|\varepsilon_m u\|_2^2$. Integrating these bounds against the positive rank-one decomposition and using Eq. (84) gives $\operatorname{Tr} T_\delta < \infty$.

A positive trace-class covariance on a real Hilbert space determines a centered Radon Gaussian measure $\mathbb{P}_\delta$ [15]. Let X denote the real one-particle random variable with this law. The finite form trace $\operatorname{Tr}_{\mathbb{R}}(\varepsilon^{1/2}T_\delta\varepsilon^{1/2}) < \infty$ also implies $X \in \operatorname{Dom}(\varepsilon^{1/2})$ almost surely. Averaging the corresponding coherent vacuum vectors gives the Bochner integral in the trace class

$$\rho_\delta = \int |D(X)\Omega\rangle\langle D(X)\Omega|\,\mathrm{d}\mathbb{P}_\delta(X). \tag{93}$$

It is positive, has trace one, and is globally normal by construction. Its Weyl characteristic functional is

$$\begin{aligned}\operatorname{Tr}(\rho_\delta W(F)) &= e^{-C_0(F,F)/2}\int e^{2i\,\mathrm{Im}\langle X, KF\rangle}\,\mathrm{d}\mathbb{P}_\delta(X) \\ &= \exp\left[-\frac{1}{2}(C_0+Q_\delta)(F,F)\right]. \end{aligned} \tag{94}$$

Here

$$Q_\delta(F,F') = \mathbb{E}\left[2\,\mathrm{Im}\langle X, KF\rangle\,2\,\mathrm{Im}\langle X, KF'\rangle\right], \tag{95}$$

which follows from Eqs. (67), (88), and (91). This proves directly that the mixture is the centered quasifree state with the completed covariance. No separate quasi-equivalence argument is needed for normality because Eq. (93) is already a trace-class density operator in the vacuum representation.

For test data F_A, F_B supported in the two regions, the completed cross covariance vanishes. Locality also gives $\sigma(F_A, F_B) = 0$. Therefore

$$\operatorname{Tr}(\rho_\delta W(F_A)W(F_B)) = \operatorname{Tr}(\rho_\delta W(F_A))\operatorname{Tr}(\rho_\delta W(F_B)). \tag{96}$$

Finite linear combinations of regional Weyl generators are norm dense in the regional Weyl C^*-algebras. Kaplansky density and the normality of ρ_δ extend Eq. (96) separately in each variable to the two regional von Neumann algebras. Thus the regional state is exactly product in the sense of Eq. (3).

Finally, Tonelli's theorem and Eq. (90) give

$$\rho_\delta(H) = \operatorname{Tr}_{\mathbb{R}}(\varepsilon^{1/2}T_\delta\varepsilon^{1/2}) = \mathcal{E}(Q_\delta) \leq C\mathcal{H}^{d-1}(S)\delta^{-d}. \tag{97}$$

The smooth-kernel argument preceding the proposition proves the Hadamard claim. This completes the proof. □

Proposition 5 proves the upper half of the main theorem and the nonemptiness of the product class at every positive collar.

Completion of the proof of Theorem 1. In Minkowski spacetime, Proposition 3, the flat frame inequality (44), and the geometric cell count (40) give the PPT lower bound in Eq. (48). Proposition 4 and the curved tangent witness give the corresponding ultrastatic lower bound (64). Proposition 5 supplies a globally normal exact-product competitor with the matching upper order in both settings. The inclusions prod $\subset$ Sep $\subset$ PPT give the two middle inequalities. The compact-mass construction in Proposition 3 and the uniform constants in the two analytic propositions give precisely the mass and bounded-geometry uniformity stated in the theorem. The count law is Eq. (50) at $\epsilon = 0$. □

8 Coefficient laws for covariance completion

8.1 Exact planar product completion

For an infinite planar interface, a translation-invariant completion yields an exact energy density. It is not the finite-total-energy competitor used in Theorem 1, but it exposes the coefficient as a function of $m\delta$. Take parallel half-spaces $z < 0$ and $z > \delta$ and Fourier transform in the $d-1$ tangential directions. In $d = 3$, put

$$\kappa = \sqrt{|\mathbf{k}|^2 + m^2}, \qquad u = -z_A \geq 0, \qquad v = z_B - \delta \geq 0. \tag{98}$$

The two vacuum cross kernels have the following Laplace decompositions. Here K_0 and K_1 are modified Bessel functions of the second kind.

$$\frac{1}{2\pi} K_0(\kappa(u+v+\delta)) = \int_1^\infty a_\phi(t) e^{-\kappa t u} e^{-\kappa t v}\, \mathrm{d}t, \tag{99}$$

$$\frac{\kappa K_1(\kappa(u+v+\delta))}{2\pi(u+v+\delta)} = \int_1^\infty a_\pi(t) e^{-\kappa t u} e^{-\kappa t v}\, \mathrm{d}t, \tag{100}$$

where

$$a_\phi(t) = \frac{e^{-\kappa\delta t}}{2\pi\sqrt{t^2-1}}, \qquad a_\pi(t) = \frac{\kappa^2 e^{-\kappa\delta t}\sqrt{t^2-1}}{2\pi}. \tag{101}$$

Each rank-one term is completed by equal positive diagonal blocks. The field profile takes opposite boundary values on the two sides. Extending it by zero would create a jump and infinite gradient energy. Instead, continue it through the collar by the minimum-energy solution of $(-\partial_z^2 + \kappa^2)\phi = 0$. Its collar energy, before the Hamiltonian factor $1/2$, is

$$2\kappa \coth(\kappa\delta/2). \tag{102}$$

The momentum profile has only an L^2 energy and needs no derivative matching.

Adding the field, collar, and momentum contributions gives the exact energy of this chosen fiber competitor,

$$C_{\mathrm{fiber}}(\kappa, \delta) = \frac{\kappa}{2\pi}\left[K_1(\kappa\delta) + \coth(\kappa\delta/2) K_0(\kappa\delta)\right]. \tag{103}$$

Define the tangential energy density by Fourier decomposition, equivalently as the energy per area in a large tangential torus followed by the thermodynamic limit. Integration over tangential momentum yields

$$\boxed{e^{\mathrm{flat}}_{\mathrm{comp}}(m, \delta) = \frac{1}{4\pi^2\delta^3}\int_{m\delta}^\infty u^2\left[K_1(u) + \coth(u/2) K_0(u)\right] \mathrm{d}u.} \tag{104}$$

The integral is finite. At $m = 0$, its coefficient is

$$\frac{1}{4\pi^2}\int_0^\infty u^2\,[K_1(u) + \coth(u/2)K_0(u)]\,\mathrm{d}u \simeq 0.11545646. \tag{105}$$

Equation (104) is the energy density of this translation-invariant half-space completion. Its total energy is infinite because the interface is infinite. The heat-kernel construction above supplies the globally normal finite-interface product states used in the theorem. The flat calculation independently displays the same cubic scale and gives an explicit massive coefficient.

8.2 Sharp momentum-block erasure and its curved limit

A related constrained minimization can be solved exactly: the least positive diagonal covariance required to erase the sharp momentum cross block. Its coefficient is distinct from the smooth product completion above and from the unrestricted separability optimum.

Take parallel half-spaces and let P_A, P_B be the sharp spatial projections onto $z < 0$ and $z > \delta$. For mass $m \geq 0$, define the energy-weighted momentum cross operator

$$K_{\pi,\delta}^{E,\sharp} = \frac{1}{2}P_A\varepsilon_m P_B. \tag{106}$$

Because P_B need not preserve $\mathrm{Dom}(\varepsilon_m)$, Eq. (106) denotes the restriction of the off-diagonal distribution kernel, not a literal composition of unbounded operators. It is defined first on $C_c^\infty(B)$. Tangential Fourier transform gives a decomposable operator whose normal fibers are trace class; the full planar trace norm is infinite only because the interface area is infinite. Let R reflect the B half-line onto the A half-line. After tangential Fourier transform, the fiber of $K_{\pi,\delta}^{E,\sharp}R$ is negative. Its negative is the positive Hankel operator with $\kappa = (|\mathbf{k}|^2 + m^2)^{1/2}$ and kernel

$$h_\kappa(u + v + \delta) = \frac{1}{2\pi}\int_\kappa^\infty e^{-t(u+v+\delta)}\sqrt{t^2 - \kappa^2}\,\mathrm{d}t. \tag{107}$$

The factor $1/2$ in Eq. (106) is included in this kernel. Positivity follows because the right side is a positive integral of the rank-one kernels $e^{-tu}e^{-tv}$. Hence its trace norm equals its trace,

$$\|h_\kappa\|_1 = \int_0^\infty h_\kappa(2u + \delta)\,\mathrm{d}u = \frac{1}{4\pi}\int_\kappa^\infty e^{-t\delta}\frac{\sqrt{t^2 - \kappa^2}}{t}\,\mathrm{d}t. \tag{108}$$

The ordinary trace norm is infinite because the interface is infinite. Define its tangential trace density by

$$\mathrm{tr}_\parallel\,|K| = \frac{1}{(2\pi)^{d-1}}\int_{\mathbb{R}^{d-1}}\mathrm{Tr}_{L^2(\mathbb{R}_+)}\,|K(\mathbf{k})|\,\mathrm{d}^{d-1}\mathbf{k}. \tag{109}$$

Rescaling $\mathbf{k} = \boldsymbol{\xi}/\delta$ and $t = s/\delta$ gives the exact scaling form

$$\mathrm{tr}_\parallel\,|K_{\pi,\delta}^{E,\sharp}| = \delta^{-d}c_d(m\delta), \tag{110}$$

$$c_d(\mu) = \frac{1}{4\pi(2\pi)^{d-1}}\int_{\mathbb{R}^{d-1}}\mathrm{d}^{d-1}\boldsymbol{\xi}\int_{\sqrt{|\boldsymbol{\xi}|^2+\mu^2}}^\infty e^{-s}\frac{\sqrt{s^2 - |\boldsymbol{\xi}|^2 - \mu^2}}{s}\,\mathrm{d}s.$$

Interchanging the two positive integrals and evaluating the tangential ball gives the one-dimensional formula

$$\boxed{c_d(\mu) = \frac{1}{2^{d+2}\pi^{d/2}\Gamma((d+2)/2)}\int_\mu^\infty e^{-s}\frac{(s^2 - \mu^2)^{d/2}}{s}\,\mathrm{d}s.} \tag{111}$$

Thus Eq. (110) is exact throughout the scaling regime $m\delta \in [0,\infty)$. The function c_d is continuous and strictly decreasing on $(0,\infty)$: increasing μ both shrinks the integration interval and strictly decreases $(s^2-\mu^2)^{d/2}$ on the surviving interval. At $\mu = 0$, the gamma duplication formula gives

$$c_d(0) = \frac{\Gamma((d+1)/2)}{4d\pi^{(d+1)/2}}. \tag{112}$$

Consequently, for fixed m, or more generally whenever $m\delta \to 0$,

$$\operatorname{tr}_{\|} |K^{E,\sharp}_{\pi,\delta}| = \frac{\Gamma((d+1)/2)}{4d\pi^{(d+1)/2}} \frac{1}{\delta^d} + o(\delta^{-d}). \tag{113}$$

In $d = 3$, the leading coefficient is

$$\frac{1}{12\pi^2\delta^3}. \tag{114}$$

The trace norm has an exact positive-completion meaning. For a Hilbert space $\mathcal{H}$, write $\mathcal{S}_1^+(\mathcal{H})$ for the cone of positive trace-class operators on $\mathcal{H}$. If $K : \mathcal{H}_B \to \mathcal{H}_A$ is trace class, then

$$\inf_{\substack{A\in\mathcal{S}_1^+(\mathcal{H}_A),\quad B\in\mathcal{S}_1^+(\mathcal{H}_B)\\ \left(\begin{smallmatrix} A & -K \\ -K^* & B\end{smallmatrix}\right)\geq 0}} \frac{\operatorname{Tr} A + \operatorname{Tr} B}{2} = \|K\|_1 . \tag{115}$$

Indeed, positivity gives $K = A^{1/2}CB^{1/2}$ with $\|C\| \le 1$, so $\|K\|_1 \le (\operatorname{Tr} A \operatorname{Tr} B)^{1/2} \le (\operatorname{Tr} A + \operatorname{Tr} B)/2$. Equality follows from the polar decomposition $K = U|K|$ by choosing $A = U|K|U^*$ and $B = |K|$. Applied fiberwise, this proves that Eq. (113) is the least positive diagonal noise needed to erase this one sharp momentum cross block.

The flat coefficient is the local model for a curved sharp interface.

Theorem 6 (Sharp mass-collar crossover for momentum completion)**.** *Let S and the collar geometry satisfy the ultrastatic hypotheses of Theorem 1. For this result, assume in addition that the collar is C^∞. Fix an integer $N > 4(d-1)+4$. In every admissible Fermi chart, impose uniform derivative bounds through order $2N+8$ on the Fermi metric, its inverse, and the Fermi Jacobian, and through order $2N+7$ on the second fundamental form. The fixed-distance pieces of the sharp regions have volume at most $C\operatorname{Area}(S)$ and uniform local geometry. A fixed smooth interface satisfies these conditions automatically. Put*

$$L_m = -\Delta_h + V + m^2, \qquad \varepsilon_m = L_m^{1/2}. \tag{116}$$

Fixed curvature coupling is included by absorbing it into the smooth potential V. There is $\delta_0 > 0$, smaller than one quarter of the tubular radius, with the following property. For $0 < \delta < \delta_0$, define the sharp cross operator first on $f \in C_c^\infty(B_\delta)$ by

$$(K^{E,\sharp}_{\pi,\delta,m} f)(x) = \frac{1}{2}\int_{B_\delta} \varepsilon_m(x,y) f(y)\, \mathrm{dvol}_h(y), \qquad x \in A_\delta. \tag{117}$$

Positive separation gives a unique trace-class extension $L^2(B_\delta) \to L^2(A_\delta)$. It is denoted schematically by $\frac{1}{2}P_{A_\delta}\varepsilon_m P_{B_\delta}$, without asserting that the sharp projection preserves $\operatorname{Dom}(\varepsilon_m)$. Writing $K = K^{E,\sharp}_{\pi,\delta,m}$, its least positive diagonal completion cost is

$$\mathcal{C}^\sharp_{\pi,+}(\delta,m) = \inf_{\substack{A\in\mathcal{S}_1^+(L^2(A_\delta)),\quad B\in\mathcal{S}_1^+(L^2(B_\delta))\\ \left(\begin{smallmatrix} A & -K \\ -K^* & B\end{smallmatrix}\right)\geq 0}} \frac{\operatorname{Tr} A + \operatorname{Tr} B}{2} = \left\| K^{E,\sharp}_{\pi,\delta,m} \right\|_1 . \tag{118}$$

For every $M < \infty$, there are $C_M > 0$ and $0 < \delta_M \le \delta_0$ such that, whenever $0 < \delta < \delta_M$ and $0 \le m\delta \le M$,

$$\boxed{\left| \mathcal{C}^{\sharp}_{\pi,+}(\delta, m) - \frac{\mathcal{H}^{d-1}(S)}{\delta^d} c_d(m\delta) \right| \le C_M \mathcal{H}^{d-1}(S) \delta^{-(d-1)}.} \tag{119}$$

Consequently, if $m_\delta \delta \to \mu \in [0, \infty)$, then

$$\frac{\delta^d}{\mathcal{H}^{d-1}(S)} \mathcal{C}^{\sharp}_{\pi,+}(\delta, m_\delta) \longrightarrow c_d(\mu). \tag{120}$$

For fixed mass, or whenever $m_\delta \delta \to 0$, the leading coefficient is the universal value in Eq. (112). In $d = 3$ it is $1/(12\pi^2)$. It is independent of the interface curvature, fixed curvature coupling, and smooth fixed potential.

Proof. Equation (115) gives the exact completion identity. The trace-class boundary-symbol estimate in Appendix D compares the curved collar operator, in trace norm, with the positive flat normal Hankel model at every point of S. Its phase-space trace is $\mathcal{H}^{d-1}(S)\delta^{-d} c_d(m\delta)$, and its remainder is the right side of Eq. (119). □

The remainder order in Eq. (119) cannot be improved in general. For two concentric spherical shells in flat three-dimensional space, with the area measured on the inner face of radius R, the Gegenbauer expansion of the massless cross kernel $1/(2\pi^2|x-y|^4)$ gives

$$\|K_{\delta,0}\|_1 = \frac{R^2}{3\pi\delta^3} + \frac{R}{3\pi\delta^2} + O_R\left(\delta^{-1}\right). \tag{121}$$

The first term is $4\pi R^2 \delta^{-3}/(12\pi^2)$. The second is the change of interface area across the collar, since $4\pi(R+\delta/2)^2 = 4\pi R^2(1+\delta/R) + O(\delta^2)$. The order $\delta^{-(d-1)}$ is therefore attained by a curved interface, while for a product interface with vanishing second fundamental form the area does not change across the collar. On $S^2_R \times \mathbb{R}$, for example, the exact heat-trace formula gives a remainder of order δ^{-1}, one power smaller.

Theorem 6 exactly solves positive diagonal momentum-covariance erasure, and $1/(12\pi^2)$ is the sharp coefficient of that restricted problem only. The planar value in Eq. (105) is the energy density of one explicit product completion, an achievable value rather than an optimum. The unrestricted separability problem is broader, because a competing state's covariance difference from the vacuum need not be positive. Its leading coefficient is unknown. Theorem 1 settles the power δ^{-d} and the existence of matching constants, not their values, and the sharp unrestricted coefficient remains a separate interface-tension problem.

9 What the law means

9.1 Separability is an ultraviolet preparation condition

At every fixed $\delta > 0$, exact product states exist with finite energy. The divergence concerns the limit $\delta \downarrow 0$, not existence at a fixed positive collar. The number of independent interface cells grows like area divided by δ^{d-1}. Neutralizing the witness in one cell forces an energy charge of at least c/δ. The heat completion supplies the matching global upper law. The cell count and the per-cell charge are independent parts of the lower proof, and their product gives the sharp power.

For $d = 3$, restoring constants gives the scale

$$E_{\mathrm{Sep}}(\delta) \asymp \hbar c \frac{\mathrm{Area}(S)}{\delta^3}. \tag{122}$$

The constants hidden in $\asymp$ depend on the selected witness and on uniform geometric bounds. The power does not.

The robust form shows that this law is not created by demanding exact factorization. For every fixed $\epsilon < \gamma_M$, a state that comes within state-norm distance ϵ of the PPT or separable set pays the same $\mathrm{Area}(S)/\delta^d$ power. Conversely, any family satisfying

$$\frac{\delta^d \rho_\delta(H_m)}{\mathcal{H}^{d-1}(S)} \longrightarrow 0 \tag{123}$$

remains a fixed positive state-norm distance from every PPT and separable state as the collar closes. The statement concerns the physical regional algebras. It does not declare abstract tensor products or finite-dimensional product density matrices inconsistent.

9.2 The local statement is stronger than the global divergence

The tolerant count law Eq. (50) is the operational core. For a fixed finite-energy state, its right side tends to zero. The state therefore violates the same pulled-back PPT witness by more than any fixed margin $\epsilon < \gamma_M$ in every sufficiently small cell of the chosen full-density interface family. Each violation uses only finitely many Weyl observables in two spacelike separated cells. The local CP maps in the construction convert the violation into a two-qubit NPT output, so the certified correlation is one-copy distillable.

Dividing the count law by the number of available cells gives an operational fraction bound, measurable once the local witness has been calibrated:

$$\boxed{\frac{N_{\mathrm{wit}}^{(\epsilon)}(\delta)}{N_\delta} \le \frac{4 B_M \alpha_0}{c_{\mathrm{pack}} (\gamma_M - \epsilon)^2} \frac{\delta^d \rho(H_m)}{\mathcal{H}^{d-1}(S)}.} \tag{124}$$

Here $N_{\mathrm{wit}}^{(\epsilon)}(\delta) = \#\{k : \rho(X_{\delta,k}) \ge -\epsilon\}$. In three spatial dimensions, the right side is proportional to $\delta^3 \rho(H_m)$ divided by the interface area. A preparation of fixed energy can therefore bring only a vanishing fraction of the witness array within the chosen PPT tolerance. Every cell outside that fraction has $D_{\mathrm{PPT},k}(\rho;\delta) > \epsilon$ and is certified NPT by the same finite measurement protocol.

The conclusion is not obtained from a nonzero two-point function alone. A separable mixed state may have connected correlations. The witness crosses the convex PPT boundary, which is why it proves nonseparability.

9.3 A nonflat example

Consider the Einstein static spacetime

$$\mathbb{R} \times S_a^3, \qquad L = -\Delta_{S_a^3} + a^{-2} + m^2. \tag{125}$$

The curvature term gives a spectral gap even at $m = 0$. Let S be a latitude two-sphere at polar angle χ. Its area is

$$\mathrm{Area}(S) = 4\pi a^2 \sin^2 \chi. \tag{126}$$

The theorem gives, uniformly for $0 \le m \le m_0$,

$$E_{\mathrm{Sep}}(\delta) \asymp \frac{4\pi a^2 \sin^2 \chi}{\delta^3}. \tag{127}$$

This example shows that the mechanism is local to the interface and does not depend on global translation or dilation symmetry.

9.4 Finite free multiplets

The scalar law adds component by component.

Corollary 7 (Finite free multiplet law)**.** *Let $n_{\mathrm{f}} \ge 1$ and consider n_{f} independent real scalar fields with masses $\boldsymbol{m} = (m_1, \dots, m_{n_{\mathrm{f}}})$ and total Hamiltonian*

$$H_{\boldsymbol{m}} = \sum_{a=1}^{n_{\mathrm{f}}} H_{m_a}. \tag{128}$$

Under the geometric hypotheses of Theorem 1, with $d \ge 3$ and $0 < \delta < \delta_M$, if $m_a \delta \le M$ for every component, then

$$\boxed{n_{\mathrm{f}} c_M \frac{\mathcal{H}^{d-1}(S)}{\delta^d} \le E^{(\boldsymbol{m})}_{\mathrm{PPT}}(\delta) \le E^{(\boldsymbol{m})}_{\mathrm{Sep}}(\delta) \le E^{(\boldsymbol{m})}_{\mathrm{prod}}(\delta) \le n_{\mathrm{f}} C_M \frac{\mathcal{H}^{d-1}(S)}{\delta^d}.} \tag{129}$$

A constant positive semidefinite internal mass matrix is included by taking m_a^2 to be its eigenvalues. The same conclusion holds for a finite direct sum of the ultrastatic scalar operators in Eq. (15) when all components belong to one common uniform family.

Proof. If the aggregate regional state is PPT, its restriction ρ_a to the ath field remains normal and PPT. The scalar lower bound applies to every component, and summing $\rho_a(H_{m_a})$ gives the left inequality without assuming independence between species. For the upper bound, take one scalar product competitor per component and regroup their tensor product by spatial region. Energies add. A constant internal mass matrix is diagonalized by a pointwise orthogonal field transformation, which preserves localization and the total Hamiltonian. The ultrastatic direct-sum statement is identical. □

9.5 Finite regulators and the continuum limit

Absolute separability remains a genuine spectral property in finite dimensions [16, 17]. There is no contradiction. A regulator contains only finitely many interface modes. As the cutoff is removed, the theorem says that keeping all of their regional PPT constraints forces the Hamiltonian expectation to follow the ultraviolet cell count. A sequence of regulated separable states can approach a normal continuum state only if it fails the regional PPT condition on an increasing set of ultraviolet cells, or if its energy diverges at the stated rate.

For a finite-mode bosonic regulator, the three energy problems have exact covariance formulations. Finite-mode means finitely many genuine canonical oscillator modes. Their Fock space is still infinite-dimensional. This is not a truncation to finitely many number states, which would destroy the exact canonical commutation relations.

Appendix E gives an exact covariance reduction for these regulated problems. For every finite collection of canonical bosonic modes and every positive quadratic Hamiltonian, the unrestricted PPT, separable, and product energy minima are semidefinite programs.

Gaussianity is a conclusion of the reduction, not an assumption on the optimizing state. The continuum frame theorem fixes the scale these programs must obey under geometric refinement, but it does not prove convergence of their normalized optimal values. That separate limit is the interface-tension problem.

10 Relation to observable field correlations

Electro-optic sampling has measured vacuum-field cross correlations between spacetime sampling volumes, including contributions dominated by spacelike separation [18]. A later experiment resolved different optical quadratures and separated vacuum fluctuations from source radiation [19]. These measurements demonstrate that correlation and causal response can be resolved independently. They motivate an analogous scalar-mode test, but they do not establish nonseparability or test the scalar theorem.

The distinction is visible in the simplest two-mode inequality. For output modes a_A, a_B, write $\Delta a_j = a_j - \langle a_j \rangle$ and define

$$M_c = \langle \Delta a_A \Delta a_B \rangle, \qquad n_A^c = \langle \Delta a_A^\dagger \Delta a_A \rangle, \qquad n_B^c = \langle \Delta a_B^\dagger \Delta a_B \rangle. \tag{130}$$

Every separable two-mode state with finite second moments obeys

$$|M_c|^2 \le n_A^c n_B^c \tag{131}$$

[20, 21]. Cross-correlation and response traces do not determine the two local occupations on the right. The published electro-optic data therefore motivate the regional observable problem but do not already test Eq. (131) or the field witness $X_{\delta,k}$.

The theorem fixes the experimental architecture; the remaining task is a finite calibration for the chosen apparatus. The existence proof passes through Reeh–Schlieder, one-copy distillability, and a finite approximation argument. Once explicit derivative-weighted mode profiles, Weyl-polynomial coefficients, and a certified witness gap have been computed, one can measure the Weyl moments entering each local witness and average the results over uniformly sampled interface cells. This samples $\overline{X}_\delta$ without requiring one collective field measurement. If an experiment obtains an estimate $\widehat{w}$ with additive uncertainty Δw and

$$\widehat{w} + \Delta w < 0, \tag{132}$$

then every PPT state is separated from the measured state by at least

$$D_{\mathrm{PPT}}(\rho; \delta) \ge -(\widehat{w} + \Delta w). \tag{133}$$

The witness therefore certifies a quantitative distance from PPT without full state tomography.

An independently calibrated upper bound $E_{\max}$ on the global mean energy gives a second check. Equation (18) guarantees a negative witness margin whenever

$$\gamma_M > K_M \sqrt{\frac{\delta^d E_{\max}}{\mathcal{H}^{d-1}(S)}}. \tag{134}$$

Repeating the same mode geometry while reducing δ would then test the predicted interface scaling. Producing the explicit witness and calibrated constants γ_M, K_M is the missing practical step. It is a finite witness-construction problem rather than full field tomography. The proposed test concerns correlation and preparation energy; it requires no signal to propagate between the regions.

11 Conclusion and scope

The proof uses four ingredients. The first three are a uniform finite local PPT gap, a positive local excitation counter which pays for neutralizing that gap, and an interface frame inequality

$$\sum_k N_k \le B\delta H. \tag{135}$$

The fourth is a regular interface which supplies at least $c_{\mathrm{pack}}\,\mathrm{Area}(S)\delta^{-(d-1)}$ cells. Any continuum theory with analogues of all four ingredients obeys the same $\mathrm{Area}(S)/\delta^d$ lower law, tolerant cell count, and state-norm exclusion from PPT. In the scalar field, the counter is a finite-mode number operator and Eq. (135) follows from the one-particle Bessel estimate. This is the structural part of the result and the precise target for extensions beyond quasifree fields.

For positive diagonal erasure of the sharp momentum cross block, the boundary-Hankel theorem goes further: it gives the exact crossover function $c_d(m\delta)$ and a curvature-independent fixed-mass leading coefficient. That coefficient belongs to the constrained covariance-completion problem; the unrestricted separability coefficient is the remaining interface-tension problem.

The law concerns global preparation energy in globally normal states of the stated vacuum representation. Positive-collar product states coexist with exact microcausality; the theorem neither identifies nonproductness with entanglement nor converts correlation into causal response. Other superselection sectors and states not normal to that representation require separate analysis.

The curved theorem is stated for ultrastatic scalar fields with a controlled positive one-particle operator. A general nonstationary spacetime has no preferred global Hamiltonian, so the quantity optimized here is not defined without further physical structure. Interacting theories do not generally possess the Fock identity used to second-quantize the frame bound. The abstract criterion Eq. (135) isolates the exact replacement an interacting proof must supply.

Exact product states exist at every positive collar. The global mean-energy infima over the larger PPT class, the separable class, and the product class all have the same sharp interface exponent. The lower bound is robust under fixed-tolerance approximation. A state that enters any fixed state-norm neighborhood of PPT below the witness gap pays the same power, while a bounded-energy state remains a fixed operational distance outside PPT as the collar closes. In three spatial dimensions, the continuum price is area divided by resolution cubed.

A Finite PPT witnesses in finite Weyl mode spaces

This appendix proves Proposition 3. The proof has four logically distinct parts: the derivative subnet retains the Reeh–Schlieder property; that property supplies a two-qubit NPT output; the corresponding field witness can be kept completely positive while it is approximated by finite Weyl polynomials; and the resulting finite data are uniform on compact mass intervals and controlled curved families.

A.1 The derivative subnet retains the required local structure

Let $\mathcal{D}_r(O)$ be the real symplectic space of Cauchy data in Eq. (23), and let $\mathcal{A}^{(r)}(O)$ be the unital Weyl C^*-algebra which it generates in the vacuum representation. Locality

and spatial translation covariance are immediate because $(-\Delta)^r$ preserves support and commutes with spatial translations. Covariance under all spacetime translations also holds. The Klein–Gordon Cauchy evolution is a matrix of functions of $\varepsilon_m = (-\Delta+m^2)^{1/2}$, so it commutes with $(-\Delta)^r$. Finite propagation speed carries the underlying compactly supported data into the causal translate of the original region. Thus a time-translated derivative datum is again a derivative datum in the translated local region. The subnet consequently inherits the full positive-energy translation representation and its spectrum condition.

The subnet is nonabelian. Choose a nonzero $h \in C_c^\infty(O \cap \Sigma)$. The two data $F = ((-\Delta)^r h, 0)$ and $G = (0, (-\Delta)^r h)$ belong to $\mathcal{D}_r(O)$ and have

$$\sigma(F, G) = \int_{\mathbb{R}^d} |(-\Delta)^r h(x)|^2 \, \mathrm{d}^d x > 0. \tag{136}$$

Their Weyl operators therefore do not commute for a suitable real rescaling.

It remains to check weak additivity without moving an unbounded multiplier through a closure. Choose one nonzero $h \in C_c^\infty(O \cap \Sigma)$ and let e_h be the one-particle vector of the derivative datum $((-\Delta)^r h, 0)$. Its momentum wavefunction is a nonzero constant times

$$|\mathbf{p}|^{2r}(|\mathbf{p}|^2 + m^2)^{-1/4}\widehat{h}(\mathbf{p}). \tag{137}$$

The Fourier transform of a nonzero compactly supported smooth function is real analytic and nonzero almost everywhere; the remaining multiplier vanishes only at $\mathbf{p} = 0$. Hence $e_h(\mathbf{p}) \neq 0$ almost everywhere. If ψ is orthogonal to every spatial translate of e_h, then the Fourier transform of the L^1 function $\overline{\psi}(\mathbf{p})e_h(\mathbf{p})$ vanishes identically. Thus $\overline{\psi}e_h = 0$ almost everywhere and $\psi = 0$. The translates of this single derivative vector, and therefore those of the derivative local spaces, have dense span. This is density of the complex span. Weyl operators take only real multiples of their arguments, and for the Weyl algebra generated by a real subspace K of one-particle vectors the vacuum is cyclic exactly when $K + iK$ is dense [22]. The complex span of the translates is $K + iK$ for their real span K, so this condition holds for the translated derivative spaces. The standard positive-energy analytic-continuation argument for Reeh–Schlieder now applies to this subnet: the vacuum is cyclic for $\mathcal{A}^{(r)}(O)$ whenever O is nonempty [23, Chap. II.5]. Notice that the argument uses translations, locality, weak additivity, and the spectrum condition. It does not assign a Lorentz-scalar transformation law to the differentiated generator.

A.2 From Reeh–Schlieder to a bounded PPT witness

Put $\mathcal{A}_A^{(r)} = \mathcal{A}^{(r)}(O_A)$ and $\mathcal{A}_B^{(r)} = \mathcal{A}^{(r)}(O_B)$. These are commuting, unital, nonabelian algebras, and the vacuum is cyclic for the first. The hypotheses of Theorem 6.2 of Verch and Werner are therefore all satisfied. In the Heisenberg picture, their theorem supplies completely positive maps

$$\Theta_A : M_2 \longrightarrow \mathcal{A}_A^{(r)}, \qquad \Theta_B : M_2 \longrightarrow \mathcal{A}_B^{(r)} \tag{138}$$

such that the two-qubit functional

$$\omega_2(C \otimes D) = \omega(\Theta_A(C)\Theta_B(D)) \tag{139}$$

is NPT [7, Def. 5.2 and Thm. 6.2]. Since $M_2 \otimes M_2$ is finite dimensional, there is a positive matrix $Y \geq 0$ for which

$$\omega_2(Y^{\top_A}) = -\beta < 0. \tag{140}$$

Expand $Y^{\top_A}$ in matrix units and replace every elementary tensor by the product of the two corresponding CP-map entries. The resulting bounded field operator will be denoted by $X^{(0)}$.

If η is algebraically PPT on $(\mathcal{A}_A^{(r)}, \mathcal{A}_B^{(r)})$, then the functional obtained from η by Eq. (139) is PPT. This is precisely the stability of the Verch–Werner PPT inequalities under local completely positive maps. Hence

$$\eta(X^{(0)}) \geq 0, \qquad \omega(X^{(0)}) = -\beta. \tag{141}$$

After division by $\max\{1, \|X^{(0)}\|\}$, this already gives a bounded selfadjoint PPT witness. What remains is to make its local content finite in Weyl modes without losing either inequality.

A.3 Finite Weyl approximation with exact complete positivity

The point of approximating Choi square roots, rather than the entries of the maps independently, is that complete positivity then survives at every finite stage. If E_{ij} are the standard matrix units, a positive Choi matrix admits a finite factorization of the form

$$\Theta_A(E_{ij}) = \sum_{\nu=1}^{r_A} a_{\nu i}^* a_{\nu j}. \tag{142}$$

The maps supplied by Theorem 6.2 take values in the chosen local C^*-algebras. Thus their Choi matrices $C_A = [\Theta_A(E_{ij})]$ and $C_B = [\Theta_B(E_{ij})]$ are positive elements of $M_2(\mathcal{A}_A^{(r)})$ and $M_2(\mathcal{A}_B^{(r)})$. Take their positive square roots $R_A = C_A^{1/2}$ and $R_B = C_B^{1/2}$. Finite derivative-Weyl polynomials are norm dense in the two Weyl C^*-algebras, so matrices $R_{A,n}, R_{B,n}$ with finite Weyl-polynomial entries can be chosen to converge in norm. Set

$$C_{A,n} = R_{A,n}^* R_{A,n}, \qquad C_{B,n} = R_{B,n}^* R_{B,n}. \tag{143}$$

These are positive Choi matrices and hence define CP maps $\Theta_{A,n}, \Theta_{B,n}$ exactly. Their norm convergence and the finite expansion of $Y^{\top_A}$ give

$$\left\| X_n - X^{(0)} \right\| \longrightarrow 0, \qquad \omega(X_n) \longrightarrow -\beta. \tag{144}$$

Every X_n is nonnegative on every algebraically PPT state because it is built from CP maps. It is selfadjoint because $Y^{\top_A}$ is selfadjoint and the product map is star preserving. Choose n with $\omega(X_n) \leq -\beta/2$ and divide by $\max\{1, \|X_n\|\}$. This proves Eq. (24) at one mass. It also explains why an entrywise approximation would be insufficient: small errors in Choi entries need not preserve positivity, whereas $R_n^* R_n$ does.

Every finite Weyl polynomial acts on the Fock space over the finite complex span E of its Weyl arguments. This is the finite mode space used in Section 4; no truncation in occupation number is made.

A.4 Uniformity in mass and in curved tangent cells

The time-zero symplectic form, and therefore the abstract Weyl relations, are independent of the mass. For a fixed finite derivative-Weyl polynomial, its vacuum expectation is a finite combination of phases and exponentials of one-particle Gram entries. Those entries are continuous at every $\mu \in [0, M]$. At $\mu = 0$, the powers $|\mathbf{p}|^{2r}$ in Eq. (23) provide an integrable mass-independent majorant for the inverse-energy factors.

At a chosen mass μ_0, perturb the finitely many derivative test data in the Choi square-root polynomials, if necessary, so that the resulting one-particle vectors are complex-linearly independent. Rebuild the Choi matrices as $R_A^* R_A$ and $R_B^* R_B$ after the perturbation; complete positivity and PPT nonnegativity therefore remain exact. Weyl generators are not norm continuous in their arguments. What is used here is their strong continuity in the regular vacuum representation: finite vacuum matrix elements converge as the test data converge. The finite coefficient ℓ^1 norms give a common operator-norm bound for normalization. The strict negative vacuum expectation consequently survives a sufficiently small perturbation without any claim of norm continuity for $W(F)$.

The perturbed Gram determinant is nonzero at μ_0. Continuity of its finite entries then gives an open mass neighborhood U_{μ_0} on which both the negative gap and a lower Gram-eigenvalue bound persist. Finitely many such neighborhoods cover $[0, M]$. Taking the minimum gap and the maximum of the finitely many inverse-Gram constants gives γ_M and G_M in Eq. (25). The uniform Sobolev estimate in Section 5 then produces B_M. For a mass in an overlap, choose either admissible member of the finite witness family.

Lemma 8 (Uniform finite-matrix tangent limit)**.** *Fix finitely many flat derivative Cauchy data and rescale them into Fermi cells by Eqs.* (59) *and* (60)*. On the controlled ultrastatic families in the main theorem, their finite symplectic matrices and ground-state symmetric covariance matrices converge uniformly, as $\delta \downarrow 0$, to the corresponding flat massless matrices.*

Proof. The symplectic statement follows by change of variables, the $J_{k,\delta}^{-1/2}$ factors, and convergence of the pulled-back metric coefficients. For the symmetric covariance, define the density isometry from the reference chart to the Fermi cell by

$$(V_{k,\delta}\varphi)(\Phi_{k,\delta}(y)) = \delta^{-d/2} J_{k,\delta}(y)^{-1/2} \varphi(y). \tag{145}$$

If $u_{k,i,\delta} = (-\Delta_h)^r f_{k,i,\delta}$, $v_{k,i,\delta} = (-\Delta_h)^r g_{k,i,\delta}$, $a_i = (-\Delta)^r f_{0,i}$, and $b_i = (-\Delta)^r g_{0,i}$, then

$$V_{k,\delta}^* u_{k,i,\delta} = \delta^{-1/2}(a_i + O(\delta)), \qquad V_{k,\delta}^* v_{k,i,\delta} = \delta^{1/2}(b_i + O(\delta)). \tag{146}$$

The remainders are uniform in every fixed Sobolev norm required by the finite profile family.

Insert nested cutoffs around the profile supports. Uniform complex-power calculus decomposes the doubly localized $L^{\pm 1/2}$ into a properly supported pseudodifferential part and a smooth remainder [24, 25, Sec. 5]. Write $|D| = (-\Delta)^{1/2}$ on the flat reference chart. On the finite profile space, the normalized operators satisfy

$$\delta^{-1} V_{k,\delta}^* L^{-1/2} V_{k,\delta} \longrightarrow |D|^{-1}, \qquad \delta V_{k,\delta}^* L^{1/2} V_{k,\delta} \longrightarrow |D|. \tag{147}$$

Indeed, after $\xi \mapsto \xi/\delta$, the normalized principal symbols are $|\xi|^{\pm 1}$. Coefficient variation and every lower symbol order gain a positive power of δ. For the negative power, the derivative profiles contribute $|\xi|^{2r}$ at zero frequency and give an integrable majorant for $|\xi|^{-1}$. After moving the derivatives onto the complementary smooth kernel, its uu and vv matrix elements are respectively $O(\delta^{4r+d-1})$ and $O(\delta^{4r+d+1})$. All seminorms are uniform under the metric, potential, reach, and gap bounds of the theorem. Equations (146) and (147) cancel their powers of δ pairwise and prove the finite-matrix convergence. At each fixed point the result agrees with the Hadamard scaling limit of Sahlmann and Verch [10, Sec. 5.3]. □

For the ultrastatic statement, fix the flat massless member of the witness family, write its vacuum gap as $\gamma_0 = -\omega_0(X) > 0$, and insert its data into each Fermi cell. The finite Choi square formulas are algebraic formulas in local Weyl generators. They therefore define CP maps, and hence a PPT-nonnegative witness, for the curved CCR without comparing curved and flat Hilbert spaces. Finite Weyl expectations are continuous functions of the finite symplectic and symmetric covariance matrices. Lemma 8 transports the strict flat negative expectation uniformly. For sufficiently small δ, every curved expectation is at most half the flat negative gap. To close the norm estimate, write the fixed finite polynomial as $X = \sum_{\ell=1}^{L} c_\ell W(F_\ell)$ after combining repeated Weyl terms and put

$$C_X = \sum_{\ell=1}^{L} |c_\ell|, \qquad \widehat{X}_{\delta,k} = \frac{X_{\delta,k}}{\max\{1, C_X\}}. \tag{148}$$

Every Weyl generator is unitary in every CCR representation, so $\left\|\widehat{X}_{\delta,k}\right\| \le 1$ uniformly. Positive scalar rescaling preserves PPT positivity, while the transported gap becomes $\gamma_0/(2\max\{1, C_X\}) > 0$. Relabeling $\widehat{X}_{\delta,k}$ as $X_{\delta,k}$ proves the final assertion of Proposition 3 and supplies the contractions used in Theorem 1.

B Proof of the frame inequalities

B.1 Flat space

Let $G_{ab}^{(\mu)}(\mathbf{x})$ be the translation Gram kernel

$$G_{ab}^{(\mu)}(\mathbf{x}) = \langle q_{a,\mu}, \tau(\mathbf{x}) q_{b,\mu} \rangle = \int_{\mathbb{R}^d} e^{-i\mathbf{p}\cdot\mathbf{x}} \overline{q}_{a,\mu}(\mathbf{p}) q_{b,\mu}(\mathbf{p}) \, \mathrm{d}^d p. \tag{149}$$

Choose r so that the number of uniformly integrable derivatives of every product $\overline{q}_{a,\mu} q_{b,\mu}$ exceeds $d-1$. The potentially most singular product multiplier obeys

$$\partial^\alpha \left(\frac{|\mathbf{p}|^{4r}}{|\mathbf{p}|^2 + \mu^2} \right) = O\Big(\mu^{4r-2-|\alpha|}\Big) \tag{150}$$

on $|\mathbf{p}| \lesssim \mu$. The volume element supplies a further factor μ^d, so the exponent is positive for the derivatives used. On $|\mathbf{p}| \ge 2\mu$, the expression is bounded by a fixed symbol of order $4r-2$. Multiplication by the Schwartz test functions supplies an integrable majorant. Terms containing one or no inverse-energy factor are less singular. This proves Eq. (35).

Let the point set obey

$$\#\{y \in \Lambda : n \le |x-y| < n+1\} \le C_\Lambda (1+n)^{d-1}. \tag{151}$$

This estimate is automatic for a relatively separated subset of $\mathbb{R}^d$. Together with the $(1+n)^{-(d+1)}$ Gram decay, it leaves the summable row series $\sum_n O((1+n)^{-2})$. Thus

$$\sup_{x\in\Lambda} \sum_{y\in\Lambda} \sum_{a,b=1}^{n_{\mathsf{E}}} |G_{ab}^{(\mu)}(x-y)| < \infty. \tag{152}$$

Schur's test bounds the Gram operator on $\ell^2(\Lambda \times \{1, \dots, n_{\mathsf{E}}\})$. If

$$Qc = \sum_{x,a} c_{x,a} \tau(x) q_{a,\mu}, \tag{153}$$

this says $QQ^* \le B_M \mathbf{1}$, which is Eq. (37). Finally,

$$\sum_{x\in\Lambda} P_{\tau(x)\mathsf{E}_\mu} = \varepsilon_\mu^{1/2} Q Q^* \varepsilon_\mu^{1/2} \le B_M \varepsilon_\mu \tag{154}$$

on the form domain of $\varepsilon_\mu^{1/2}$.

For positive one-particle forms $A \le B$, second quantization gives $\mathrm{d}\Gamma(A) \le \mathrm{d}\Gamma(B)$ on every n-particle sector and hence on Fock space. This proves Eq. (39) without a domain interchange.

B.2 Bounded geometry

We prove Proposition 4 in two steps. Let $L = -\Delta_h + V \ge \nu_0^2$ on a bounded-geometry manifold and $\varepsilon = L^{1/2}$.

B.2.1 Finite-profile tangent covariance

Fractional powers of L are nonlocal, so they cannot be replaced by powers of a chart-restricted differential operator. Only finitely many covariance pairings are needed here. Put

$$u_{k,i,\delta} = (-\Delta_h)^r f_{k,i,\delta}, \qquad v_{k,i,\delta} = (-\Delta_h)^r g_{k,i,\delta}, \tag{155}$$

and let $a_i = (-\Delta)^r f_{0,i}$ and $b_i = (-\Delta)^r g_{0,i}$ be the corresponding flat profiles.

Lemma 8 specializes to

$$\langle u_{k,i,\delta}, L^{-1/2} u_{k,j,\delta}\rangle \longrightarrow \langle a_i, |D|^{-1} a_j\rangle, \tag{156}$$

$$\langle v_{k,i,\delta}, L^{1/2} v_{k,j,\delta}\rangle \longrightarrow \langle b_i, |D| b_j\rangle, \tag{157}$$

$$\langle u_{k,i,\delta}, v_{k,j,\delta}\rangle \longrightarrow \langle a_i, b_j\rangle. \tag{158}$$

The one-particle Gram matrix is written directly as

$$\begin{aligned} G_{k,\delta;ij} &= \langle s_{k,i,\delta}, s_{k,j,\delta}\rangle \\ &= \frac{1}{2}\big(\langle u_{k,i,\delta}, L^{-1/2} u_{k,j,\delta}\rangle + \langle v_{k,i,\delta}, L^{1/2} v_{k,j,\delta}\rangle \\ &\qquad + i\langle u_{k,i,\delta}, v_{k,j,\delta}\rangle - i\langle v_{k,i,\delta}, u_{k,j,\delta}\rangle\big). \end{aligned} \tag{159}$$

Lemma 8 makes this matrix converge uniformly to the nondegenerate flat massless Gram matrix selected in Appendix A. Hence, after reducing δ_0, there are uniform constants $0 < g_- < C$ such that

$$g_- \mathbf{1} \le G_{k,\delta} \le C \mathbf{1}. \tag{160}$$

This proves the first part of Proposition 4 and shows that multiplication by $G_{k,\delta}^{-1/2}$ in the subsequent orthonormalization is uniformly bounded.

B.2.2 Direct synthesis and the form inequality

Before orthonormalization, the energy-divided profiles are

$$\widetilde{q}_{k,i,\delta} = \varepsilon^{-1/2} s_{k,i,\delta} = \frac{1}{\sqrt{2}} \left(L^{-1/2} (-\Delta_h)^r f_{k,i,\delta} + i (-\Delta_h)^r g_{k,i,\delta} \right). \tag{161}$$

The separated cell centers and fixed compact reference supports imply that all rescaled primitive supports have uniformly bounded overlap.

Let $\mathbf{c} = (c_{k,i})$ be a finitely supported coefficient array and set

$$F_{\mathbf{c}} = \sum_{k,i} c_{k,i} f_{k,i,\delta}, \qquad G_{\mathbf{c}} = \sum_{k,i} c_{k,i} g_{k,i,\delta}. \tag{162}$$

The two rescalings in Eqs. (59) and (60) give

$$\left\| \nabla^j f_{k,i,\delta} \right\|_2 \le C\delta^{2r-1/2-j}, \qquad 0 \le j \le 2r-1, \tag{163}$$

$$\left\| \nabla^j g_{k,i,\delta} \right\|_2 \le C\delta^{2r+1/2-j}, \qquad 0 \le j \le 2r. \tag{164}$$

Bounded overlap turns these cellwise estimates into the simultaneous bounds

$$\|F_{\mathbf{c}}\|^2_{H^{2r-1}} + \|G_{\mathbf{c}}\|^2_{H^{2r}} \le C\delta \sum_{k,i} |c_{k,i}|^2. \tag{165}$$

This is the full-density step: its constant is independent of the number of interface cells.

Uniform coercive elliptic regularity gives $L : H^1(\Sigma) \to H^{-1}(\Sigma)$ as a uniformly coercive isomorphism. Put $w = (-\Delta_h)^r F_{\mathbf{c}}$. The identity below uses the H^{-1}, H^1 dual pairing:

$$\begin{aligned} \left\| L^{-1/2} w \right\|_2^2 &= \langle w, L^{-1} w \rangle_{H^{-1},H^1} \\ &\le C \|w\|^2_{H^{-1}} \le C \|F_{\mathbf{c}}\|^2_{H^{2r-1}}, \end{aligned} \tag{166}$$

$$\|(-\Delta_h)^r G_{\mathbf{c}}\|_2^2 \le C \|G_{\mathbf{c}}\|^2_{H^{2r}}. \tag{167}$$

Equations (165)– (167) show that the raw synthesis operator

$$\widetilde{\mathcal{T}}_\delta \mathbf{c} = \sum_{k,i} c_{k,i} \widetilde{q}_{k,i,\delta} \tag{168}$$

satisfies $\left\| \widetilde{\mathcal{T}}_\delta \right\|^2 \le C\delta$.

Let $\mathcal{R}_\delta$ be block diagonal on the coefficient space, with block $G_{k,\delta}^{-1/2}$ in cell k. Explicitly,

$$e_{k,a,\delta} = \sum_i s_{k,i,\delta} (G_{k,\delta}^{-1/2})_{ia}, \qquad q_{k,a,\delta} = \sum_i \widetilde{q}_{k,i,\delta} (G_{k,\delta}^{-1/2})_{ia}. \tag{169}$$

Equation (160) gives $\|\mathcal{R}_\delta\| \le g_-^{-1/2}$. Orthonormalization within each cell therefore changes the synthesis operator to

$$\mathcal{T}_\delta = \widetilde{\mathcal{T}}_\delta \mathcal{R}_\delta, \qquad \mathcal{T}_\delta \mathcal{T}_\delta^* \le B\delta \mathbf{1}. \tag{170}$$

For $\psi \in \mathrm{Dom}(\varepsilon^{1/2})$,

$$\begin{aligned} \left\langle \psi, \sum_k P_{\mathsf{E}_{k,\delta}} \psi \right\rangle &= \sum_{k,a} |\langle e_{k,a,\delta}, \psi \rangle|^2 \\ &= \sum_{k,a} |\langle q_{k,a,\delta}, \varepsilon^{1/2} \psi \rangle|^2 \\ &= \langle \varepsilon^{1/2} \psi, \mathcal{T}_\delta \mathcal{T}_\delta^* \varepsilon^{1/2} \psi \rangle. \end{aligned} \tag{171}$$

Consequently,

$$\sum_k P_{\mathsf{E}_{k,\delta}} = \varepsilon^{1/2} \mathcal{T}_\delta \mathcal{T}_\delta^* \varepsilon^{1/2} \le B\delta\varepsilon \tag{172}$$

as quadratic forms on $\mathrm{Dom}(\varepsilon^{1/2})$. Applying this inequality sector by sector gives

$$\sum_k \mathrm{d}\Gamma(P_{\mathsf{E}_{k,\delta}}) \le B\delta H. \tag{173}$$

This proves Eq. (63) without a global shell estimate. The uniform tangent-gap statement needed to combine the frame with the PPT witness was proved separately in Appendix A.

C The heat-kernel product completion

C.1 Balanced heat estimates

We prove Eqs. (82) and (83). For $0 < t \le 1$, the bounded-geometry heat calculus gives, with one fixed smaller Gaussian exponent for all j, k,

$$|\nabla_x^j \nabla_z^k p_t(x,z)| \le C_{j,k} t^{-(d+j+k)/2} \exp\left[-\frac{\operatorname{dist}(x,z)^2}{ct}\right]. \tag{174}$$

The constants are uniform on the bounded coefficient families in the theorem [11, Thms. 6.1 and 6.5] [12, Thm. 4.1]. The cited zero-order theorem permits uniformly bounded lower-order terms. In normal charts, scaled interior parabolic estimates extend the derivative bounds to the real potential V, with constants depending on the displayed bounds for its derivatives. One common exponent may be used for all j, k. Apply the interior estimate on $B(x, \sqrt{t})$ at a slightly earlier heat time. The zeroth-order Gaussian bound on that ball retains a fixed fraction of $\operatorname{dist}(x,z)^2/t$, while only the power prefactor depends on j, k. Symmetry gives the same argument in the second variable. Use the support distances $a(z)$ and $b(z)$ from Eq. (79). Combining Eq. (174) with Eq. (80) gives

$$r(t,z)^2 e^{-\eta a(z)^2/t} = r(t,z)^{-2} e^{-\eta b(z)^2/t} = e^{-\eta(a(z)^2+b(z)^2)/(2t)}, \tag{175}$$

and therefore

$$\sum_{J=A,B} \left\| \widetilde{u}^J_{t,z} \right\|_2^2 \le C t^{-d/2} \exp\left[-\frac{\eta}{2t}(a(z)^2 + b(z)^2)\right]. \tag{176}$$

Inside the tubular collar, write $z = (s, u)$. The support conditions imply

$$a(z)^2 + b(z)^2 \ge c_2 \left[(u - \delta/2)^2 + \delta^2\right]. \tag{177}$$

The normal integral is at most $C\sqrt{t}\, e^{-c\delta^2/t}$. The Fermi Jacobian turns the tangential integral into $C\mathcal{H}^{d-1}(S)$. Hence

$$\int_\Sigma \sum_{J=A,B} \left\| \widetilde{u}^J_{t,z} \right\|_2^2 \operatorname{dvol}_h(z) \le C\mathcal{H}^{d-1}(S) t^{-(d-1)/2} e^{-c\delta^2/t} + R_0(t). \tag{178}$$

Outside the tube, the two compact cutoff supports cannot be simultaneously close. Gaussian integration gives a fixed-distance remainder R_0 which is independent of δ and integrable against $t^{-3/2}\mathrm{d}t$ after the large-t estimate is included.

The quadratic form of L is bounded by the H^1 norm on each bounded geometry chart. Leibniz' rule, Eq. (174), and $|\nabla \chi_J| \le C\delta^{-1}$ give

$$\int_\Sigma \sum_{J=A,B} \left\| L^{1/2} \widetilde{u}^J_{t,z} \right\|_2^2 \operatorname{dvol}_h(z) \le C(t^{-1} + \delta^{-2} + 1)\mathcal{H}^{d-1}(S) t^{-(d-1)/2} e^{-c\delta^2/t} + R_1(t). \tag{179}$$

The fixed-distance remainder is integrable against $t^{-1/2}\mathrm{d}t$. For the bounded interface families defined in Section 2, the fixed-distance support and cutoff bounds give the estimate

$$\int_0^1 t^{-3/2} R_0(t)\, \mathrm{d}t + \int_0^1 t^{-1/2} R_1(t)\, \mathrm{d}t + R_\infty \le C\mathcal{H}^{d-1}(S)(1 + \delta^{-1}), \tag{180}$$

where R_∞ denotes the corresponding large-time weighted integrals. The factor δ^{-1} comes from the large-time energy integral. For $t \ge 1$ the balancing factor equals one, so the collar

gradients contribute $\sum_J \|\nabla\chi_J\|_2^2 \le C\mathcal{H}^{d-1}(S)\delta^{-1}$ against the integrable weight $t^{-(d+1)/2}$, $d \ge 3$. Since $\delta^{-1} \le \delta_M^{d-1}\delta^{-d}$ for $0 < \delta < \delta_M$, this remainder is absorbed into the order of Eq. (74).

Because $r(t,z)$ is independent of the kernel variables, no derivative of it appears in the local energy form. The added field and momentum covariances are the monotone positive integrals

$$Q_{\phi\phi,\delta} = \frac{1}{2\sqrt{\pi}} \int_0^\infty t^{-1/2} \int_\Sigma (\widetilde{u}^A_{t,z} - \widetilde{u}^B_{t,z})^{\otimes 2} \,\mathrm{dvol}_h(z)\mathrm{d}t, \tag{181}$$

$$Q_{\pi\pi,\delta} = \frac{1}{4\sqrt{\pi}} \int_0^\infty t^{-3/2} \int_\Sigma (\widetilde{u}^A_{t,z} + \widetilde{u}^B_{t,z})^{\otimes 2} \,\mathrm{dvol}_h(z)\mathrm{d}t. \tag{182}$$

For $t \ge 1$, the balancing factor is one. The spectral gap gives exponential large-time decay in the curved theory. In massless Minkowski space, the explicit Euclidean heat kernel and compact cutoff supports give integrable polynomial decay. Equations (178) and (179), followed by the gamma integrals in Section 7, prove Eq. (74).

Differentiation in either kernel variable adds only powers of $t^{-1/2}$. Equation (177) retains $e^{-c\delta^2/t}$, so every differentiated small-time integral converges at fixed $\delta > 0$. The completed Cauchy-data kernel is smooth. The spacetime two-point correction obtained by evolving either entry is therefore a smooth bisolution. Its wavefront set is empty, so adding it to the vacuum two-point function preserves the Hadamard condition.

C.2 The trace-class operator integral

We now complete the operator argument behind Proposition 5. Regard the complex one-particle Hilbert space as a real Hilbert space and use the displacement map Z in Eq. (88). The field completion in Eq. (181) and the momentum completion in Eq. (182) induce the positive operator

$$\begin{aligned} T_\delta = {} & \frac{1}{2\sqrt{\pi}} \int_0^\infty t^{-1/2} \int_\Sigma |Z(\widetilde{u}^A_{t,z} - \widetilde{u}^B_{t,z}, 0)\rangle\langle Z(\widetilde{u}^A_{t,z} - \widetilde{u}^B_{t,z}, 0)|_{\mathbb{R}} \,\mathrm{dvol}_h(z)\mathrm{d}t \\ & + \frac{1}{4\sqrt{\pi}} \int_0^\infty t^{-3/2} \int_\Sigma |Z(0, \widetilde{u}^A_{t,z} + \widetilde{u}^B_{t,z})\rangle\langle Z(0, \widetilde{u}^A_{t,z} + \widetilde{u}^B_{t,z})|_{\mathbb{R}} \,\mathrm{dvol}_h(z)\mathrm{d}t. \end{aligned} \tag{183}$$

The notation $Z(u_\phi, u_\pi)$ makes the two components of the classical mean configuration explicit.

For a gapped theory,

$$\|Z(u)\|^2 \le \frac{1}{2\nu_0} \|u\|_E^2 . \tag{184}$$

For Minkowski space, $\varepsilon_m^{-1} \le |\mathbf{p}|^{-1}$, so the compact-support estimate in Eq. (92) controls the momentum component uniformly for all $m \ge 0$. The field component obeys

$$\langle u_\phi, \varepsilon_m u_\phi\rangle \le \|u_\phi\|_2 \, \|\varepsilon_m u_\phi\|_2 \le C_{\mathcal{K}_0} \, \|\varepsilon_m u_\phi\|_2^2 , \tag{185}$$

where compact-support Poincaré gives $\|u_\phi\|_2 \le C_{\mathcal{K}_0} \, \||\mathbf{p}| u_\phi\|_2 \le C_{\mathcal{K}_0} \, \|\varepsilon_m u_\phi\|_2$. Thus the same support-dependent constant works throughout the uniform massive regime. Taking the trace of the monotone positive integrals in Eq. (183) and applying these estimates shows that T_δ is trace class. Taking the energy trace instead gives

$$\mathrm{Tr}_{\mathbb{R}}(\varepsilon^{1/2} T_\delta \varepsilon^{1/2}) = \mathcal{E}(Q_\delta) < \infty. \tag{186}$$

Both identities follow first for finite subintegrals and then by monotone convergence. Thus no interchange involving a conditionally convergent operator integral is used. This explicit

integral establishes the two trace claims used in Proposition 5. Its proof then constructs the Radon Gaussian measure, computes the characteristic functional, and derives exact regional productness without repeating these operator estimates.

D Boundary-Hankel trace asymptotics

This appendix proves the local trace law used in Theorem 6. The argument has two scales. The normal coordinate is dilated by the collar width, which turns the two curved faces into parallel half-lines one unit apart. The tangential variables are treated semiclassically. The order-one principal symbol then produces the same positive Hankel fiber as the exact planar calculation.

Choose a collar cutoff equal to one for $|u| < u_0/4$ and supported in $|u| < u_0/2$. Denote the resulting collar-localized part of $K^{E,\sharp}_{\pi,\delta,m}$ by $K^{\mathrm{coll}}_{\delta,m}$. Write the Fermi metric and volume element as

$$h = \mathrm{d}u^2 + h_u(s), \qquad \mathrm{dvol}_h = J(s,u)\mathrm{d}A_h(s)\mathrm{d}u, \qquad J(s,0) = 1. \tag{187}$$

Define the normal dilations

$$(\mathcal{U}^A_\delta f)(s,\rho) = \delta^{1/2} J(s,-\delta\rho)^{1/2} f(s,-\delta\rho), \tag{188}$$

$$(\mathcal{U}^B_\delta f)(s,v) = \delta^{1/2} J(s,\delta(1+v))^{1/2} f(s,\delta(1+v)), \tag{189}$$

for $\rho, v > 0$. These maps identify the collar pieces with subspaces of $L^2(S \times \mathbb{R}_+, \mathrm{d}A_h \mathrm{d}\rho)$.

Put $n = d-1$, $\mu = m\delta$, and

$$\kappa_\mu(s,\xi) = \sqrt{|\xi|^2_{h_0(s)} + \mu^2}, \qquad (s,\xi) \in T^*S. \tag{190}$$

For $\kappa \geq 0$, let H_κ be the normal Hankel operator on $L^2(\mathbb{R}_+)$ with kernel

$$\mathsf{H}_\kappa(\rho,v) = \frac{1}{2\pi} \int_\kappa^\infty e^{-t(\rho+v+1)} \sqrt{t^2-\kappa^2}\, \mathrm{d}t. \tag{191}$$

It is positive because

$$\mathsf{H}_\kappa = \frac{1}{2\pi} \int_\kappa^\infty e^{-t} \sqrt{t^2-\kappa^2} |e_t\rangle\langle e_t|\, \mathrm{d}t, \qquad e_t(\rho) = e^{-t\rho}. \tag{192}$$

We now fix the tangential quantizations used below. Let $\{\chi_j\}$ be a finite smooth partition with $\sum_j \chi_j^2 = 1$, subordinate to bounded cotangent charts of S. In one such chart, the operator-valued semiclassical Weyl quantization is

$$(\mathrm{Op}^w_{\delta,j}(b)f)(s) = \frac{1}{(2\pi\delta)^n} \int e^{i(s-s')\cdot\xi/\delta} b\Big(\frac{s+s'}{2}, \xi\Big) f(s')\, \mathrm{d}s'\mathrm{d}\xi. \tag{193}$$

Here $b(s,\xi)$ acts on the normal factor $L^2(\mathbb{R}_+)$, and chart cutoffs are understood on both sides. We define $\mathrm{Op}^w_{\delta,S}$ by summing these localized operators; changes of coordinates contribute the usual one-lower-order symbol terms. For the anti-Wick quantization, let $\psi^\delta_{s,\xi}$ be the normalized standard Gaussian coherent state at phase-space point (s,ξ). For a trace-class-valued symbol $b(s,\xi)$ on $L^2(\mathbb{R}_+)$ with integrable trace norm, set

$$\mathrm{Op}^{\mathrm{AW}}_{\delta,S}(b) = \sum_j \frac{1}{(2\pi\delta)^n} \int |\chi_j\psi^\delta_{s,\xi}\rangle\langle\chi_j\psi^\delta_{s,\xi}| \otimes b_j(s,\xi)\, \mathrm{d}s\mathrm{d}\xi. \tag{194}$$

The integral acts on $L^2(S)\widehat{\otimes}L^2(\mathbb{R}_+)$. It is positive term by term when $b(s,\xi) \geq 0$. Changing charts changes the operator only by one lower semiclassical order.

Lemma 9 (Operator-valued trace-ideal estimate)**.** *Let $N > 4n + 4$. Suppose a symbol $b_\delta(s, \xi)$, compactly supported in the base variable, takes values in $\mathcal{S}_1(L^2(\mathbb{R}_+))$ and*

$$\mathfrak{m}_N(b_\delta) := \sum_{|\alpha|+|\beta|\le N} \int_{T^*S} \langle \xi \rangle^N \left\| \partial_s^\alpha \partial_\xi^\beta b_\delta(s, \xi) \right\|_1 \, \mathrm{d}\lambda_{T^*S} < \infty. \tag{195}$$

Then, uniformly for $0 < \delta \le 1$,

$$\left\| \mathrm{Op}^w_{\delta,S}(b_\delta) \right\|_1 + \left\| \mathrm{Op}^{\mathrm{AW}}_{\delta,S}(b_\delta) \right\|_1 \le C\delta^{-n} \mathfrak{m}_N(b_\delta). \tag{196}$$

Here $\mathrm{d}\lambda_{T^*S}$ is Liouville measure, equal to $\mathrm{d}A_h(s)\mathrm{d}\xi_s$ in an $h_0(s)$-orthonormal cotangent frame. The estimate is scale covariant. If $\delta \le r \le 1$, the symbol is supported where either $\kappa_\mu \le 2r$ or $r/2 \le \kappa_\mu \le 2r$, and

$$\sum_{|\alpha|+|\beta|\le N} \int \langle \eta \rangle^N \left\| \partial_s^\alpha \partial_\eta^\beta b_{\delta,r}(s, r\eta) \right\|_1 \, \mathrm{d}A_h(s)\mathrm{d}\eta \le A, \tag{197}$$

then both trace norms in Eq. (196) are at most $C(r/\delta)^n A$.

There is a separate compact-frequency form which does not require derivatives through the conic point. If $b_{\delta,r}(s, r\eta)$ is supported in a fixed η-ball and

$$\sum_{|\alpha|\le 2n+2} \int \|\partial_s^\alpha b_{\delta,r}(s, r\eta)\|_1 \, \mathrm{d}A_h(s)\mathrm{d}\eta \le A_0, \tag{198}$$

then its Weyl and anti-Wick trace norms are at most $C(r/\delta)^n A_0$.

Proof. For anti-Wick quantization, take the trace norm inside the Bochner coherent-state integral, using the ideal triangle inequality, in Eq. (194). The coherent-state resolution gives the factor $(2\pi\delta)^{-n}$ and the zeroth symbol seminorm.

For Weyl quantization, localize to one bounded chart, extend it to a fixed flat torus, and put $\mathcal{K} = L^2(\mathbb{R}_+)$. Choose an even integer q such that

$$\frac{n}{2} < q, \qquad 2q + 2n + 2 < N. \tag{199}$$

Such a choice exists because $N > 4n + 4$. On the torus set

$$P_\delta = (1 - \delta^2 \Delta_{\mathbb{T}^n})^{q/2}, \qquad J_\delta = P_\delta^{-1}. \tag{200}$$

Since q is even, P_δ is a semiclassical differential operator with the s-independent polynomial symbol $p(\xi) = (1 + |\xi|^2)^{q/2}$. On $C^\infty(\mathbb{T}^n) \otimes \mathcal{K}$, define

$$Y_\delta = (P_\delta \otimes \mathbf{1}_\mathcal{K}) \, \mathrm{Op}^w_\delta(b_\delta)(P_\delta \otimes \mathbf{1}_\mathcal{K}). \tag{201}$$

Once Y_δ is shown to extend boundedly, the common-core identity $P_\delta J_\delta = J_\delta P_\delta = \mathbf{1}$ and continuity give

$$\mathrm{Op}^w_\delta(b_\delta) = (J_\delta \otimes \mathbf{1}_\mathcal{K}) Y_\delta (J_\delta \otimes \mathbf{1}_\mathcal{K}). \tag{202}$$

We now prove that extension and its projective estimate. Because p is a polynomial, both Moyal products terminate and the exact Weyl symbol of Y_δ is

$$y_\delta = p \, \#_\delta \, b_\delta \, \#_\delta \, p. \tag{203}$$

Writing $D = (\partial_s, \partial_\xi)$, the finite composition formula gives, for $|\gamma| \le 2n+2$,

$$\|D^\gamma y_\delta(s,\xi)\|_1 \le C_\gamma \sum_{|\rho|\le 2q+|\gamma|} \langle\xi\rangle^{2q} \|D^\rho b_\delta(s,\xi)\|_1 . \tag{204}$$

Therefore

$$\sum_{|\gamma|\le 2n+2} \int \|D^\gamma y_\delta(s,\xi)\|_1 \, \mathrm{d}s\mathrm{d}\xi \le C\mathfrak{m}_N(b_\delta). \tag{205}$$

Indeed, Eq. (199) keeps the largest derivative strictly below N, and the factor $\langle\xi\rangle^{2q}$ is absorbed by the weight in Eq. (195).

Choose a smooth finite-overlap partition of the phase chart into unit boxes. For each box cutoff θ_λ, fix an enlarged cutoff θ_λ^+ which equals one on $\operatorname{supp}\theta_\lambda$. On each box Fourier-expand the localized $\mathcal{S}_1(\mathcal{K})$-valued symbol,

$$\theta_\lambda y_\delta = \sum_{k\in\mathbb{Z}^{2n}} a_{\lambda k} C_{\lambda k}, \qquad C_{\lambda k} \in \mathcal{S}_1(\mathcal{K}). \tag{206}$$

Applying $(1-\Delta_{s,\xi})^{n+1}$ before taking each Fourier coefficient gives

$$\|C_{\lambda k}\|_1 \le C\langle k\rangle^{-2n-2} \sum_{|\gamma|\le 2n+2} \int_{\operatorname{supp}\theta_\lambda^+} \|D^\gamma y_\delta(s,\xi)\|_1 \, \mathrm{d}s\mathrm{d}\xi. \tag{207}$$

The scalar operators $B_{\lambda k} = \mathrm{Op}_\delta^w(a_{\lambda k})$ have uniformly bounded norm: after fixed window factors they are phase-space translations and modulations of one unit-box operator. Since $2n+2 > \dim(T^*S) = 2n$, the coefficient series is absolutely summable. Summing first in k, then over the finite-overlap boxes, and using Eq. (205) yields the convergent projective decomposition

$$Y_\delta = \sum_{\lambda,k} B_{\lambda k} \otimes C_{\lambda k}, \qquad \sum_{\lambda,k} \|B_{\lambda k}\| \, \|C_{\lambda k}\|_1 \le C\mathfrak{m}_N(b_\delta), \tag{208}$$

which proves the required bounded extension. Approximation in the Bochner $W^{N,1}$ norm passes the argument from finite-rank-valued symbols to general $\mathcal{S}_1(\mathcal{K})$-valued symbols. This is the point at which the normal trace ideal enters; no identity on $\mathcal{K}$ is treated as Hilbert–Schmidt. It is also the projective-tensor form of the operator-valued trace criterion in [26, Thm. 3.4.40]. For each term,

$$\|(J_\delta B_{\lambda k} J_\delta) \otimes C_{\lambda k}\|_1 \le \|J_\delta\|_2^2 \, \|B_{\lambda k}\| \, \|C_{\lambda k}\|_1 . \tag{209}$$

Since $\|J_\delta\|_2^2 = \mathrm{Tr}(1-\delta^2\Delta_{\mathbb{T}^n})^{-q} = O(\delta^{-n})$, summing Eq. (209) proves the Weyl estimate. The finite quadratic partition of unity globalizes it; the elementary ideal estimates are collected in Ref. [27].

For the scale-covariant statement, substitute $\xi = r\eta$. The effective semiclassical parameter is δ/r, while Eq. (197) supplies the same projective norm on a fixed phase region. Repeating the preceding proof replaces δ^{-n} by $(r/\delta)^n$.

For the compact-frequency form, Fourier-expand only in the bounded base variable. In a chart, the Weyl kernel of one base Fourier mode and one fixed η is a rank-one tangential kernel tensored with its $\mathcal{S}_1(\mathcal{K})$-valued Fourier coefficient. Integration by parts in the base variable makes the coefficient trace norms absolutely summable; the $2n+2$ derivatives in Eq. (198) are more than enough. Integrating the rank-one nuclear decomposition over the fixed η-ball gives $(2\pi\delta/r)^{-n}A_0$. The anti-Wick estimate follows directly from its coherent-state integral. Thus neither estimate differentiates the symbol in η at the conic point. □

Lemma 10 (Curved boundary-symbol estimate)**.** *For every $M < \infty$, there are $C_M, \delta_M > 0$ such that, for $0 < \delta < \delta_M$ and $0 \le \mu \le M$, there is a positive trace-class operator*

$$\mathsf{B}_{\delta,\mu} = \delta^{-1}\,\mathrm{Op}^{\mathrm{AW}}_{\delta,S}\left(\mathsf{H}_{\kappa_\mu(s,\xi)}\right) \tag{210}$$

such that

$$\left\| -\mathcal{U}^A_\delta K^{\mathrm{coll}}_{\delta,m} (\mathcal{U}^B_\delta)^* - \mathsf{B}_{\delta,\mu} \right\|_1 \le C_M \mathcal{H}^{d-1}(S)\delta^{-n}. \tag{211}$$

Its phase-space trace is

$$\mathrm{Tr}\,\mathsf{B}_{\delta,\mu} = \frac{\delta^{-d}}{(2\pi)^n} \int_{T^*S} \mathrm{Tr}_{L^2(\mathbb{R}_+)}\,\mathsf{H}_{\kappa_\mu(s,\xi)}\,\mathrm{d}\lambda_{T^*S}(s,\xi) + O_M\left(\mathcal{H}^{d-1}(S)\delta^{-n}\right). \tag{212}$$

The complement of the collar satisfies

$$\left\| K^{E,\sharp}_{\pi,\delta,m} - K^{\mathrm{coll}}_{\delta,m} \right\|_1 \le C_M \mathcal{H}^{d-1}(S). \tag{213}$$

Proof. For fixed $\delta > 0$, the two sharp regions are positively separated. The kernel of $L_m^{1/2}$ is smooth off the diagonal, so its restriction to the compact product $\overline{A}_\delta \times \overline{B}_\delta$ maps $L^2(B_\delta)$ into $H^q(A_\delta)$ for every q. For $q > d$, the embedding $H^q(A_\delta) \hookrightarrow L^2(A_\delta)$ is trace class. This factorization proves the asserted trace-class extension. A partition of unity separates the collar-collar block from terms with at least one variable in the fixed-distance region. Pseudolocality makes every such complementary kernel uniformly smooth. The sharp-region volume and local-geometry bounds in Theorem 6 then give Eq. (213).

In a Fermi chart, parameter-dependent complex-power calculus gives the principal symbol

$$\sigma_1(L_m^{1/2})(s,u;\xi,\tau) = \sqrt{\tau^2 + |\xi|^2_{h_u(s)} + m^2} \tag{214}$$

[28–30]. Under

$$u = -\delta\rho, \quad u' = \delta(1+v), \quad \xi_{\mathrm{phys}} = \xi/\delta, \quad \tau_{\mathrm{phys}} = \sigma/\delta, \tag{215}$$

the leading symbol becomes

$$\delta^{-1}\sqrt{\sigma^2 + |\xi|^2_{h_0(s)} + \mu^2}. \tag{216}$$

For $\rho + v + 1 > 0$, its inverse normal Fourier transform is

$$-\frac{1}{2}\mathcal{F}^{-1}_{\sigma\to\rho+v+1}\sqrt{\sigma^2+\kappa^2} = \mathsf{H}_\kappa(\rho,v). \tag{217}$$

This fixes both the sign and the factor $1/(2\pi)$ in Eq. (191).

The boundary calculus gives the following trace-class-valued symbol expansion:

$$-\mathcal{U}^A_\delta K^{\mathrm{coll}}_{\delta,m} (\mathcal{U}^B_\delta)^* = \delta^{-1}\,\mathrm{Op}^w_{\delta,S}(\mathsf{H}_{\kappa_\mu}) + \mathrm{Op}^w_{\delta,S}(b_{\delta,\mu}) + R_{\delta,\mu}. \tag{218}$$

Here is the symbol accounting behind this expansion. Put $\mathcal{K} = L^2(\mathbb{R}_+)$ and

$$q_\mu(s,\xi,\sigma) = \sqrt{\sigma^2 + |\xi|^2_{h_0(s)} + \mu^2}. \tag{219}$$

After the Fermi half-density conjugations and normal dilation, the complete tangential symbol has the form

$$a_{\delta,\mu}(s,\xi) = \delta^{-1}\mathsf{H}_{\kappa_\mu(s,\xi)} + b_{\delta,\mu}(s,\xi), \qquad b_{\delta,\mu} \in S^0_{\mathrm{wc}}(T^*S;\mathcal{S}_1(\mathcal{K})), \tag{220}$$

modulo a remainder of trace norm $O_M(\mathrm{Area}(S))$. The notation S^0_{wc} means the weak-conic estimates displayed in Eq. (224) below.

Seeley's complex-power theorem supplies the parameter-dependent scalar expansion $q_1 + q_0 + q_{-1} + \cdots$ for $L_m^{1/2}$. Under Eq. (215), a term homogeneous of degree $1-j$ carries the factor δ^{j-1}. The weakly polyhomogeneous boundary composition of Grubb and Seiler permits normal truncation, Fermi-Jacobian conjugation, and chart composition term by term [28–30]. We retain the terms $j = 1, \dots, n+1$ in $b_{\delta,\mu}$, multiplied by a cutoff $1 - \vartheta_{\mathrm{core}}$ that vanishes on the conic core $\kappa_\mu \le 4\delta$. In physical variables this core is the fixed ball $\left(|\xi_{\mathrm{phys}}|^2_{h_0} + m^2\right)^{1/2} \le 4$, so the corresponding part of $L_m^{1/2}$ is removed by a low-energy cutoff in physical frequency. The removed part has a uniformly smooth kernel on the compact collar pair and contributes $O_M(\mathrm{Area}(S))$ to $R_{\delta,\mu}$. On an annulus $\kappa_\mu \asymp r$, the first omitted term has rescaled trace-class amplitude $O((\delta/r)^{n+1})$. The tangential trace estimate contributes $(r/\delta)^n$, so that annulus costs $O(\delta/r)$. The dyadic sum over $\delta \le r \le 1$ is bounded, and the $r = \delta$ core is $O(1)$ by the compact-frequency clause of Lemma 9. Thus the first omitted term, the truncated normal tails, and the uniformly smoothing chart terms together form an $O_M(\mathrm{Area}(S))$ trace remainder.

A representative first correction shows how the scalar calculus becomes a trace-class-valued estimate. Let $G^{ij}(s) = \partial_u h^{ij}_u(s)|_{u=0}$. On the left face,

$$\delta^{-1}\left[\sqrt{\sigma^2 + |\xi|^2_{h_{-\delta\rho}(s)} + \mu^2} - q_\mu(s,\xi,\sigma)\right] = -\rho\frac{G^{ij}(s)\xi_i\xi_j}{2q_\mu} + O(\delta\rho^2 q_\mu). \tag{221}$$

Since $\mathcal{F}^{-1}_{\sigma\to a}(q_\mu^{-1}) = \pi^{-1}K_0(\kappa_\mu a)$, the displayed first term contributes, up to the fixed left-symbol sign convention,

$$b^{(1)}_{\mathrm{met}}(s,\xi;\rho,v) = \frac{\rho G^{ij}(s)\xi_i\xi_j}{4\pi} K_0(\kappa_\mu(s,\xi)(\rho+v+1))\,. \tag{222}$$

Using $K_0(\kappa a) = \int_1^\infty e^{-\kappa t a}(t^2-1)^{-1/2}\,\mathrm{d}t$, this is a Bochner integral of rank-one normal kernels. In particular,

$$\left\|b^{(1)}_{\mathrm{met}}(s,\xi)\right\|_1 \le C|\xi|^2 \int_1^\infty \frac{e^{-\kappa_\mu t}}{\sqrt{t^2-1}} \left\|\rho e^{-\kappa_\mu t\rho}\right\|_2 \left\|e^{-\kappa_\mu t v}\right\|_2 \,\mathrm{d}t \le C. \tag{223}$$

Each tangential derivative costs at most one factor κ_μ^{-1}. Base derivatives preserve the estimate because $\partial_s^k \kappa_\mu = O(\kappa_\mu)$. The subprincipal, right-face, Jacobian, potential, and chart-composition corrections have the same finite Laplace-rank-one form. On $\kappa_\mu \asymp r \ge \delta$, their Taylor remainders are $O(\delta/r)$ in rescaled seminorms. The core $\kappa_\mu \lesssim \delta$ is left unexpanded, and this is necessary rather than convenient. The terms with $j \ge 2$, such as the potential term q_{-1}, carry no factor $|\xi|^2$ to offset the normal norms, so their trace norms grow like $(\delta/\kappa_\mu)^{j-1}$ as $\kappa_\mu \downarrow 0$. They are bounded on every annulus $\kappa_\mu \asymp r \ge \delta$, where $\delta/r \le 1$, but not uniformly on the core, which is why the core is removed from $b_{\delta,\mu}$.

The exponential $e^{-t(\rho+v+1)}$ makes every normal trace-class seminorm integrable. The only loss of ordinary symbol smoothness occurs at the conic point $(\xi,\mu) = (0,0)$. Termwise composition, including chart transitions, now gives the uniform weak-conic bounds

$$\int_S e^{c|\xi|_{h_0(s)}}\kappa_\mu(s,\xi)^{|\beta|}\left\|\partial_s^\alpha\partial_\xi^\beta b_{\delta,\mu}(s,\xi)\right\|_1 \,\mathrm{d}A_h(s) \le C_{\alpha\beta M}\,\mathrm{Area}(S), \tag{224}$$

$$\|R_{\delta,\mu}\|_1 \le C_M\,\mathrm{Area}(S). \tag{225}$$

The first inequality is asserted on the annuli $\kappa_\mu \asymp r$ with $r \ge \delta$, uniformly in r. This is all the argument uses, because $b_{\delta,\mu}$ vanishes on the core.

To quantize it uniformly, choose a dyadic partition with ϑ_r supported where $r/2 \le \kappa_\mu \le 2r$ and a buffer $\widetilde{\vartheta}_r = 1$ on that support, vanishing outside $r/4 \le \kappa_\mu \le 4r$. Use a core cutoff supported in $\kappa_\mu \le 4\delta$. All Moyal compositions are made with the buffered symbols. Cutoff derivatives cost exactly r^{-1}, commutators meet only finitely many adjacent annuli, and the finite overlap preserves the estimates. Thus the region $\kappa_\mu \le 1$ splits into annuli $\kappa_\mu \asymp r$, with $r = 2^{-j} \ge \delta$, and one conic core. On an annulus, the substitution $\xi = r\eta$ cancels the factor $\kappa_\mu^{-|\beta|}$ in every rescaled derivative. The scale-covariant part of Lemma 9 therefore gives

$$\left\|\mathrm{Op}^w_{\delta,S}(b^{(r)}_{\delta,\mu})\right\|_1 \le C_M \,\mathrm{Area}(S)\delta^{-n} r^n. \tag{226}$$

On the core, $b_{\delta,\mu}$ vanishes, and the removed low-frequency part is already in $R_{\delta,\mu}$ with trace norm $O_M(\mathrm{Area}(S))$. No estimate differentiates through $\xi = 0$. The exponentially decaying region $\kappa_\mu \ge 1$ obeys the ordinary seminorm estimate. Since $\sum_{j\ge 0} 2^{-jn} < \infty$, summing the annuli gives

$$\left\|\mathrm{Op}^w_{\delta,S}(b_{\delta,\mu})\right\|_1 \le C_M \mathcal{H}^{d-1}(S)\delta^{-n}. \tag{227}$$

It remains to replace the leading Weyl quantization by the positive anti-Wick model. Away from $(\xi,\mu) = (0,0)$, the identity $\mathrm{Op}^{\mathrm{AW}}_\delta(a) = \mathrm{Op}^w_\delta(e^{\delta\Delta/4}a)$ shows that, after the outer factor δ^{-1}, the difference has order zero. Uniformly in $0 \le \mu \le M$, split the conic zone at $\kappa_\mu(s,\xi) \le 2\delta^{1/2}$. The compact-frequency clause of Lemma 9, applied separately to the Weyl and anti-Wick symbols with $r = \delta^{1/2}$, uses only their uniformly bounded base derivatives. Including the outer factor δ^{-1}, it bounds each core trace norm by

$$O_M\Big(\mathrm{Area}(S)\delta^{-d+n/2}\Big) = O_M\big(\mathrm{Area}(S)\delta^{-n}\big)\,, \qquad n = d-1 \ge 2, \tag{228}$$

On the complement, use the normal-fiber derivative estimate

$$\int_S e^{c\kappa_\mu(s,\xi)}\kappa_\mu(s,\xi)^{|\beta|-1}\left\|\partial_s^\alpha\partial_\xi^\beta \mathsf{H}_{\kappa_\mu(s,\xi)}\right\|_1 \mathrm{d}A_h(s) \le C_{\alpha\beta M}\,\mathrm{Area}(S), \qquad |\beta| \ge 1. \tag{229}$$

It follows directly by differentiating the positive Laplace integral Eq. (192). Duhamel's identity

$$\delta^{-1}(e^{\delta\Delta/4} - 1)\mathsf{H}_{\kappa_\mu} = \frac{1}{4}\int_0^1 e^{t\delta\Delta/4}\Delta\mathsf{H}_{\kappa_\mu}\,\mathrm{d}t \tag{230}$$

then shows that on a dyadic annulus $\kappa_\mu \asymp r \ge \delta^{1/2}$, the scale-covariant trace estimate is

$$C_M\,\mathrm{Area}(S)\delta^{-n}r^{n-1}. \tag{231}$$

The sum over these annuli is uniform because $\sum r^{n-1} < \infty$ for $n = d-1 \ge 2$; the region $\kappa_\mu \ge 1$ is exponentially smaller. Hence

$$\left\|\delta^{-1}\,\mathrm{Op}^w_{\delta,S}(\mathsf{H}_{\kappa_\mu}) - \mathsf{B}_{\delta,\mu}\right\|_1 \le C_M\mathcal{H}^{d-1}(S)\delta^{-n}. \tag{232}$$

Equations (218), (227), and (232) prove Eq. (211). Taking the trace in the coherent-state definition Eq. (194), including the one-lower-order chart terms already bounded above, gives Eq. (212). □

It remains to evaluate that trace. Equation (192) gives

$$\mathrm{Tr}\,\mathsf{H}_\kappa = \frac{1}{4\pi}\int_\kappa^\infty e^{-t}\frac{\sqrt{t^2-\kappa^2}}{t}\,\mathrm{d}t. \tag{233}$$

At each s, choose an $h_0(s)$-orthonormal cotangent frame. In that frame, Liouville measure is $\mathrm{d}A_h(s)\mathrm{d}\xi_s$, so the fiber integral is the same at every point. Equations (111) and (212) therefore give

$$\operatorname{Tr} \mathsf{B}_{\delta,\mu} = \mathcal{H}^{d-1}(S)\delta^{-d}c_d(\mu) + O_M\Big(\mathcal{H}^{d-1}(S)\delta^{-(d-1)}\Big)\,. \tag{234}$$

Because $\mathsf{B}_{\delta,\mu} \geq 0$,

$$\Big|\Big\|K^{\mathrm{coll}}_{\delta,m}\Big\|_1 - \operatorname{Tr} \mathsf{B}_{\delta,\mu}\Big| \leq \Big\|-\mathcal{U}^A_\delta K^{\mathrm{coll}}_{\delta,m}(\mathcal{U}^B_\delta)^* - \mathsf{B}_{\delta,\mu}\Big\|_1\,. \tag{235}$$

Combining Eqs. (213), (211), and (234) proves Eq. (119).

E Finite-mode covariance programs

This appendix supplies the regulated computational statement cited in Section 9. Finite-mode means finitely many genuine canonical oscillator modes; their Fock space remains infinite dimensional.

Proposition 11 (Exact finite-mode covariance reduction)**.** *Let* $\mathbf{r} = (q_1, p_1, \ldots, q_n, p_n)^{\mathsf{T}}$ *contain finitely many canonical variables satisfying*

$$[r_j, r_k] = i\Omega_{jk}\mathbf{1}. \tag{236}$$

Suppose their symplectic space splits as $\mathsf{V}_A \oplus \mathsf{V}_B \oplus \mathsf{V}_C$*, where* C *contains any modes outside the tested pair. Let*

$$H_{\mathrm{reg}} = \frac{1}{2}\mathbf{r}^{\mathsf{T}} G \mathbf{r} - E_0, \qquad G = G^{\mathsf{T}} > 0. \tag{237}$$

For a state with mean $\bar{\mathbf{r}}$*, define the centered covariance*

$$\Gamma_{jk} = \frac{1}{2}\rho(\{r_j - \bar{r}_j\mathbf{1}, r_k - \bar{r}_k\mathbf{1}\})\,. \tag{238}$$

Write the AB principal block as

$$\Gamma_{AB} = \begin{pmatrix} \Gamma^{(0)}_A & Z \\ Z^{\mathsf{T}} & \Gamma^{(0)}_B \end{pmatrix}. \tag{239}$$

The full covariance is physical when

$$\Gamma + \frac{i}{2}\Omega \succeq 0. \tag{240}$$

Within that physical set, impose one of the following three regional constraints:

$$\text{PPT}: \quad \Gamma_{AB} + \frac{i}{2}(\Omega_A \oplus -\Omega_B) \succeq 0, \tag{241}$$

$$\text{Sep}: \quad \Gamma_{AB} \succeq \Gamma'_A \oplus \Gamma'_B, \quad \Gamma'_A + \frac{i}{2}\Omega_A \succeq 0, \quad \Gamma'_B + \frac{i}{2}\Omega_B \succeq 0 \quad \text{for some real symmetric } \Gamma'_A, \Gamma'_B, \tag{242}$$

$$\text{prod}: \quad Z = 0. \tag{243}$$

Let $\mathfrak{F}_{\mathrm{PPT}}$*,* $\mathfrak{F}_{\mathrm{Sep}}$*, and* $\mathfrak{F}_{\mathrm{prod}}$ *denote the resulting semidefinite-representable feasible sets. For any one of these classes* $\mathcal{C}$*,*

$$\boxed{\inf_{\rho:\,\rho_{AB}\in\mathcal{C}} \rho(H_{\mathrm{reg}}) = \min_{\Gamma\in\mathfrak{F}_{\mathcal{C}}} \left\{\frac{1}{2}\operatorname{Tr}(G\Gamma) - E_0\right\}.} \tag{244}$$

Thus all three unrestricted finite-mode energy minima are semidefinite programs. An optimizer may be chosen Gaussian although no Gaussian restriction was imposed.

Proof. Every finite-energy state has a covariance satisfying Eq. (240). A Weyl displacement that removes $\bar{\mathbf{r}}$ preserves each regional state class and lowers the energy by $\bar{\mathbf{r}}^{\mathsf{T}} G \bar{\mathbf{r}}/2$. It is therefore enough to optimize over centered states.

If ρ_{AB} is PPT, partial transposition reverses the symplectic form on the B modes, giving Eq. (241). A separable covariance obeys Eq. (242), while a product covariance has $Z = 0$ [31].

Conversely, every covariance satisfying Eq. (240) is realized by a centered Gaussian density state. The PPT constraint makes its Gaussian AB marginal PPT. The separability constraint invokes the converse Werner–Wolf construction. If $Z = 0$, the Gaussian characteristic function factorizes, so the marginal is product. In all three cases the realizing Gaussian state has the same quadratic energy as the supplied covariance.

It remains only to show that the infima are attained. Since $G > 0$, every bounded objective sublevel has bounded Γ. In the separable problem,

$$0 \preceq \Gamma'_A \preceq \Gamma_A^{(0)}, \qquad 0 \preceq \Gamma'_B \preceq \Gamma_B^{(0)}. \tag{245}$$

The auxiliary covariances are bounded on the same sublevel. A convergent subsequence preserves every semidefinite constraint, including the existential separability constraints. Hence all three minima are attained. □

F Flat-interface integral

The Fourier identities used in Section 8.1 are

$$K_0(\kappa s) = \int_1^\infty \frac{e^{-\kappa s t}}{\sqrt{t^2-1}}\,\mathrm{d}t, \tag{246}$$

$$\frac{\kappa K_1(\kappa s)}{s} = \kappa^2 \int_1^\infty e^{-\kappa s t}\sqrt{t^2-1}\,\mathrm{d}t. \tag{247}$$

For $e_t(u) = e^{-\kappa t u}$,

$$\|e_t\|_2^2 = \frac{1}{2\kappa t}, \qquad \langle e_t, (-\partial_u^2 + \kappa^2) e_t\rangle = \frac{\kappa(t^2+1)}{2t}. \tag{248}$$

There are two half-lines and a Hamiltonian factor $1/2$. Combining their field energy with the momentum energy reduces the outside-collar integrand to

$$\frac{\kappa}{2\pi}\frac{t e^{-\kappa\delta t}}{\sqrt{t^2-1}}, \tag{249}$$

whose integral is $\kappa K_1(\kappa\delta)/(2\pi)$.

The minimum-energy field interpolation between boundary values $+1$ and -1 is odd about the collar midpoint. Direct solution of $(-\partial_z^2+\kappa^2)\phi = 0$ gives the energy in Eq. (102). Multiplying it by the field Laplace weight and the Hamiltonian factor produces

$$\frac{\kappa}{2\pi}\coth(\kappa\delta/2)K_0(\kappa\delta). \tag{250}$$

This proves Eq. (103). Finally, $\mathrm{d}^2\mathbf{k} = 2\pi\kappa\,\mathrm{d}\kappa$ and $u = \kappa\delta$ give Eq. (104).

The quoted massless coefficient can be checked without a special-function quadrature. The standard Mellin integral gives $\int_0^\infty u^2 K_1(u)\mathrm{d}u = 2$. Expanding $\coth(u/2) = 1 + 2\sum_{n\ge1} e^{-nu}$ gives

$$\int_0^\infty u^2 \coth(u/2)K_0(u)\mathrm{d}u = \frac{\pi}{2} + 2\sum_{n=1}^\infty J_n, \tag{251}$$

where

$$J_1 = \frac{4}{15}, \qquad J_n = \frac{(2n^2+1)\,\operatorname{arcosh} n}{(n^2-1)^{5/2}} - \frac{3n}{(n^2-1)^2} \quad (n>1). \tag{252}$$

This follows from $J_n = 2\int_0^\infty (n+\cosh t)^{-3}\mathrm{d}t$. The summand is $O(\log n/n^3)$, so the series converges absolutely and evaluates the coefficient in Eq. (105).

Data availability

No data were generated or analyzed in this work.

Code availability

No code is required for the proofs. The numerical coefficient in Eq. (105) can be reproduced from the absolutely convergent series in Eqs. (251) and (252).

Author contributions

M.C. formulated the problems (arising from a broader motivating framework), set their hypotheses and claim boundaries, directed the proof strategy, and verified the results.

AI use disclosure

OpenAI GPT-5.6 Sol, used through Codex, assisted with literature and method selection and carried out proof construction, calculations, and drafting under the author's direction. Anthropic's Claude Fable 5 reviewed drafts. Google's Gemini 4 Carbon independently checked Theorem 6 numerically and identified a gap in Appendix D, repaired in this version. The author takes full responsibility for the content.

Competing interests

The author declares no competing interests.

References

[1] Detlev Buchholz. Product states for local algebras. *Communications in Mathematical Physics*, 36(4): 287–304, 1974. DOI: 10.1007/BF01646201.

[2] Christopher J. Fewster. The split property for quantum field theories in flat and curved spacetimes. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 86(2):153–175, 2016. DOI: 10.1007/s12188-016-0130-9.

[3] Ko Sanders. On separable states in relativistic quantum field theory. *Journal of Physics A: Mathematical and Theoretical*, 56:505201, 2023. DOI: 10.1088/1751-8121/ad0bca.

[4] Stephen J. Summers and Reinhard Werner. Maximal violation of Bell's inequalities for algebras of observables in tangent spacetime regions. *Annales de l'Institut Henri Poincaré Physique Théorique*, 49(2):215–243, 1988. URL https://www.numdam.org/item/AIHPA_1988__49_2_215_0/.

[5] Stefan Hollands and Ko Sanders. *Entanglement Measures and Their Properties in Quantum Field Theory*, volume 34 of *SpringerBriefs in Mathematical Physics*. Springer, 2018. DOI: 10.1007/978-3-319-94902-4.

[6] Cédric Bény, Christopher T. Chubb, Terry Farrelly, and Tobias J. Osborne. Energy cost of entanglement extraction in complex quantum systems. *Nature Communications*, 9:3792, 2018. DOI: 10.1038/s41467-018-06153-w.

[7] Rainer Verch and Reinhard F. Werner. Distillability and positivity of partial transposes in general quantum field systems. *Reviews in Mathematical Physics*, 17(5):545–576, 2005. DOI: 10.1142/S0129055X05002364.

[8] Wojciech Dybalski. A sharpened nuclearity condition and the uniqueness of the vacuum in QFT. *Communications in Mathematical Physics*, 283:523–542, 2008. DOI: 10.1007/s00220-008-0514-5.

[9] Wojciech Dybalski. A sharpened nuclearity condition for massless fields. *Letters in Mathematical Physics*, 84:217–230, 2008. DOI: 10.1007/s11005-008-0244-9.

[10] Hanno Sahlmann and Rainer Verch. Microlocal spectrum condition and Hadamard form for vector-valued quantum fields in curved spacetime. *Reviews in Mathematical Physics*, 13(10):1203–1246, 2001. DOI: 10.1142/S0129055X01001010.

[11] Laurent Saloff-Coste. Uniformly elliptic operators on Riemannian manifolds. *Journal of Differential Geometry*, 36(2):417–450, 1992. DOI: 10.4310/jdg/1214448748.

[12] Ingolf Buttig and Jürgen Eichhorn. The heat kernel for p-forms on manifolds of bounded geometry. *Acta Scientiarum Mathematicarum*, 55(1–2):33–51, 1991. URL https://acta.bibl.u-szeged.hu/15294/.

[13] Marek J. Radzikowski. Micro-local approach to the Hadamard condition in quantum field theory on curved space-time. *Communications in Mathematical Physics*, 179(3):529–553, 1996. DOI: 10.1007/BF02100096.

[14] Elliott H. Lieb. Sharp constants in the Hardy–Littlewood–Sobolev and related inequalities. *Annals of Mathematics*, 118(2):349–374, 1983. DOI: 10.2307/2007032.

[15] Vladimir I. Bogachev. *Gaussian Measures*, volume 62 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1998. DOI: 10.1090/surv/062.

[16] Marek Kuś and Karol Życzkowski. Geometry of entangled states. *Physical Review A*, 63:032307, 2001. DOI: 10.1103/PhysRevA.63.032307.

[17] Jennifer Ahiable, Naga Bhavya Teja Kothakonda, and Andreas Winter. The geometry of absolute separability and other convex matrix properties from spectrum, 2026. arXiv:2608.03390 [quant-ph].

[18] Francesca Fabiana Settembrini, Frieder Lindel, Alexa Marina Herter, Stefan Yoshi Buhmann, and Jérôme Faist. Detection of quantum-vacuum field correlations outside the light cone. *Nature Communications*, 13(1):3383, 2022. DOI: 10.1038/s41467-022-31081-1.

[19] Alexa Herter, Frieder Lindel, Laura Gabriel, Stefan Yoshi Buhmann, and Jérôme Faist. Experimentally separating vacuum fluctuations from source radiation. *Nature Communications*, 17(1):2863, 2026. DOI: 10.1038/s41467-026-69142-4.

[20] Mark Hillery and M. Suhail Zubairy. Entanglement conditions for two-mode states. *Physical Review Letters*, 96:050503, 2006. DOI: 10.1103/PhysRevLett.96.050503.

[21] Frieder Lindel, Alexa Herter, Valentin Gebhart, Jérôme Faist, and Stefan Yoshi Buhmann. Entanglement harvesting from electromagnetic quantum fields. *Physical Review A*, 110(2):022414, 2024. DOI: 10.1103/PhysRevA.110.022414.

[22] Huzihiro Araki. A lattice of von Neumann algebras associated with the quantum theory of a free Bose field. *Journal of Mathematical Physics*, 4(11):1343–1362, 1963. DOI: 10.1063/1.1703912.

[23] Rudolf Haag. *Local Quantum Physics: Fields, Particles, Algebras*. Springer, Berlin, 2 edition, 1996. DOI: 10.1007/978-3-642-61458-3.

[24] Christian Gérard, Omar Oulghazi, and Michał Wrochna. Hadamard states for the Klein–Gordon equation on Lorentzian manifolds of bounded geometry. *Communications in Mathematical Physics*, 352(2):519–583, 2017. DOI: 10.1007/s00220-017-2847-4.

[25] Bernd Ammann, Robert Lauter, Victor Nistor, and András Vasy. Complex powers and non-compact manifolds. *Communications in Partial Differential Equations*, 29(5–6):671–705, 2004. DOI: 10.1081/PDE-120037329.

[26] Giuseppe De Nittis, Max Lein, and Marcello Seri. A magnetic pseudodifferential calculus for operator-valued and equivariant operator-valued symbols, 2022. arXiv:2210.05731.

[27] Barry Simon. *Trace Ideals and Their Applications*, volume 120 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2 edition, 2005. DOI: 10.1090/surv/120.

[28] Robert T. Seeley. Complex powers of an elliptic operator. In *Singular Integrals*, volume 10 of *Proceedings of Symposia in Pure Mathematics*, pages 288–307. American Mathematical Society, Providence, RI, 1967. DOI: 10.1090/pspum/010/0237943.

[29] Gerd Grubb. A weakly polyhomogeneous calculus for pseudodifferential boundary problems. *Journal of Functional Analysis*, 184(1):19–76, 2001. DOI: 10.1006/jfan.2000.3649.

[30] Jörg Seiler. Singular Green operators in the edge algebra formalism. *Journal of Mathematical Analysis and Applications*, 511:126041, 2022. DOI: 10.1016/j.jmaa.2022.126041.

[31] Reinhard F. Werner and Michael M. Wolf. Bound entangled Gaussian states. *Physical Review Letters*, 86(16):3658–3661, 2001. DOI: 10.1103/PhysRevLett.86.3658.