Perfect play in a binary-scoring nonlocal game requires infinite dimension but only finite energy

Seth Douglas Nidhal Mghirbi

Version 2.0.0, October 2026

Abstract

A finite rational binary-scoring game can require infinite dimension for perfect play while admitting finite energy and finite entanglement entropy. Our game has ten Alice questions, eleven Bob questions, at most six answers, and 52 uniform question pairs. It is spatially perfect and never finitely perfect, with sharp dimension profile $D_\Gamma(\varepsilon) = 10\log_2(1/\varepsilon) + O(\log\log(1/\varepsilon))$. Among pure perfect strategies, the displayed strategy attains the least entanglement entropy, 5.362740278498934 ebits, and least mean occupation, 275/18 per oscillator. We classify all perfect normal type-I marginal spectra and prove that this spectrum is the unique majorization maximum. These constants are also the small-loss limits of the entropy and one-sided oscillator energy prices. In the massive free scalar field we construct a pure perfect state in two strictly separated bounded regions, equal to the vacuum outside a bounded containing region and with every polynomial energy moment finite. In contrast, every fixed faithful split state has a positive loss floor; the complementary-wedge vacuum has supremum one without attainment. Energy-analytic states, bounded-region invertible vacuum preparations and product-Kraus preparations cannot be perfect. Every perfect pure spatial strategy has infinite Schmidt rank; finite entropy is attained but not forced. In every type-II_1 factor the normalized trace has a loss floor, while bounded faithful prepared densities win perfectly with sharp trace-entropy deficit $\Delta_* = \ln(7/6) - \ln(2)/7$. Type-$\mathrm{II}_\infty$ permits every finite trace entropy. A graded modular ladder gives a normal semifinite hosting criterion. Its scalar fingerprint has sharp capacity and entropy-gap bounds; the capacity floor is approached but not attained in type I, and attained in type II.

1 The separation

A finite nonlocal game is specified by finite question and answer sets, a probability distribution π on question pairs, and a predicate V taking values in $\{0, 1\}$. A pure strategy consists of a unit vector $\psi \in H_A \otimes H_B$ and local positive operator-valued measurements (POVMs). It is *perfect* if its winning probability is one. We call it finite-dimensional when both local spaces are finite-dimensional, and spatial when they may be separable infinite-dimensional Hilbert spaces. The associated correlation sets are C_q and C_{qs}, respectively. Answer sets may depend on the question; padding by zero effects puts them in the usual fixed-alphabet convention.

Let $D_G(\varepsilon)$ be the least integer d for which a strategy on $\mathbb{C}^d \otimes \mathbb{C}^d$ with arbitrary POVMs has loss at most ε. Smaller spaces can be padded. Mixed shared states are also allowed in this definition. Throughout, dimension refers to this POVM convention.

For this game finite dimension cannot pay for perfection, whereas finite oscillator energy and finite ordinary entanglement entropy suffice. Among perfect pure spatial strategies, the displayed

The authors developed this work with AI models from more than one developer and take responsibility for every claim. How the work was done, and what the machine checks cover, is described at the end of the paper.

strategy minimizes entanglement entropy at $\log_2 24 + 7/9$ ebits and mean oscillator occupation at $275/18$ per side. The reason is an integer spectral classification with a unique majorization maximum, not merely passivity at a fixed spectrum. Diffuse type-II algebras permit perfect prepared play, although their normalized trace states cannot win. The field construction uses arbitrary bounded local PVMs and an ideal joint preparation, with energy depending on the chosen field modes.

Theorem 1.1. *There is a rational binary-scoring game Γ with ten Alice questions, eleven Bob questions, at most six answers per question, and a uniform distribution on fifty-two supported pairs, such that:*

1. *Γ has a normalized perfect strategy on $\ell^2 \otimes \ell^2$;*
2. *no finite-dimensional strategy is perfect;*
3. *for each integer $K \geq 1$, an explicit strategy of local dimension $20K$ has loss*

$$\ell_{20K}^{\mathrm{cmp}} = \frac{1}{624(4^K - 1)}. \tag{1}$$

Both the finite-dimensional supremum and spatial value are one, but only the latter is attained, and

$$D_\Gamma(\varepsilon) \leq 20 \left\lceil \log_4 \left(1 + \frac{1}{624\varepsilon} \right) \right\rceil \qquad (\varepsilon > 0). \tag{2}$$

For $0 < \varepsilon < 1$ one also has $D_\Gamma(\varepsilon) \leq 10 \log_2(1/\varepsilon) + 20$.

Theorem 1.2 (Sharp dimension profile)**.** *There are constants $C > 0$ and $0 < \varepsilon_0 < 1/2$ such that every finite-dimensional strategy for Γ with loss at most ε, $0 < \varepsilon \leq \varepsilon_0$, using arbitrary local POVMs and a mixed or pure shared state, has each local dimension at least*

$$10 \log_2(1/\varepsilon) - 10 \log_2 \log_2(1/\varepsilon) - C. \tag{3}$$

Consequently,

$$D_\Gamma(\varepsilon) = 10 \log_2(1/\varepsilon) + O\big(\log \log(1/\varepsilon)\big) \qquad (\varepsilon \downarrow 0). \tag{4}$$

The lower bound controls all approximate strategies, through a robust sector law and a count of Schmidt multiplicities. The key analytic input is a dimension-independent square-root extraction for the row/column magic-square game, proved in Section 6. Sections 7 and 8 establish the robust sector law, gluing estimates and rank count. Local measurement dilation preserves the Schmidt rank used in this count, even when it enlarges the ambient Hilbert spaces.

Context and credit. Coladangelo and Stark separated finite-dimensional from spatial correlations with finite question and answer sets [10], using offset decompositions and Schmidt multisets. Beigi obtained such a separation with four binary measurements per party [3]. These are separations of exact correlations. Here every spatial correlation winning Γ perfectly lies outside C_q, so a finite game with a 0/1 winning predicate detects the distinction at perfect attainment. This answers the possibility raised in the second-page remark of Coladangelo–Stark [9] that a binary or integer-valued game score might not witness the separation. Binary scoring refers to the predicate, not to the number of answers.

Slofstra's nonclosure game [27] has a perfect approximate correlation but no perfect spatial strategy. His Theorem 3.2 equates finite and spatial perfect existence for linear-system games.

Kim–Paulsen–Schafhauser prove $C^s_q(n,m) = C^s_{qs}(n,m)$ for synchronous correlations in their Theorem 3.10 [20]. These restricted classes do not include the present construction.

Our building block is Mančinska's Flagged Magic-Square game [23], including its calibrated frustrated bridge. Her Theorem 11 is a finite-dimensional exact self-test. We prove the sector identity on arbitrary Schmidt supports without appealing to a spatial extension of that theorem. The exact magic-square rigidity used below is Proposition 2 and Appendix A.2 of that paper, also proved here in the operator form we need. Lalonde's earlier game [21] requires nonmaximal entanglement for perfect play, but has a perfect strategy in local dimension six. The additional step here is the offset gluing that rules out every finite dimension.

Quantitative dimension witnesses predate this construction; for example, Slofstra–Vidick [28] prove polynomial dimension lower bounds for finite games. Douglas [15, Theorem 6] states a matching logarithmic law for the optimal-value deficit of I_{3322}. We do not use that paper's lower-bound certificates. The present theorem gives the sharp leading coefficient for this explicit spatially perfect, non-finitely-perfect game of value one. No absolute historical-priority claim is made.

2 The game and its finite building block

2.1 The flagged magic square

Write $\mathcal{R} = \{s \in \{\pm1\}^3 : s_1s_2s_3 = 1\}$ and $\mathcal{C}_j = \{t \in \{\pm1\}^3 : t_1t_2t_3 = \eta_j\}$, where $(\eta_1,\eta_2,\eta_3) = (1,1,-1)$. Alice's questions are R_1,R_2,R_3,T_0,T_1; Bob's are C_1,C_2,C_3,U_0,U_1. Put

$$c_0 = r_0 = (1,1,1), \qquad c_1 = r_1 = (1,-1,-1), \qquad \sigma_{xy} = (-1)^{xy}.$$

The answers to R_i are $\mathcal{R}\sqcup\{f\}$, and those to C_j are $\mathcal{C}_j\sqcup\{f\}$. The answers to T_x are $\{+,-\}\sqcup(\mathcal{C}_1\setminus\{c_x\})$, and those to U_y are $\{+,-\}\sqcup(\mathcal{R}\setminus\{r_y\})$. A sign answer is distinct from a sign triple. All twenty-five pairs have equal positive probability. The rules are:

1. At (R_i,C_j), (f,f) wins, a single flag loses, and nonflag answers (s,t) win exactly when $s_j = t_i$.
2. At (T_x,C_1), a sign wins exactly against f or c_x. A triple wins exactly against the identical triple.
3. At (R_1,U_y), a sign for Bob wins exactly against f or r_y. A triple wins exactly against the identical triple.
4. At (T_x,U_y), two signs $a,b \in \{\pm1\}$ win exactly when $ab = \sigma_{xy}$. If at least one answer is a triple, they win.
5. All other pairs accept every answer pair.

This is G_{FMS}, in the conventions of [23].

Here is its ideal strategy. Let X,Y,Z be the Pauli matrices and use the square

$$(M_{ij}) = \begin{pmatrix} X\otimes I & I\otimes X & X\otimes X \\ I\otimes Z & Z\otimes I & Z\otimes Z \\ X\otimes Z & Z\otimes X & Y\otimes Y \end{pmatrix}. \tag{5}$$

Each row and column commutes; row products are I and column products are $I,I,-I$. Denote the rank-one joint spectral projections by $\pi^{\mathrm{row}}_{i,s}$ and $\pi^{\mathrm{col}}_{j,t}$. On $\mathbb{C}f\oplus\mathbb{C}^4$ take

$$\psi_* = \tfrac{1}{3}|ff\rangle + \tfrac{\sqrt{2}}{3}\sum_{z=0}^{3}|zz\rangle. \tag{6}$$

The row and column measurements consist of their spectral projections and the flag projection. Choose the distinguished eigenvectors as

$$|c_0\rangle = (1,0,1,0)^T/\sqrt{2}, \qquad |c_1\rangle = (0,1,0,1)^T/\sqrt{2},$$
$$|r_0\rangle = (1,1,1,1)^T/2, \qquad |r_1\rangle = (1,-1,1,-1)^T/2.$$

They satisfy $\langle c_x|r_y\rangle = \sigma_{xy}/\sqrt{2}$. The sign projections for T_x are onto $(|f\rangle \pm |c_x\rangle)/\sqrt{2}$, with the other column-one projections as the remaining outcomes. Bob uses $(|f\rangle \pm |r_y\rangle)/\sqrt{2}$ and the other row-one projections. All these projections are real. The row/column and calibration rules follow from their supports. For two signs the amplitude is proportional to $1 + ab\sigma_{xy}$, so the forbidden products vanish. Thus (6) wins perfectly.

2.2 Two copies and one label

Form Γ with Alice questions (h, q) and Bob questions (h, r), $h \in \{1, 2\}$, for gadget questions q, r. Bob has an additional question Lab with answers A, B. Copy one retains the gadget answers. Add a boundary answer $\perp$ to every question of copy two.

There are twenty-five internal pairs per copy and two label pairs:

$$((1, R_1), \mathsf{Lab}), \qquad ((2, R_1), \mathsf{Lab}).$$

Each of the fifty-two pairs has probability $1/52$. Copy one uses the gadget predicate. In copy two, $(\perp, \perp)$ wins, exactly one boundary answer loses, and otherwise use the gadget predicate. At the links accept exactly when

$$\begin{aligned} [a \neq f] &= [b = A] && \text{in copy one,} \\ [a \in \{f, \perp\}] &= [b = A] && \text{in copy two.} \end{aligned} \tag{7}$$

The first line concerns R_1, so $a \neq f$ means a magic-square answer. Unsupported pairs have probability zero. These rules specify a rational game with ten and eleven questions, five answers in copy one, six in copy two, and two label answers. The ancillary `game.json` records all supported pairs and forbidden answer pairs.

3 An inverse-free sector identity

Fix bases and identify ψ with its Hilbert–Schmidt coefficient operator $K : H_B \to H_A$. Bob's transposed effects are denoted by F. Then

$$p(E, F) = \|E^{1/2} K F^{1/2}\|_2^2, \qquad p(E, F) = 0 \implies EKF = 0. \tag{8}$$

Compress to the supports of the reduced densities. Now K is injective with dense range, and KK^*, K^*K are positive trace-class operators. Compression preserves correlations. No inverse of K will be used.

Lemma 3.1 (Checked classes)**.** *At a supported pair suppose the answers split into matching classes, with different classes always losing. In a perfect full-support strategy their coarse effects E_c, F_c are orthogonal projections, $E_c K = K F_c$, and E_c commutes with KK^*.*

Proof. The forbidden products and completeness give $E_c K = E_c K F_c = K F_c$. For $c \neq d$, $E_c E_d K = E_c K F_d = 0$. Dense range gives $E_c E_d = 0$, and completeness gives $E_c^2 = E_c$. Apply the same argument to K^* on Bob's side. Finally $E_c K K^* = K F_c K^*$ is self-adjoint, proving commutation. □

For orthogonal complete sectors satisfying $P_s K = K Q_s$, the compact operator K is the direct sum of its sector maps. Hence

$$\mathrm{Sch}(K) = \bigsqcup_s \mathrm{Sch}(P_s K Q_s). \tag{9}$$

Here Sch is the multiset of nonzero singular values with multiplicities. All sector vectors are unnormalized. Formula (9) holds also for countably infinite square-summable multisets, by the compact spectral theorem.

Lemma 3.2 (Exact magic-square algebra)**.** *Let reflections O_{ij} commute within rows and columns, with row products I and column products $I, I, -I$. There is a unitary identification with $\mathbb{C}^4 \otimes J$ carrying them to $M_{ij} \otimes I_J$ in* (5)*, for an arbitrary multiplicity space J.*

Proof. This is the operator rigidity underlying [23, Proposition 2 and Appendix A.2]. Put $a = O_{11}$, $b = O_{12}$, $c = O_{21}$, $d = O_{22}$. The cross commutations are $[a, b] = [a, c] = [b, d] = [c, d] = 0$. Parity gives $O_{13} = ab$, $O_{23} = cd$, $O_{31} = ac$, $O_{32} = bd$, and $acbd = O_{33} = -abcd$. Multiplication on the left by a and on the right by d gives $cb = -bc$. Using $[ac, bd] = 0$ and the cross commutations then gives $ad = -da$.

An anticommuting reflection pair is unitarily X, Z on $\mathbb{C}^2$ tensored with a multiplicity space: the second reflection exchanges the first's two eigenspaces. The second pair commutes with the first, so acts on its multiplicity space and permits the same construction. Taking (a, d) as the first-qubit (X, Z) and (b, c) as the second-qubit (X, Z) recovers (5). This works with infinite multiplicity. □

Proposition 3.3 (Flagged sector law)**.** *In every perfect finite-dimensional or spatial strategy for G_{FMS}, the flag and magic-square sector maps satisfy*

$$\mathrm{Sch}(K_m) = 4 \cdot (\sqrt{2}\ \mathrm{Sch}(K_f)). \tag{10}$$

Here $m \cdot S$ means m copies of S. The two sectors vanish together.

Proof. Flag consistency gives a common flag projection on each side, independent of the row or column, and $K = K_f \oplus K_m$. In the other sector, Lemma 3.1 makes the checked cell classes projections. A joint POVM with projection marginals is their joint spectral measurement: each fine effect is supported in the ranges specified by its marginals, and completeness recovers their products. Thus the cell reflections satisfy the row relations. Intertwining through K_m transfers Bob's column relations to Alice by dense-range cancellation, and conversely by injectivity. Lemma 3.2 on each side yields

$$K_m = I_4 \otimes D_m.$$

An intertwiner of the two full M_4 actions has this form by comparison of its sixteen matrix blocks. On its multiplicity supports D_m is injective with dense range.

Calibration identifies the sign supports as

$$S_x = P_f \oplus (|c_x\rangle\langle c_x| \otimes I), \qquad S'_y = Q_f \oplus (|r_y\rangle\langle r_y| \otimes I).$$

Put $X_x = E_x^+ - E_x^-$ and $Y_y = F_y^+ - F_y^-$, self-adjoint contractions on their sign supports. Identify each support with flag space direct sum multiplicity space and write

$$W_{xy} = S_x K S'_y = K_f \oplus \frac{\sigma_{xy}}{\sqrt{2}} D_m, \qquad H = K_f \oplus \frac{D_m}{\sqrt{2}}. \tag{11}$$

Each W_{xy} is injective with dense range. The two forbidden sign events imply

$$(I + X_x)W_{xy}(I - \sigma_{xy}Y_y) = 0,$$
$$(I - X_x)W_{xy}(I + \sigma_{xy}Y_y) = 0.$$

Their sum and difference give both identities

$$X_x W_{xy} = \sigma_{xy} W_{xy} Y_y, \qquad X_x W_{xy} Y_y = \sigma_{xy} W_{xy}. \tag{12}$$

Consequently $X_x^2 W_{xy} = W_{xy}$ and $W_{xy} Y_y^2 = W_{xy}$. Dense range and injectivity give $X_x^2 = Y_y^2 = I$. This proves projectivity of the unchecked signs with bounded operators alone.

For $y = 0$, $W_{00} = W_{10} = H$, so $X_0 = X_1 =: X$. For $y = 1$, $W_{01} = H$ and $W_{11} = \Delta_A H = H\Delta_B$, where $\Delta = I \oplus (-I)$ on each side. The intertwining equations imply $X = -\Delta_A X \Delta_A$. Thus

$$X = \begin{pmatrix} 0 & \beta^* \\ \beta & 0 \end{pmatrix}, \qquad \beta^*\beta = I, \quad \beta\beta^* = I.$$

The map β is unitary between flag and multiplicity spaces. From $XH = HY_0$ and self-adjointness follows $XHH^* = HH^*X$. Its off-diagonal blocks give

$$\beta K_f K_f^* = \tfrac{1}{2} D_m D_m^* \beta.$$

Their positive eigenvalues have this ratio. Taking square roots and counting the four copies in K_m proves (10). Unitarity of β also makes the sectors vanish together. □

This directly proves the spatial law without extending a finite-dimensional self-test by assumption. It makes no assertion about general commuting-operator strategies.

4 Gluing and the finite-dimensional obstruction

The support graph has Alice and Bob questions as its two vertex sets and supported pairs as edges; delete unused questions. Its connectedness is the hypothesis needed to propagate a common boundary projection.

Theorem 4.1 (Connected-support gluing)**.** *Let G be a finite rational game with connected support. At a supported pair (z, z'), suppose its answers split into consistent classes U, L. Assume every finite-dimensional perfect full-support strategy satisfies*

$$\mathrm{Sch}(K_U) = m \cdot (\kappa\, \mathrm{Sch}(K_L)), \qquad m \in \mathbb{N}_{\geq 1}, \quad \kappa > 1, \tag{13}$$

with both sectors nonzero. Suppose G has a perfect strategy with one Schmidt coefficient l in L and m coefficients κl in U.

Take two copies of the supported predicate, add the boundary answer to copy two, and link $(1, z), (2, z)$ to one Bob label question by $[a \in U] = [b = A]$ and $[a \in L \text{ or } a = \perp] = [b = A]$. Give all these pairs positive rational probabilities. The resulting game is spatially perfect, never finitely perfect, and has finite-dimensional supremum one. If (13) holds spatially, the multiset identities below hold for every perfect spatial strategy.

Proof. In copy two, boundary/nonboundary is a checked decomposition on every edge. Lemma 3.1 and connectedness propagate common projections $P_\perp, Q_\perp$ with $P_\perp K = K Q_\perp$. All other copy-two effects lie in their complements; conditioning on this complement when nonzero gives a perfect strategy for G.

Let $S_1 = \mathrm{Sch}(K_L^1)$ and $S_2 = \mathrm{Sch}(K_L^2)$ be the unnormalized lower sectors, and $T = \mathrm{Sch}(P_\perp K)$. The links give equalities of projections,

$$P_U^1 = P_L^2 + P_\perp, \qquad P_L^1 = P_U^2. \tag{14}$$

For example both sides of the first equality intertwine K with Bob's label-A projection, so dense range gives equality. Equations (9) and (13) imply

$$S_1 = m \cdot (\kappa S_2), \qquad m^2 \cdot (\kappa^2 S_2) = S_2 \sqcup T. \tag{15}$$

The nonboundary sector cannot vanish: the second link would kill copy one's lower sector, and its sector law would kill the whole state. Thus S_2 is nonempty. In finite dimension its least element $\mu > 0$ belongs to the left side of the second identity, so $\mu = \kappa^2 s$ for some $s \in S_2$. This contradicts $s = \mu/\kappa^2 < \mu$.

For spatial existence use basis vectors $A(i,j,k)$, $1 \le i,j \le m$, and $B(j,k)$, $1 \le j \le m$, $k \ge 0$. Give them respective Schmidt coefficients a_k, b_k, where

$$a_k = \kappa b_k, \qquad b_{k+1} = b_k/\kappa^2, \qquad b_0^2 = \frac{1-\kappa^{-4}}{m^2\kappa^2 + m}.$$

The squared coefficients sum to one. Copy one uses ideal blocks with $B(j,k)$ as L and $A(i,j,k)$ as U. Copy two uses $A(1,1,k+1)$ as L and $B(j,k)$ as U, since $b_k = \kappa a_{k+1}$. All other A vectors are boundary in copy two. Direct sums of ideal measurements, with the boundary projection, are bounded complete POVMs. Bob's label distinguishes A and B. Each internal pair wins on ideal blocks or boundary/boundary, and the two linked Alice projections equal the A projection. The diagonal state therefore wins both links. Finite truncation and compressed measurements give loss tending to zero by Lemma 5.1 below. □

Connectedness is explicit: on disconnected support, consistency gives one boundary projection per component and does not alone give a global projection. The concrete gadget has full pair support, so this issue does not arise for Γ.

Apply Theorem 4.1 with $m = 4$, $\kappa = \sqrt{2}$, $z = R_1$, $z' = C_1$, $U = \mathcal{R}$, $L = \{f\}$. The sector law is Proposition 3.3 and the ideal strategy is (6). In particular

$$S_1 = 4 \cdot (\sqrt{2} S_2), \qquad 16 \cdot (2 S_2) = S_2 \sqcup T. \tag{16}$$

This proves finite impossibility without assuming positive boundary weight. In fact the weight is forced: copy one's law gives $\|S_1\|_2^2 = 1/9$, the second link gives $\|S_2\|_2^2 = 1/72$, and the first gives

$$\|T\|_2^2 = 8/9 - 1/72 = 7/8. \tag{17}$$

This mass calculation does not replace the minimum-coefficient argument.

5 Perfect spatial play and finite compression

5.1 The normalized state

For the concrete game index i, j by $0, 1, 2, 3$. On each side use the basis $A(i, j, k), B(j, k)$, $k \geq 0$, and put

$$\begin{aligned} \psi &= \sum_{k\geq 0} \sum_{i,j=0}^{3} a_k |A(i,j,k)A(i,j,k)\rangle \\ &\quad + \sum_{k\geq 0} \sum_{j=0}^{3} b_k |B(j,k)B(j,k)\rangle, \\ b_k &= \frac{2^{-k}}{4\sqrt{3}}, \qquad a_k = \sqrt{2} b_k. \end{aligned} \tag{18}$$

The squared mass of level k is $16a_k^2 + 4b_k^2 = (3/4)4^{-k}$, so $\|\psi\| = 1$. Each local space is isomorphic to ℓ^2; there are two Hilbert tensor factors.

Copy one applies the ideal PVMs on each block

$$\big(B(j,k), A(0,j,k), A(1,j,k), A(2,j,k), A(3,j,k)\big).$$

Copy two applies them on

$$\big(A(0,0,k+1), B(0,k), B(1,k), B(2,k), B(3,k)\big),$$

and outputs $\perp$ on every remaining A vector. The coefficient ratio in each block is $(1, \sqrt{2}, \sqrt{2}, \sqrt{2}, \sqrt{2})$ since $b_k = \sqrt{2} a_{k+1}$. Bob labels the A and B subspaces. Alice's copy-one nonflag and copy-two flag-plus-boundary projections both equal the A projection; the diagonal state satisfies the links.

This covers all supported pairs. In copy one the nine row/column pairs, four calibration pairs and four sign pairs win by the ideal strategy; the remaining eight are free. In copy two all twenty-five pairs win on ideal blocks or boundary/boundary, with no mixed amplitude. Both links win as just shown. Thus (18) is a normalized perfect spatial strategy.

5.2 Loss under compression

Lemma 5.1 (Tail-to-loss bound)**.** *A perfect spatial strategy with Schmidt tail mass t outside a finite retained set has a normalized truncation with compressed measurements of loss at most t.*

Proof. Its loss operator satisfies $0 \leq L \leq I$ and $L\psi = 0$. The normalized truncation ϕ has overlap $\sqrt{1-t}$ with ψ. For $0 < t < 1$ write $\phi = \sqrt{1-t}\psi + \sqrt{t}\eta$ with $\eta \perp \psi$ unit. Then $\langle \phi, L\phi\rangle = t\langle \eta, L\eta\rangle \leq t$. Compression of the local effects preserves this expectation and completeness. The zero-tail case is immediate. □

Retaining levels $0, \ldots, K-1$ keeps $20K$ coefficients and mass $1 - 4^{-K}$. Compressed effects are POVMs; the sign effects at the incomplete last block need not be projections. The lemma gives

$$D_\Gamma(\varepsilon) \leq 20\lceil \log_4(1/\varepsilon)\rceil \leq 10 \log_2(1/\varepsilon) + 20 \quad (0 < \varepsilon < 1).$$

The constant 10 uses arbitrary POVMs. A PVM-only convention requires a dilation and its dimension cost.

The compressed loss can be calculated exactly. All copy-one blocks remain intact. The only incomplete copy-two block is the last B quadruple, whose next-level flag is removed. Its unnormalized coefficient matrix is $\mathrm{diag}(0, b_{K-1}, b_{K-1}, b_{K-1}, b_{K-1})$. Only the four sign pairs lose. For forbidden answer signs u, v, $uv = -\sigma_{xy}$, and the probability is

$$\left| \frac{uv\, b_{K-1}}{2} \langle c_x | r_y \rangle \right|^2 = \frac{b_{K-1}^2}{8}.$$

Each of four pairs has two forbidden sign choices, giving total unweighted mass b_{K-1}^2. All row/column, calibration, boundary and link events still have zero loss. Dividing by 52 and the retained mass gives

$$\ell_{20K}^{\mathrm{cmp}} = \frac{b_{K-1}^2}{52(1-4^{-K})} = \frac{4^{-(K-1)}}{48 \cdot 52(1-4^{-K})} = \frac{1}{624(4^K-1)}.$$

Solving for K gives (2) and completes Theorem 1.1. This is the loss of the specified compression; no optimality among all strategies of dimension $20K$ is asserted.

6 Quantitative consistency and magic-square extraction

In this and the next section Bob's operators are the physical operators on H_B, without transposition. Operators on opposite parties commute. For a vector v, $\mathrm{Sch}(v)$ now denotes its nonincreasing Schmidt sequence, padded with zeros. Multiset unions are sorted before comparison. All vectors other than the original strategy vector may be unnormalized. Singular-value perturbation gives

$$\| \mathrm{Sch}(v) - \mathrm{Sch}(w) \|_2 \le \|v - w\|. \tag{19}$$

Indeed, the matrix singular-value inequality follows by applying the Hermitian eigenvalue inequality to the self-adjoint block matrices with off-diagonal blocks K, K^*; finite-rank approximation gives the compact Hilbert–Schmidt case. Sorting nonnegative sequences also cannot increase their Euclidean distance.

Lemma 6.1 (Approximate checked classes)**.** *For projective measurements, suppose mismatched answer classes at a pair always lose, with absolute losing mass at most η on a vector v. Their class projections satisfy $\|(E_c - F_c)v\| \le \sqrt{\eta}$. For the sign tests of G_{FMS}, writing $S_x = E_+^{T_x} + E_-^{T_x}$, $X_x = E_+^{T_x} - E_-^{T_x}$ and similarly S'_y, Y_y for Bob, one has*

$$\|(X_x Y_y - \sigma_{xy} S_x S'_y)v\| \le 2\sqrt{\eta}.$$

Proof. The squared norm of $(E_c - F_c)v$ is the sum of the probabilities of the disjoint mismatched-class events involving c. For signs, expand in the orthogonal vectors $E_a^{T_x} F_b^{U_y} v$ with $a, b \in \{+, -\}$. The coefficient $ab - \sigma_{xy}$ is zero on winning events and has magnitude two on losing events. □

Left multiplication by a contraction preserves a vector estimate. Right multiplication needs additional control. The following rule will provide it. If s are Alice reflections with Bob mirrors $\bar{s}$ satisfying $\|(s - \bar{s})v\| \le e$, then a word w of length l mirrors to the reversed Bob word with error at most le. Consequently, for an Alice operator R with $\|R\| \le 2$ and $\|Rv\| \le re$,

$$\|Rwv\| \le (r + 2l)e. \tag{20}$$

To see this, replace the rightmost letter by its mirror, commute that mirror past the remaining Alice letters, and continue. The resulting Bob word commutes with R and is unitary. The rule holds with the parties interchanged; the mirroring assertion also holds for contractions.

Theorem 6.2 (Square-root word extraction)**.** *Let $\|\phi\| \le 1$ and let local PVMs have valid magic-square row and column answers, with absolute losing mass at most η on each of the nine pairs. There are local isometries $V = V_A \otimes V_B$ and an auxiliary vector j such that, for every word w of length at most six in Bob's cell reflections and cell projections, including the empty word,*

$$\|Vw\phi - (I \otimes \widehat{w})(\Phi_4 \otimes j)\| \le C_{\mathrm{MS}}\sqrt{\eta}, \qquad C_{\mathrm{MS}} = 2^{24}. \tag{21}$$

Here $\Phi_4 = \frac{1}{2}\sum_{z=0}^{3} |zz\rangle$, and $\widehat{w}$ uses the ideal cells in (5), in the same order. The vector j need not be normalized. The constant is independent of dimension, also on separable Hilbert spaces.

This statement is proved below for the actual row/column game. For comparison, the analytic estimate of Wu et al. [34, Appendix C, Eq. (C35)] has a fourth-root rate in a correlation-error parameter; per-pair loss η gives checked correlation errors at most 2η. Coladangelo–Stark's magic-square statement is [11, Theorem 6.9], applying their general Theorem 4.16 to an equation/variable game. Their Definition 4.14 uses trace norm for the state and squared state-dependent operator distance. We do not use its asserted linear trace-norm rate: Lemma A.3 there fails already for $v_t = \sqrt{1-t}|0\rangle + \sqrt{t}|1\rangle$ and $P = |0\rangle\langle 0|$, since the squared projector defect is t but $\||v_t\rangle\langle v_t| - P\|_1 = 2\sqrt{t} > 6t$ for $0 < t < 1/9$.

Proof of Theorem 6.2. Let A_{ij}, B_{ij} be the cell reflections and put $e = 2\sqrt{\eta}$. Consistency gives $\|(A_{ij} - B_{ij})\phi\| \le e$. Alice's rows commute and have product I; Bob's columns commute and have products $I, I, -I$. Write $T \simeq_r U$ for $\|(T-U)\phi\| \le re$. Put

$$a = A_{11}, \quad b = A_{12}, \quad c = A_{21}, \quad d = A_{22}, \quad g = A_{31}, \quad h = A_{32}, \quad i = A_{33}.$$

Mirroring two letters to a Bob column, using its parity and mirroring back, gives

$$[a,c] \simeq_4 0, \qquad [b,d] \simeq_4 0,$$
$$g \simeq_3 ac, \qquad h \simeq_3 bd, \qquad i \simeq_3 -abcd.$$

Also $[a,b] = [c,d] = [g,h] = 0$ and $gh = i$ exactly. Substitution of h on the right, then of g through its two-letter suffix using (20), gives $gh \simeq_{10} acbd$. Thus $acbd \simeq_{13} -abcd$. Remove the left a and the right d, the latter by (20), to obtain $\{b,c\} \simeq_{15} 0$.

Similarly $hg \simeq_{10} bdac$, so $acbd \simeq_{20} bdac$. Commuting b, d through the suffix ac gives $bdac \simeq_8 dabc$; exact $[c,d] = 0$ followed by commuting b, d through c gives $abcd \simeq_6 adbc$. Combining these with the $13e$ estimate yields $(ad + da)bc \simeq_{47} 0$. Removing bc costs at most $4e$, so $\{a,d\} \simeq_{51} 0$. Transferring two-letter relations to Bob costs at most $4e$ more. Therefore $55e$ bounds all the corresponding commutation and anticommutation defects on both sides.

Consider the order-32 real Pauli group

$$G = \langle J, a, b, c, d : J^2 = a^2 = b^2 = c^2 = d^2 = 1,\ J \text{ central},$$
$$ad = Jda,\ bc = Jcb,\ \text{all other pairs commute}\rangle.$$

Every element has a unique form $J^z a^u b^v c^w d^x$ with binary exponents. Let f_A, f_B evaluate these forms in the physical generators with $J = -I$. Sorting a product involves at most eight letters and at most 32 adjacent exchanges. Each exchange costs at most $(55 + 16)e$ by (20), and cancellations are exact. In particular, on either party,

$$\|(f(g)f(h) - f(gh))\phi\| \le 4096e \qquad (g, h \in G). \tag{22}$$

Define isometries

$$W_A v = 32^{-1/2} \sum_{g \in G} |g\rangle \otimes f_A(g) v, \qquad W_B v = 32^{-1/2} \sum_{g \in G} |g\rangle \otimes f_B(g) v.$$

On the register use the right regular representation $(R_h F)(g) = F(gh)$. Squaring and averaging (22) over g gives

$$\|W_A f_A(h)\phi - R_h W_A \phi\| \le H, \qquad H = 4096e, \tag{23}$$

and likewise for Bob. An opposite-party isometry preserves these errors.

Since $f(Jg) = -f(g)$ exactly, each isometry has range in the register's $J = -1$ subspace. Every representation on this subspace is a multiple of the real four-dimensional representation $a = XI, b = IX, c = IZ, d = ZI$. For completeness, diagonalize the commuting c, d. The involution a exchanges the two d eigenspaces and preserves c; b exchanges the two c eigenspaces and preserves d. The four joint eigenspaces have equal multiplicity, and their identifications by a, b give precisely these Pauli actions. This also works for infinite multiplicity. Compose W_A, W_B with these identifications and call the result V.

Put $v = V\phi$. For each of the four generators, consistency and (23) give $\|(M_s^A - M_s^B)v\| \le 2H + e$. The four involutions $M_s^A M_s^B$ commute, with common $+1$ projection $P_0 = |\Phi_4\rangle\langle\Phi_4| \otimes I$. Hence

$$\|(I - P_0)v\|^2 \le \frac{1}{4} \sum_s \|(I - M_s^A M_s^B)v\|^2 \le (2H + e)^2.$$

Taking $\Phi_4 \otimes j = P_0 v$ proves the empty-word bound $\|v - \Phi_4 \otimes j\| \le 8193e$.

Every other Alice cell equals, or is within $3e$ of, its normal-form word of degree at most four. Its extraction error is thus at most $H + 3e + 8193e$. Bob's cells follow from consistency and $(M \otimes I)\Phi_4 = (I \otimes M^T)\Phi_4$, at a further cost of at most e. The common bound $2^{14}e$ covers a single reflection or projection on either side; a cell projection is $(I \pm M)/2$.

For a Bob word, mirror its last letter to Alice at cost at most e, commute it past the Bob prefix, apply the prefix estimate, and use the Alice single-letter estimate on $\Phi_4 \otimes j$. Passing from the physical isometric action to that target costs at most another $8193e$. Thus each added letter costs at most $e + 8193e + 2^{14}e$. For six letters, including the empty-word error, this is less than $2^{18}e$, which is covered by $2^{24}\sqrt{\eta}$. All estimates used absolute losing masses and $\|\phi\| \le 1$ throughout; there is no normalization factor or dimension dependence. □

7 The robust flagged sector law

We write $v \approx w$ in this section when $\|v - w\| \le C\sqrt{\eta}$, for a fixed constant independent of the strategy and its dimension. The constant may increase at successive steps: only finitely many contractions, bounded sums and words of bounded length occur.

Lemma 7.1 (Two-way spectral domination)**.** *Suppose local contractions satisfy* $\|v - (A \otimes B)w\| \le d_1$ *and* $\|w - (A' \otimes B')v\| \le d_2$. *Then*

$$\|\operatorname{Sch}(v) - \operatorname{Sch}(w)\|_2 \le \sqrt{d_1^2 + d_2^2}.$$

Proof. Let $u = \operatorname{Sch}(v)$, $t = \operatorname{Sch}(w)$, and let s, s' be the Schmidt sequences after the two contractions, respectively. Singular values decrease entrywise under contractions, so $s \le t$, $s' \le u$. Equation (19) gives $\|(u - t)_+\|_2 \le d_1$ and $\|(t - u)_+\|_2 \le d_2$. These positive parts have disjoint supports. The same argument for compact Hilbert–Schmidt maps follows by finite-rank approximation. □

Theorem 7.2 (Robust flagged sector law)**.** *There is an absolute constant C_{F} such that every projective G_{FMS} strategy on a vector ψ of norm at most one, with absolute losing mass at most η on each of its twenty-five pairs, satisfies*

$$\left\|\mathrm{Sch}(m) - 4\cdot(\sqrt{2}\,\mathrm{Sch}(f))\right\|_2 \le C_{\mathrm{F}}\sqrt{\eta}, \quad f = E_f^{R_1}\psi, \quad m = \psi - f. \tag{24}$$

This holds in finite dimension and on separable spatial Hilbert spaces.

Proof. Put $P = E_f^{R_1}$, $Q = F_f^{C_1}$, and write P^i, Q^j for the other row and column flags. Let $H_y = E_{r_y}^{R_1}$, $G_x = F_{c_x}^{C_1}$. Cell projections, denoted $E_{i;j,s}, F_{j;i,s}$, sum the nonflag answers with cell sign s. Lemma 6.1 gives

$$\begin{aligned} P^i\psi \approx Q^j\psi \approx P\psi, \qquad & E_{i;j,s}\psi \approx F_{j;i,s}\psi, \\ Qm \approx P^i m \approx Q^j m \approx 0, \qquad & Pm = 0. \end{aligned} \tag{25}$$

The calibrations and signs give

$$S_x\psi \approx (Q + G_x)\psi, \quad S'_y\psi \approx (P + H_y)\psi, \quad X_xY_y\psi \approx \sigma_{xy}S_xS'_y\psi. \tag{26}$$

Mirrors and the nonflag game. If $L = F_{j;i,s}$ and $M = E_{i;j,s}$, then $L\psi \approx M\psi$. Moreover

$$PM\psi \approx PL\psi = LP\psi \approx LQ^j\psi = 0, \qquad MP\psi \approx MP^i\psi = 0.$$

Thus $Lm \approx Mm$. Replacing rightmost letters and commuting the opposite-party mirrors proves this assertion for every bounded-length Bob word, in reversed Alice order. The corresponding reflection estimate follows by bounded linear combinations.

Merge each row or column flag into a fixed valid nonflag answer, with the row-one target outside $\{r_0, r_1\}$ and the column-one target outside $\{c_0, c_1\}$. The merged measurements are PVMs with valid parities. For a nonflag losing event Λ, its amplitude on m differs from its amplitude on ψ by $\Lambda P\psi \approx \Lambda P^i\psi = 0$. Events involving merged flags have amplitudes $O(\sqrt{\eta})$ by (7). There are a fixed number of outcomes, so the losing *mass* on m is $O(\eta)$ on every pair. Theorem 6.2 therefore applies at square-root rate. An unmerged word differs from the merged word by a bounded sum of terms containing a column flag. Mirror the suffix to Alice, commute it past that flag and use $Q^j m \approx 0$. Each term is $O(\sqrt{\eta})$. Consequently the extraction applies to the unmerged words below as well.

The needed ideal words. Let

$$\widetilde{B}_y = F_{3;1,(r_y)_3}F_{2;1,(r_y)_2}F_{1;1,(r_y)_1}.$$

It mirrors H_y on m. Define the physical reflections

$$A = E_{1;2,+} - E_{1;2,-} + P, \quad B_{12} = F_{2;1,+} - F_{2;1,-} + Q^2, \quad B = F_{1;2,+} - F_{1;2,-} + Q.$$

In the ideal square $A = IX$, $B = IZ$. The exact ideal identities are

$$\begin{aligned} \widetilde{B}_0G_1 &= \widetilde{B}_0G_0B_{12}, & \widetilde{B}_1G_1 &= -\widetilde{B}_1G_0B_{12}, \\ \widetilde{B}_1G_0 &= B\widetilde{B}_0G_0. \end{aligned} \tag{27}$$

Indeed, $IXc_0 = c_1$, $IXr_y = (-1)^y r_y$, $IZr_0 = r_1$ and $IZc_0 = c_0$. The reversed order under Bob's transposed action on Φ_4 is essential. A column answer projector uses two cell projections because of parity, so the longest word displayed has length six. Theorem 6.2 transports these identities to m. For its auxiliary vector j the same theorem and singular-value perturbation give

$$\mathrm{Sch}(m) \approx 4\cdot(\mathrm{Sch}(j)/2), \qquad \mathrm{Sch}(\widetilde{B}_0G_0m) \approx \mathrm{Sch}(j)/(2\sqrt{2}). \tag{28}$$

For the second equality, the product of the rank-one row and column projections has singular value $|\langle c_0|r_0\rangle| = 1/\sqrt{2}$; the factor $1/2$ comes from Φ_4.

Frustration on the sign supports. Put $z_{xy} = S_x S'_y m$, $u = S_0 S'_0 \psi$ and $\Delta_A = 2P - I$, $\Delta_B = 2Q - I$. Calibration implies $S_x f \approx S'_y f \approx f$. For example, $S_x P\psi \approx S_x Q\psi = QS_x\psi \approx Q(Q + G_x)\psi \approx f$. Using mirrors on m then gives

$$z_{xy} \approx \widetilde{B}_y G_x m, \qquad Pz_{xy} \approx Qz_{xy} \approx 0, \qquad S_x S'_y \psi \approx f + z_{xy}. \tag{29}$$

For P annihilation commute P through the Bob word and use $Pm = 0$. For Q annihilation first mirror the *whole* Bob word to Alice, then commute Q through that word and use $Qm \approx 0$. The three identities (27), consistency $Am \approx B_{12} m$, and $Af = f$, $Bf \approx f$ yield

$$S_1 S'_0 \psi \approx Au, \quad S_0 S'_1 \psi \approx Bu, \quad S_1 S'_1 \psi \approx AB\Delta_A u. \tag{30}$$

Here A commutes with P exactly, and $\Delta_A u \approx \Delta_B u$ by both annihilations in (29).

Set $X' = AX_1A$, $Y' = BY_1B$. Substituting (30) in the four sign tests and conjugating by the indicated reflections gives

$$X_0 Y_0 u \approx u, \quad X' Y_0 u \approx u, \quad X_0 Y' u \approx u, \quad \Delta_A X' \Delta_A Y' u \approx -u. \tag{31}$$

Subtract the first two equations and apply Y_0, using $Y_0^2 u = u$, to get $(X_0 - X')u \approx 0$. This also holds on $\Delta_A u$: replace $\Delta_A u$ by $\Delta_B u$ and commute Δ_B past the Alice operators. Replace X' in the fourth equation and add the third. To remove Y', note that

$$Y'^2 u = BS'_1 Bu \approx BS'_1 (S_0 S'_1 \psi) = B(S_0 S'_1 \psi) \approx u.$$

It follows that

$$(X_0 + \Delta_A X_0 \Delta_A)u = 2(PX_0P + P^\perp X_0 P^\perp)u \approx 0.$$

Since $u \approx f + z_{00}$ and $Pz_{00} \approx 0$, the two orthogonal range blocks imply $PX_0 f \approx 0$ and $P^\perp X_0 z_{00} \approx 0$. Projecting $X_0 Y_0 u \approx u$ onto these ranges, and using that Y_0 commutes with P, proves the two directions

$$X_0 Y_0 f \approx z_{00}, \qquad X_0 Y_0 z_{00} \approx f. \tag{32}$$

For example the P projection identifies $PX_0Y_0z_{00}$ with f, and its complementary component is small by the preceding block estimate; the other direction follows from the complementary projection.

Spectra. The operators X_0, Y_0 are contractions, so Lemma 7.1 gives $\mathrm{Sch}(f) \approx \mathrm{Sch}(z_{00})$. Combine (29) and (28): four copies of $\mathrm{Sch}(j)/2$ are four copies of $\sqrt{2}\,\mathrm{Sch}(j)/(2\sqrt{2})$. This proves (24). All operations have dimension-independent bounds and bounded word lengths. They establish a finite absolute C_{F}; no useful numerical value for this constant is asserted. □

8 Gluing estimates and logarithmic rank

Lemma 8.1 (Approximate gluing)**.** *There is a constant C_{G} such that a unit-vector PVM strategy for Γ of loss at most ε has the following properties. Let L_A, L_B be Bob's two label-answer projections, put*

$$\alpha = L_A\psi, \quad \beta = L_B\psi, \quad \phi_2 = E_f^{(2,R_1)}\psi, \qquad c = \mathrm{Sch}(\alpha), \quad b = \mathrm{Sch}(\beta), \quad a = \mathrm{Sch}(\phi_2).$$

The subscripts on L_A, L_B denote answers, not tensor factors. With $P_\perp = E_\perp^{(2,R_1)}$, $Q_\perp = F_\perp^{(2,C_1)}$, $t = \mathrm{Sch}(P_\perp Q_\perp \psi)$ and $\delta = C_{\mathrm{G}}\sqrt{\varepsilon}$,

$$\|c - (a \sqcup t)\|_2 \le \delta, \tag{i}$$

$$\|c - 4 \cdot (\sqrt{2}b)\|_2 \le \delta, \tag{ii}$$

$$\|b - 4 \cdot (\sqrt{2}a)\|_2 \le \delta, \tag{iii}$$

$$\|\,\mathrm{Sch}(\psi) - (c \sqcup b)\|_2 \le \delta, \tag{iv}$$

$$\|\alpha\|^2 + \|\beta\|^2 = 1. \tag{v}$$

Proof. Every supported pair has losing mass at most 52ε. Link one identifies the first-copy magic-square vector with α and its flag vector with β, with vector error $O(\sqrt{\varepsilon})$. Apply Theorem 7.2 and singular-value perturbation to obtain (ii).

For (iii) remove the boundary coherently: put $\psi' = (I - P_\perp)\psi$. Boundary consistency on all twenty-five copy-two pairs identifies each boundary vector with $P_\perp\psi$ within $O(\sqrt{\varepsilon})$. In particular every row or column boundary projection applied to ψ' is $O(\sqrt{\varepsilon})$. Merge the boundary into a valid retained answer, choosing a magic-square answer at a row or column and an ordinary triple at a sign question. For a nonboundary losing projector Λ, $\Lambda P_\perp \psi$ is small: replace $P_\perp$ by the boundary projection at the relevant opposite-party question, where its product with Λ is zero. Thus every losing amplitude on ψ' is $O(\sqrt{\varepsilon})$ and the merged FMS losing mass is $O(\varepsilon)$. At R_1 its flag vector is exactly ϕ_2, and its nonflag vector is exactly $E_{\mathrm{MS}}^{(2,R_1)}\psi$, which link two identifies with β. The robust sector law proves (iii).

For (i), link two identifies α with $(E_f^{(2,R_1)} + P_\perp)\psi$. At the pair $((2,R_1),(2,C_1))$ the three classes flag, boundary and magic square are consistent, so this vector is within $O(\sqrt{\varepsilon})$ of

$$w = E_f^{(2,R_1)} F_f^{(2,C_1)} \psi + P_\perp Q_\perp \psi.$$

Its two summands are orthogonal blocks on *both* parties, hence their Schmidt lists form the union $\mathrm{Sch}(w)$. The first block is within $O(\sqrt{\varepsilon})$ of ϕ_2. Singular-value perturbation and sorting prove (i). For (iv) use the two-sided block vector

$$E_{\mathrm{MS}}^{(1,R_1)} L_A \psi + E_f^{(1,R_1)} L_B \psi.$$

Consistency makes this vector close to ψ and its blocks close to α, β, respectively. Finally (v) is exact completeness of the label PVM. A single enlarged constant covers all four errors. □

Lemma 8.2 (Descent with multiplicities)**.** *Suppose finitely supported nonnegative sorted sequences a, b, c, t satisfy (i)–(iii) in the infinity norm with error at most $\delta > 0$. Let $\mu_0 = \max a$. For every integer $k \ge 0$ with $\mu_0 2^{-k} \ge 64\delta$ there are distinct positive entries $\mu_k = a_{i_k}$ such that*

$$|\mu_k - \mu_0 2^{-k}| \le (2 + \sqrt{2})\delta. \tag{33}$$

If K is the number of these integers, then $K \ge \log_2(\mu_0/(64\delta))$. At each index i_k, four entries of b are within δ of $\sqrt{2}\mu_k$ and sixteen entries of c are within $(1+\sqrt{2})\delta$ of $2\mu_k$. These are disjoint index blocks for distinct k.

Proof. For a positive entry μ of a, select its occurrence p in the sorted union $a \sqcup t$, retaining occurrences when values coincide. Read (i) and (ii) at p to obtain $|b_{\lceil p/4 \rceil} - \mu/\sqrt{2}| \le \sqrt{2}\delta$. Read (iii) at $\lceil p/4 \rceil$ to find an entry of the zero-padded a within $g_0 = (1 + 1/\sqrt{2})\delta$ of $\mu/2$. Iteration contracts

the preceding error by $1/2$, so its total is at most $g_1 = 2g_0 = (2+\sqrt{2})\delta$. Retain exactly the integers in the statement. Their entries are positive, so none is a padded zero. Consecutive entries differ by at least $\mu_0 2^{-k}/2 - 2g_1 > 0$. Thus their a indices are distinct. If $\mu_0 \geq 64\delta$, the number retained is $\lfloor \log_2(\mu_0/(64\delta)) \rfloor + 1$; if not, take $K = 0$. Both conventions give the stated inequality.

Relation (iii) supplies the four b indices $4(i_k - 1) + 1, \ldots, 4i_k$. Relation (ii) supplies the sixteen c indices $16(i_k - 1) + 1, \ldots, 16i_k$, with the stated accumulated error. Distinct i_k give disjoint blocks, regardless of equal spectral values. □

Proof of Theorem 1.2. First consider a finite-dimensional mixed-state POVM strategy on local spaces of dimensions d_A, d_B. Loss is linear in the state, so some pure component has loss at most ε using the same measurements. Let its Schmidt rank be $r \leq \min(d_A, d_B)$. Give the finitely many local POVMs common Naimark dilations with fixed pure local ancillas. The resulting PVM strategy preserves the correlations and the Schmidt rank r. The dilated ambient dimensions may be larger and are never used as the original dimensions. Every one-sided projection and local contraction used below has Schmidt rank at most r.

Apply Lemma 8.1. Write $x = \|a\|_2$, $y = \|b\|_2$, $z = \|c\|_2$. Its relations give

$$z \leq \sqrt{8}y + \delta, \quad y \leq \sqrt{8}x + \delta, \quad y^2 + z^2 = 1.$$

Since $y^2 + (\sqrt{8}y + \delta)^2 \leq (3y + \delta)^2$,

$$y \geq (1 - \delta)/3, \qquad x \geq (1 - 4\delta)/\sqrt{72}. \tag{34}$$

Choose $\varepsilon_0 > 0$ so that $C_{\mathrm{G}}\sqrt{\varepsilon_0} \leq 1/100$. Then $x > 0.11$. As $\operatorname{rank}(\phi_2) \leq r$, $\mu_0 = \max a \geq 0.11/\sqrt{r}$.

The 2-norm estimates imply infinity-norm estimates, so Lemma 8.2 applies. All its selected b and c entries exceed δ: even the smallest chain entry is at least $(64 - 2 - \sqrt{2})\delta$. Relation (iv) then forces at least $20K$ positive entries in $\operatorname{Sch}(\psi)$. Indeed, match occurrences in the sorted union $c \sqcup b$; an occurrence larger than δ cannot match a zero with coordinate error at most δ. Multiplicity cannot be lost by coincident values. Consequently

$$\begin{aligned} r \geq 20K &\geq 20 \log_2 \frac{\mu_0}{64 C_{\mathrm{G}} \sqrt{\varepsilon}} \\ &\geq 10 \log_2(1/\varepsilon) - 10 \log_2 r - 20 \log_2(64 C_{\mathrm{G}}/0.11). \end{aligned} \tag{35}$$

Put $L = \log_2(1/\varepsilon)$. If $r \geq 10L$ the desired lower bound is immediate. Otherwise $\log_2 r \leq \log_2 L + \log_2 10$ in (35). Absorb the fixed constants to get $r \geq 10L - 10 \log_2 L - C$. This bounds each original local dimension. Together with Theorem 1.1 it proves the assertion for D_Γ and all the stated quantifiers. At zero loss the same descent has arbitrarily many positive entries; one treats that case directly, without taking a logarithm of zero. □

9 Exact tails and remaining questions

Proposition 9.1 (Exact perfect-state tails)**.** *Every normalized perfect spatial strategy has a nonempty sector multiset S_2 closed under halving. If $s_0 = \max S_2$, its Schmidt tail after deleting the r largest coefficients satisfies*

$$t_\psi(r) \geq 96 s_0^2 \, 4^{-\lceil r/20 \rceil} \geq 24 s_0^2 \, 2^{-r/10}. \tag{36}$$

Proof. The spatial sector law gives (16) for compact sector maps. If $s \in S_2$, the second identity puts s in $16 \cdot (2S_2)$, hence $s/2 \in S_2$. Square summability gives a largest coefficient $s_0 > 0$ and thus the distinct chain $s_0 2^{-k}$.

The full Schmidt multiset contains sixteen copies of $2s_0 2^{-k}$ from copy one's upper sector and four copies of $\sqrt{2}s_0 2^{-k}$ from its lower sector, for every k. This selected submultiset has squared mass $72s_0^2 4^{-k}$ per level, and tail $96s_0^2 4^{-K}$ after K full levels. Deleting r coefficients from the full multiset cannot leave less squared mass than optimally deleting r from this submultiset. Deleting $20\lceil r/20\rceil$ selected coefficients can only decrease the latter tail. This proves the first inequality; $\lceil r/20\rceil \le r/20 + 1$ gives the second. □

No nonzero finite-rank projection can commute with every Alice effect in a perfect full-support strategy. Projecting the vector onto it preserves all losing zeros and gives a nonzero finite-Schmidt-rank perfect vector. Bob can then be compressed to its finite Schmidt support, contradicting Theorem 1.1.

Proposition 9.1 concerns exact perfect states. Corollary 13.7 below determines the least discarded mass among perfect spectra. Neither tail statement identifies optimal game loss at a fixed finite dimension or compares every approximate strategy with a truncation of a perfect one.

Corollary 9.2 (Profile at dyadic accuracy)**.** *For $A_\Gamma(n) = \log_2 D_\Gamma(2^{-n})$ one has*

$$A_\Gamma(n) = \log_2 n + \log_2 10 + O\big((\log n)/n\big) \qquad (n \to \infty).$$

Proof. Theorems 1.1 and 1.2 give $D_\Gamma(2^{-n}) = 10n + O(\log n)$. Take logarithms. □

Finite energy and operator-algebraic play

10 Two oscillators and finite energy

All logarithms in entanglement entropies are base two; trace entropies and modular variables below use natural logarithms. Oscillator energies are occupation numbers, with the zero point omitted. Arbitrary bounded local measurements are allowed. In particular, these statements impose no Gaussian restriction or finite control-time bound.

Theorem 10.1 (An explicit finite-energy perfect strategy)**.** *The perfect strategy of* (18) *is a strategy on two harmonic oscillators with*

$$\langle N_A\rangle = \langle N_B\rangle = \frac{275}{18}, \qquad \langle N_A + N_B\rangle = \frac{275}{9}.$$

Every polynomial number moment is finite. These are the least possible occupations among perfect pure spatial strategies on two oscillators, and the passive minima for the displayed Schmidt spectrum. Retaining K levels gives the loss (1) *and mean occupation per side*

$$\langle N\rangle_K = \frac{275}{18} - \frac{20K}{4^K - 1}.$$

Proof. Identify $A(i,j,k)$ with $|20k + 4j + i\rangle$ and $B(j,k)$ with $|20k + 16 + j\rangle$. This is a bijection onto the oscillator basis. The probabilities, in decreasing order, are sixteen copies of $4^{-k}/24$ and then four of $4^{-k}/48$, at each $k \geq 0$. They are decreasing even across adjacent levels. Their level mass is $(3/4)4^{-k}$, with expected level $1/3$. The unnormalized mean offset in a level is

$$\frac{0 + \cdots + 15}{24} + \frac{16 + \cdots + 19}{48} = \frac{155}{24}.$$

Hence the mean is $20/3 + (155/24)(4/3) = 275/18$. Decreasing probabilities assigned to increasing energies minimize their mean: finite-dimensional rearrangement, followed by spectral cutoffs and monotone convergence, gives this passive lower bound for every placement. The displayed placement attains it. Global optimality over perfect pure strategies follows from the integer-spectrum argument in Theorem 10.4.

The generating function is

$$\langle z^N\rangle = \frac{\sum_{r=0}^{15} z^r/24 + \sum_{r=16}^{19} z^r/48}{1 - z^{20}/4}.$$

Its radius exceeds one, giving every polynomial moment. The local measurements are the orthogonal direct sums of the finite ideal PVMs already described in Section 5, plus the boundary and label projections. They are everywhere-defined projections of norm at most one, and their finite outcome sums are the identity in the strong topology. Thus the oscillator identification introduces no unbounded measurement-domain assumption. The block and link proof of Theorem 1.1 verifies all fifty-two pairs. Finally, conditioning the geometric level distribution on $k < K$ gives mean level $1/3 - K/(4^K - 1)$; the offset distribution is unchanged, proving the truncation formula. □

10.1 Attainment of bounded resource prices

For a game G, let $S_G(\varepsilon)$ be the infimum of pure-state entanglement entropy among spatial strategies of loss at most ε. Let $E_G(\varepsilon)$ be the corresponding total mean energy on a specified pair of nonnegative

local Hamiltonians, allowing normal mixed states and arbitrary POVMs. Empty infima are infinite. These definitions specify the state domain; a finite-dimensional infimum at loss zero would be a different quantity.

Theorem 10.2 (Three compactness criteria)**.** *For a finite game, bounded dimension as loss tends to zero is equivalent to finite-dimensional perfect attainment. A finite limit $s = \lim_{\varepsilon\downarrow 0} S_G(\varepsilon)$ is attained by a perfect spatial pure strategy of entropy s. If both local Hamiltonians have compact resolvent, a finite limit $e = \lim_{\varepsilon\downarrow 0} E_G(\varepsilon)$ is attained by a perfect strategy, and a perfect pure strategy has energy e as well.*

Proof. In a fixed finite dimension, density operators and the finitely many POVMs form compact sets and the loss is continuous. A convergent subsequence proves the dimension statement.

For entropy choose pure strategies with losses tending to zero and entropies tending to s. Transport their Schmidt supports to one common countable basis, with decreasing probabilities $p_i^{(n)}$. Since $p_i^{(n)} \le 1/i$, their tails past $r > 1$ are at most $(s+1)/\log_2 r$, after discarding a finite prefix. Coordinatewise subselection and this tightness give convergence in ℓ^1 to a probability vector p. The canonical pure vectors converge in norm because $(\sqrt{x} - \sqrt{y})^2 \le |x - y|$. Extract simultaneous ultraweakly convergent subsequences of all local effects. Positivity and completeness persist. Tensor products converge weakly: verify matrix elements on product vectors, then finite sums, and use uniform operator bounds and density. Combining this with norm convergence of vectors makes the limiting loss zero. Fatou gives $S(p) \le s$. Every perfect strategy bounds every approximate price from above, so equality holds. Transport to Schmidt supports avoids an unjustified extension of an infinite-dimensional basis map to a unitary on its original kernel.

For energy put $H = h_A \otimes I + I \otimes h_B$. Its bounded-energy spectral projections have finite rank. If $\operatorname{tr}(\rho H) \le e + 1$, then $\operatorname{tr}(\rho 1_{(R,\infty)}(H)) \le (e+1)/R$. Finite-rank compression and the square-root tail estimate give trace-norm precompactness of these density operators. Extract a trace-norm limit and weak limits of the measurements as above. The limit is perfect, and lower semicontinuity of energy, obtained by bounded spectral cutoffs, gives energy at most e; it must equal e. In a pure decomposition of this state, positivity of the loss makes every positive-weight component perfect. Each has energy at least e, so some, indeed almost every component of finite weight, has energy e. This proves the pure assertion without assuming the minimizing state was pure. □

These criteria separate attainment from computation. For Γ, the exact entropy and one-sided oscillator endpoints are computed in Theorem 10.5, using the spectral classification of Section 13.3.

Proposition 10.3 (A comparison with matched logarithmic games)**.** *Suppose a rational value-one game has matched finite-dimensional laws with $\alpha > 0$:*

$$\log_2 D_G(\varepsilon) = \alpha \log_2(1/\varepsilon) + o(\log(1/\varepsilon)), \qquad d_G^{\rightarrow}(\varepsilon) = \alpha \log_2(1/\varepsilon) + o(\log(1/\varepsilon)),$$

where $d_G^{\rightarrow}$ is the original-state one-way distillable-entanglement demand, and maximally entangled states attain the dimension upper law. With m oscillator modes per party its optimal energy has exponent α/m:

$$E_G(\varepsilon) = \varepsilon^{-\alpha/m + o(1)}.$$

Proof. The maximum entropy of m oscillators at total mean occupation E is $mg(E/m)$, where $g(x) = (x+1)\log_2(x+1) - x\log_2 x$. Relative entropy against the product thermal state proves this bound, and $g(x) \le \log_2(\mathrm{e}(x+1))$. Since directional distillation is at most either marginal entropy, each mean occupation is at least $m(2^{d_G^{\rightarrow}(\varepsilon)/m}/\mathrm{e} - 1)$. For normal spatial strategies, finite

local truncations approximate their correlations with arbitrarily small relative loss slack; local monotonicity of distillation gives the same lower bound with ε replaced by $(1 + o(1))\varepsilon$.

A rank-D upper strategy fits into the m-mode grid with each coordinate less than $\lceil D^{1/m} \rceil$. Its mean occupation is at most $m\lceil D^{1/m} \rceil$. The two laws squeeze the exponent. The matched games in [16] supply examples. This proposition uses both matched laws; entropy divergence alone does not imply this power law, and Γ has the separate bounded-energy behavior above. □

10.2 Global oscillator optima and small-loss limits

Let $\Lambda_* = (p_0^*, p_1^*, \ldots)$ be the decreasing probabilities of Theorem 10.1: sixteen copies of $4^{-k}/24$, then four of $4^{-k}/48$, at every level. Section 13.3 proves that this is the unique majorization maximum among perfect type-I marginal spectra.

Theorem 10.4 (The least oscillator occupation)**.** *For any perfect strategy with a normal type-I Alice marginal and any nonnegative compact-resolvent local Hamiltonian with ordered eigenvalues* $e_0 \le e_1 \le \cdots$,

$$\langle H_A \rangle \ge \sum_{i \ge 0} e_i p_i^*.$$

When the series is finite it is attained by preparing the displayed strategy in the ordered energy eigenbasis. If the series is infinite, every perfect strategy has infinite energy on that platform. In particular $\langle N_A \rangle \ge 275/18$. *Equality for the oscillator requires spectrum* Λ_* *and passive placement. For pure spatial strategies, each side has mean occupation at least* $275/18$, *and the total at least* $275/9$, *all attained by the displayed strategy.*

Proof. For a rank-r projection, the variational principle bounds its probability by the sum of the first r marginal eigenvalues. Expand the Hamiltonian as a monotone sum of positive steps at its energy-prefix projections. With p the ordered marginal spectrum this gives

$$\langle H_A \rangle \ge e_0 + \sum_{r \ge 1} (e_r - e_{r-1}) T_r(p) \ge e_0 + \sum_{r \ge 1} (e_r - e_{r-1}) T_r(\Lambda_*) = \sum_i e_i p_i^*.$$

Theorem 13.5 proves the second inequality. All sums are nonnegative, so no finite-expectation assumption is needed. Ordered preparation attains the expression. For N_A every increment is one. Equality forces every spectral tail to equal the released tail, hence $p = \Lambda_*$; equality in the energy-prefix inequalities is passive placement. The computed value is $275/18$. Pure strategies have the same nonzero Schmidt spectrum on Bob, giving the other side and the total. □

Attainment here permits preparation and transport on the Schmidt support. It is not an assertion that every entire-space unitary orbit reaches a passive minimum: a prescribed marginal with an extra kernel retains it on its unitary orbit, while the oscillator's full-support passive optimum has none. Nor does compact resolvent imply a finite bound; for instance $e_i = 2^i - 1$ makes the displayed series infinite. The pure Bob-side conclusion is the only Bob-side energy assertion used here.

Define $E_{A,\Gamma}(\varepsilon)$ as the infimum of Alice's mean oscillator occupation over normal spatial strategies of loss at most ε, allowing mixed states and unrestricted Bob dimension and POVMs. The entropy price S_Γ retains the pure spatial convention above. Finite-only prices at positive loss are defined by the same infima restricted to finite-dimensional strategies.

Theorem 10.5 (Exact limiting resource prices)**.**

$$\lim_{\varepsilon \downarrow 0} S_\Gamma(\varepsilon) = \log_2 24 + 7/9, \qquad \lim_{\varepsilon \downarrow 0} E_{A,\Gamma}(\varepsilon) = 275/18.$$

The same limits hold for the finite-only prices at positive loss. These are limits, not quantitative convergence rates; the finite-only feasible set at zero loss is empty.

Proof. Monotonicity and the perfect released strategy give finite upper limits. The entropy compactness argument in Theorem 10.2 produces a perfect pure limit at the limiting entropy. Theorem 13.6 bounds that entropy below by H_Γ, so equality follows.

For one-sided energy, purify any mixed strategy into Bob without changing Alice's marginal, cost or correlations. Order its Schmidt probabilities p_i and transport both Schmidt supports to the canonical vector $\sum_{i\ge 0}\sqrt{p_i}|ii\rangle$, placing Alice passively. Her new mean $\sum_i ip_i$ is no greater than the original energy, and Bob is unconstrained. Along a sequence approaching a finite limiting cost e,

$$\sum_{i\ge r} p_i \le (e + o(1))/r.$$

Coordinatewise subselection and this tightness give ℓ^1 convergence of the lists and norm convergence of canonical pure vectors. Simultaneous ultraweak subsequences of the finitely many local effects remain POVMs. Product effects converge weakly first on product vectors and then by density, so the limiting loss is zero. Fatou bounds its Alice occupation by e. Theorem 10.4 forces $e \ge 275/18$. This controls Bob by Schmidt transport, without applying two-sided Hamiltonian compactness to an unconstrained Bob system.

For finite-only prices the original K-level compressions have loss $1/[624(4^K - 1)]$, energy tending to $275/18$ and entropy

$$H_K = H_\Gamma - \frac{2K}{4^K - 1} + \log_2(1 - 4^{-K}) \longrightarrow H_\Gamma.$$

They supply the upper limits. Restriction to a smaller feasible class supplies the lower limits already proved. □

11 A localized finite-energy field state

We use the vacuum Fock representation of the free real scalar field of mass $m_0 > 0$ in $d+1$ Minkowski dimensions, $d \ge 1$. The Hamiltonian is $H = \mathrm{d}\Gamma(h)$, $h = \sqrt{-\Delta + m_0^2}$, with vacuum energy zero. Regions below have a positive spacelike separation collar; in particular their local algebras form a split pair. The split input for this field is standard nuclearity theory [7]. Energy bounds depend on the mass and chosen smearing functions.

Theorem 11.1 (Localized perfect play with all energy moments)**.** *For two nonempty strictly spacelike separated bounded open regions there are local PVMs and a normalized pure field vector Ψ winning Γ perfectly, with $\Psi \in D(H^s)$ for every $s > 0$. There is a bounded containing region R such that the state agrees with the vacuum on its causal exterior. It is obtained by a joint algebra-local isometry in $\mathcal{A}(R)$, not by a field-only local unitary or a product of independent local operations. Its nonzero Lüders outcome vectors retain every polynomial energy moment.*

11.1 Local oscillator factors and energy forms

Choose small double cones inside the regions with simultaneous interior points, in a fixed Lorentz frame. On that slice choose real smooth compactly supported functions f_A, f_C on disjoint small balls on Alice's side and f_B on Bob's, with $\langle f_i, f_j\rangle = \delta_{ij}$. Fix $\mu > 0$ and set

$$q_i = \sqrt{\mu}\,\phi_0(f_i), \qquad p_i = \pi_0(f_i)/\sqrt{\mu}.$$

The regular Weyl relations have $[q_i, p_j] = i\delta_{ij}$, with other commutators zero. Their bounded Weyl operators belong to the indicated local algebras. To justify time-zero localization, approximate

$\phi_0(f)$ and $\pi_0(f)$ by spacetime smears $\chi_\eta(t)f(x)$ and $-\chi'_\eta(t)f(x)$, with a smooth shrinking time mollifier. Their one-particle vectors converge, by rapid Fourier decay and dominated convergence, to the time-zero data. The thin supports remain in the chosen cones. Strong continuity of Weyl operators and von Neumann closure prove membership, rather than a formal commutator argument alone.

Write

$$u_i = \sqrt{\mu}\, h^{-1/2} f_i, \quad v_i = h^{1/2} f_i/\sqrt{\mu}, \quad F = \mathrm{span}_{\mathbb{C}}\{u_i, v_i : i = A, B, C\}.$$

This finite complex space lies in every $D(h^s)$ and has dimension between three and six. Its underlying real phase space contains the three specified nondegenerate canonical pairs. Complete them symplectically: subtract from each candidate its components along these pairs, then pair a nonzero complement vector with one of nonzero symplectic pairing and normalize. Iteration exhausts the finite complement. A finite-dimensional metaplectic implementer [14, Section 2.4] gives a factorization

$$\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C \otimes \mathcal{H}_D \otimes \Gamma_s(F^\perp).$$

The first three factors are oscillator spaces, with their full bounded operator algebras in the corresponding local algebras. The completed spectator coordinates need not be local. No entire type-III local algebra is identified with a type-I algebra.

Let N_F be physical particle number in F and $N_c = N_A + N_B + N_C + N_D$ the completed canonical number. For each integer $s \geq 1$ put

$$A_s = \max\{1, \max_{1 \leq r \leq s} \|h^r|_F\|^{1/r}\}.$$

On the k-particle space over F, expand $H^s = (\sum_{j=1}^k h_j)^s$. There are k^s ordered terms, each a tensor product of maps $h^{r_j}|_F$ of product norm at most A_s^s. Their ranges may leave F, which causes no problem. Orthogonality of total particle-number sectors gives

$$\|H^s \xi\| \leq A_s^s \|N_F^s \xi\|. \tag{37}$$

Finite-sector approximation and closedness extend this to its stated domain on $\Gamma_s(F) \otimes \Omega_{F^\perp}$. This proof does not assume that F is invariant under h.

The finite Bogoliubov relations imply, on the oscillator Schwartz core,

$$\|(N_F + 1)^s \xi\| \leq C_s \|(N_c + 1)^s \xi\|, \qquad \|(N_c + 1)^s \xi\| \leq C'_s \|(N_F + 1)^s \xi\|. \tag{38}$$

Indeed each number power is a finite sum of words of length at most $2s$ in the other creation and annihilation operators. A word of length l has norm at most a fixed multiple of $(k+1)^{l/2}$ on number sector k, and moves it by at most l. Squaring and summing the finitely overlapping output sectors proves the bounds. A finite metaplectic transformation preserves Schwartz space; graph-norm approximation and closedness therefore extend the bounds. The physical vacuum has every canonical number moment, since applying a power of the canonical quadratic number operator to it produces a finite polynomial vector. Denote $v_p = \langle \Omega, (N_c + 1)^p \Omega \rangle < \infty$.

Proposition 11.2 (The sorted-index energy criterion)**.** *A specified pure spatial strategy with decreasing Schmidt probabilities p_n, $n \geq 0$, has a realization by two oscillator factors of finite total mean occupation if and only if $\sum_n n p_n < \infty$. This condition suffices to transplant it, with the same correlations, into two localized oscillator subfactors of the massive free field with finite mean field energy. This transplantation alone does not assert exterior-vacuum localization.*

Proof. Passive rearrangement gives the minimum mean occupation $\sum_n np_n$ on each side; assigning the Schmidt basis to number states attains it. The necessity also follows by the Ky Fan tail inequality and summation of the oscillator spectral tails. To obtain localized field factors repeat the finite canonical completion with just A, B. Put the unused completed modes and $F^\perp$ in their Gaussian vacua. If b_j are canonical and a_i physical annihilators, the finite Bogoliubov relations give the quadratic-form bound

$$\langle N_F\rangle \le C(1+\sum_j \langle b_j^* b_j\rangle) = C(1+2\sum_n np_n).$$

The physical one-particle density is supported in F. With the finite form matrix $(h_F)_{ij} = \langle h^{1/2}w_i, h^{1/2}w_j\rangle$ in an orthonormal basis of F, mean energy is $\mathrm{Tr}(h_F\gamma) \le \|h_F\|\langle N_F\rangle$. This uses a form matrix, not invariance of F. All measured operators remain in the localized factors. Infinite Schmidt rank and finite one-particle mode count are compatible: Fock space over one mode is already infinite-dimensional. □

11.2 An isometry that stores the original input

For $t = a+b+c$ define

$$J(a,b,c) = \binom{t+2}{3} + a(t+1) - \frac{a(a-1)}{2} + b.$$

At fixed t, the successive runs have lengths $t-a+1$ and cover exactly the integers from $\binom{t+2}{3}$ to $\binom{t+3}{3}-1$. Thus J bijects triples with nonnegative integers. Also $J+1 \le \binom{t+3}{3} \le (t+1)^3$. On the three specified modes let

$$V|a,b,c\rangle = |\psi_\Gamma\rangle_{AB} \otimes |J(a,b,c)\rangle_C,$$

where ψ_Γ is Theorem 10.1's vector; let V be the identity on spectators. The images are orthonormal and complete in $\psi_\Gamma \otimes \mathcal{H}_C$, so $V^*V = I$ and $VV^* = |\psi_\Gamma\rangle\langle\psi_\Gamma| \otimes I$. It lies in the three-factor type-I algebra, hence in the joined laboratory algebra and in $\mathcal{A}(R)$ for a bounded containing double cone.

Expand the actual vacuum as $\Omega = \sum_{a,b,c} |a,b,c\rangle\xi_{abc} \otimes \Omega_{F^\perp}$. The spectator vectors need not be orthogonal. Then

$$\Psi = V\Omega = \psi_\Gamma \otimes \chi \otimes \Omega_{F^\perp}, \qquad \chi = \sum_{a,b,c} |J(a,b,c)\rangle\xi_{abc}.$$

This keeps the original correlations in storage; it does not assume the vacuum initially factorizes across local modes. For spectator number u,

$$J(a,b,c) + u + 1 \le (a+b+c+u+1)^3.$$

Orthogonality of storage labels, the spectral resolution of N_D and monotone convergence therefore imply

$$\langle \chi, (N_C + N_D + 1)^p \chi\rangle \le v_{3p}.$$

With $X = N_A + N_B$ and $Y = N_C + N_D + 1$ on the product output,

$$\langle \Psi, (N_c+1)^p \Psi\rangle \le 2^{p-1}(\langle \psi_\Gamma, X^p \psi_\Gamma\rangle + v_{3p}) < \infty.$$

Take $p = 2s$ in (38) and (37). This proves every required energy domain, including the factor of two between moments and operator norms. For example $\langle X^2\rangle = 47530/27$, so the second canonical moment is at most $95060/27 + 2v_6$, giving an explicit finite geometry-dependent energy bound.

The game PVMs act on the first two factors and give exactly the original perfect correlation. For $A \in \mathcal{A}(R')$, locality and $V^*V = I$ give $\langle \Psi, A\Psi \rangle = \langle \Omega, A\Omega \rangle$. This is the classical local-isometry mechanism of strict localization [22, Theorem I]; see also [4, Section 4]. The addition here is its explicit game resource and energy-domain control. A finite Lorentz boost used to choose simultaneous points preserves all moments: the spectrum condition bounds the original energy by a fixed multiple of the chosen-frame energy.

11.3 Measurements and the meaning of preparation

Each copy-one oscillator block has occupation span at most sixteen, and each copy-two block at most four. On a block of span w, the largest and smallest values of $n + u + 1$ have ratio at most $w + 1$, for every spectator occupation u. An outcome projection P therefore satisfies

$$\|(N_c + 1)^s P\xi\| \leq (w + 1)^s \|(N_c + 1)^s \xi\|.$$

Square and sum this estimate over the orthogonal blocks; diagonal labels and boundaries have ratio one. Products of Alice and Bob outcomes retain the bound successively. Equations (37)–(38) prove the outcome assertion of Theorem 11.1. It concerns finiteness of moments for this smooth finite-span class and Lüders instruments, not preservation of their numerical values or every possible instrument.

The normal UCP map $A \mapsto V^*AV$ fixes exterior observables and sends the vacuum state to Ψ deterministically. It is a joint entangling operation, not LOCC. An ideal pure ancillary qubit gives a bounded algebra-local dilation: for $\gamma > 0$,

$$G = \gamma(V \otimes |1\rangle\langle 0| + V^* \otimes |0\rangle\langle 1|)$$

is self-adjoint of norm γ, and $e^{-itG}\xi|0\rangle = \cos(\gamma t)\xi|0\rangle - i \sin(\gamma t)V\xi|1\rangle$. At $t = \pi/(2\gamma)$ it prepares the isometric output. This ancillary-qubit localization mechanism is known; see [17, Section V.2]. With free Hamiltonian H_0, the interaction $W(t) = e^{-itH_0}Ge^{itH_0}$ gives propagator $e^{-itH_0}e^{-itG}$; the displayed vacuum trajectory is in every required domain, so differentiating this vector product justifies the evolution without a bounded $[H, V]$ assumption. Finite time translates lie in an enlarged bounded region. Transporting the final measurements absorbs the free evolution. This is an ideal bounded algebra-local control model; no microscopic interaction-density synthesis is asserted.

For fixed split regions let $\delta_0 > 0$ be their vacuum loss floor proved in Theorem 12.1. The mass gap $H \geq m_0(I - |\Omega\rangle\langle\Omega|)$ and the vacuum fidelity bound give, for any normal state of mean energy E, half trace distance at most $\sqrt{E/m_0}$. Comparing the same loss effect in that state and the vacuum yields

$$m_0(\delta_0 - \varepsilon)_+^2 \leq E_\Gamma(\varepsilon) \leq E_* < \infty,$$

where E_* is the preceding prepared state's energy. Thus this unrestricted normal-state energy price is bounded above and bounded away from zero near perfect play. This argument asserts no value for its infimum, no field-energy optimizer, and no universal conversion of 275/18 into physical units.

12 Faithful states, the vacuum, and preparation

The mechanisms in this section are established operator-algebraic background. Joint faithfulness and independent local algebras already exclude perfect nontrivial correlations in the work of Rédei–Summers [25, proof of Lemma 3, equations (18)–(19)]. Summers–Werner's faithful-state argument [30, Theorem 2.1] supplies the uniform-gap mechanism for CHSH. We give their finite-game formulations and the explicit consequences for Γ. Local faithfulness of the two marginals alone is insufficient: the ideal entangled gadget has faithful marginals and a nonfaithful joint state.

Theorem 12.1 (Faithfulness and a uniform fixed-state floor)**.** *Let M, N commute and satisfy Schlieder independence: the product of nonzero positive local operators is nonzero. In a state faithful on $M \vee N$, perfect finite-game play exists exactly when the game has a classical perfect strategy. If the state ω is normal and faithful and there is a normal product state σ on $M \vee N$, every classically imperfect finite game has a positive loss floor in this fixed state, uniform over local POVMs.*

More generally, a faithful state on a commuting pair admitting a product state cannot attain a commuting Bell maximum strictly above its classical maximum. For a split pair, the analogous statement applies to its spatial maximum.

Proof. Every supported losing product $E_a^x F_b^y$ is positive. Zero total loss and faithfulness make each such product zero. Independence then makes one factor zero. Choose a nonzero answer effect for each question; completeness ensures that one exists. Every supported pair of selected effects is nonzero, hence winning. These choices give a deterministic classical perfect strategy. The converse uses effects zero and identity.

For the uniform assertion, write $\delta_{\mathrm{cl}} = 1 - \omega_{\mathrm{cl}}(G) > 0$ and define, for $0 < t \leq 1$,

$$\beta_t(\sigma\|\omega) = \min\{\omega(T) : 0 \leq T \leq I,\ \sigma(T) \geq t\}.$$

The feasible effect set is ultraweakly compact, and both functionals are continuous. Its minimum is positive: a zero minimizer would vanish by faithfulness, contradicting its σ value. The game loss effect L is a positive contraction and $\sigma(L) \geq \delta_{\mathrm{cl}}$, because product responses are classical. Thus $\omega(L) \geq \beta_{\delta_{\mathrm{cl}}}(\sigma\|\omega) > 0$. The floor is for the specified state and pair, not uniform over varying faithful states.

For a Bell operator B with universal commuting upper bound β_{qc}, the operator $\beta_{qc} I - B$ is positive, since every vector with those measurements is a commuting strategy. Faithful equality forces $B = \beta_{qc} I$. Its expectation in the product state is classical, contradicting $\beta_{qc} > \beta_{\mathrm{cl}}$. On a split pair the same argument uses its tensor factorization and the spatial upper bound. Normality of the product state is needed for the uniform argument, not this qualitative one. □

Proposition 12.2 (The classical gap and a testing price)**.** *The classical value of Γ is $51/52$. For a fixed faithful normal split state put*

$$I_t = \inf_{\sigma \text{ normal product}} [-\ln \beta_t(\sigma\|\omega)].$$

Then $-\ln t \leq I_t < \infty$ and every Γ strategy in that state has loss at least $e^{-I_{1/52}} > 0$. A family attaining loss at most ε must have $I_{1/52} \geq \ln(1/\varepsilon)$. This is a testing quantity, not an energy or geometric rate.

Proof. A deterministic perfect first FMS copy is impossible. Row/column flag consistency would make all six flag statuses the same. All flags force four inconsistent sign parities. No flags force the magic-square row/column parity contradiction. Thus at least one of the fifty-two pairs loses. For attainment of this bound choose in copy one Alice's rows all $(+\,+\,+)$, $T_0 = +$, $T_1 = (+\,+\,+)$; Bob's first two columns $(+\,+\,+)$, third $(+\,+\,-)$, $U_0 = +$, $U_1 = (+\,+\,+)$. Choose boundary answers throughout copy two and Bob's label A. Exactly (R_3, C_3) of copy one loses. This proves the value.

The testing minimum is positive by Theorem 12.1 and at most t, using $T = tI$. One normal product comparison gives a finite upper bound on the infimum. Taking the supremum of the individual testing lower bounds gives the stated exponential expression, with no optimizer assumption. An algebra-pair isomorphism transports all states and effects, so the quantity is invariant under

such transport. If reverse relative entropy is finite, data processing under the binary measurement $\{L, I - L\}$ also gives

$$\omega(L) \geq \exp[-(S(\sigma\|\omega) + \ln 2)/\delta_{\mathrm{cl}}],$$

since $S(\sigma\|\omega) \geq \delta_{\mathrm{cl}} \ln(1/\omega(L)) - \ln 2$. This uses the reverse order shown, rather than ordinary mutual information [2]. □

12.1 Correlation closures and the full commutant

Proposition 12.3 (Normal approximation and fixed-vector access)**.** *If one local algebra is the strong closure of increasing finite-dimensional unital algebras, all normal finite correlations lie in C_{qa}. Normal split correlations are spatial. For a type III_1 factor M on a separable Hilbert space with Bob's full commutant M', the closure of the correlations of any fixed unit vector contains C_{qa}; it equals C_{qa} if M is hyperfinite.*

Proof. For each Alice POVM form the column $(\sqrt{E_a})_a$, a contraction column. Matrix-amplified Kaplansky density gives contraction columns in a common finite-dimensional approximating algebra, converging strongly-star. Their squared entries sum to at most identity; add the positive deficit to the first answer. These are exact POVMs and converge strongly. Normality makes their finitely many correlations with fixed Bob effects converge. A representation of $\bigoplus_k M_{n_k}$ decomposes into $\bigoplus_k (\mathbb{C}^{n_k} \otimes K_k)$, with commuting Bob acting on the multiplicity spaces. Purify the state and compress Bob in each block to Schmidt rank at most n_k. A shared direct-sum label realizes the full correlation with finite local dimension at most $\sum_k n_k$. This proves approximation; it is the standard hyperfinite mechanism [26, Section 4]. A split tensor factorization directly gives the spatial statement.

For the last assertion, proper infiniteness embeds normally represented unital copies of $B(\ell^2)$ in M and M'. Splitting off these commuting type-I subfactors gives $\ell^2 \otimes \ell^2 \otimes K$, in which every finite target strategy can be represented. The local-unitary clause of [32, Theorem 51, equation (97), arXiv version 2] approximates any target vector by $uv\Omega$, $u \in M$, $v \in M'$. Transport the measurements by these unitaries. Every effect probability changes by at most twice the vector error. This gives the inclusion; hyperfinite approximation gives the reverse. The full commutant hypothesis is essential: it cannot be replaced by a strictly separated bounded Bob region, where the faithful-vacuum floor applies. □

12.2 Wedge supremum one without attainment

Theorem 12.4 (The massive wedge vacuum)**.** *For complementary wedges in the massive free scalar field, the vacuum has Γ supremum one but no perfect strategy, even with local POVMs.*

Proof. The modular group of a wedge vacuum is the boost group, and wedge duality identifies the opposite algebra with the full commutant [5, Theorem 4]. Its centralizer is scalar. Here is the latter step explicitly. On the positive mass shell use longitudinal rapidity θ, with $p_0 = m_\perp \cosh\theta$, $p_1 = m_\perp \sinh\theta$; invariant measure is $(1/2)d\theta\, d^{d-1}p_\perp$. On an n-particle sector a boost translates all rapidities together. Taking $u = \theta_1$ and $v_j = \theta_j - \theta_1$ gives an $L^2(\mathbb{R}, du)$ regular-translation factor. No nonzero vector in it is invariant: Fourier support would lie at a measure-zero point. Symmetrization adds no invariant vectors, and the Fock direct sum leaves only the vacuum. If A is in the vacuum centralizer, modular invariance makes $A\Omega$ boost-invariant, hence scalar times Ω; separatingness gives $A \in \mathbb{C}I$.

If self-adjoint $A \in M$, $B \in M'$ satisfy $A\Omega = B\Omega$, then for $x \in M$, $\omega(Ax) = \langle B\Omega, x\Omega\rangle = \omega(xA)$. Thus A lies in the centralizer, the doubles mechanism of Werner [33, Section II]. Self-adjointness, or both adjoint mirror identities, matters here.

For commuting effects P, Q, binary mismatch equals

$$P(I-Q)+(I-P)Q=(P-Q)^2+(P-P^2)+(Q-Q^2).$$

Zero expectation makes each positive summand vanish on Ω. Local separatingness gives projections P, Q and the mirror identity. Apply this to first-copy flag checks: all flag projections coincide on the vacuum and hence locally, and their self-adjoint mirrors put them in the scalar centralizer. They are either all zero or all identity.

If identity, calibration eliminates all triples at the four bridge questions. Their binary sign reflections satisfy $X_0\Omega = Y_0\Omega = Y_1\Omega$ and $X_1\Omega = Y_0\Omega = -Y_1\Omega$, impossible for a reflection. If zero, every magic-square cell marginal has a self-adjoint mirror and is a scalar projection. Their deterministic signs would have row products $+1$ and column products $+1, +1, -1$, again impossible. Thus the first FMS copy, and hence Γ, cannot be perfect.

Massive free wedge algebras are hyperfinite type III_1 in this separable vacuum representation [6, 18]. Proposition 12.3, applied to the finite compressions of Theorem 1.1, gives vacuum strategies approaching one. This proves the supremum assertion without importing a spatial Schmidt decomposition into the wedge. Summers–Werner already discuss wedge CHSH nonattainment via trivial centralizers [30, p. 221]; the addition here is the explicit FMS/Γ consequence. □

12.3 Analytic energy and restricted preparation

Theorem 12.5 (What cannot prepare perfect play)**.** *In the massive free-field vacuum representation, no nonzero vector with bounded spectral energy support, or in $D(e^{aH})$ for some $a > 0$, wins Γ perfectly in two bounded strictly separated regions. Every perfect normal state therefore has infinite positive exponential energy moments. No bounded-region invertible field operation, in particular no such unitary, prepares perfect play from the vacuum. Nor does a finite LOCC protocol or normal operation with a product-Kraus decomposition across the laboratories.*

Proof. The energy-analytic Reeh–Schlieder theorem, with locality, positive-energy covariance and weak additivity, makes these vectors cyclic for every nonempty open local algebra [8, Section 3], [29, Section 3]. An open region in the common causal complement is therefore cyclic; commutation makes the vector separating for the joined game algebra. Theorem 12.1 excludes perfection. If $\langle e^{tH}\rangle < \infty$, the vector belongs to $D(e^{tH/2})$. For a normal mixed state with finite exponential moment, every positive-weight vector in its spectral decomposition has that property and its mixture is locally faithful. Thus mixed states are excluded too. Finite mean energy is not bounded spectral support: Theorem 11.1 exhibits the difference. The assertion concerns moments, not a pointwise regular form of spectral tails.

Enlarge a bounded region to contain both the preparation and game algebras. If U is invertible there and $AU\Omega = 0$, then $(U^{-1}AU)\Omega = 0$; vacuum separatingness gives $A = 0$. Thus the output is still separating for the game algebra. Covariant time translations obey the same argument in transported regions.

For a nonzero product branch AB, a losing event of zero conditional probability gives

$$\omega_\Omega((A^*E_a^xA)(B^*F_b^yB)) = 0.$$

Vacuum faithfulness and independence force one filtered factor to vanish. Their local sums are A^*A and B^*B, both nonzero. Selecting nonzero filtered answers for every question therefore gives a classical perfect strategy, a contradiction. A sum or integral of branches with zero total loss has zero loss in each positive-weight branch, almost everywhere in an integral. This proves the stated class, not an unspecified closure of infinite protocols. The isometry in Theorem 11.1 is joint and noninvertible, so it avoids both obstructions. □

Faithfulness, analytic-vector cyclicity and their preservation are known mechanisms, not new principles about all possible physical preparations. Halvorson–Clifton [19] establish generic Bell correlation, which is compatible with failure of perfect attainment. Neither a marginal's faithfulness nor the type-III label alone fixes the answer. Our claims concern the finite-game formulations and these specified vacuum and preparation classes.

13 Entropy is different from dimension

Theorem 13.1 (Infinite rank and minimum pure entropy)**.** *Every perfect pure spatial strategy for* Γ *has infinite Schmidt rank. The displayed perfect strategy attains the minimum pure-state entanglement entropy*

$$H_\Gamma = \log_2 24 + \frac{7}{9} = 5.362740278498934\ldots \quad \textit{ebits.}$$

Finite entropy is attained but not forced on all perfect strategies. The minimum concerns pure-state entanglement; the corresponding lower bound also holds for Alice marginal entropy of normal mixed extensions.

Proof. A finite-Schmidt-rank perfect vector could be compressed to its two finite Schmidt supports, retaining complete POVMs and all correlations. This contradicts Theorem 1.1. Equivalently the exact sector identity forces $s/2$ whenever $s > 0$ occurs in its nonempty second flag sector, hence infinitely many nonzero coefficients.

For the displayed state, sum the probabilities in Theorem 10.1. The expected $2k$ contribution is $2/3$, and the extra binary logarithm in the four B probabilities has total weight $1/9$. This gives $\log_2 24 + 7/9$. The lower bound for every perfect pure strategy follows from Theorem 13.6 below.

For a counterexample to universal finite entropy, tensor this state with a countable pure state whose jth block has 2^j equal probabilities $c/(j^2 2^j)$, $j \geq 1$, $c = 6/\pi^2$. Its mass is one and its entropy is at least $c\sum_j 1/j = \infty$. Tensor every game effect with identity. Correlations remain perfect and the local spaces remain separable, but the entropy is infinite. Thus a large enough ordinary entropy budget admits the exhibited state, whereas every finite dimension budget excludes all perfect states. No area, gravity or control-cost interpretation follows without additional assumptions. □

13.1 Modular ladders and the role of the boundary

Let M carry a faithful normal state φ, with modular automorphisms σ_t. A modular eigen-operator of eigenvalue $\lambda > 0$ satisfies $\sigma_t(u) = \lambda^{it}u$. In the semifinite case write $\varphi = \mathrm{Tr}(h\,\cdot)$, $\mathrm{Tr}\,h = 1$, after compression to the support of h. The eigen-operator relation means $hu = \lambda uh$ as an affiliated-operator identity, equivalently $g(h)u = ug(\lambda h)$ for every bounded Borel g.

For an integer $n \geq 1$, an (n, λ) *ladder* consists of a boundary projection B and modular eigen-operators $u_1, \ldots, u_n$ such that

$$u_i^* u_j = \delta_{ij}(I - B), \qquad \sum_{i=1}^{n} u_i u_i^* = I. \tag{39}$$

All source and range projections belong to the centralizer. Put $\rho = n\lambda$. For a semifinite density define the state law $\mu_e(g) = \mathrm{Tr}(he\, g(-\ln h))$ for centralizer projections e, and the trace measure $\eta_e(A) = \mathrm{Tr}(e1_A(h))$. Let S_a translate real measures by a, and D_c multiply positive spectral values by c.

Theorem 13.2 (The ladder law)**.** *Every ladder satisfies $\varphi(B) = 1 - 1/\rho$, so $\rho \geq 1$, with $B = 0$ exactly at $\rho = 1$. In the semifinite case, if $\rho > 1$,*

$$\mu_I = \sum_{j \geq 0} \rho^{-j} S_{j \ln \lambda} \mu_B, \qquad \eta_I = \sum_{j \geq 0} n^{-j} (D_{\lambda^{-j}})_* \eta_B. \tag{40}$$

Thus the scalar modular law is that of $K_B + (\ln \lambda) J$, with independent variables, normalized boundary law K_B, and $\Pr(J = j) = (1 - 1/\rho)\rho^{-j}$.

No normal semifinite critical ladder exists for $n \geq 2$. The exceptional critical case $n = \lambda = 1$ has $B = 0$ and a centralizer unitary; the boundary series does not apply there. Finite-dimensional ladders require $\lambda = 1$. In a type-I factor, $\lambda < 1$ is impossible, $\lambda = 1$ has finite-rank examples, and $\lambda > 1$ forces infinite rank and has examples. Every II_1 factor admits a normal ladder for $n \geq 2$, $\rho > 1$; its density is unbounded when $\lambda < 1$.

In the semifinite case, with finite second modular moment,

$$\operatorname{Var} K = \operatorname{Var} K_B + (\ln \lambda)^2 \frac{\rho}{(\rho - 1)^2}.$$

If $T = \operatorname{Tr} I < \infty$ and $n \geq 2$, trace entropy may equal $-\infty$, and

$$S_{\mathrm{Tr}}(h) \leq \ln T - \Delta(n, \lambda), \qquad \Delta(n, \lambda) = \ln \frac{\rho - 1}{\lambda(n - 1)} - \frac{\ln \lambda}{\rho - 1}. \tag{41}$$

Equality holds exactly when h is scalar on B, and is attained in every finite type-II factor.

Proof. In the standard representation, $\Delta_\varphi^{it} u\Omega = \lambda^{it} u\Omega$. Tomita's identity $Su\Omega = u^*\Omega$, $S = J_\varphi \Delta_\varphi^{1/2}$, gives $\varphi(uu^*) = \lambda\varphi(u^*u)$. Summing (39) proves the boundary mass and, by faithfulness, the critical equivalence.

For $a = \ln \lambda$, bounded spectral calculus and trace cyclicity give

$$\mu_I = \rho S_{-a} \mu_{I-B}, \qquad \mu_I = \mu_B + \rho^{-1} S_a \mu_I.$$

After L substitutions the remaining measure has mass ρ^{-L}. This proves the first series for $\rho > 1$ and its normalized independent-law description. Likewise $\eta_I = \eta_B + n^{-1}(D_{1/\lambda})_* \eta_I$. On a set contained in $[c, \infty)$ its Lth remainder is at most $1/(c\rho^L)$, because $\eta_I([y, \infty)) \leq 1/y$. This proves the trace series, also when total support trace is infinite.

At $\rho = 1$, $n \geq 2$, the first identity would make a probability measure invariant under the nonzero translation $\ln \lambda = -\ln n$. An interval shorter than this translation with positive mass would have disjoint translates of infinite total mass. For $n = \lambda = 1$, the relations instead say that u_1 is a centralizer unitary and $B = 0$; for example any faithful matrix density with $u_1 = I$ is legal. Its nonzero state law cannot equal a zero-boundary series.

In finite dimension the positive spectrum of h is λ times its positive spectrum on $I - B$. The smallest eigenvalue contradicts $\lambda > 1$, and the largest contradicts $\lambda < 1$. In type I, let ζ count eigenvalues on $I - B$. Since $\zeta \leq n(D_\lambda)_* \zeta$, integer multiplicity forces y/λ^k for every k whenever y occurs. For $\lambda < 1$ these eigenvalues are unbounded, impossible for a density; for $\lambda > 1$ they give infinite rank. To construct the latter, take ζ with one value $c\lambda^{-k}$ for each $k \geq 0$, where $c = (\lambda - 1)/(n\lambda^2)$, and let h have n copies of $\lambda\zeta$. Choose one vector of each level after the top as source and map it to the n preceding-level vectors. This gives the relations and eigenvalue λ. A flat density in dimension n gives the $\lambda = 1$ example. These constructions concern the support corner.

For II_1, with normalized trace τ, choose orthogonal projections E_j summing strongly to identity, with

$$\tau(E_j) = \frac{n - 1}{n} n^{-j}, \qquad h|_{E_j} = c\lambda^{-j}, \qquad c = \frac{\rho - 1}{\lambda(n - 1)}.$$

The geometric sums give $\tau(h) = 1$. Put $B = E_0$, split each E_j into n pieces equivalent to E_{j+1}, and sum the corresponding partial isometries strongly along the levels. They have the required common source, orthogonal full ranges and density ratio. If $\lambda < 1$, the affiliated density is unbounded but trace integrable and faithful. More generally the renewal law then has arbitrarily negative modular values on sets of positive mass, so no bounded density can realize it.

The independent geometric variable has mean $1/(\rho - 1)$ and variance $\rho/(\rho - 1)^2$, proving the variance identity. Without a finite second moment we use it only when a finite mean and an extended variance have been specified; normality alone does not supply logarithmic moments.

In finite trace, internal equivalence gives $t_B = \operatorname{Tr} B = T(n-1)/n$, and $w_B = \varphi(B) = (\rho - 1)/\rho$. With $g_B = (t_B/w_B)h_B$ and $\tau_B = \operatorname{Tr}/t_B$ on the boundary, the renewal law gives

$$\ln T - \Delta(n, \lambda) - S_{\mathrm{Tr}}(h) = \tau_B(g_B \ln g_B) \geq 0.$$

This is relative entropy against the uniform boundary. The positive part of the modular mean has finite integral in finite trace, so the formula remains meaningful at entropy $-\infty$. Equality is precisely $g_B = B$. The displayed type-II construction has scalar boundary and attains it; trace rescaling gives any finite T. □

The deficit is also the relative entropy of the geometric law of ratio $1/\rho$ against that of ratio $1/n$, so it is nonnegative and vanishes exactly at $\lambda = 1$. The scalar renewal uses a standard resolvent mechanism. The one-vertex specialization of [1, Theorem 3.1] has the same scalar ratio under $ne^{-\beta} = 1/\rho$, or $\beta = \ln(n^2\lambda)$. This correspondence does not identify the operator representations or recover the full graph-algebra theorem: its generators have full source and deficient ranges, unlike (39).

For perspective, the geometric oscillator density $(1 - q)\operatorname{diag}(1, q, q^2, \ldots)$ has a $(1, 1/q)$ ladder given by the backward shift and vacuum boundary. A boundaryless $(n, 1/n)$ ladder instead gives the dimension-changing isomorphism $x \mapsto [u_i^* x u_j]$ from M to $M_n(M)$, with inverse $[x_{ij}] \mapsto \sum u_i x_{ij} u_j^*$. The modular KMS identity gives $\varphi(u_i x u_j^*) = \varphi(x u_j^* u_i)/n = \delta_{ij}\varphi(x)/n$, splitting off a maximally mixed n-level marginal. This is algebraic state splitting, not fixed-pure-ancilla unitary extraction from a faithful state; that convention has a support obstruction. The known semifinite embezzlement obstruction is credited in Section 14. No assertion that all ladder resources diverge near criticality is needed: at $n = 1$ the mean and variance contributions both tend to one.

13.2 The graded hosting criterion

A *graded* $(4, 2)$ *ladder* for a semifinite density h consists of orthogonal centralizer projections F, P, B summing to its support and partial isometries w_i, v_i, $1 \leq i \leq 4$, with

$$\begin{aligned} w_i^* w_j &= \delta_{ij} F, & \sum_i w_i w_i^* &= P + B, \\ v_i^* v_j &= \delta_{ij} P, & \sum_i v_i v_i^* &= F, \qquad h w_i = 2 w_i h, \quad h v_i = 2 v_i h. \end{aligned} \tag{42}$$

The equations for unbounded h mean bounded spectral intertwining, as above.

Theorem 13.3 (Hosting perfect play)**.** *A normal semifinite Alice marginal admits perfect* Γ *play with a suitable commuting partner if and only if its support corner carries a graded* $(4, 2)$ *ladder. This is an existence criterion for measurements, not a claim that arbitrary measurements on the state win.*

Proof. For necessity apply the standard-partner reduction of Lemma 14.1. Zero probabilities give $Eh^{1/2}Q = 0$. The internal sector argument of Lemma 14.2, including its support-cutoff cancellations, supplies a flag f, four magic corners with internal matrix units commuting with h, and β such that $\beta^*\beta = f$, $\beta\beta^* = q$, and $d\beta = 2\beta h_f$. Thus the four maps $e_{i1}\beta$ from the flag to magic pieces satisfy $hw_i = 2w_i h$. Apply this to both copies. The links identify first flag with second magic, $F_1 = M_2$, and first magic with second flag plus boundary, $M_1 = P + B$. These are exactly (42). The argument uses internal partial isometries, not just equal Hilbert multiplicities, and remains valid at infinite support trace.

Conversely the first five equivalent corners $\{F, w_i w_i^*\}$ have square-root density

$$\operatorname{diag}(1, \sqrt{2}, \sqrt{2}, \sqrt{2}, \sqrt{2}) \otimes h_F^{1/2}$$

in their matrix-unit representation. Put the ideal FMS PVMs there. Every forbidden product is the corresponding zero five-dimensional product tensored with $h_F^{1/2}$. Do the same on $\{P, v_i v_i^*\}$ for copy two and give B its boundary answer. Right-hand measurements give the standard commuting partner. Bob's label is $(P + B, F)$; both checked links identify these same centralizer projections and win. The sum of each PVM is identity on support; complete it off support arbitrarily. The square-root matrix, rather than the density itself, has ratio $\sqrt{2}$. □

Putting $u_i = w_i + v_i$ gives an ungraded $(4, 2)$ ladder: all cross terms vanish because sources and ranges lie in the orthogonal grading. Its $\rho = 8$ law has $\Pr(J = j) = (7/8)8^{-j}$. This is exactly the law written as $2G + Q$ in Theorem 14.5. Also $\Delta(4, 2) = \ln(7/6) - \ln 2/7$. The ladder packages their common algebraic mechanism; Section 14 gives direct game-specific proofs. The extra constraint in type I is integer spectral multiplicity.

13.3 Integer spectra and the majorization maximum

For a positive summable multiset σ, write $m_\sigma(y)$ for the multiplicity of y, and let

$$\Lambda(\sigma) = 16 \cdot (4\sigma) \sqcup 4 \cdot (2\sigma).$$

Call σ admissible if its multiplicities are nonnegative integers,

$$\sum_{\mu \in \sigma} \mu = 1/72, \qquad m_\sigma(\mu) \le 16 m_\sigma(\mu/4) \quad (\mu > 0). \tag{43}$$

Multiplicity is included in the sum. Each positive value has finite multiplicity by summability.

Theorem 13.4 (Classification of perfect type-I marginals)**.** *A prescribed trace-class Alice marginal on a type-I factor admits perfect Γ play, with arbitrary POVMs and a suitable commuting partner, exactly when its nonzero spectrum is $\Lambda(\sigma)$ for an admissible σ. Zero eigenvalues on the ambient Hilbert space are irrelevant. The criterion covers initially mixed commuting extensions; sufficiency can be realized by a pure spatial strategy.*

Proof. Theorem 13.3 gives a graded ladder on support. Count eigenvalues on P by σ. Its four density-doubling equivalences give

$$\eta_F = 4(D_2)_*\sigma, \qquad \eta_{P+B} = 16(D_4)_*\sigma, \qquad \eta_B = 16(D_4)_*\sigma - \sigma \ge 0.$$

Thus the full spectrum is $\Lambda(\sigma)$, its total is $72 \sum \sigma = 1$, and the last positive measure is precisely the multiplicity inequality.

Conversely construct a magic space of spectrum $16(D_4)_*\sigma$. The inequality chooses within each eigenspace a subspace P of spectrum σ; let the remainder be B. On an orthogonal space F put spectrum $4(D_2)_*\sigma$. Eigenbasis bijections define four maps $P \to F$ and four maps $F \to P + B$, each doubling density, with the required common sources and orthogonal ranges. They give the graded ladder and hence perfect play. Purifying the type-I density gives a spatial pure realization. Transport on the prescribed marginal's support and complete measurements on its kernel. □

For a decreasing nonnegative list x, put $\mathrm{KF}_r(x) = \sum_{i=1}^{r} x_i$, $T_r(x) = \sum_{i>r} x_i$, and $G_x(t) = \sum_i (x_i - t)_+$. Set

$$\sigma_* = \{4^{-k}/96 : k \geq 0\}, \qquad \Lambda_* = \Lambda(\sigma_*).$$

This is the displayed sixteen/four spectrum of Theorem 10.1.

Theorem 13.5 (A unique majorization maximum spectrum)**.** *For every admissible σ and every rank r,*

$$\mathrm{KF}_r(\Lambda(\sigma)) \leq \mathrm{KF}_r(\Lambda_*).$$

Equality at all ranks holds exactly when $\sigma = \sigma_*$. Uniqueness concerns the spectrum, not the measurements or the whole strategy.

Proof. Integrality in (43) forces $\mu/4^k$ for every $k \geq 1$ whenever μ occurs. Order $\sigma_1 \geq \sigma_2 \geq \cdots$. The distinct forced values $\sigma_j/4, \sigma_j/16, \ldots$ are strictly smaller than σ_j and hence occur after position j, also in the presence of ties. Therefore

$$T_j(\sigma) \geq \sigma_j/3, \quad T_{j-1}(\sigma) = \sigma_j + T_j(\sigma) \leq 4T_j(\sigma), \quad T_j(\sigma) \geq 4^{-j}/72.$$

These are the tails of σ_*, proving its Ky Fan domination. Equality at every rank forces the same decreasing list.

To transfer the order without assumptions about interlacing, use

$$G_x(t) = \sup_r [\mathrm{KF}_r(x) - rt], \qquad \mathrm{KF}_r(x) = \inf_{t \geq 0} [G_x(t) + rt].$$

The first crossing index and $t = x_r$ prove these identities; at $t > 0$ summability leaves finitely many positive terms. The transform satisfies

$$G_{\Lambda(\sigma)}(t) = 64G_\sigma(t/4) + 8G_\sigma(t/2).$$

Both coefficients are positive, so the Ky Fan order of σ transfers to Λ. If both Lambda spectra agree, the two nonnegative G gaps in this formula have zero sum for every t. Each vanishes, whence $G_\sigma = G_{\sigma_*}$ and $\sigma = \sigma_*$. This proves the stated uniqueness. □

13.4 Entropy and every truncation rank

Majorization is an established entanglement ordering [24]. Here the game-specific content is the admissible integer multiset and its forced quarter-chain. The following infinite-list argument supplies the required entropy conclusions directly.

Theorem 13.6 (Sharp positive-order entropy bounds)**.** *Every normal perfect type-I marginal has von Neumann entropy at least H_Γ. For $0 < \alpha < 1$ or $1 < \alpha < \infty$, its Rényi entropy obeys*

$$H_\alpha \geq \frac{1}{1-\alpha} \log_2 \frac{16\,24^{-\alpha} + 4\,48^{-\alpha}}{1 - 4^{-\alpha}}.$$

Equality in the von Neumann bound or at any such finite positive order holds exactly for spectrum Λ_*. Every bound is attained by the displayed pure strategy. For pure strategies these are entanglement entropies; for initially mixed extensions the assertion concerns Alice's marginal entropy.

Proof. Let x be the marginal spectrum, $y = \Lambda_*$, and $D(t) = G_y(t) - G_x(t) \geq 0$. Both lists have mass one and entries at most one. Integrating the identity

$$\int_0^1 \frac{x - (x-t)_+}{t}\, dt = x - x \ln x$$

and using cutoffs gives $S(x) - S(y) = \int_0^1 D(t)\, dt/t \geq 0$ in nats, with a well-defined extended value because $S(y) < \infty$. For $\alpha > 1$, nonnegative integration of

$$x^\alpha = \alpha(\alpha - 1) \int_0^1 t^{\alpha-2} (x-t)_+\, dt$$

gives $\sum x_i^\alpha \leq \sum y_i^\alpha$. For $0 < \alpha < 1$, use instead

$$x^\alpha = \alpha(1-\alpha) \int_0^\infty t^{\alpha-2} \min(x,t)\, dt, \qquad \sum_i \min(x_i, t) = 1 - G_x(t).$$

This reverses the power-sum inequality. The definition of Rényi entropy gives the claimed directions. If the spectra differ, some continuous G gap is positive on an interval; each strictly positive integration kernel makes its finite equality impossible. Thus equality is precisely $x = y$. Summing the sixteen/four geometric probabilities gives the displayed power sum, and the order-one value is $\log_2 24 + 7/9$. □

Corollary 13.7 (The sharp discarded-mass profile)**.** *For a perfect type-I marginal and any rank-r projection Q, $\operatorname{tr}(h(I-Q)) \geq T_r(\Lambda_*)$. With $r = 20k + j$, $0 \leq j < 20$,*

$$T_r(\Lambda_*) = 4^{-k} \begin{cases} 1 - j/24, & 0 \leq j \leq 16, \\ 1/3 - (j-16)/48, & 16 \leq j < 20. \end{cases}$$

In particular $T_{20k} = 4^{-k}$. No perfect pure strategy discards less Schmidt probability at a fixed truncation rank. This is not an optimal game-loss assertion for the compressed measurements.

Proof. The Ky Fan variational principle bounds $\operatorname{tr}(hQ)$ by the first r eigenvalues, and Theorem 13.5 bounds those by the released prefix. Sum the remaining sixteen/four level probabilities to obtain the formula. The displayed strategy attains each retained-weight bound. □

13.5 A capacity infimum attained only beyond type I

Theorem 13.8 (Sharp but unattained type-I capacity)**.** *Among perfect normal type-I marginals with finite second modular moment,*

$$\inf C_E = \frac{8}{49} (\ln 2)^2,$$

and no such marginal attains it. The same bound is attained by the scalar-boundary II_1 witness of Theorem 14.3.

Proof. The graded criterion and Theorem 13.2 give the lower bound. Equality would make the boundary density scalar, $h_B = bB$. In type I its counting measure is $m\delta_b$ for a finite positive integer m: its probability mass bounds mb by one. The ungraded trace renewal would then put multiplicity $m/4^j$ at $b/2^j$ for every j. Eventually this lies strictly between zero and one, contrary to integer multiplicity.

For an explicit approach choose admissible values $\mu_k = c_N 4^{-k}$ with multiplicities

$$m_k = \begin{cases} 16^{N-k}, & 0 \le k \le N, \\ 1, & k > N, \end{cases} \qquad c_N = \left(72 \sum_{k \ge 0} m_k 4^{-k} \right)^{-1}.$$

Each inequality $m_k \le 16 m_{k+1}$ holds, so Theorem 13.4 gives a perfect infinite-rank strategy. Its boundary consists of 16^{N+1} copies at $4c_N$ and fifteen copies at $c_N 4^{-k}$ for every $k \ge N$. The normalized off-top boundary mass is $t_N = 5/(16\,64^N + 5)$. Conditional on this tail its modular displacement, in units of $\ln 2$, is $2(N + 1 + G)$ with G geometric of ratio $1/4$. Its mean is $2N + 8/3$ and variance $16/9$. Consequently

$$\frac{C_E}{(\ln 2)^2} = \frac{8}{49} + \frac{16}{9} t_N + t_N(1 - t_N)(2N + 8/3)^2. \tag{44}$$

The positive excess is $O(N^2 64^{-N})$ and tends to zero. At $N = 0$ the expression is $152/81$, the released state's value. These are actual perfect infinite strategies, not finite approximants confused with perfect ones. The type-II scalar boundary attains equality by the internal construction in Section 14. □

14 Prepared play in type II and its spectral fingerprint

Let M have a faithful normal semifinite trace Tr. A normal Alice state has density $h \ge 0$, $\operatorname{Tr} h = 1$, support p, and square root $K = h^{1/2} \in L^2(M, \operatorname{Tr})$. The standard partner is the full right multiplication commutant on this L^2 space. Trace entropy is $S_{\operatorname{Tr}}(h) = -\operatorname{Tr}(h \ln h)$, which can be negative. Support trace is not Hilbert-space dimension. Existence statements below specify the standard partner; the necessary bounds permit arbitrary commuting extensions with the given normal Alice marginal.

14.1 Compression to the standard partner

Lemma 14.1 (Marginal reduction and the tracial obstruction)**.** *Every commuting strategy with Alice marginal $a \mapsto \operatorname{Tr}(ha)$ has the same correlations as left and right POVMs on $L^2(pMp, \operatorname{Tr})$, with vector K. If $h = p/T$, $T = \operatorname{Tr} p < \infty$, its Γ loss is at least $(25/52)\delta_{\mathrm{FMS}} > 0$, where $\delta_{\mathrm{FMS}} = 1 - \omega_{\mathrm{tr}}(G_{\mathrm{FMS}})$.*

Proof. In the strategy's GNS representation define $V(aK) = \pi(a)\xi$. Equality of the two squared norms proves this map is isometric. Its initial closure is $L^2(M)p$: spectral cutoffs $e_n = 1_{[1/n,n]}(K)$ and bounded inverses on e_n approximate bounded square-integrable elements of that right support. The range reduces Alice's algebra. Compress a commuting Bob effect by V; the compression is a positive contraction commuting with left multiplication, and the finite sum of compressed effects is identity. The commutant on $L^2(M)p$ is right pMp, giving

$$p(a, b|x, y) = \operatorname{Tr}(K E_a^x K F_b^y).$$

Replace Alice effects by $pE_a^x p$ and work in pMp. This uses compression, not invariance of Bob's original effects on Alice's cyclic subspace. The standard density realization is established background [12, Section 2].

For $h = p/T$ the correlation is tracial in the normalized corner. Mančinska's strict tracial FMS gap [23, Theorem 3] gives $\delta_{\mathrm{FMS}} > 0$; discarding all losses except copy one's twenty-five equally weighted pairs proves the claimed floor. Compactness of the universal tracial POVM class ensures

the strict gap: the tracial states of the universal finite-POVM C^*-algebra form a weak-star compact set, and the score is continuous. Thus this is stronger than nonattainment by finite maximally entangled vectors. □

Lemma 14.2 (The internal density-sector law)**.** *In a perfect FMS strategy after this reduction, the magic corner consists of four internally equivalent copies of a flag corner with doubled density. In particular, for flag f and magic m,*

$$\begin{aligned} \operatorname{Tr} m &= 4 \operatorname{Tr} f, & \operatorname{Tr}(h_m) &= 8 \operatorname{Tr}(h_f), \\ \operatorname{Tr}(h_m \ln h_m) &= 8 \operatorname{Tr}(h_f \ln h_f) + 8 \ln 2 \operatorname{Tr}(h_f). \end{aligned} \tag{45}$$

The spectral-measure version requires no entropy moment or finite support trace.

Proof. Compress to the faithful support of h. For positive effects, zero correlation is $\|E^{1/2} K F^{1/2}\|_2^2 = 0$, hence $EKF = 0$. Checked matching classes give $E_c K = K F_c$ by summing zero cross terms. Orthogonality gives $E_c E_d K = 0$ for $c \neq d$, so the class effects are projections. Cancellation is valid even for unbounded K: if bounded $AK = 0$, then $Ae_n = 0$ on the spectral cutoffs above, whence $A = 0$. Taking adjoints and products in L^1 gives $E_c h = h E_c$, and thus $E_c = F_c$.

The checked magic-square reflections commute with h. Their elementary algebra from Lemma 3.2 generates matrix units of an internal $M_4 \subset mMm$: two anticommuting reflection pairs commute with each other. Let q be a minimal diagonal corner of this matrix algebra and $d = qhq$. The magic density is four equivalent copies of d. The distinguished column and row isometries from q have overlaps $\sigma_{xy} q/\sqrt{2}$. Pulling the calibrated sign supports back to $f + q$ gives

$$W_{xy} = K_f + \sigma_{xy} \sqrt{d}/\sqrt{2}.$$

The two forbidden sign equations yield both $X_x W_{xy} = \sigma_{xy} W_{xy} Y_y$ and $X_x W_{xy} Y_y = \sigma_{xy} W_{xy}$. Full-support cancellation makes these sign differences reflections. With $H = K_f + \sqrt{d}/\sqrt{2}$ and $Z = f - q$, the three positive cells identify $X_0 = X_1 = X$, and the frustrated cell gives $X = -ZXZ$. Its off-diagonal block $\beta \in qMf$ satisfies $\beta^*\beta = f$, $\beta\beta^* = q$. From $XH = HY_0$ and its adjoint one gets $XH^2 = H^2 X$, hence

$$d = 2\beta h_f \beta^*.$$

These are internal partial isometries, so they preserve the trace and spectral calculus. They prove the trace, probability and logarithmic identities, and the corresponding bounded-Borel functional identities. If either sector vanishes, the equivalence forces both to vanish. For affiliated unbounded densities the same cutoff cancellation and $L^2 \cdot L^2 \subset L^1$ justify all products used. One does not truncate the full strategy and assume it remains perfect. On finite support, the negative part of $u \ln u$ is bounded and a finite total entropy makes all corner moments finite; otherwise the entropy upper inequality below is trivial at $-\infty$. □

14.2 A perfect density in every finite factor

Theorem 14.3 (Prepared type-II play and a sharp entropy deficit)**.** *Every* II_1 *factor with its standard partner admits perfect* Γ *play in a bounded faithful normal density. For every perfect normal semifinite strategy with support trace* $T < \infty$*,*

$$S_{\mathrm{Tr}}(h) \leq \ln T - \Delta_*, \qquad \Delta_* = \ln(7/6) - \frac{\ln 2}{7} = 0.05512965403298\ldots.$$

Equality is attained in every type-II finite corner of trace T. In a II_1 *factor the normalized trace has a positive loss floor, whereas a full-support prepared density attains trace entropy $-\Delta_*$ and wins perfectly.*

Proof. In a II_1 factor use continuous projection dimension and equivalence of equal-trace projections to choose orthogonal U_k, V_k, $k \geq 0$, with

$$\tau(U_k) = \tfrac{3}{4}16^{-k}, \quad \tau(V_k) = \tfrac{3}{16}16^{-k}, \quad h = \sum_{k\geq 0} \left(\tfrac{7}{6}4^{-k}U_k + \tfrac{7}{12}4^{-k}V_k\right).$$

Their traces sum to one, as does $\tau(h)$. The density is bounded by $7/6$ and has support one, though no bounded inverse. Split U_k into four pieces equivalent to V_k and use the ideal FMS PVMs with flag V_k. For copy two split V_k into four pieces equivalent to U_{k+1} and use flag U_{k+1}. Indeed their traces satisfy $\tau(U_k) = 4\tau(V_k)$ and $\tau(V_k) = 4\tau(U_{k+1})$, while the paired density ratios are two. Use boundary answer on U_0. All ideal blocks have the required square-root coefficient ratio. Direct sums give bounded complete PVMs; cross-copy compatibility is not required by unsupported question pairs. Bob labels $U = \sum U_k$ and $V = \sum V_k$. Both linked Alice projections are U, commuting with h, so every supported pair wins. The level state masses are $(7/8)64^{-k}$ and $(7/64)64^{-k}$, which sum to one. Direct summation gives $\tau(h \ln h) = \Delta_*$.

For the universal bound denote copy-one flag and magic by F_1, M_1, and copy-two flag, magic and boundary by P, M_2, B. Connected matching and the two links give $F_1 = M_2$ and $M_1 = P + B$. Two applications of (45) imply

$$\operatorname{Tr} P = T/20, \quad \operatorname{Tr} B = 3T/4, \qquad \operatorname{Tr}(hP) = 1/72, \quad \operatorname{Tr}(hB) = 7/8.$$

The nonboundary sector cannot vanish, by the same linked-sector argument as before. Write $A = \operatorname{Tr}(h_P \ln h_P)$, $J = \operatorname{Tr}(h_B \ln h_B)$ and $D = \operatorname{Tr}(h \ln h)$. The logarithmic sector identities give

$$D = 72A + \tfrac{17}{9}\ln 2, \qquad J = 63A + \tfrac{16}{9}\ln 2, \qquad D = \tfrac{8}{7}J - \tfrac{1}{7}\ln 2.$$

Jensen on the spectral trace measure of the boundary gives $J \geq (7/8)\ln(7/(6T))$, proving the bound. More precisely, set $g_B = (6T/7)h_B$ and $\tau_B = \operatorname{Tr}/(3T/4)$ on BMB. Then

$$\ln T - \Delta_* - S_{\mathrm{Tr}}(h) = \tau_B(g_B \ln g_B) \geq 0.$$

Equality is exactly a scalar boundary density $h_B = 7B/(6T)$. In any finite type-II corner perform the preceding construction with its normalized trace and divide the resulting density by T relative to Tr. Its boundary is scalar, attaining equality. Completion of measurements off the support changes no correlation. The trace-floor assertion is Lemma 14.1. □

Corollary 14.4 (Infinite semifinite support and finite entropies)**.** *A II_∞ factor with its standard partner admits perfect play at every finite real value of trace entropy. If it is countably decomposable, these values also have bounded faithful witnesses.*

Proof. Choose a finite corner of trace $T = \exp(s + \Delta_*)$. Theorem 14.3 gives entropy s. For the faithful version partition identity into countably many equal-trace projections of trace T_0, put the normalized-corner equality density $h_{0,j}$ in each, and set $h = \sum_{j\geq 0} 2^{-j-1}h_{0,j}/T_0$. Its direct-sum measurements are perfect, its support is identity, and its entropy is $2\ln 2 + \ln T_0 - \Delta_*$. Vary T_0. Countable decomposability is needed for this full-support normal-state assertion, not the finite-corner one. □

For the displayed finite-factor density, retaining $0 \leq k < K$ discards trace 16^{-K} and probability 64^{-K}. The only incomplete block is the final V_{K-1} block. The ideal coefficient matrix $\operatorname{diag}(0, b, b, b, b)$ loses unweighted probability b^2 in its four sign pairs, so the normalized loss is

$$\ell_K = \frac{7}{208(64^K - 1)}.$$

The retained corners have integral ranks at trace denominator $4\,16^K$ and yield matrix strategies of dimension $d_K = 4(16^K - 1)$. These are finite approximants of the trace construction, with compressed POVMs at the last block. They do not improve the dimension tradeoff: the original compression with $L = 3K - 2$ has smaller dimension $20(3K - 2)$ and no greater loss. Indeed with $X = 64^K \geq 64$, $16/(X - 16) \leq 21/(X - 1)$ is equivalent to $X \geq 64$, comparing the two exact losses, and $20(3K - 2) < 4(16^K - 1)$ holds at one and increases thereafter. Entropy relative to a rescaled trace must also be rescaled; approximate strategies are not subject to the exact-perfect entropy inequality.

14.3 A necessary modular spectral law

Put $K_h = -\ln h$ on its support, and define its state-weighted spectral law $\mu(g) = \mathrm{Tr}(hg(K_h))$. When defined, write $S_0 = \ln \mathrm{Tr}\, p$, $S_1 = \mu(K_h)$, $S_2 = -\ln \mathrm{Tr}\, h^2$ and $C_E = \mathrm{Var}_\mu K_h$. This capacity is established terminology [13, Sections 2–3]. Trace rescaling shifts the entropies and K_h together and leaves gaps and variance unchanged.

Theorem 14.5 (A logarithm-of-two fingerprint)**.** *For every normal semifinite perfect strategy there is a boundary law ν_B such that*

$$\mu = \nu_B * \mathrm{law}((\ln 2)J), \qquad \Pr(J = 2k) = \tfrac{7}{8}64^{-k}, \quad \Pr(J = 2k+1) = \tfrac{7}{64}64^{-k} \quad (k \geq 0).$$

If the modular mean and variance are finite,

$$C_E \geq \tfrac{8}{49}(\ln 2)^2.$$

Whenever the entropy differences are defined, without subtracting two infinities,

$$S_0 - S_1 \geq \Delta_*, \qquad S_1 - S_2 \geq \frac{\ln 2}{7} + \ln(14/15).$$

The three constants are sharp in type II, attained simultaneously by Theorem 14.3's scalar-boundary witness. This is a necessary scalar spectral law, not a sufficient condition for perfect play or a physical tensor decomposition.

Proof. For a density-reducing projection e, let $\mu_e(g) = \mathrm{Tr}(heg(K_h))$, and let S_t translate measures by t. The internal sector law gives, for bounded Borel functions and $a = \ln 2$,

$$\mu_m = 8S_{-a}\mu_f.$$

The proof of Lemma 14.2 used no finite-support-trace assumption before its entropy estimate, so this measure identity also holds at infinite support trace. The links imply

$$\mu_P + \mu_B = 64S_{-2a}\mu_P, \quad \mu_P = \sum_{k\geq 1} 64^{-k} S_{2ka}\mu_B, \quad \mu_{M_2} = \sum_{k\geq 1} 8\,64^{-k} S_{(2k-1)a}\mu_B.$$

Iterate the first equation to prove the second; its remaining positive measure has mass $64^{-n}\mu_P(\mathbb{R}) \to 0$. The full measure is $\mu_B + \mu_P + \mu_{M_2}$, of mass $(8/7)\mu_B(\mathbb{R}) = 1$. Hence $\nu_B = (8/7)\mu_B$ is a probability law and gives the stated convolution.

Realize its two summands as independent classical random variables. Here $J = 2G + Q$, with G geometric of ratio $1/64$ and Q Bernoulli of probability $1/9$. Elementary sums give

$$\mathbb{E}J = 1/7, \quad \mathrm{Var}\, J = 8/49, \quad \mathbb{E}2^J = 7/6, \quad \mathbb{E}2^{-J} = 14/15.$$

Variance adds. Also $\operatorname{Tr} p = \mathbb{E}_\mu e^{K_h}$ and $\operatorname{Tr} h^2 = \mathbb{E}_\mu e^{-K_h}$. Independence and scalar Jensen give

$$S_0 - S_1 = \ln \mathbb{E}_{\nu_B} e^X - \mathbb{E}_{\nu_B} X + \Delta_*,$$
$$S_1 - S_2 = \mathbb{E}_{\nu_B} X + \ln \mathbb{E}_{\nu_B} e^{-X} + \tfrac{1}{7} \ln 2 + \ln(14/15).$$

These identities prove the inequalities in their defined ranges and the unambiguous extended cases. With finite terms equality means ν_B is a point mass; the same condition minimizes variance. The type-II witness has constant boundary density, and therefore attains all three bounds. No finite-matrix perfect witness is claimed. □

For comparison the explicit oscillator state's actual boundary has sixteen A vectors at level zero and fifteen at every later level. Its conditional modular law is

$$\Pr(K_B = \ln 24) = 16/21, \qquad \Pr(K_B = \ln 24 + 2k \ln 2) = \tfrac{5}{7} 4^{-k} \ (k \geq 1).$$

Convolution gives its full sixteen/four level law; exact moments give $\operatorname{Var}((K_h - \ln 24)/\ln 2) = 152/81 = 8/49 + 6800/3969$. Its capacity is thus strictly above the type-II minimum. The fingerprint alone omits the internal measurement algebra: perfect play also requires nonzero matrix corners in the density centralizer.

Normal semifinite algebras have no normal universal embezzling states [31, Corollary 54, arXiv version 4]. This known obstruction concerns an entanglement-transformation task, and does not forbid the prepared strategies for one fixed game constructed here. The stronger quoted invariant assertion uses a separable-predual hypothesis [31, Corollary 55]. No gravity or conserved-charge interpretation is part of the preceding theorems.

Remaining questions

The leading logarithmic dimension law and the exact limiting entropy and oscillator prices are established. What are the quantitative rates at which the resource prices approach their limits? Can the dimension law's $O(\log\log(1/\varepsilon))$ remainder be reduced to $O(1)$, and which finite strategies minimize loss at each dimension? The sharp discarded-mass profile of perfect spectra does not settle optimal loss after compression. The unique maximizing spectrum does not assert uniqueness of measurements or of the whole strategy. Quantitative stability of the exact integer, density and fingerprint laws remains open. The field construction also leaves microscopic implementation costs unresolved; its field-energy optimum is not fixed by the oscillator number 275/18.

How this work was done

The authors developed the construction and proofs with AI models from more than one developer. The proofs were checked against the cited primary statements and by finite symbolic and numerical calculations. The square-root word extraction is proved here directly, with its norm convention, constant and row/column hypotheses stated explicitly. Nonexistence in finite dimension follows from the sector-and-minimum proof; numerical optimization is not a certificate of nonexistence.

The ancillary package contains the rational predicate, a deterministic generator, symbolic checks, finite compression checks, recorded outputs, a sequential runner, and SHA-256 checksums. The symbolic checks verify all forbidden gadget products exactly, projectivity, completeness, overlap signs, normalization and the incomplete-block loss. Numerical checks cover every supported pair for the specified compressions and compare their loss with (1). Positivity tolerances are supplemental;

positivity follows analytically by compression of projections. Additional exact checks cover the three bridge words, the order-32 real Pauli group and the trace-norm counterexample. Numerical checks test regular-representation extraction, complex local perturbations and sector defects, with one solver thread. A separate descent check verifies thresholds and multiplicities. These finite tests do not establish a uniform lower bound; the analytic sector, gluing and rank proofs do. The infinite strategy is justified by the normalized direct-sum construction. No proof-assistant certification or exhaustive historical-priority search is claimed.

The finite-energy checks recompute geometric occupation and entropy sums, passive truncations, exact storage enumeration and moment inequalities, local infinite-matrix rows, finite symplectic completion and noninvariant smooth-mode energy bounds. The vacuum checks independently reconstruct the classical value, testing inequalities and finite-algebra compression. The type-II checks reconstruct its matrix models, internal trace ratios, sharp boundary entropy identity and fingerprint convolution. All are finite checks of the stated interfaces; none is a continuum field simulation, a proof of a universal optimum, or a substitute for the written domain and operator-algebra arguments. The external faithfulness, localization and approximation inputs are credited where used. All numerical-library thread counts are one; recorded outputs, source adaptations and hash manifests accompany the runner.

The ladder checks represent infinite admissible multisets by finite integer data and an exact geometric tail. They compare 329 spectra using exact rational Ky Fan sums and full-tail passive energies, and test logarithmic entropy and positive-order Rényi evaluations separately. Exact rational capacity gaps test the perfect infinite approach sequence, and boundary checks retain the exceptional critical identity ladder. These computations supplement the classification, majorization and one-sided compactness proofs; they do not optimize arbitrary approximate strategies or prove a convergence rate.

References

[1] Astrid an Huef, Marcelo Laca, Iain Raeburn, and Aidan Sims. KMS states on the C^*-algebras of finite graphs. `https://arxiv.org/abs/1205.2194v1`, 2012. arXiv:1205.2194v1.

[2] Huzihiro Araki. Relative entropy of states of von Neumann algebras. *Publications of the Research Institute for Mathematical Sciences*, 11:809–833, 1976.

[3] Salman Beigi. Separation of quantum, spatial quantum, and approximate quantum correlations. *Quantum*, 5:389, 2021. arXiv:2004.11103.

[4] Stephan De Bièvre. Local states of free bose fields. `https://arxiv.org/abs/math-ph/0511037`, 2005. arXiv:math-ph/0511037.

[5] Joseph J. Bisognano and Eyvind H. Wichmann. On the duality condition for a Hermitian scalar field. *Journal of Mathematical Physics*, 16:985–1007, 1975.

[6] Detlev Buchholz, Claudio D'Antoni, and Klaus Fredenhagen. The universal structure of local algebras. *Communications in Mathematical Physics*, 111:123–135, 1987.

[7] Detlev Buchholz and Eyvind H. Wichmann. Causal independence and the energy-level density of states in local quantum field theory. *Communications in Mathematical Physics*, 106:321–344, 1986.

[8] Rob Clifton and Hans Halvorson. Entanglement and open systems in algebraic quantum field theory. *Studies in History and Philosophy of Modern Physics*, 32:1–31, 2001. arXiv:quant-ph/0001107.

[9] Andrea Coladangelo and Jalex Stark. Separation of finite and infinite-dimensional quantum correlations, with infinite question or answer sets. `https://arxiv.org/abs/1708.06522`, 2017. arXiv:1708.06522.

[10] Andrea Coladangelo and Jalex Stark. Unconditional separation of finite and infinite-dimensional quantum correlations. `https://arxiv.org/abs/1804.05116`, 2018. arXiv:1804.05116.

[11] Andrea Coladangelo and Jalex Stark. Robust self-testing for linear constraint system games. `https://arxiv.org/abs/1709.09267v2`, 2019. arXiv:1709.09267v2.

[12] Jason Crann, David W. Kribs, Rupert H. Levene, and Ivan G. Todorov. State convertibility in the von Neumann algebra framework. `https://arxiv.org/abs/1904.12664`, 2019. arXiv:1904.12664.

[13] Jan de Boer, Jarkko Järvelä, and Esko Keski-Vakkuri. Aspects of capacity of entanglement. *Physical Review D*, 99:066012, 2019. arXiv:1807.07357.

[14] Jan Dereziński. Bosonic quadratic Hamiltonians. `https://arxiv.org/abs/1608.03289v2`, 2017. arXiv:1608.03289v2.

[15] Seth Douglas. The I_{3322} quantum value is attained spatially but not in finite dimension. `https://arxiv.org/abs/2609.05555v2`, 2026. arXiv:2609.05555v2.

[16] Seth Douglas and Nidhal Mghirbi. The uncomputable price of precision. `https://doi.org/10.5281/zenodo.23753741`, 2026. Version 4.0.0. Preprint.

[17] Jan Gulla and Johannes Skaar. Photon localization: Local measurements, localized states, and causality. `https://arxiv.org/abs/2610.30473v1`, 2026. arXiv:2610.30473v1.

[18] Uffe Haagerup. Connes' bicentralizer problem and uniqueness of the injective factor of type III_1. *Acta Mathematica*, 158:95–148, 1987.

[19] Hans Halvorson and Rob Clifton. Generic Bell correlation between arbitrary local algebras in quantum field theory. *Journal of Mathematical Physics*, 41:1711–1717, 2000. arXiv:math-ph/9909013.

[20] Se-Jin Kim, Vern I. Paulsen, and Christopher Schafhauser. A synchronous game for binary constraint systems. `https://arxiv.org/abs/1707.01016`, 2017. arXiv:1707.01016.

[21] Olivier Lalonde. Maximally entangled states are not complete for pseudo-telepathy. `https://arxiv.org/abs/2608.05378`, 2026. arXiv:2608.05378.

[22] A. L. Licht. Strict localization. *Journal of Mathematical Physics*, 4:1443–1447, 1963.

[23] Laura Mančinska. The right entanglement for perfect nonlocality and zero-error communication. `https://arxiv.org/abs/2610.06512`, 2026. arXiv:2610.06512v1.

[24] Michael A. Nielsen. Conditions for a class of entanglement transformations. *Physical Review Letters*, 83:436–439, 1999. arXiv:quant-ph/9811053.

[25] Miklós Rédei and Stephen J. Summers. Local primitive causality and the common cause principle in quantum field theory. *Foundations of Physics*, 32:335–355, 2002. arXiv:quant-ph/0108023v1.

[26] Volkher B. Scholz and Reinhard F. Werner. Tsirelson's problem. `https://arxiv.org/abs/0812.4305v1`, 2008. arXiv:0812.4305v1.

[27] William Slofstra. The set of quantum correlations is not closed. `https://arxiv.org/abs/1703.08618`, 2017. arXiv:1703.08618v2.

[28] William Slofstra and Thomas Vidick. Entanglement in non-local games and the hyperlinear profile of groups. `https://arxiv.org/abs/1711.10676v2`, 2018. arXiv:1711.10676v2.

[29] Stephen J. Summers. Yet more ado about nothing: the remarkable relativistic vacuum state. `https://arxiv.org/abs/0802.1854v2`, 2009. arXiv:0802.1854v2.

[30] 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.

[31] Lauritz van Luijk, Alexander Stottmeister, Reinhard F. Werner, and Henrik Wilming. Embezzlement of entanglement, quantum fields, and the classification of von Neumann algebras. `https://arxiv.org/abs/2401.07299v4`, 2025. arXiv:2401.07299v4.

[32] Lauritz van Luijk, Alexander Stottmeister, Reinhard F. Werner, and Henrik Wilming. Pure state entanglement and von Neumann algebras. *Communications in Mathematical Physics*, 406:296, 2025. Numbered references use arXiv:2409.17739v2.

[33] Reinhard F. Werner. EPR states for von Neumann algebras. `https://arxiv.org/abs/quant-ph/9910077v1`, 1999. arXiv:quant-ph/9910077v1.

[34] Xingyao Wu, Jean-Daniel Bancal, Matthew McKague, and Valerio Scarani. Device-independent parallel self-testing of two singlets. *Physical Review A*, 93:062121, 2016. arXiv:1512.02074v2.