Perfect play with spatial quantum strategies is Σ_2^0-complete

Seth Douglas Nidhal Mghirbi

Version 1.0.0, November 2026

Abstract

When does a finite nonlocal game admit perfect spatial quantum play? We prove that this decision problem is complete for the second existential level of the arithmetical hierarchy. The same classification holds when a perfect strategy must have finite marginal entropy. Hardness persists among games whose finite-dimensional value is one and which have no finite-dimensional perfect strategy. A finite clock-phase order supplies the existential witness: normality forces a finite order on a nonzero sector, while the strategy chooses the order.

For the upper bound, we normalize the largest Schmidt coefficient and use compact operator sandwiches. Fixing a lower bound on the largest Schmidt probability gives a co-recursively enumerable stratum. No computable bound selects a sufficient stratum for every game. Exact attainment of a rational Bell bound, of an unspecified supremum, and of a specified scalar value is also Σ_2^0-complete. The hardness proof extends the modular-shift and rational-row toolkit of our earlier entropy-game paper to four calibrated bridges and a deterministic computation on rooted finite squares.

AI systems assisted with this work. The authors take responsibility for its claims; see "How this work was done." This preprint has not been peer reviewed.

Contents

1 Introduction **3**
1.1 Prior work and contribution . . . 4
1.2 Proof route . . . 5

2 Preliminaries **5**

3 The upper bound **7**
3.1 Normalize one Schmidt coefficient . . . 7
3.2 An effective finite predicate . . . 7
3.3 No effective choice of stratum; vanishing largest weight . . . 8
3.4 Finite entropy and rational slack . . . 8

4 A finite compiler with four bridges **9**
4.1 The normal flagged magic-square bridge . . . 9
4.1.1 Questions, calibration and exact bindings . . . 9
4.1.2 Normal bridge theorem . . . 10
4.1.3 A shared question fixes transported labels . . . 12
4.2 Rational projection equations that certify products . . . 12
4.2.1 Graphs and triangles . . . 12
4.2.2 Binding prescribed parameters . . . 13
4.3 Rational rows as Boolean winning conditions . . . 13
4.3.1 A certified scalar matrix algebra . . . 14
4.3.2 Signed reflections and column graphs . . . 14
4.4 The admissible presentations . . . 15
4.5 A dimension-independent completeness estimate . . . 19

5 The rooted tree and computation squares **19**
5.1 Normality removes every rootless sector . . . 19
5.2 Each depth is a finite square . . . 20
5.3 A finite deterministic tableau . . . 20
5.4 A worked front and mask example . . . 22

6 A finite-order clock and its period reader **22**
6.1 A determinant supplies a finite witness . . . 22
6.2 The clock and a counter without a positive lower bound . . . 23
6.3 A bounded inverse on every prime clock . . . 23

7 The reduction **24**

8 Value one for every constructed game **25**

9 Exact Bell attainment **26**
9.1 A specified global bound . . . 26
9.2 No bound supplied . . . 27
9.3 A specified scalar value . . . 27

10 Discussion and open problems **27**
10.1 Prescribed tables and escaping mass . . . 28
10.2 The Fu–Miller–Slofstra remark . . . 28
10.3 A possible von Neumann type ladder . . . 29
10.4 Other questions . . . 29

How this work was done **29**

A Reproducibility **29**

1 Introduction

When does a finite game admit a strategy that wins every round? Approaching certainty and attaining certainty are different questions. They also depend on the model of quantum strategy. Table 1 puts the distinction first. The spatial entry is the result of this paper.

Decision problem	Complete class	Source
Perfect finite-dimensional strategy	Σ_1^0	[8]
Perfect commuting strategy	Π_1^0	[18, 14]
Finite-dimensional value equals one	Π_2^0	[14]
Perfect spatial strategy	Σ_2^0	This paper

Table 1: Finite rational Boolean games; computable many-one reductions. A value is a supremum. A perfect strategy attains one.

The finite-dimensional and spatial models use tensor products. The latter allows an arbitrary normal state on infinite-dimensional local Hilbert spaces. The commuting model instead allows the two measurement algebras to act on one Hilbert space and commute. Its compact state-space description makes perfect commuting play equivalent to commuting value one. Spatial play has no such compactness in normalized states.

Our upper bound keeps one Schmidt coefficient from disappearing. After that coefficient is scaled to one, the amplitude operator belongs to a norm-compact family of compact operators. Losing probabilities become homogeneous zero equations. These survive a limit even when some squared Schmidt mass escapes. The surviving state can be normalized again.

The hardness mechanism goes in the other direction. A finite Weyl clock commutes with a trace-class density. Each positive density eigenspace is finite-dimensional, so a determinant forces the clock phase to have finite order. The strategy chooses that order. A deterministic computation reads it and checks an infinite sequence of decidable conditions. This is the existential quantifier, and no computable bound on the game can capture its size. The modular-shift toolkit comes from our earlier paper [5]; here four calibrated bridges connect it to the computation.

Theorem A (spatial perfection). *Deciding whether a finite rational* $0/1$ *game has a perfect normal spatial strategy is* Σ_2^0*-complete. The same holds when a perfect strategy with finite marginal entropy is required. Hardness holds under the promise* $\omega_q(G) = 1$*, for a family in which no instance has a perfect finite-dimensional strategy.*

Finite marginal entropy can be required on both sides. For pure strategies the two marginal entropies agree. For a perfect mixed strategy with one finite marginal entropy, a pure component has finite entropy and is still perfect; see Section 3.

Theorem B (Bell attainment). *For rational Bell functionals, each of the following is* Σ_2^0*-complete in the normal spatial model: attainment of a supplied rational global upper bound; attainment of the supremum, with no bound supplied; and attainment of a supplied rational scalar value. In the first problem, validity of the supplied bound can either be promised or required as part of the decision problem.*

Theorem C (compactness and Schmidt mass). *Uniformly for integers $\kappa \geq 1$, the existence of a perfect pure spatial strategy with largest Schmidt probability at least $1/\kappa$ is a Π^0_1 property. Their union is $\mathrm{PERF}_{\mathrm{qs}} \in \Sigma^0_2$. No total computable function $\kappa(G)$ bounds a sufficient stratum for every spatially perfect game. If $\omega_q(G) = 1$ and no perfect spatial strategy exists, the largest Schmidt probability tends to zero along every sequence of pure spatial strategies whose winning probabilities tend to one.*

1.1 Prior work and contribution

The finite-dimensional entry in Table 1 uses perfect completeness in Theorem 12.6 of Ji, Natarajan, Vidick, Wright and Yuen [8], not just a gap statement about a supremum. Upper membership follows by a dimension search and real quantifier elimination. Slofstra's embedding theorem [18] and the account of Mousavi, Nezhadi and Yuen [14] give the commuting entry. Their value-one theorem gives the third entry. The commuting upper bound uses the convergent moment hierarchy of Navascués, Pironio and Acín [15]. The operator-solution framework for binary constraint system games originates with Cleve and Mittal [3]; Cleve, Liu and Slofstra gave the commuting-operator extension [2]. Slofstra also established nonclosure of finite-dimensional quantum correlations [17]. These are distinct earlier results. Lin's $\mathsf{MIP}^{co} = \mathsf{coRE}$ theorem concerns gap approximation in the commuting model [10]; it does not classify spatial attainment.

Infinite-dimensional spatial correlations were constructed by Coladangelo and Stark [4]. Kim, Paulsen and Schafhauser prove that perfect spatial play collapses to perfect finite-dimensional play for synchronous games [9, Corollary 3.11]. Lupini and coauthors establish the corresponding equivalence for imitation games [11, Corollary 5.4]. Our games retain nonsynchronous bridge questions. No YES instance is synchronous or an imitation game, since either condition would imply finite-dimensional perfect play.

Mančinska's flagged magic-square construction supplies the bridge geometry [12]. The normal bridge, shared-label refinement, graph products and Boolean rational-row tests were developed in [5, Lemmas 4.1 and 4.2, Theorem B, Eq. (32)]. The versions needed here are restated and proved in Section 4. That section extends the assembly from a particular presentation to four bridges and arbitrary finite lists of endpoint labels. The rooted four-ary computation and the finite-period reader then give the hardness reduction. The compact-sandwich upper bound does not use this compiler.

Fu, Miller and Slofstra study fixed-scenario correlation membership [7]. A remark there asserts, without proof, that fixed-scenario membership in C_q and C_{qs} is in RE; Section 10 records a conditional tension with Theorem A. We do not identify prescribed-table membership with game perfection.

To our knowledge, Theorems A and B give the first Σ^0_2-completeness classifications of these spatial attainment problems. This priority statement is provisional. It rests on the cited hierarchy and correlation papers, their nearby results, and a targeted literature search. It is not an exhaustive bibliographic certificate.

$$\begin{gathered}
\text{Presentation} \xrightarrow{\text{Lemma 4.4}} \text{finite game } G_e \\
\text{Perfect normal strategy} \longrightarrow \text{normal realization} \\
\longrightarrow \text{Fock tree} \longrightarrow \text{finite clock sector} \\
\longrightarrow \text{deterministic tableau} \longrightarrow \forall t\ R(e,n,t) \\
\text{Witness } n \longrightarrow \text{Fock/Weyl model} \\
\longrightarrow \text{scaffold, frame, row ancilla} \longrightarrow \text{perfect strategy}
\end{gathered}$$

Figure 1: Dependencies of the hardness proof. Compactness and Bell attainment have separate upper-bound arguments.

1.2 Proof route

Figure 1 summarizes the hardness proof. Section 3 proves Theorem C and the finite-entropy upper bound. Section 4 gives the finite game compiler. Sections 5–7 encode a chosen finite witness and all its checks. Section 8 proves the value-one promise with four explicit boundary defects. Section 9 proves Theorem B. The verification appendix records what the finite programs check.

2 Preliminaries

A finite rational Boolean game is a finite list of questions and answers, a rational distribution $\pi(x, y)$, and a predicate $V(a, b \mid x, y) \in \{0, 1\}$. A spatial strategy consists of positive operator-valued measurements (POVMs) $E_{a|x}$ and $F_{b|y}$ and a normal state, represented by a positive trace-class density ρ of trace one, on $\mathcal{H}_A \otimes \mathcal{H}_B$. Its winning probability is

$$\sum_{xyab} \pi(x, y)V(a, b \mid x, y) \operatorname{Tr} \rho(E_{a|x} \otimes F_{b|y}).$$

We write C_q, C_{qs}, C_{qc} for finite-dimensional, normal spatial, and commuting correlations. For a strategy $\mathcal{S}$, write $\omega(G, \mathcal{S})$ for the displayed winning probability and $\mathrm{loss}_G(\mathcal{S}) = 1 - \omega(G, \mathcal{S})$. The set $\mathrm{PERF}_{\mathrm{qs}}$ consists of games with a perfect normal spatial strategy. Its stratum $\mathrm{PERF}_{\mathrm{qs}}(\kappa)$ requires a perfect pure strategy whose largest Schmidt probability is at least $1/\kappa$. The set $\mathrm{PERF}_{\mathrm{finS}}$ requires a perfect strategy with finite marginal entropy. The value ω_q is the supremum over C_q. Finite Schmidt truncations show that C_q and C_{qs} have the same closure and the same supremum for every Bell functional. Entropy uses logarithms to base two: $S(h) = -\operatorname{Tr}(h \log_2 h)$ and $H(p) = -\sum_i p_i \log_2 p_i$, with $0 \log_2 0 = 0$. A projection-valued measurement (PVM) is a POVM consisting of orthogonal projections. Write $|\Phi_d\rangle = d^{-1/2} \sum_{j=0}^{d-1} |j, j\rangle$ for a finite maximally entangled vector. Vectorization of an operator T is $|T\rangle\rangle = \sum_{ij} T_{ij}|i, j\rangle$ when this sum is square summable. A purification is a unit vector with the given state as its partial trace over an added environment.

The set Σ_2^0 consists of predicates $\exists n \, \forall t \, R(e, n, t)$ with decidable R. The class Π_1^0 consists of universal decidable predicates. All encodings and reductions in this paper are effective finite encodings and computable many-one reductions. Completeness is for an unrestricted finite scenario, not one fixed question-and-answer size.

Fact 2.1 (pure components). If a mixed normal strategy is perfect, write $\rho = \sum_j \lambda_j |\psi_j\rangle\langle\psi_j|$. Each losing event is positive. Its expectation vanishes in every component with $\lambda_j > 0$, so every such component is perfect with the same measurements. Trace-class marginals have separable support. Thus all arguments may be made on separable state supports.

Fact 2.2 (matching and support). For effects on opposite sides, perfect binary matching gives, on a purification Ω,

$$0 = \langle E + F - 2EF\rangle = \langle (E - F)^2 + (E - E^2) + (F - F^2)\rangle.$$

Each positive term annihilates Ω. If T is the coefficient operator of the pure vector on the chosen Schmidt cut, matching gives $ET = TF^{\mathsf{T}}$; dense coefficient range first makes the local state support reducing. On that support the effects are projections and commute with the marginal density. A positive operator B with $\operatorname{Tr}(hB) = 0$ for faithful h is zero, since $B^{1/2}h^{1/2} = 0$ and $h^{1/2}$ has dense range.

Fact 2.3 (sharp compressions and joint measurements). If PEP is a projection for an effect E and support projection P, positivity shows that E acts as identity on its range and zero on its kernel inside $P\mathcal{H}$. Hence E reduces the support. A POVM with sharp coarse marginals is their joint spectral measurement: each fine effect is supported in the corresponding marginal ranges, and summation gives the commuting products. These facts apply after compression before projectivity is used.

Fact 2.4 (synchronization and forbidden events). We call a projection synchronized when its binary question is perfectly matched to the corresponding opposite question. If synchronized projections satisfy $P_A^{\mathrm{sync}}\Omega = P_B^{\mathrm{sync}}\Omega$ and a forbidden event gives $Q_A P_B^{\mathrm{sync}}\Omega = 0$, dense Schmidt range gives the local identity $Q_A P_A^{\mathrm{sync}} = 0$. The arguments work on either party and on purifications of mixed states.

Table 2 collects the notation.

Table 2: Notation used throughout the proof. Distinct fonts and subscripts separate question names, state spaces and operators.

Symbol	Meaning
κ, D_{Sch}	Schmidt stratum integer and normalized amplitude operator
$D_{\mathrm{esc},n}$	escaping-mass example in Section 10
$\mathsf{R}_i, \mathsf{C}_j, \mathsf{T}_x, \mathsf{B}_y$	flagged magic-square questions
$P_\circ, \mathcal{K}$	root projection and root multiplicity space
$P_{\mathrm{mag}}, h_{\mathrm{mag}}$	magic support and its density block
$T_f, T_{\mathrm{mag}}, T_{\mathrm{cal}}, \Delta$	flag/magic coefficient blocks, their calibrated diagonal and its sign reflection
$U_x, U_y; P_x, P_y$	cyclic coordinate shifts; zero-column and zero-row projections
$J_{\mathrm{W}}, K_{\mathrm{W}}, Z_{\mathrm{clk}}$	Weyl pair and central clock phase
$W, C_{\mathrm{ctr}}, P_{\mathrm{ret}}, T_{\mathrm{inv}}$	coordinate clock, counter, return pulse and bounded reader inverse
P_c	tableau symbol projections, $c \in \Gamma$
e_g, M_g, w_g, v	scaffold corners, magic blocks, single bridges and their chain
L_r	decoded copy of a typed label L at corner e_r
$\mathcal{H}_{\mathrm{base}}, F_i, G_{ij}$	graph base space, frame projections and graph projections
B_{ij}, s_{par}	parameter-binding projections and its reflection label
$P_\triangle, \Pi_t, \mathcal{I}_t$	triangle projection, row-column partition and column isometries
$h_{\mathrm{old}}, h^{[m]}$	marginal before row ancilla; truncated Fock/Weyl density
C_{abs}	sum of absolute Bell coefficients
$\mathcal{S}, \omega(G, \mathcal{S})$	strategy and its winning probability

3 The upper bound

3.1 Normalize one Schmidt coefficient

For a pure strategy write $\psi = \sum_i s_i |i,i\rangle$ with $s_1 \geq s_2 \geq \cdots$ and $\sum_i s_i^2 = 1$. Set $d_i = s_i/s_1$ and $D_{\mathrm{Sch}} = \mathrm{diag}(d_i)$. For some integer κ,

$$d_1 = 1, \qquad d_i \downarrow 0, \qquad \operatorname{Tr} D_{\mathrm{Sch}}^2 \leq \kappa. \tag{1}$$

For positive effects E, F, vectorization gives

$$p(E,F) = \frac{\|\sqrt{E}\, D_{\mathrm{Sch}} \sqrt{F^{\mathsf{T}}}\|_{\mathrm{HS}}^2}{\operatorname{Tr} D_{\mathrm{Sch}}^2}, \qquad p(E,F) = 0 \iff E D_{\mathrm{Sch}} F^{\mathsf{T}} = 0.$$

To prove the reverse implication, $E D_{\mathrm{Sch}} F^{\mathsf{T}} = 0$ sends $D_{\mathrm{Sch}} \operatorname{ran} F^{\mathsf{T}}$ into $\ker E$. The closures of the ranges of F^{T} and $\sqrt{F^{\mathsf{T}}}$ agree, and $\ker E = \ker \sqrt{E}$. Continuity gives the square-root sandwich identity. The forward implication follows by multiplying the zero square-root sandwich by the square roots.

Let P_n be the first n coordinate projection. Monotonicity in (1) gives

$$\|D_{\mathrm{Sch}} - P_n D_{\mathrm{Sch}} P_n\| = d_{n+1} \leq \sqrt{\kappa/(n+1)}. \tag{2}$$

Consequently the family in (1) is compact in operator norm: take coordinatewise limits and use this common tail estimate; Fatou preserves the squared-sum bound. Finite families of POVMs on ℓ^2 are compact in the weak operator topology. Positivity and the finite sum-to-identity relations survive these limits.

Lemma 3.1 (Compact sandwiches)**.** *If $D_{\mathrm{Sch},j} \to D_{\mathrm{Sch}}$ in norm with compact D_{Sch}, and uniformly bounded E_j, F_j converge weakly to E, F, then $E_j D_{\mathrm{Sch},j} F_j$ converges weakly to $E D_{\mathrm{Sch}} F$.*

Proof. For fixed v, compactness makes $D_{\mathrm{Sch}} F_j v \to D_{\mathrm{Sch}} F v$ in norm. Pair this with the weakly convergent vectors $E_j^* u$. The replacement of D_{Sch} by $D_{\mathrm{Sch},j}$ has norm error tending to zero. □

3.2 An effective finite predicate

Let $\mathcal{L}_G$ be the losing tuples with $\pi(x,y) > 0$. Define $R(G,\kappa,n)$ by the existence of n-dimensional POVMs and real numbers $1 = d_1 \geq \cdots \geq d_n \geq 0$, $\sum_i d_i^2 \leq \kappa$, such that

$$|(E_{a|x} D_{\mathrm{Sch},n} F_{b|y}^{\mathsf{T}})_{ij}|^2 \leq \kappa/(n+1) \quad ((x,y,a,b) \in \mathcal{L}_G;\ i,j \leq n). \tag{3}$$

This is semialgebraic over the rationals. Splitting complex entries into real and imaginary parts and using principal minors for positivity makes real quantifier elimination a decision procedure.

A perfect strategy with $s_1^2 \geq 1/\kappa$ supplies all these witnesses. Compress its POVMs. Since P_n commutes with D_{Sch} and $E D_{\mathrm{Sch}} F^{\mathsf{T}} = 0$, the finite sandwich is $-P_n E (D_{\mathrm{Sch}} - D_{\mathrm{Sch},n}) F^{\mathsf{T}} P_n$, with norm bounded by (2).

Conversely choose a witness for every n. Extend $D_{\mathrm{Sch},n}$ by zero and each POVM by a fixed deterministic POVM on the complement. Pass to a subsequence with $D_{\mathrm{Sch},n} \to D_{\mathrm{Sch}}$ in norm and all effects weakly convergent. For every fixed matrix entry, (3) and Lemma 3.1 give $E D_{\mathrm{Sch}} F^{\mathsf{T}} = 0$. We retain

$$1 \leq \operatorname{Tr} D_{\mathrm{Sch}}^2 \leq \kappa, \qquad d_1 = 1.$$

The normalized vectorization of D_{Sch} is a perfect normal strategy with largest Schmidt probability $1/\operatorname{Tr} D_{\mathrm{Sch}}^2 \geq 1/\kappa$. Therefore

$$G \in \mathrm{PERF}_{\mathrm{qs}}(\kappa) \iff \forall n\ R(G,\kappa,n), \qquad \mathrm{PERF}_{\mathrm{qs}} = \bigcup_{\kappa\geq 1} \mathrm{PERF}_{\mathrm{qs}}(\kappa). \tag{4}$$

This proves the upper membership in Theorems A and C.

3.3 No effective choice of stratum; vanishing largest weight

If a total computable $\kappa(G)$ were sufficient for every YES game, (4) would put spatial perfection in Π_1^0. Perfect completeness of [8, Theorem 12.6] already makes it Σ_1^0-hard. NO instances of that map have spatial value at most $1/2$, because the finite-dimensional and spatial suprema agree. Thus the same map reduces the halting set to spatial perfection. The halting set cannot be Π_1^0, so such a function does not exist.

Now suppose pure spatial strategies approach perfect play while $p_{\max} \geq 1/\kappa$ on a subsequence. Normalize their largest Schmidt amplitudes. For a losing event,

$$\|ED_{\mathrm{Sch}}F^{\mathsf{T}}\|_{\mathrm{HS}} \leq \|\sqrt{E}D_{\mathrm{Sch}}\sqrt{F^{\mathsf{T}}}\|_{\mathrm{HS}} = \sqrt{\operatorname{Tr} D_{\mathrm{Sch}}^2\, p(E,F)}.$$

Positive question weights force every such probability to zero. The same compact limit gives a perfect spatial strategy. Contraposition proves the final assertion of Theorem C, including sequences of finite-dimensional pure strategies. Equivalently their Schmidt min-entropy tends to infinity.

3.4 Finite entropy and rational slack

For sorted Schmidt probabilities p_i, $p_i \leq 1/i$. If $H(p) \leq B$, then

$$\delta_n = \sum_{i>n} p_i \leq B/\log_2(n+1). \tag{5}$$

The normalized first-n truncation majorizes p, so has entropy at most B. Its pure-state trace distance from the original state is $\sqrt{\delta_n}$. This bounds the change in any winning probability. Sorted distributions with entropy at most B are compact in ℓ^1: (5) gives tightness and entropy is lower semicontinuous. Their square roots are compact in ℓ^2. Coupled with weak compactness of POVMs, this gives compactness of the resulting correlations.

For effectivity use slack rather than treating entropy as a polynomial. Fix integers $B, n \geq 1$. Choose computably a rational $\eta_{B,n} > 0$ converging to zero with $\eta_{B,n} > \sqrt{B/\log_2(n+1)}$. For every $k \geq 1$ compute a rational polynomial $H_{n,k}$ such that

$$\sup_{p\in\Delta_n} |H_{n,k}(p) - H(p)| \leq 1/k, \qquad \Delta_n = \{p_i \geq 0 : \sum_i p_i = 1\}.$$

Such polynomials are effective, for example by Bernstein approximation with rationally rounded coefficients to the computable continuous function $-t\log_2 t$ on $[0,1]$, with error $1/(nk)$ per summand. At precision k require sorted $p \in \Delta_n$, real amplitudes $s_i \geq 0$ with $s_i^2 = p_i$, finite POVMs, and

$$\mathrm{loss}_G(s,E,F) \leq \eta_{B,n}, \qquad H_{n,k}(p) \leq B + 1/n + 1/k.$$

The loss is a rational polynomial in the amplitudes and real and imaginary parts of the effects. Real quantifier elimination therefore decides this finite relaxation. Exact entropy feasibility implies every relaxation. Conversely, if every k-relaxation is feasible, compactness

of this finite variable domain yields a subsequential limit; the uniform approximation error gives $H(p) \le B + 1/n$ and the closed loss constraint survives. Thus exact stage feasibility is equivalent to feasibility of all these rational relaxations, even though their sets need not be nested. Flattening $\forall n \forall k$ gives a uniform Π^0_1 predicate for the stratum $\mathrm{PERF}_{\mathrm{finS}}(B)$ with supplied entropy budget B, defined by a perfect pure strategy with $H(p) \le B$.

A strategy of finite entropy satisfies the stages for an integer B above its entropy. Conversely stage witnesses have a common entropy bound $B+1$, hence the tail compactness just proved. A limit is perfect and has entropy at most B. Thus $\mathrm{PERF}_{\mathrm{finS}} \in \Sigma^0_2$. For mixed states, concavity of marginal entropy gives $\sum_j \lambda_j S((\psi_j)_A) \le S(\rho_A)$. If the right side is finite, a positive-weight component has finite entropy and is perfect. This proves the claimed equivalence of the entropy conventions.

4 A finite compiler with four bridges

The modular operators below are not synchronized. A synchronized operator commutes with the marginal density, so it cannot itself change a density eigenvalue. We retain conditional flagged magic-square questions to supply that change. The density centralizer is the set of bounded operators commuting with the density. Everything submitted to the graph compiler is in it.

We first restate the bridge, label and rational-row tools from [5, Lemmas 4.1 and 4.2, Theorem B, Eq. (32)], with their proofs. These arguments concern arbitrary normal supports and finite word lists. We then give the four-bridge assembly. No product of two modular transports is submitted as a centralizer word.

4.1 The normal flagged magic-square bridge

The magic-square relations originate with Mermin and Peres [13, 16]. We use Mančinska's flagged magic-square (FMS) geometry [12]. We spell out the predicate and prove its normal extension because the published finite-dimensional statement does not suffice here. All FMS question and answer indices start at zero; frame indices below start at one.

4.1.1 Questions, calibration and exact bindings

Alice has $\mathsf{R}_0, \mathsf{R}_1, \mathsf{R}_2, \mathsf{T}_0, \mathsf{T}_1$ and Bob has $\mathsf{C}_0, \mathsf{C}_1, \mathsf{C}_2, \mathsf{B}_0, \mathsf{B}_1$. A row question has a flag answer f or a triple of product $+1$. A column question has f or a triple of product $+1, +1, -1$ according to its column. The magic square is

$$\begin{pmatrix} XI & IX & XX \\ IZ & ZI & ZZ \\ XZ & ZX & YY \end{pmatrix}. \tag{6}$$

Here X, Y, Z denote the Pauli matrices, I the 2×2 identity, and adjacent letters their tensor product. The displayed matrices describe a canonical realization; the predicate only checks the corresponding signs. Select column-zero outcomes $c_0 = (+\,+\,+)$, $c_1 = (+\,-\,-)$ and row-zero outcomes $r_0 = (+\,+\,+)$, $r_1 = (+\,-\,-)$. Question T_x has answers $+, -$ and the three column-zero triples other than c_x; B_y analogously retains all row-zero triples other than r_y.

Table 3 and Figure 2 summarize the rules. On $\mathsf{R}_i, \mathsf{C}_j$, reject flag/nonflag mismatch and otherwise check the intersecting cell; two flags are accepted. On $\mathsf{T}_x, \mathsf{C}_0$, a sign is allowed only against f or c_x; a retained triple must equal the opposite triple. On $\mathsf{R}_0, \mathsf{B}_y$ use the analogous rule with r_y. On $\mathsf{T}_x, \mathsf{B}_y$, reject two signs $a, b \in \{+1, -1\}$ unless $ab = (-1)^{xy}$; all pairs with a nonsign answer are accepted. Other question pairs are accepted. For a conditional instance

add an inactive answer $\perp$, accept two inactive answers, and reject exactly one inactive answer. Bind each noninactive event to its synchronized active projection and all row/column flags to the synchronized flag.

Let P_{mag} be the nonflag magic support. Its four column-zero projections $m_{00}, m_{01}, m_{10}, m_{11}$ correspond respectively to $(+++), (+--), (-+-), (--+)$. On Alice, m_{00} is read from the retained $(+++)$ answer of T_1 and the other three from T_0. On Bob all four come from C_0. Each of these chosen occurrences is matched to its opposite synchronized binary question. The positive nonflag ZI event is read from Alice's R_1 component 1 and Bob's C_1 component 1. The positive IX event comes from Alice's R_0 component 1 and Bob's C_1 component 0. Put

$$\mathfrak{x} = I - 2P_{\mathrm{mag}} + 2E_{ZI,+}, \quad \mathfrak{z} = I - 2P_{\mathrm{mag}} + 2E_{IX,+}, \quad t_{ab} = \mathfrak{x}^a \mathfrak{z}^b m_{00}.$$

These are reflections extended by identity outside P_{mag}, and $t_{ab}^* t_{ab} = m_{00}$, $t_{ab} t_{ab}^* = m_{ab}$.

Questions	Answers	Tested rule
$\mathsf{R}_i, \mathsf{C}_j$	flag or parity triple	flags match; nonflags agree at (i, j)
$\mathsf{T}_x, \mathsf{C}_0$	sign or retained triple	signs pair only with f, c_x; retained triples match
$\mathsf{R}_0, \mathsf{B}_y$	triple or sign	symmetric calibration at r_y
$\mathsf{T}_x, \mathsf{B}_y$	signs or triples	two signs satisfy $ab = (-1)^{xy}$
Conditional copies	also inactive $\perp$	activity matches; flags and events bind to named projections

Table 3: Flagged magic-square questions, answers and rules. Unspecified pairs are accepted.

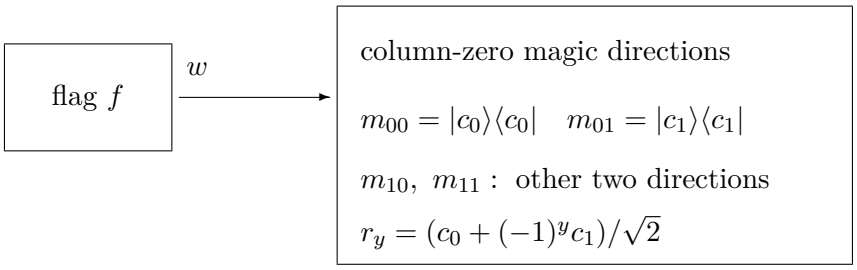


Figure 2: Selected column and row directions. The flag transports to c_0; m_{ab} denote projections.

4.1.2 Normal bridge theorem

Lemma 4.1. *A perfect normal conditional FMS strategy has on its active supports the decomposition*

$$\mathcal{H}_{\mathrm{active}} = \mathcal{H}_f \oplus (\mathbb{C}^4 \otimes \mathcal{H}_{\mathrm{mag}}), \qquad h = h_f \oplus (I_4 \otimes h_{\mathrm{mag}}).$$

The calibrated transports w_A, w_B from the flag to the selected c_0 corner satisfy $w^* w = f$, $w w^* = m_{00}$ and $hw = 2wh$. On a purification their matched components satisfy

$$\Omega_{c_0 c_0} = \sqrt{2}(w_A \otimes w_B \otimes I_{\mathrm{env}})\Omega_{ff}. \tag{7}$$

The flag answer and its projection are both denoted f. The other three magic directions obey the transported versions of this equation.

Proof. First apply Section 2 to the activity, flags and checked cells. They reduce the original state supports and are sharp there. On the compressed nonflag space each compressed fine effect lies below its assigned sharp cell marginals. Thus it is their joint spectral projection.

The projective-compression fact then makes the original fine effects reduce the support. Alice's row relations and Bob's column relations transfer through the dense coefficient range, in reversed product order; within a context the products commute.

For completeness, write $a = \mathsf{M}_{00}, b = \mathsf{M}_{01}, c = \mathsf{M}_{10}, d = \mathsf{M}_{11}$ for the four top-left cell reflections of the square, where M_{ij} denotes the checked cell observable. Context relations give

$$[a,b] = [c,d] = [a,c] = [b,d] = 0, \qquad (ac)(bd) = -(ab)(cd).$$

Cancel the outside a, d to obtain $cb = -bc$. Commutation of ac and bd then gives $da = -ad$. The pairs (a,d) and (b,c) are anticommuting Pauli pairs commuting across pairs. Each pair exchanges its two eigenspaces unitarily. Hence they generate $M_4 \otimes I_{\mathcal{H}_{\mathrm{mag}}}$. Marginal commutation gives the stated density form. The synchronized magic units identify its coefficient block as $I_4 \otimes T_{\mathrm{mag}}$, with T_{mag} positive, injective and dense-range. Write the flag coefficient block as T_f with the same properties.

Calibration makes Alice's sign support $S_x^A = f + \Pi_{c_x}$ and Bob's $S_y^B = f + \Pi_{r_y}$. The selected directions have overlap $\langle c_x, r_y\rangle = \sigma_{xy}/\sqrt{2}$, where $\sigma_{xy} = (-1)^{xy}$. Their compressed coefficient operator is

$$W_{xy} = \mathrm{diag}(T_f, \sigma_{xy} T_{\mathrm{mag}}/\sqrt{2}).$$

For compressed sign effects E_a, B_b, a forbidden pair gives $E_a^{1/2} W_{xy} (B_b^T)^{1/2} = 0$, and therefore $E_a W_{xy} B_b^T = 0$. Summing the two complete binary effects gives

$$E_a W_{xy} = W_{xy} B_{\sigma_{xy} a}^T = E_a W_{xy} B_{\sigma_{xy} a}^T.$$

Injectivity and dense range imply $B_b^2 = B_b$ and $E_a^2 = E_a$ on the sign supports. Their zero extensions give projective compressions; thus the original sign effects reduce the state supports. Their differences $\mathcal{X}_x, \mathcal{Y}_y$ obey

$$\mathcal{X}_x W_{xy} = \sigma_{xy} W_{xy} \mathcal{Y}_y^T. \tag{8}$$

Use the matrix units $t_{ab} t_{cd}^*$ to identify each selected magic direction with the same multiplicity space, and use $\sqrt{2}\Pi_{c_x}\Pi_{r_y}$ with its sign σ_{xy} for the selected row directions. Set $T_{\mathrm{cal}} = \mathrm{diag}(T_f, T_{\mathrm{mag}}/\sqrt{2})$ and $\Delta = \mathrm{diag}(I, -I)$. For $(x,y) = (0,0), (0,1), (1,0)$, Eq. (8) reads respectively $\mathcal{X}_0 T_{\mathrm{cal}} = T_{\mathrm{cal}} \mathcal{Y}_0^T$, $\mathcal{X}_0 T_{\mathrm{cal}} = T_{\mathrm{cal}} \mathcal{Y}_1^T$, and $\mathcal{X}_1 T_{\mathrm{cal}} = T_{\mathrm{cal}} \mathcal{Y}_0^T$. Dense range of T_{cal} gives $\mathcal{X}_0 = \mathcal{X}_1 = \mathcal{X}$; injectivity gives $\mathcal{Y}_0 = \mathcal{Y}_1 = \mathcal{Y}$. For $(1,1)$ it reads $\mathcal{X}\Delta T_{\mathrm{cal}} = -\Delta T_{\mathrm{cal}} \mathcal{Y}^T = -\Delta \mathcal{X} T_{\mathrm{cal}}$, so dense range gives $\mathcal{X}\Delta = -\Delta\mathcal{X}$. In blocks $\mathcal{X} = (\mathcal{X}_{ij})$, this says $\mathcal{X}_{11} = \mathcal{X}_{22} = 0$. Since $\mathcal{X} = \mathcal{X}^*$ and $\mathcal{X}^2 = I$, its remaining block $\beta = \mathcal{X}_{21}$ is unitary. Thus

$$\mathcal{X} = \begin{pmatrix} 0 & \beta^* \\ \beta & 0 \end{pmatrix}, \qquad \beta : \mathcal{H}_f \longrightarrow \mathcal{H}_{\mathrm{mag}} \text{ unitary}.$$

Self-adjointness and $\mathcal{X} T_{\mathrm{cal}} = T_{\mathrm{cal}} \mathcal{Y}^T$ imply $[\mathcal{X}, T_{\mathrm{cal}}^2] = 0$. Here $h_f = T_f^2$ and $h_{\mathrm{mag}} = T_{\mathrm{mag}}^2$. Its off-diagonal block gives $h_{\mathrm{mag}} \beta = 2\beta h_f$. The corresponding corner transport has the required ratio.

The actual Alice extraction is $w_A = \Pi_{c_0} \mathcal{X}_0 f$. Bob's r_0 projection is the retained $(+++)$ answer of B_1, equivalently the sign support of B_0 minus its flag. His extraction to the same c_0 direction is

$$w_B = \sqrt{2}\Pi_{c_0}\Pi_{r_0}\mathcal{Y}_0 f.$$

The coefficient $\sqrt{2}$ occurs in extraction, not in the predicate. To see the phase in (7) directly, put $\Omega' = (S_0^A \otimes S_0^B)\Omega$. The zero sign events give $\mathcal{X}_0\Omega' = \mathcal{Y}_0\Omega'$. Project onto $c_0 \otimes f$. The left side is $(w_A \otimes I)\Omega_{ff}$; the right side is

$$2^{-1/2}(I \otimes f\mathcal{Y}_0 |r_0\rangle\langle c_0|)\Omega_{c_0 c_0}.$$

Apply Bob's flag-to-r_0 block and then $|c_0\rangle\langle r_0| = \sqrt{2}\Pi_{c_0}\Pi_{r_0}$. This gives (7).

For a mixed state, first apply the pure-vector calculation on the cut $A \mid B\,\mathrm{env}$. The synchronized opposite effects reduce the Schmidt support, and the sign effects do so after their compressed sharpness is proved. Thus every compressed product used above is the restriction of an original Bob-only product tensored with I_{env}. The resulting Alice identities hold on her original marginal support. Repeat on $B \mid A\,\mathrm{env}$ for Bob's identities. The component calculation then gives Eq. (7) on the original purification, using $w_A \otimes w_B \otimes I_{\mathrm{env}}$. Tracing out the environment preserves the displayed component identity. □

4.1.3 A shared question fixes transported labels

Lemma 4.2 (Label refinement)**.** *Retain the original FMS questions. For a synchronized projection P^{lab}, add Alice's $\mathsf{T}_0[P^{\mathrm{lab}}]$ with answers (a, z). Ignore z against the unchanged $\mathsf{B}_0, \mathsf{B}_1, \mathsf{C}_0$ and use the original predicates; match z to Bob's binary P^{lab}. Add the symmetric $\mathsf{B}_0[P^{\mathrm{lab}}]$ against unchanged $\mathsf{T}_0, \mathsf{T}_1, \mathsf{R}_0$. Also retain activity matching. In every perfect normal strategy the coarse signs equal the original signs and commute with P^{lab} on their state supports. Conversely commuting labels give these refinements. Distinct labels require no mutual commutation.*

Proof. Write P_A for the state support and R_A for the sign support. For an arbitrary new fine effect $Q_{a,z}$ first form $q_{a,z} = P_A Q_{a,z} P_A$. Calibration identifies its compressed retained effects and the coarse sign sum R_A. Let $\mathcal{X}'$ be its compressed sign difference. The unchanged opposite sign gives $\mathcal{X}' W_{0y} = W_{0y}\mathcal{Y}_y^T$, with exactly the same dense-range W_{0y} as for $\mathcal{X}_0$. Hence $\mathcal{X}' = \mathcal{X}_0$. Its coarse effects are $(R_A \pm \mathcal{X}_0)/2$, hence projections. Projective compression now makes the original coarse effects reduce P_A.

Label matching gives $P_A(\sum_a Q_{a,1})P_A = P_A^{\mathrm{lab}}$. Positivity gives $0 \le q_{a,1} \le P_A^{\mathrm{lab}}$ and $0 \le q_{a,0} \le I - P_A^{\mathrm{lab}}$. Thus the compressed coarse effects commute with P_A^{lab}, and so does $\mathcal{X}_0$. The symmetric argument applies on Bob. Conversely use $E_a P^{\mathrm{lab}}, E_a(I - P^{\mathrm{lab}})$ when the original coarse effects commute with P^{lab}. A separate duplicated question for each label always returns the same original sign. This proves the final assertion. □

If a label is block diagonal with flag block P_f^{lab} and identical magic blocks P_m^{lab}, the commutator says $\beta P_f^{\mathrm{lab}} = P_m^{\mathrm{lab}}\beta$. In particular, for a graph projection of a block unitary it certifies transport of that same unitary by the calibrated w. It does not supply a fresh bridge with an independent gauge.

4.2 Rational projection equations that certify products

All projections in this section are synchronized. We give the finite rules used in the compiler. The projection-partition viewpoint is standard [6]; the formulas here provide the needed actual multiplication identities. See also the graph-projection encoding in [1, Section 4].

4.2.1 Graphs and triangles

Let F_i, F_j be orthogonal, $S = F_i + F_j$, and introduce projections $G, H, G^\perp, H^\perp$ with

$$G + G^\perp = S, \quad H + H^\perp = S, \quad 5H + S = 3F_i + 4G. \tag{9}$$

On S, put $\Delta_{ij} = F_i - F_j$, $\Xi_{ij} = 2G - S$. These are reflections, and $2H - S = (3\Delta_{ij} + 4\Xi_{ij})/5$. Squaring forces $\Delta_{ij}\Xi_{ij} + \Xi_{ij}\Delta_{ij} = 0$. Thus

$$\Xi_{ij} = \begin{pmatrix} 0 & u_{ij} \\ u_{ij}^* & 0 \end{pmatrix}, \qquad u_{ij}^* u_{ij} = F_j, \quad u_{ij} u_{ij}^* = F_i.$$

The graph blocks are certified, rather than assumed. Conversely any such unitary corner map gives the four projections in (9).

For a partition $F_1 + F_2 + F_3 = I$ impose the graph tests on all three pairs and add a projection $P_\triangle$ with

$$3P_\triangle = 2(G_{12} + G_{23} + G_{13}) - I. \tag{10}$$

Its block matrix has diagonal $I/3$ and off-diagonal $u_{ij}/3$. The $(1,3)$ block of $P_\triangle^2 = P_\triangle$ gives $2u_{13} + u_{12}u_{23} = 3u_{13}$, hence $u_{13} = u_{12}u_{23}$. Conversely this product identity makes the entire matrix a projection.

Start with a reference triangle for identity. Its scalar matrix algebra is called the *frame*. It gives scalar matrix units and $\mathcal{H} = \mathbb{C}^3 \otimes \mathcal{H}_{\text{base}}$. For every unitary label U, use graphs in all three positions. Triangles $(\mathsf{U}, 1, \mathsf{U})$ and $(1, \mathsf{U}, \mathsf{U})$ identify the placements. A triangle $(\mathsf{U}, \mathsf{V}, \mathsf{W})$ then certifies $\mathsf{W} = \mathsf{UV}$; a triangle with product identity certifies inverse labels. Build each word from its proper prefix and final token. Finally a real rational equation $\sum_w q_w \mathsf{U}_w = 0$ is precisely

$$2\sum_w q_w G_{12}(\mathsf{U}_w) - \Big(\sum_w q_w\Big)(F_1 + F_2) = 0. \tag{11}$$

Its two off-diagonal blocks are the equation and its adjoint. All products therefore refer to actual common operators. Synchronization makes the extracted unitaries commute with the base marginal density.

4.2.2 Binding prescribed parameters

A prescribed synchronized projection P^{par} is linked to the frame by projections $P^{\text{par}}F_i, NF_i$ with

$$P^{\text{par}}F_i + NF_i = F_i, \quad \sum_i P^{\text{par}}F_i = P^{\text{par}}, \quad \sum_i NF_i = I - P^{\text{par}}.$$

These symbols are names for projections before the equations identify their products; the names NF_i denote complements of $P^{\text{par}}F_i$ inside F_i. For every pair ij, introduce $B_{11}, B_{10}, B_{01}, B_{00}$ and impose

$$\begin{aligned} B_{11} + B_{10} &= P^{\text{par}}F_i + P^{\text{par}}F_j, & B_{00} + B_{01} &= NF_i + NF_j, \\ B_{11} + B_{01} &= G_{ij}(1), & G_{ij}(s_{\text{par}}) &= B_{11} + B_{00}. \end{aligned}$$

Projection sums give the indicated joint partitions. The symbol s_{par} is a unitary parameter label. These equations force P^{par} to commute with the frame and identify the extracted s_{par} with $2e - I$, where $P^{\text{par}} = I_3 \otimes e$. In particular these parameters are the specified FMS events, not independent replacements. The four-bridge assembly below uses twenty-five such parameters.

These rules are a deterministic finite compiler for the finite word list. Cancellation of adjacent inverse tokens is performed before taking proper prefixes. All coefficients stay rational. Conversely a centralizer unitary assignment with the prescribed parameter projections gives every graph, complement, H, triangle and joint projection simultaneously by these formulas.

4.3 Rational rows as Boolean winning conditions

We implement each rational row by a subspace orthogonality test. This is the general rational-row theorem of [5, Theorem B], restated here to specify the extension and its quantitative loss.

Lemma 4.3 (Boolean rational rows)**.** *Let a finite strategy specification name finitely many synchronized binary projections and finitely many rational linear identities among them. There is a finite extension by Boolean tests with the following properties.*

Soundness: *in every perfect normal spatial strategy for the extended tests, the identities hold as operator identities for the named support-restricted projections on each local state support.*

Completeness: *any normal spatial strategy with these synchronized projections, satisfying the identities on its supports, extends to satisfy all the added tests perfectly after adjoining a finite maximally entangled state. The old measurements act trivially on the new factor. The rest of the old strategy, including nonsynchronous questions, is retained; if its old tests were perfect, the whole extension is perfect.*

4.3.1 A certified scalar matrix algebra

Choose $d_{\max} = 2^{2r}$ large enough for every row construction below. Use r synchronized magic squares, binding each context's cell marginals to synchronized binary cells. The Pauli-pair calculation in Lemma 4.1 gives two anticommuting pairs per square. Four-answer joint questions for all pairs of axes from distinct squares force cross-square commutation. Hence their algebra is $M_{d_{\max}} \otimes I$ on the support. Synchronization makes the marginal $I_{d_{\max}}/d_{\max} \otimes h$.

Add joint questions between every named old projection and every scalar axis. Their sharp marginals force the old projections into the commutant: $E_i = I_{d_{\max}} \otimes e_i$. This step concerns only synchronized projections, never an FMS sign. Add global Z and global X measurements with all bit outcomes, binding each bit marginal to its axis. They are the actual joint spectral measurements.

For a row dimension $d = 2^n \le d_{\max}$ use the first n qubits; the rest are spectators. For each nonzero $j \in \mathbb{F}_2^n$, choose a basis $\ell_1, \dots, \ell_{n-1}$ of $j^\perp$. Add a joint context for $X^j, Z^{\ell_1}, \dots, Z^{\ell_{n-1}}$, binding its marginals to the corresponding global parity events. These commute because $\ell \cdot j = 0$. The zero-Z-syndrome outcomes with X^j sign $\pm$ are

$$R_j^{\pm} = \tfrac{1}{2}(|0\rangle \pm |j\rangle)(\langle 0| \pm \langle j|) \otimes I_{\text{spectator}}.$$

Write $|+_d\rangle = d^{-1/2} \sum_{j=0}^{d-1} |j\rangle$. The reference $R_d = |+_d\rangle\langle +_d| \otimes I_{\text{spectator}}$ is the event that the first n global X bits are positive. It includes all spectator outcomes.

4.3.2 Signed reflections and column graphs

For one row $S = \alpha_0 I + \sum_i \alpha_i e_i$, let L clear all coefficient denominators and set

$$s_i = 2e_i - I, \quad c_i = L\alpha_i, \quad c_0 = 2L\alpha_0 + \sum_i L\alpha_i.$$

Then $2LS = c_0 I + \sum_i c_i s_i$. Expand this into signed reflections with integer repetitions, add a canceling $I, -I$ pair with $V_0 = I$, and pad by further such pairs to a power of four d. The list has even length because its coefficient sum is even. Thus

$$\sum_{j=0}^{d-1} V_j = 2LS, \qquad V_0 = I, \qquad V_j \in \{\pm I,\ \pm(2e_i - I)\}. \tag{12}$$

For each $j > 0$ use a six-answer joint question of e_i and the three-way scalar measurement $R_j^+, R_j^-, I - R_j^+ - R_j^-$. For $V_j = \pm I$, bind its binary bit to a synchronized deterministic anchor: its 1-projection is I and its 0-projection is zero. Matching both marginals identifies all joint effects. Appropriate coarse events are the two-coordinate graph projections

$$G_j^{\pm} = \frac{1}{2}\begin{pmatrix} I & \pm V_j \\ \pm V_j & I \end{pmatrix}$$

on coordinates $0, j$, zero elsewhere. A negative reflection sign exchanges the two coarse events.

Add a synchronized d-answer question with projections Π_t, indexed by $t \in \mathbb{F}_2^n$. Put $\chi_t(j) = (-1)^{t \cdot j}$. Forbid Π_t together with G_j^- when $\chi_t(j) = 1$, or G_j^+ when $\chi_t(j) = -1$. The forbidden-event support rule gives actual orthogonality. A vector in $\operatorname{ran} \Pi_t$ therefore has coordinates $v_j = \chi_t(j) V_j v_0$.

Define

$$\mathcal{I}_t \xi = d^{-1/2} \big(\chi_t(j) V_j \xi\big)_{j=0}^{d-1}, \qquad \mathcal{P}_t = \mathcal{I}_t \mathcal{I}_t^*.$$

Unitarity and character orthogonality give $\mathcal{I}_t^* \mathcal{I}_u = \delta_{tu} I$ and $\sum_t \mathcal{P}_t = I$. Neither equation requires different V_j to commute. The kernel conditions give $\Pi_t \le \mathcal{P}_t$. Since both families sum to identity, $\Pi_t = \mathcal{P}_t$ for every t.

Finally forbid Π_0 jointly with the opposite R_d. Then $\Pi_0 R_d = 0$, equivalently $R_d \mathcal{I}_0 = 0$. The column overlap is

$$\mathcal{I}_0^* R_d \mathcal{I}_0 = d^{-2} \Big(\sum_j V_j\Big)^* \Big(\sum_j V_j\Big).$$

It vanishes exactly when $S = 0$. This proves soundness on arbitrary normal supports.

For completeness append $\Phi_{d_{\max}}$, use the scalar Paulis and the actual joint products, and define the column partition by the $\mathcal{I}_t$. Synchronized original projections commute with the marginal and transfer across a purification. Products transfer in reversed order. Define Bob's column effects by those transferred products, equivalently by transpose in canonical Schmidt coordinates. This order matters for noncommuting e_i. All column and synchrony tests hold. The row-zero test holds precisely because $S = 0$. These constructions act trivially on every old measurement's new factor and prove completeness, including for mixed original states.

Write h_{old} for the trace-one marginal before the row ancilla is appended (including the qutrit frame when present). For the honest column construction, without assuming $S = 0$, the sole row-zero event has probability

$$p(\Pi_0, R_d) = \frac{4L^2}{d^3} \operatorname{Tr}(h_{\mathrm{old}} S^2). \tag{13}$$

Indeed the scalar density is I_d/d, the column isometry is normalized by $d^{-1/2}$, and the reference vector by $d^{-1/2}$. The overlap computed above supplies $(2LS)^*(2LS)$. Spectator factors have normalized trace one. All column-graph and matching tests remain exact even when the row is nonzero. This formula is [5, Eq. (32)] in the general row setting.

4.4 The admissible presentations

Definition (admissible presentation; normal realization). An admissible presentation is a finite rational list with four contraction symbols A_α ($\alpha \in \{0,1\}^2$), a centralizer projection label $P_\circ$ for the root, and finitely many centralizer labels, each typed as a unitary, projection, positive contraction ($0 \le T \le I$), or general contraction ($\|T\| \le 1$). Its application equations (the identities enforced by the final row tests) are rational $*$-polynomials in the centralizer labels, the creation and range equations

$$A_\alpha^* A_\beta = \delta_{\alpha\beta} I, \qquad \sum_\alpha A_\alpha A_\alpha^* = I - P_\circ,$$

and endpoint equations $L A_\alpha = A_\beta B$ or $L A_\alpha = 0$. Finite sums $\sum_\nu B_\nu A_{\alpha_\nu} C_\nu = 0$ are allowed. Here L, B, B_ν, C_ν are centralizer polynomial expressions. Apart from the displayed creation and range equations, no monomial contains more than one creation symbol. Every label used on a left endpoint is included in the transported unitary list Λ defined in the proof of Lemma 4.4.

A normal realization consists of a nonzero Hilbert space $\mathcal{H}_0$, a faithful positive trace-class operator h_0 with $\operatorname{Tr} h_0 = 1$, and bounded operators of the declared types on $\mathcal{H}_0$. Every centralizer label, including $P_\circ$, commutes with h_0. All application equations hold as operator identities on $\mathcal{H}_0$, and

$$h_0 A_\alpha = A_\alpha h_0/16.$$

Thus all identities are imposed on the state support, where the density is faithful. No positive lower spectral bound is required for a positive contraction. Finite entropy of a realization means $S(h_0) < \infty$.

Lemma 4.4 (Four-bridge compiler)**.** *Every admissible presentation effectively yields a finite rational Boolean game. Every perfect normal spatial strategy yields a normal realization in the sense just defined, on the nonzero source corner with faithful trace-one density $h_0 = 121 e_0 h e_0$. Every such normal realization extends to a perfect normal spatial strategy after adjoining the scaffold and finite ancillary factors. Finite marginal entropy is preserved by the completeness construction.*

Proof. We give the assembly and both directions.

Corners and parameters. Use seventeen orthogonal synchronized corners

$$e_0; \quad m_{g,00}, m_{g,01}, m_{g,10}, m_{g,11} \quad (1 \le g \le 4).$$

Put $e_g = m_{g,00}$ and $M_g = \sum_{ab} m_{g,ab}$, and impose the cover $e_0 + \sum_g M_g = I$. Conditional bridge g has flag e_{g-1}, magic support M_g and activity $e_{g-1} + M_g$. Use all event bindings described above in both orientations, including the flag/four-magic/inactive partition. Bind to the qutrit frame the seventeen corners and the eight positive ZI, IX events. These are the twenty-five frame-bound parameters. Introduce eight axis labels $\mathfrak{x}_g, \mathfrak{z}_g$ with the explicit equations

$$\mathfrak{x}_g = I - 2M_g + 2E_{g,ZI,+}, \qquad \mathfrak{z}_g = I - 2M_g + 2E_{g,IX,+}.$$

The calibrated partial isometries satisfy $h w_g = 2 w_g h$. Thus

$$v = w_4 w_3 w_2 w_1, \quad v^* v = e_0, \quad v v^* = e_4, \quad hv = 16 vh. \tag{14}$$

Trace cyclicity gives four magic masses twice the preceding flag mass in each bridge. The cover fixes

$$\operatorname{Tr}(h e_0) = 1/121, \qquad 121 = 1 + 4(2 + 4 + 8 + 16). \tag{15}$$

In particular the source corner is nonzero.

$$e_0 \xrightarrow{w_1} \boxed{M_1 \supset e_1} \xrightarrow{w_2} \boxed{M_2 \supset e_2} \xrightarrow{w_3} \boxed{M_3 \supset e_3} \xrightarrow{w_4} \boxed{M_4 \supset e_4}$$

corner	e_0	e_1	e_2	e_3	e_4
$121 \operatorname{Tr}(h e_g)$	1	2	4	8	16
$121 \operatorname{Tr}(h M_g)$	–	8	16	32	64

Figure 3: Four calibrated bridges. Each M_g has four equal-mass magic corners; its selected e_g is the next flag.

Figure 3 displays the four-bridge chain.

Every transported label. Encode a unitary directly, a projection P by $2P - I$, a positive contraction T_+ by a unitary u_+ with $T_+ = (2I + u_+ + u_+^*)/4$, and a general contraction by an average of two unitaries. Let Λ be the resulting finite unitary label list that must be transported. For every $k \in \Lambda$, $r = 0, \dots, 4$, introduce a unitary $u_{r,k}$ and impose

$$(u_{r,k} - I)(I - e_r) = 0. \tag{16}$$

The unitary reduces e_r and is identity on its complement. For a typed label L, its corner copy L_r is explicitly decoded as

type of L	L_r
unitary	$e_r u_{r,L}$
projection	$e_r(I + u_{r,L})/2, \quad u_{r,L}^2 = I$
positive contraction	$e_r(2I + u_{r,L} + u_{r,L}^*)/4$
general contraction	$e_r(u_{r,L,+} + u_{r,L,-})/2$

Thus $(P_\circ)_4$ denotes the decoded root projection at e_4. The two encodings of a general contraction are transported separately when that label occurs at a left endpoint.

For each $g = 1, \dots, 4$ and every $k \in \Lambda$ introduce a unitary $v_{g,k}$ with

$$v_{g,k} = I - (e_{g-1} + M_g) + e_{g-1}u_{g-1,k} + \sum_{a,b} \mathfrak{x}_g^a \mathfrak{z}_g^b e_g u_{g,k} \mathfrak{z}_g^{-b} \mathfrak{x}_g^{-a}. \tag{17}$$

For this bridge, apply Lemma 4.2 to the six projections

$$F_1, F_2, F_3, G_{12}(1), G_{13}(1), G_{23}(1)$$

and the $|\Lambda|$ graph projections $G_{12}(v_{g,k})$. Thus there are $6 + |\Lambda|$ refinements of each retained sign question per bridge, on each party. Each label has its own duplicated question; the unchanged opposite questions are always retained.

The six frame commutators make $w_g = I_3 \otimes \widetilde{w}_g$. The commutator of the original sign with $G_{12}(v_{g,k})$, in its (c_0, f) block, is precisely

$$w_g u_{g-1,k} = u_{g,k} w_g. \tag{18}$$

For clarity the sign has off-diagonal block w_g between flag and selected magic direction, while the block unitary has the respective diagonal blocks $u_{g-1,k}, u_{g,k}$. Their commutator equates these two products. Consequently $L_4 v = v L_0$ for every decoded transported label, and for every expression built from them. Here L_r means its decoded copy at e_r, whose corner identity is e_r. The same actual w_g works for every label.

In the application the transported typed list is exactly

$$P_\circ,\ U_x,\ U_y,\ P_x,\ P_y,\ W,\ C_{\mathrm{ctr}},\ Z_{\mathrm{clk}},\ J_{\mathrm{W}},\ K_{\mathrm{W}}. \tag{19}$$

Table 4 records the label types. Here $P_\circ, P_x, P_y$ are projections, C_{ctr} is positive, and the other six are unitary. Thus one unitary encoding per typed label suffices, so $|\Lambda| = 10$. There are fifty support equations, forty block equations and sixteen refinements per bridge per party. Source-only labels $P_{\mathrm{ret}}, T_{\mathrm{inv}}, P_c$ do not occur at left endpoints of an A_α relation and need no transport.

Extraction and exact translation. Introduce four centralizer contractions $a_\alpha = (a_{\alpha,+} + a_{\alpha,-})/2$ and impose $a_\alpha = e_4 a_\alpha e_0$. On the root corner extract $A_\alpha = v^* a_\alpha$. The adjoint of (14) gives $hv^* = v^* h/16$, so $h_0 A_\alpha = A_\alpha h_0/16$, where $h_0 = 121 e_0 h e_0$. Supply the following polynomials to the graph compiler:

Label	Type	Transported?	Use
$P_\circ$	projection	yes	Fock root
U_x, U_y	unitary	yes	coordinate increments
P_x, P_y	projection	yes	boundaries
W	unitary	yes	clock
C_{ctr}	positive contraction	yes	counter
$Z_{\mathrm{clk}}, J_{\mathrm{W}}, K_{\mathrm{W}}$	unitary	yes	Weyl clock
P_{ret}	projection	no	pulse
T_{inv}	general contraction	no	reader inverse
P_c	projection	no	tableau
a_α	general contraction	scaffold only	lifted creation

Table 4: Typed application labels. Transported labels are copied through all bridges for use at a left endpoint.

Root identity	Centralizer identity
$A_\alpha^* A_\beta = \delta_{\alpha\beta} I$	$a_\alpha^* a_\beta = \delta_{\alpha\beta} e_0$
$\sum_\alpha A_\alpha A_\alpha^* = I - P_\circ$	$\sum_\alpha a_\alpha a_\alpha^* = e_4 - (P_\circ)_4$
$L A_\alpha = A_\beta B$	$L_4 a_\alpha = a_\beta B_0$
$L A_\alpha = 0$	$L_4 a_\alpha = 0$
Root-only polynomial	Same polynomial in source copies

Both implications hold. The range condition and $vv^* = e_4$ identify $A_\alpha^* A_\beta$ with $a_\alpha^* a_\beta$. Conjugating the range-sum equation by v uses the transported $P_\circ$. Transport of decoded adjoints also gives $v^* L_4 = L_0 v^*$. For an endpoint relation multiply by v and use $L_4 v = v L_0$; conversely use injectivity of v^* on e_4. No polynomial submitted to graph compilation contains v.

Expand the finite words, bind their parameters, apply graphs and triangles, and then Lemma 4.3. Retain all FMS and refinement questions. Pad the two question sets if necessary and sample uniformly from all pairs, accepting every unspecified pair. The result has rational question distribution and a Boolean predicate. The support arguments prove soundness on a nonzero corner, using Alice alone; mirror equations are unnecessary.

One simultaneous realization. Given a root realization with density h_0, take

$$h_{\mathrm{sc}} = \frac{1}{121} \operatorname{diag}(h_0;\, 2h_0^{\oplus 4};\, 4h_0^{\oplus 4};\, 8h_0^{\oplus 4};\, 16h_0^{\oplus 4}). \tag{20}$$

Use identity corner identifications in the bridges and $a_\alpha = |e_4\rangle\langle e_0| \otimes A_\alpha$. Its commutation with h_{sc} is exactly the modular relation. Every positive eigenspace of a faithful trace-class density is finite-dimensional. In each such block the polar partial isometry of a contraction T extends to a unitary V_{pol}, and

$$T = (U_+ + U_-)/2, \qquad U_\pm = V_{\mathrm{pol}}(|T| \pm i\sqrt{I - |T|^2}).$$

Direct sums give centralizer encodings for all four a_α at once. The root isometries A_α are never themselves written as averages of centralizer unitaries. The contraction completion takes place on the scaffold. Encode each source contraction inside the source density blocks, choose $2P - I$ for projections, and use

$$u_+ = 2T_+ - I + 2i\sqrt{T_+(I - T_+)}$$

for a positive contraction T_+. Copy these same choices through the four corner identifications and extend by identity outside the corners. The block equations now hold simultaneously.

Append the scalar qutrit frame, use the graph formulas, and take actual joint products for all refinements. Their block unitaries commute with the coarse FMS effects. Separate labels need not commute. Finally append the finite maximally entangled row factor. All tests hold with one state and one assignment. The scaffold and ancillary dimensions are finite, so they add only finite entropy. □

4.5 A dimension-independent completeness estimate

The exact construction gives a useful quantitative statement. Keep the label types, centralizer conditions, contractivity and modular covariance exact, but allow residuals r_j in the application equations of Section 4.4. Write $\|T\|_{2,h}^2 = \mathrm{Tr}(hT^*T)$. The constructed strategy satisfies

$$1 - \omega(G, \mathcal{S}) \le C_{\mathrm{comp}} \sum_j \|r_j\|_{2,h}^2, \tag{21}$$

where C_{comp} depends only on the finite presentation and the compiler.

Here is the proof, including the density factors. The same scaffold, exact unitary encodings and actual products make all structural tests perfect. Only the final application row-zero tests may lose. Let p_j be the lifted polynomial residual. For a source or root-polynomial row it is the source copy of r_j; for an endpoint row it is vr_j; for a range row it is vr_jv^*. Consequently

$$\|p_j\|_{2,h_{\mathrm{sc}}}^2 = \frac{\beta_j}{121}\|r_j\|_{2,h}^2, \qquad \beta_j = \begin{cases} 16, & \text{range row at } e_4, \\ 1, & \text{otherwise.} \end{cases}$$

For the endpoint case use $v^*v = e_0$. For the range case use $v^*h_{\mathrm{sc}}v = 16h/121$. These are equalities with no density-gap estimate.

The qutrit graph row S_j has off-diagonal blocks p_j, p_j^*. Since p_j commutes with h_{sc},

$$\mathrm{Tr}(h_{\mathrm{old}}S_j^2) = \frac{1}{3}\,\mathrm{Tr}[h_{\mathrm{sc}}(p_j^*p_j + p_jp_j^*)] = \frac{2\beta_j}{363}\|r_j\|_{2,h}^2.$$

Apply (13). For uniform Q-by-Q sampling one may take

$$C_{\mathrm{comp}} = \max_j \frac{8L_j^2\beta_j}{363d_j^3Q^2}. \tag{22}$$

Summing event probabilities gives an upper bound even if rejecting rules overlap. The dimension, clock order, density spectrum and spectator dimensions do not enter this constant. No continuity of polar completions or approximate dense-range cancellation is used: the structural realization is built exactly. This is a completeness estimate for constructed strategies under the stated exact structural hypotheses. It is not a robust extraction claim for arbitrary nearly perfect strategies.

5 The rooted tree and computation squares

5.1 Normality removes every rootless sector

Require four isometries and a root projection with

$$A_{ab}^*A_{cd} = \delta_{ac}\delta_{bd}I, \qquad \sum_{ab} A_{ab}A_{ab}^* = I - P_\circ, \qquad h_0A_{ab} = A_{ab}h_0/16. \tag{23}$$

Here h_0 is a faithful density on the extracted root-corner Hilbert space. For $w = w_1 \cdots w_m$ set $A_w = A_{w_1} \cdots A_{w_m}$ (the empty product is I) and put $Q_m = \sum_{|w|=m} A_wA_w^*$. Orthogonality and

trace cyclicity give $\mathrm{Tr}(h_0 Q_m) = 4^{-m}$. The projections Q_m decrease. Their strong limit has zero h_0-weight, hence is zero by faithfulness. Also $Q_m - Q_{m+1} = \sum_{|w|=m} A_w P_\circ A_w^*$. Therefore

$$\mathcal{H} = \bigoplus_w A_w \mathcal{K} \simeq \ell^2(\mathcal{W}) \otimes \mathcal{K}, \qquad \mathcal{K} = P_\circ \mathcal{H} \neq 0. \tag{24}$$

Here $\mathcal{W}$ denotes finite words in the four-letter alphabet $\{0,1\}^2$. The modular equation makes h_0 commute with every $A_w P_\circ A_w^*$. Since $\mathrm{Tr}(h_0 P_\circ) = 3/4$, transporting its root block gives

$$h_0 = h_F \otimes \sigma, \qquad h_F = \frac{3}{4} \sum_w 16^{-|w|} |w\rangle\langle w|, \qquad \mathrm{Tr}\,\sigma = 1. \tag{25}$$

In particular

$$S(h_F) = \log_2(4/3) + 4/3. \tag{26}$$

Indeed the depth-m mass is $(3/4)4^{-m}$ and its mean depth is $1/3$; the log probability at each word is $\log_2(4/3) + 4m$. This proves the formula and excludes a rootless normal summand.

5.2 Each depth is a finite square

For $w = (a_1, b_1) \cdots (a_m, b_m)$ define

$$x(w) = \sum_{j=1}^m a_j 2^{j-1}, \qquad y(w) = \sum_{j=1}^m b_j 2^{j-1}.$$

The first letter is the least significant bit. Prefixing A_{ab} maps (x, y) to $(2x + a, 2y + b)$. Depth m is a $2^m \times 2^m$ square. Require centralizer unitaries U_x, U_y and projections P_x, P_y satisfying, for $a, b \in \{0, 1\}$,

$$U_x A_{0b} = A_{1b}, \qquad U_x A_{1b} = A_{0b} U_x, \qquad U_x P_\circ = P_\circ, \tag{27}$$
$$U_y A_{a0} = A_{a1}, \qquad U_y A_{a1} = A_{a0} U_y, \qquad U_y P_\circ = P_\circ, \tag{28}$$
$$P_x A_{0b} = A_{0b} P_x, \qquad P_x A_{1b} = 0, \qquad P_x P_\circ = P_\circ, \tag{29}$$
$$P_y A_{a0} = A_{a0} P_y, \qquad P_y A_{a1} = 0, \qquad P_y P_\circ = P_\circ. \tag{30}$$

Induction on word length identifies U_x and U_y as cyclic increments of x and y respectively, acting identically on $\mathcal{K}$. The same induction makes P_x, P_y the indicators of $x = 0, y = 0$. In particular U_x, U_y commute, and

$$P_{x,\mathrm{last}} = U_x^* P_x U_x, \qquad P_{y,\mathrm{last}} = U_y^* P_y U_y$$

mark the final column and row. These are conclusions on the full root multiplicity, not just selected vectors.

5.3 A finite deterministic tableau

Choose a deterministic radius-one cellular automaton on the half-line. Its finite alphabet Γ includes start s, blank b and halt H, and its update is $\Phi(\ell, c, r, \varepsilon)$ with left boundary symbol $\#$ and external pulse $\varepsilon \in \{0, 1\}$. The initial row is $s, b, b, \ldots$. We use these conventions:

- $\Phi(b, b, b, \varepsilon) = b$ for both pulse values.
- The initial start state ignores the pulse at time zero. Later waiting-front states respond to positive-time pulses.
- Cells passed by the waiting front keep their unary marks. Frozen cells and all states of the later input-checking computation ignore subsequent pulses.

- A transition using the pulse at time y produces row $y+1$. Thus the pulse at time q records the front at coordinate q in that next row.

Here Φ names the update function; f in the following table is a front-state symbol. Here is an explicit front layer, with alphabet $\{s, b, f, z, H\}$. The right neighbour is unused in these six rules:

condition	next centre
$c = s$	b
$c = f$	z if $\varepsilon = 1$, b if $\varepsilon = 0$
$c = b,\ \ell = s$	f
$c = b,\ \ell = f$	f if $\varepsilon = 0$, b if $\varepsilon = 1$
$c = z$	z
$c = H$	H

All remaining blank centres stay blank. At time $y \geq 1$, until a return, the unique front is at $x = y$. The first positive pulse at q writes z at $(q, q+1)$ and suppresses the front's successor. The frozen mark's distance from the left boundary is the unary input q; blank cells between the boundary and the mark represent its unary positions. Pulse-dependence is confined to neighbourhoods containing the front. In particular a blank right of f may read the pulse.

Add a persistent unary bit $u \in \{0, 1\}$, initially zero, with $u' = 1$ when the old front centre is s or f, and $u' = u$ otherwise. Thus passed cells retain their unary marks on a separate track, even though their front-layer component is blank. Take the product alphabet $\{s, b, f, z, H\} \times \{0, 1\} \times \mathcal{T}$, with a finite TM alphabet $\mathcal{T} = \mathcal{A} \times (Q \cup \{\text{no head}\})$ for the fixed one-tape half-line program in Section 7. All TM cells start blank without a head. The transition $f \to z$ initializes its head in the start state at that cell. Thereafter a centre carrying (a, p) writes the symbol prescribed by $\delta(p, a, u)$ and retains a head only for a stationary move. A head moving right (left) is received by the right (left) neighbour, whose old tape symbol is retained. A centre receiving no head retains its symbol without a head. Invalid multiple incoming heads may be sent to a fixed error state; they never occur from this initialization. The left boundary # is read as a tape end marker and no transition moves left of it. The frozen front track is read-only input for this TM. A halting head writes a designated halt symbol H in the combined alphabet and keeps it. These prescriptions give a finite radius-one lookup table effectively from the finite TM transition table; all subsequent TM transitions ignore the pulse. Denote this finite product alphabet, with the halt convention, by Γ. In the combined alphabet, s and b abbreviate the start and fully blank product states. Collapse any product state with a halting head to the single absorbing symbol H.

Blank quiescence holds before a front or head arrives. The s rule ignores the time-zero pulse and seeds f at time one via its right neighbour. The mark is permanent and later pulses cannot restart a head. Finally, all disturbances start at the origin, or at the marked front within its radius-one cone, so no nonblank cell can occur at $x > y$. These checks give all four conventions and the blank right ghost used below.

Let $(P_c)_{c \in \Gamma}$ be a centralizer PVM and P_{ret} a centralizer projection, with $P_{\text{ret},0} = I - P_{\text{ret}}$, $P_{\text{ret},1} = P_{\text{ret}}$. Require

$$P_s P_x P_y = P_x P_y, \qquad P_b (I - P_x) P_y = (I - P_x) P_y, \qquad P_H = 0. \tag{31}$$

Define centralizer polynomials

$$L_{\#} = P_x, \qquad L_c = U_x P_c U_x^* (I - P_x),$$
$$N_r = U_x^* P_r U_x (I - P_{x,\text{last}}) + \mathbf{1}_{\{r=b\}} P_{x,\text{last}}.$$

For every $a \in \Gamma$ impose

$$\left[U_y^* P_a U_y - \sum_{\Phi(\ell,c,r,\varepsilon)=a} L_\ell P_c N_r P_{\text{ret},\varepsilon} \right] (I - P_{y,\text{last}}) = 0. \tag{32}$$

The factors have the displayed order. They need not commute globally. The factor $I - P_{y,\mathrm{last}}$ removes a spurious transition from the final row to row zero. The ghost to the right of the last column is blank.

Lemma 5.1 (Definite computation)**.** *Suppose a root vector ξ belongs to a sector where the clock pulse on every coordinate has a specified scalar value. Then every P_c acts on $A_w\xi$ as the indicator of the genuine classical computation at $(x(w), y(w))$.*

Proof. Work at a fixed depth m and induct on time, for every spatial coordinate at once. Equation (31) gives row zero. On an induction vector every factor $L_\ell, P_c, N_r, P_{\mathrm{ret},\varepsilon}$ acts by its classical zero or one. Their ordered product therefore selects exactly one update. Equation (32) gives $U_y^* P_a U_y$ with the correct scalar on that vector. Multiplying by U_y gives the next-row assertion.

The half-line and the square agree throughout this induction. At time y no disturbance has travelled beyond coordinate y. In particular the ghost cell 2^m is blank for all times $y \le 2^m - 1$, by radius one and quiescence under either pulse. The left ghost is always #. This proves the genuine-prefix assertion. A halting event at finite coordinates and time occurs in every sufficiently large square and contradicts $P_H = 0$. □

The tableaux at distinct depths need not intertwine with the creation operators. Determinism forces each of them to the same finite prefix; soundness needs only one sufficiently large square. Every depth has positive normal weight by (25).

5.4 A worked front and mask example

Table 5 is rendered by `review/generate_toy.py` from `tableau(3,3)[3]`. Its d denotes a passed cell and m a frozen mark. These persistent annotations differ from the minimal front layer above, whose passed cells are blank and whose mark is z; both put the mark at $(3, 4)$. The pulse at time zero is ignored by s. The pulse at $y = 3$ produces the mark in row 4, and the later pulse at $y = 6$ has no effect.

y	pulse	0	1	2	3	4	5	6	7	update mask
0	1	s	b	b	b	b	b	b	b	1
1	0	d	f	b	b	b	b	b	b	1
2	0	d	d	f	b	b	b	b	b	1
3	1	d	d	d	f	b	b	b	b	1
4	0	d	d	d	m	b	b	b	b	1
5	0	d	d	d	m	b	b	b	b	1
6	1	d	d	d	m	b	b	b	b	1
7	0	d	d	d	m	b	b	b	b	0

Table 5: Depth $m = 3$, period $q = 3$: the decoder's 8×8 toy tableau. P_x marks column 0, $P_{x,\mathrm{last}}$ column 7, P_y row 0 and $P_{y,\mathrm{last}}$ row 7. The last-column neighbour is replaced by blank; the last-row mask suppresses the toroidal $7 \to 0$ update.

6 A finite-order clock and its period reader

6.1 A determinant supplies a finite witness

Add centralizer unitaries $J_{\mathrm{W}}, K_{\mathrm{W}}, Z_{\mathrm{clk}}$ with

$$\begin{gathered} J_{\mathrm{W}} K_{\mathrm{W}} = Z_{\mathrm{clk}} K_{\mathrm{W}} J_{\mathrm{W}}, \qquad [Z_{\mathrm{clk}}, J_{\mathrm{W}}] = [Z_{\mathrm{clk}}, K_{\mathrm{W}}] = 0, \\ [J_{\mathrm{W}}, A_{ab}] = [K_{\mathrm{W}}, A_{ab}] = [Z_{\mathrm{clk}}, A_{ab}] = 0. \end{gathered} \tag{33}$$

A unitary commuting with an isometry also commutes with its adjoint: its inverse commutes with the isometry, and taking adjoints gives the assertion. Thus these labels reduce $\mathcal{K} = P_\circ \mathcal{H}$ and commute with σ. Their action on every word sector is their root action.

Every positive eigenspace of σ is finite-dimensional. It reduces $J_{\mathrm{W}}, K_{\mathrm{W}}, Z_{\mathrm{clk}}$ and decomposes into Z_{clk} eigenspaces. On a nonzero such eigenspace of dimension d, with phase ζ, the Weyl relation gives

$$\det(J_{\mathrm{W}} K_{\mathrm{W}}) = \zeta^d \det(K_{\mathrm{W}} J_{\mathrm{W}}), \qquad \zeta^d = 1. \tag{34}$$

Since the positive eigenspaces exhaust $\mathcal{K}$, Z_{clk} has pure point spectrum there and every phase is a root of unity. At least one sector is nonzero. Conversely a q-dimensional Weyl pair realizes every primitive qth root, with density I_q/q.

6.2 The clock and a counter without a positive lower bound

Take a centralizer unitary W and a positive centralizer contraction C_{ctr} with

$$W P_\circ = P_\circ, \qquad W A_{ab} = A_{ab} Z_{\mathrm{clk}}^b W^2, \tag{35}$$

$$C_{\mathrm{ctr}} P_\circ = P_\circ, \qquad C_{\mathrm{ctr}} A_{ab} = 2^{-b} A_{ab} C_{\mathrm{ctr}}^2. \tag{36}$$

Also require Z_{clk} to commute with $U_x, U_y, P_x, P_y, W, C_{\mathrm{ctr}}, P_{\mathrm{ret}}, T_{\mathrm{inv}}$ and all P_c. Induction on word length, using the coordinate recurrence, gives

$$W|_{(x,y)} = Z_{\mathrm{clk}}^y, \qquad C_{\mathrm{ctr}}|_{(x,y)} = 2^{-y} I. \tag{37}$$

In particular C_{ctr} is injective even though its spectrum accumulates at zero.

For a centralizer projection P_{ret} and contraction T_{inv}, impose

$$(W - I) P_{\mathrm{ret}} = 0, \qquad T_{\mathrm{inv}}(W - I) = \tfrac{1}{2} C_{\mathrm{ctr}}(I - P_{\mathrm{ret}}). \tag{38}$$

The first equation puts $\operatorname{ran} P_{\mathrm{ret}}$ in $\ker(W - I)$. For ξ in that kernel the second gives $C_{\mathrm{ctr}}(I - P_{\mathrm{ret}})\xi = 0$. Injectivity of C_{ctr} gives $P_{\mathrm{ret}}\xi = \xi$. Thus

$$P_{\mathrm{ret}} = \operatorname{proj} \ker(W - I). \tag{39}$$

On a root phase sector of order q, the pulse is exactly $\mathbf{1}_{\{q|y\}}$ at every spatial coordinate. It cannot vanish spuriously or postpone the first return.

6.3 A bounded inverse on every prime clock

Lemma 6.1 (A good primitive phase)**.** *For every prime q there is a primitive qth root ζ such that*

$$|\zeta^t - 1| \geq 2^{-t-1} \quad (t \geq 1,\ q \nmid t).$$

Proof. For $1 \leq r < q$,

$$|e^{2\pi i r/q} - 1| \geq 4 \min(r, q - r)/q.$$

Fix $1 \leq t < q$. Multiplication by t permutes nonzero residues. A phase $e^{2\pi i k/q}$ violating the required inequality must satisfy $\min(tk \bmod q, q - (tk \bmod q)) < q 2^{-t-3}$. There are at most $q 2^{-t-2}$ such residues. Summing over t excludes fewer than $q/4$ choices, while $q/4 < q - 1$. One of the $q - 1$ nontrivial phases works. For $t \geq q$ reduce to a nonzero residue $r < q$; its lower bound 2^{-r-1} is at least 2^{-t-1}. □

For this phase set

$$T_{\mathrm{inv}}|_{(x,y)} = \begin{cases} 2^{-y-1}(\zeta^y - 1)^{-1} I, & q \nmid y, \\ 0, & q \mid y. \end{cases} \tag{40}$$

Lemma 6.1 gives $\|T_{\mathrm{inv}}\| \leq 1$. Equations (37) and (39) then verify (38). The game coefficients remain rational: the phase belongs to the strategy, not the predicate.

7 The reduction

Fix any decidable predicate $R(e,n,t)$ and the set $S = \{e : \exists n\, \forall t\ R(e,n,t)\}$. We effectively choose the cellular automaton in Section 5 to run the following program, with e hardwired.

Read a positive-time return. At time y a waiting front is at position y. Ignore the time-zero pulse. At the first positive-time pulse, freeze the front and retain its coordinate as unary input q. With the transition convention above the mark is written into row $q+1$ at position q. Test q. If $q = 1$ or q is composite, halt. Otherwise compute the unique $n \geq 0$ with $q = p_{n+1}$, the $(n+1)$st prime. Then evaluate $R(e,n,0), R(e,n,1), \ldots$ and halt at the first false answer. All later computation states ignore the external pulse.

For a Σ_2^0-complete choice of S, there is no computable bound on a sufficient clock order as a function of the input game: such a bound would leave only finitely many possible n in this reduction, making $\exists n\, \forall t\ R(e,n,t)$ a Π_1^0 predicate. Indeed one tests, for every finite time cutoff, whether some one of those finitely many candidates survives all checks so far; if none survives forever, all have failed by a common finite cutoff. This contradicts Σ_2^0-hardness.

This finite program has an effective finite radius-one simulation. The initial-front rule, frozen unary track and ordinary tape/head tracks give the stated boundary conventions. We do not require a bound on its running time. Each completed false check causes the halt symbol after finitely many updates.

Include the Fock, coordinate, clock, reader and tableau equations above, all typing relations, and the stated Z_{clk} commutators. They form an admissible finite presentation. For example their transported polynomials include

$$W_4 a_{ab} = a_{ab} Z_{\mathrm{clk},0}^{b} W_0^2, \qquad C_{\mathrm{ctr},4} a_{ab} = 2^{-b} a_{ab} C_{\mathrm{ctr},0}^2, \qquad J_{\mathrm{W},4} a_{ab} = a_{ab} J_{\mathrm{W},0}.$$

The source-only return equations and tableau equations are centralizer polynomials. Lemma 4.4 therefore computes a finite rational Boolean game G_e.

Theorem 7.1. *The map $e \mapsto G_e$ satisfies*

$$G_e \in \mathrm{PERF}_{\mathrm{qs}} \iff G_e \in \mathrm{PERF}_{\mathrm{finS}} \iff \exists n\, \forall t\ R(e,n,t).$$

No G_e has a perfect finite-dimensional strategy.

Proof. A perfect normal strategy extracts a nonzero faithful root realization. The determinant argument supplies a nonzero root phase sector of finite order q. By (33) its word descendants persist at all depths. Equations (35)–(39) force the exact return sequence, with its first positive return at q. Lemma 5.1 needs only that the initial root vector lie in this Z_{clk}-eigenspace: the clock pulse is then scalar on all its word descendants. The imposed $[Z_{\mathrm{clk}}, P_c] = 0$ relations are redundant for this induction. An invalid q halts. For valid $q = p_{n+1}$ any false $R(e,n,t)$ also halts. A sufficiently large square would contain that halt, contrary to $P_H = 0$. Hence the extracted n witnesses membership in S.

Conversely take a witness n and set $q = p_{n+1}$. Use (25) with $\sigma = I_q/q$, a good primitive phase and its Weyl pair, the canonical coordinate permutations, (37), (39), (40), and the classical nonhalting tableau. Every required operator is bounded and commutes with the density when so specified. This is a simultaneous realization. On $\mathbb{C}^q$ one may take $J_{\mathrm{W}}|j\rangle = \zeta^j |j\rangle$, $K_{\mathrm{W}}|j\rangle = |j+1 \bmod q\rangle$, and $Z_{\mathrm{clk}} = \zeta I$. Its root entropy and largest Schmidt probability are

$$S(h) = \log_2 q + \log_2(4/3) + 4/3, \qquad p_{\max} = 3/(4q). \tag{41}$$

The compiler adds finite factors, giving a perfect finite-entropy strategy.

The tensor-product witness can be written without an implicit extension. Let $\mathcal{H}_F = \ell^2(\mathcal{W})$ and let $\mathcal{H}_{\mathrm{sc}} = \mathbb{C}^{17}$ index the scaffold corners. Take

$$\mathcal{H}_A = \mathcal{H}_B = \mathcal{H}_{\mathrm{sc}} \otimes \mathcal{H}_F \otimes \mathbb{C}^q \otimes \mathbb{C}^3 \otimes \mathbb{C}^{d_{\max}}.$$

If $\mu = (1,2,2,2,2,4,4,4,4,8,8,8,8,16,16,16,16)/121$, use the pure vector

$$|\psi\rangle = |\sqrt{\mu}\rangle\rangle \otimes |\sqrt{h_F}\rangle\rangle \otimes |\Phi_q\rangle \otimes |\Phi_3\rangle \otimes |\Phi_{d_{\max}}\rangle, \quad |\sqrt{\mu}\rangle\rangle = \sum_i \sqrt{\mu_i}|i,i\rangle.$$

The Fock factor is $\sum_w \sqrt{(3/4)16^{-|w|}}|w,w\rangle$. Alice uses the scaffold FMS effects, actual graph projections, refinements and column partitions of Lemma 4.4; Bob uses the prescribed opposite FMS effects and the transposed transferred centralizer products. The row factor is the finite maximally entangled ancilla of Lemma 4.3. These simultaneous effects were checked in the compiler proof. The marginal entropy is $H(\mu) + S(h_F) + \log_2 q + \log_2 3 + \log_2 d_{\max} < \infty$.

Finally finite-dimensional perfect play would extract four isometries with mutually orthogonal ranges on one nonzero finite-dimensional root space. Their range dimensions would sum to four times the dimension of that space, which is impossible. The argument applies to every input, whether or not it has a witness. □

Together with Section 3, this proves both completeness statements in Theorem A. The value-one promise is established next. For the promise version, the reduction maps every input into that promise; the unrestricted Σ_2^0 membership predicate restricts to it, giving promise-problem completeness. The construction can be made syntactically nonsynchronous and non-imitation on every input by adding a fresh binary question to each side, accepting every pair involving it. Either answer then has the same nonempty allowed-answer support against every opposite question, which violates the imitation separation condition. This addition does not change perfection or the value-one property.

Size of the finite game. Here is a coarse order-of-magnitude bound, independent of the witness q. Let $g = |\Gamma|$, let N_e count the named graph projections, rows and word prefixes after endpoint translation, and let b_e bound the bit length of their rational coefficients. The lookup table has at most $2(g+1)g^2$ entries, so unshared explicit polynomials give $N_e = \mathrm{poly}(g)$ for this fixed list of operator relations. Denominator clearing and signed repetitions give row dimension $d_{\max} \le 2^{O(N_e(b_e+\log(N_e+1)))}$. The scalar contexts, column questions and marginal bindings are polynomial in $N_e, d_{\max}$, so the padded question count Q_e and maximum answer count obey the same coarse exponential bound. An explicit predicate table is polynomial in those two counts and has the same exponential order. Thus the effective compilation depends on the hardwired program, never on its chosen prime or running time. For scale, the bundled minimal one-isometry example already has 748 rows and 27 991 questions per party; it is not a size claim for the full reduction.

8 Every constructed game has value one

The finite-dimensional approximations do not need a witness for e. Truncate the Fock tree to depths $0, \dots, m$ and choose a prime $q > 2^m$ with a good phase (Lemma 6.1). Use the normalized density

$$h^{[m]} = \frac{\sum_{|w|\le m} 16^{-|w|}|w\rangle\langle w|}{\sum_{j=0}^{m} 4^{-j}} \otimes \frac{I_q}{q}.$$

Every represented time is less than q. The decoder therefore remains in the waiting stage in every square and never halts. Let P_m denote the top-depth projection. Its mass is

$$t_m = \mathrm{Tr}(h^{[m]} P_m) = \frac{3}{4^{m+1} - 1}. \tag{42}$$

Every endpoint identity $LA_\alpha = A_\beta B$ remains exact. Below top depth it is the canonical identity; on top depth both sides vanish because the endpoint expression B preserves depth. The modular equations, the range-sum identity and all off-diagonal $A_\alpha^* A_\beta = 0$ also remain exact. The only failed original relations are the four diagonal isometry equations:

$$A_\alpha^* A_\alpha - I = -P_m \quad (\alpha = 00, 01, 10, 11). \tag{43}$$

All graph and triangle projections are defined from the actual products. Every structural graph, synchrony, binding, FMS, refinement and column test stays exact.

Here is the full row accounting. For each α let L_α be the denominator-clearing integer and d_α the column dimension of the row obtained by translating $a_\alpha^* a_\alpha - e_0 = 0$ into graph projections. These depend on the finite input presentation, not on m or q. In the scalar qutrit frame its defect S_α has off-diagonal blocks $-(P_m)_0$ and adjoint and zero diagonal blocks. The frame density is $I_3/3$ and the source scaffold mass is $1/121$, so

$$\mathrm{Tr}(h_{\mathrm{old}} S_\alpha^2) = \frac{2}{3} \frac{1}{121} t_m. \tag{44}$$

The same calculation applies separately to each of the four rows. Their row-zero events have distinct column questions and occur in no other failed rule. If the game has Q_e equally padded questions per party, (13) gives

$$1 - \omega(G_e, \mathcal{S}_m) \le C_e \frac{3}{4^{m+1} - 1}, \tag{45}$$

$$C_e = \frac{1}{Q_e^2} \frac{2}{363} \sum_{\alpha \in \{00,01,10,11\}} \frac{4 L_\alpha^2}{d_\alpha^3}. \tag{46}$$

If the four row constructions have common L, d, this is $C_e = 32L^2/(363 Q_e^2 d^3)$. We include one oriented row-zero event per row. The opposite column questions are retained for identification without a second row-zero test. This fixes the convention in the constant. Duplicating the four row-zero tests in the opposite orientation would multiply the constant by two.

Since the density factor I_q/q has trace one, no factor of q enters the loss. Thus $\omega_q(G_e) = 1$ for every e, proving the promise in Theorem A. This calculation uses four rows, rather than suppressing them inside an unspecified continuity estimate.

9 Exact Bell attainment

Let $L(p) = \sum_{xyab} c_{xyab} p(a, b \mid x, y)$ have rational coefficients and let $C_{\mathrm{abs}} = \sum |c_{xyab}|$. Its value on any strategy lies in $[-C_{\mathrm{abs}}, C_{\mathrm{abs}}]$. Finite Schmidt truncation shows that its supremum over spatial strategies equals its supremum over finite-dimensional strategies.

9.1 A specified global bound

Assume first that $\beta \in \mathbb{Q}$ is a global upper bound. For fixed measurements the Bell operator

$$B(E, F) = \sum c_{xyab} E_{a|x} \otimes F_{b|y}$$

satisfies $B \le \beta I$. A normal state attains β only if every positive-weight pure component does so. For a pure vector this is equivalent to $B\psi = \beta\psi$, because $\beta I - B$ is positive. Compress to its Schmidt supports and normalize the largest amplitude as before. Vectorization yields the homogeneous equation

$$\sum c_{xyab} E_{a|x} D_{\mathrm{Sch}} F_{b|y}^{\mathsf{T}} = \beta D_{\mathrm{Sch}}. \tag{47}$$

At stage n use the variables of (3) and require every squared entry of the residual in (47) to be at most $C_{\mathrm{abs}}^2 \kappa/(n+1)$. Compression of a true witness satisfies this bound: only the compact tails of the sandwiched terms contribute, with coefficient sum C_{abs}. Conversely the compact-sandwich lemma gives an exact nonzero limit solution. Its vectorization is an attaining strategy. Thus attainment is $\exists\kappa\,\forall n$ of a decidable predicate.

If validity is not promised, impose also $\omega_q(L) \le \beta$. This is Π_1^0: for each finite dimension, the universal assertion that all states and POVMs give value at most β is decidable by real quantifier elimination. Finite-dimensional density states suffice for this validity test. The conjunction remains Σ_2^0.

9.2 No bound supplied

The quantifier pattern is explicitly $\exists\kappa \ge 1\,\forall n \ge 1\ \mathcal{R}_{\mathrm{Bell}}(L,\kappa,n)$, with $\mathcal{R}_{\mathrm{Bell}}$ decidable as follows. At stage n introduce also a real variable $\beta \in [-C_{\mathrm{abs}}, C_{\mathrm{abs}}]$, use the same residual bound, and require $\beta \ge \omega_n(L)$, where ω_n is the maximum in local dimension n. This latter requirement is a universal semialgebraic formula in β, not a numerical oracle. Full real quantifier elimination decides the stage formula with its existential and universal real variables.

An attaining strategy supplies all stages. Conversely pass to a subsequence with $\beta_n \to \beta_*$, $D_{\mathrm{Sch},n} \to D_{\mathrm{Sch}}$ in norm and effects weakly convergent. For every fixed dimension d, all sufficiently late stages have $\beta_n \ge \omega_d(L)$. Hence $\beta_* \ge \sup_d \omega_d(L)$. The limit equation (47) gives a unit vector attaining β_*. Its expectation cannot exceed that supremum, so equality holds. This proves Σ_2^0 membership for attainment of the unspecified supremum.

9.3 A specified scalar value

Let $t \in \mathbb{Q}$ be supplied without a bound promise. A value strictly between the infimum and supremum has finite-dimensional values on either side, by equality of the finite and spatial extrema. A finite direct-sum mixture then attains t. Existence of a finite-dimensional strategy with value exactly t is Σ_1^0, by dimension search and quantifier elimination. Every remaining attainable value is an endpoint. Upper endpoints are handled above; lower endpoints are upper endpoints for $-L$. This finite union is Σ_2^0.

For hardness in all three cases use the winning-probability functional of G_e. Its coefficients are rational, its global bound and finite-dimensional supremum are both one, and attaining the specified value one is precisely perfect spatial play. Theorem 7.1 and Section 8 therefore prove Theorem B. In its promised-bound version every reduction output has valid bound 1, and the upper predicate restricts to the valid-bound domain, so the same map establishes promise-problem completeness.

10 Discussion and open problems

The hierarchy separates four ways to ask for certainty. Finite-dimensional perfection has a finite existential witness. Commuting perfection is controlled by compactness. Value one asks for arbitrarily accurate finite witnesses. Spatial perfection combines an unbounded finite choice with an infinite verification. Here the chosen object is the order of a clock phase.

The entropy restriction does not lower the complexity. Each YES instance G_e of our reduction has a finite-entropy witness. There is no computable entropy budget sufficient for every finite-entropy YES game: such a budget $B(G)$ would make $\mathrm{PERF}_{\mathrm{finS}}$ a Π^0_1 set by the uniform budget predicate in Section 3, contradicting its Σ^0_1-hardness from perfect completeness of $\mathsf{MIP}^* = \mathsf{RE}$. The analogous impossibility for a sufficient largest-Schmidt stratum is Theorem C. The fixed-κ compactness result does not give a uniform modulus over games. It also concerns zero constraints and extremal scalar eigen-equations, not a full prescribed correlation table.

10.1 Prescribed tables and escaping mass

For a concrete obstruction to the latter extension, put

$$D_{\mathrm{esc},n} = \mathrm{diag}(1, \underbrace{n^{-1/2}, \ldots, n^{-1/2}}_{n \text{ entries}}, 0, \ldots).$$

Then $\operatorname{Tr} D^2_{\mathrm{esc},n} = 2$ and $D_{\mathrm{esc},n} \to |0\rangle\langle 0|$ in operator norm. Alice measures $|0\rangle\langle 0|$ and its complement; Bob has one trivial outcome. Every normalized vectorization of $D_{\mathrm{esc},n}$ gives Alice probabilities $(1/2, 1/2)$. The normalized limit gives $(1, 0)$. The largest-amplitude normalization preserves zero equations while changing nonzero normalized probabilities.

Unproved remarks on prescribed-table bounds. The following upper-bound remarks are not proved in this paper and are not used in any theorem. For a prescribed table with rational or real-algebraic entries, direct coding of countable operators and a normal vector gives a Σ^1_1 upper bound. A supplied computable Schmidt-tail modulus makes the normalized state family effectively compact and gives a Π^0_1 upper bound. An existential choice from a uniformly total computable family of such moduli gives Σ^0_2. Finite entropy also gives Σ^0_2, by (5) with an existential integer budget and rational slack. If an arbitrary program for a computable modulus is existentially supplied, totality and the tail conditions give a Σ^0_3 upper bound. These proposed bounds concern the respective restricted realization classes; establishing their effective coding details is outside the present proofs. We do not prove that every algebraic spatial table has a computable-tail realization, or give a completeness classification for unrestricted table membership.

10.2 The Fu–Miller–Slofstra remark

After Theorem 1.1, Fu, Miller and Slofstra state that their fixed-scenario membership problem is in RE for $t = q$ or $t = qs$ [7]. We found no proof of the general qs upper assertion in that discussion. The following consequence needs a uniform interpretation over finite scenarios.

Suppose membership of real-algebraic tables in C_{qs} has a uniform RE procedure. If every spatially perfect game in our reduction had an algebraic perfect table, enumerate all algebraic tables of its scenario, test the losing zeros exactly, and dovetail the membership procedures. This would enumerate precisely the YES inputs of a Σ^0_2-complete family, a contradiction. Under that interpretation, either the RE remark fails for qs, or some spatially perfect games in this family have no perfect correlation table with all entries algebraic. The conclusion is existential over YES games; it does not say that every YES game has only nonalgebraic perfect tables.

The literal remark is per fixed scenario. The scenario of G_e grows with e, which is hardwired into the CA alphabet, so separate nonuniform RE procedures for each scenario give no contradiction. Tension arises only if there is one procedure uniform across scenarios, together with algebraic perfect tables for every YES game of the reduction. This is an open

discrepancy, not a refutation. Resolving the uniformity and algebraicity questions is a natural next problem.

10.3 A possible von Neumann type ladder

Does perfect play with normal states and type III ambient local algebras give exactly the games with perfect commuting play? Which games admit perfect play with type II ambient local algebras? Is there a useful intermediate decision classification? These are questions, not consequences proved here. The meaning of normality, the allowed ambient algebras, and whether both local algebras must have the designated type must be fixed before comparing such classes. We make no type II or type III classification claim.

10.4 Other questions

Can the game compiler be made substantially smaller? Can the period reader be realized by a simpler predicate? What are the optimal effective bounds when one supplies an entropy budget or a Schmidt-tail modulus? The finite-entropy and arbitrary-spatial decision sets have the same completeness degree, although the sets themselves differ: the entropy-game construction in [5] has spatial perfection and no finite-entropy perfect strategy. A classification of prescribed-table membership would require additional control of the normalization that escapes in the example above.

How this work was done

The authors selected the questions and proof standards, directed the research, and take responsibility for every claim. AI systems from more than one developer assisted with exploration, derivation, programming and drafting.

The mathematical evidence consists of the proofs in this paper and the explicitly cited external theorems. The ancillary programs provide reproducible finite checks. Before the rewrite, review by a different model family checked the underlying derivations of the upper bound, compiler, rooted computation, clock, reduction, value-one argument and Bell attainment. The final manuscript was audited separately; the present revision addresses that manuscript audit. The finite programs test truncated tree and clock relations, a pulse-reading front, good primitive phases, and the compiler ingredients described in Appendix A. Their arithmetic and scope are stated there; numerical checks are distinguished from exact checks.

These finite checks do not prove the infinite-dimensional compactness argument, the general reduction, or an unrestricted robustness theorem by sampling. The asymptotic conclusions rest on the written proofs. This preprint has not been peer reviewed.

A Reproducibility

The accompanying `verification/` directory contains the tree-and-clock checker, the primitive-phase checker, their execution logs, and the independent compiler replication artifacts available at packaging time. Its README states commands, dependencies, provenance and limitations. `SHA256SUMS` identifies every bundled file apart from itself. Public paths are relative. The three compiler entry points accept `-output-dir DIR`, defaulting to `compiler/generated/`; they refuse to overwrite the shipped compiler directory or toolkit. Windows memory telemetry is guarded on other platforms. Fresh revision logs and the check summary are in `logs/revision-checks/`; generated output directories and Python bytecode caches are excluded from the shipped-file manifest.

The numerical tree checker constructs the truncated Fock model, Weyl pair, coordinate increments, clock, counter and exact-return candidate. Its pulse-driven cellular automaton is a small front-and-mark model. It is not an implementation of every Turing-machine decoder used by the reduction. The checker tests the creation and range identities, modular scaling, endpoint recurrences, clock and reader equations, tableau propagation, centralizer commutators and top-depth defects. The four diagonal isometry equations fail exactly at the top depth, as predicted by (43); all other reported residuals are at floating-point precision. The analytical mass is (42).

The primitive-phase checker tests the existence of a good primitive phase for every prime below 4000 and tests the chord lower bound only for primes below 400. Its trigonometric evaluations use floating point. Lemma 6.1 is the proof for all primes. The independent exact tree suite checks ten models, depths 0 through 4 and clock orders $2, 3$, with 2000 exact checks in total. A separate minimal compiler instance has four bridges and one modular isometry, but no transported endpoint labels. It has 748 rational rows and 27 991 questions per party. Its independent structural audit checks 77 750 serialized bindings, 11 652 column-graph rules, all 748 row expansions, 333 scalar pair contexts and 2472 old forbidden events. It does not instantiate the full ten-label computation compiler. These exact replication calculations and their relative-path packaging are separately identified in the ancillary README. No proof-assistant certification is claimed.

The paper is built with `python build.py` from `paper/`. This runs `pdflatex`, BibTeX and the required repeat passes, and fails if the final log contains a warning, an undefined reference, or an overfull or underfull box. The plain-text abstract is `abstract.txt`.

References

[1] Ranyiliu Chen and Yuming Zhao. A non-robust quantum correlation self-test, 2026. arXiv:2609.25117v1.

[2] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58:012202, 2017. arXiv:1606.02278.

[3] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *Automata, Languages, and Programming (ICALP 2014)*, pages 320–331, 2014. arXiv:1209.2729.

[4] Andrea Coladangelo and Jalex Stark. An inherently infinite-dimensional quantum correlation. *Nature Communications*, 11:3335, 2020. Preprint arXiv:1804.05116.

[5] Seth Douglas and Nidhal Mghirbi. Perfect play in a finite nonlocal game requires infinite entanglement entropy. https://doi.org/10.5281/zenodo.23902656, 2026. Version 1.0.0, preprint.

[6] Tobias Fritz. Curious properties of free hypergraph C^*-algebras. *Journal of Operator Theory*, 83(2):423–445, 2020. arXiv:1808.09220.

[7] Honghao Fu, Carl A. Miller, and William Slofstra. The membership problem for constant-sized quantum correlations is undecidable. https://arxiv.org/abs/2101.11087v5, 2025. arXiv:2101.11087v5; remark after Theorem 1.1.

[8] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP* = RE. https://arxiv.org/abs/2001.04383v4, 2020. arXiv:2001.04383v4, Theorem 12.6.

[9] Se-Jin Kim, Vern I. Paulsen, and Christopher Schafhauser. A synchronous game for binary constraint systems. *Journal of Mathematical Physics*, 59:032201, 2018. arXiv:1707.01016.

[10] Junqiao Lin. $\mathsf{MIP}^{co} = \mathsf{coRE}$, 2025. Preprint arXiv:2510.07162.

[11] Martino Lupini, Laura Mančinska, Vern I. Paulsen, David E. Roberson, Giannicola Scarpa, Simone Severini, Ivan G. Todorov, and Andreas Winter. Perfect strategies for non-signalling games. *Mathematical Physics, Analysis and Geometry*, 23:7, 2020. arXiv:1804.06151.

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

[13] N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical Review Letters*, 65:3373–3376, 1990.

[14] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. Nonlocal games, compression theorems, and the arithmetical hierarchy. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1–11, 2022. Full version arXiv:2110.04651v2.

[15] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10:073013, 2008. arXiv:0803.4290.

[16] Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151:107–108, 1990.

[17] William Slofstra. The set of quantum correlations is not closed. *Forum of Mathematics, Pi*, 7:e1, 2019. arXiv:1703.08618.

[18] William Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 33:1–56, 2020. arXiv:1606.03140.