

Constant-precision undecidability for amplified determinants with capacity one

Lluis Eriksson

17 October 2026

Working paper

Abstract

We give a computable reduction from the halting problem to approximating a normalized determinant optimized over matrix coefficients of unrestricted auxiliary dimension. The coefficient matrices are real rational, orthonormal for normalized Hilbert–Schmidt inner product, and have covariance-optimized operator capacity exactly one. Nevertheless, approximating the determinant optimum to one fixed additive precision is undecidable. The proof combines the synchronous linear-game hardness theorem of Taller and Vidick with a homogeneous multiplication block, weighted rational normalization, and two quantitative rounding arguments. A positional-question normalization handles repeated queries without identifying distinct answer positions. The constructed pencils have a linear scalar dimension in the expanded relation list and an explicit constant gap. It follows that even at one fixed precision no total recursive uniform bound on the required auxiliary dimension exists. The hardness of the source games is an inherited theorem; the contribution is its transfer to this constrained determinant objective. The gap constant is very small, and no practical numerical hardness threshold or globally certified priority is claimed.

1 The objective and the result

For a tuple $B = (B_1, \dots, B_s)$ in $M_D(\mathbb{C})$, define

$$\mathcal{F}_m(B) = \max_{\sum_{\ell=1}^s \tau_m(Y_\ell^* Y_\ell) = 1} \left| \det \left(\sum_{\ell=1}^s B_\ell \otimes Y_\ell \right) \right|^{2/(Dm)}, \quad \mathcal{F}_\infty(B) = \sup_{m \geq 1} \mathcal{F}_m(B), \quad (1)$$

where $\tau_d = \text{Tr}/d$. Each fixed- m feasible set is compact. The variables Y_ℓ are arbitrary complex matrices: they are not required to be unitary, Hermitian, or mutually commuting.

For a completely positive map $\Phi : M_D(\mathbb{C}) \rightarrow M_D(\mathbb{C})$ let

$$\text{cap}(\Phi) = \inf_{X \succ 0, \det X = 1} \det \Phi(X).$$

For a density matrix $Q \in M_s(\mathbb{C})$ put

$$\Phi_Q(X) = \sum_{\ell, h} Q_{\ell h} B_\ell X B_h^*, \quad \mathcal{C}(B) = \sup_{Q \succeq 0, \text{Tr } Q = 1} \text{cap}(\Phi_Q)^{1/D}.$$

We call $\mathcal{C}$ the covariance-optimized capacity. The normalization in the following theorem makes it explicitly computable on every constructed instance.

Theorem 1. *Set*

$$\gamma = \frac{1}{36\,288\,000\,000}.$$

There is a computable map from Turing machines T to finite real rational matrix tuples $B(T)$ such that

$$\tau_D(B_\ell^* B_h) = \delta_{\ell h}, \quad \mathcal{C}(B) = 1, \quad (2)$$

and

$$T \text{ halts} \implies \mathcal{F}_\infty(B(T)) \geq 1 - \gamma/4, \quad (3)$$

$$T \text{ does not halt} \implies \mathcal{F}_\infty(B(T)) \leq 1 - \gamma. \quad (4)$$

Thus approximating $\mathcal{F}_\infty$ to additive error $\gamma/8$ is undecidable, even on the class (2).

The reduction is computable in the original machine description. It may explicitly expand a succinct game into a much larger finite list; we do not claim polynomial running time in the machine description. Once that list is given explicitly, the matrix construction has polynomial size.

Corollary 2. *For the same class there is no total recursive function of input length that bounds a dimension m satisfying*

$$\mathcal{F}_m(B) \geq \mathcal{F}_\infty(B) - \gamma/8.$$

In particular, there is no total recursive uniform auxiliary-dimension bound for arbitrary requested precision.

These are approximation statements, stronger than undecidability of exact equality with an upper bound. A previous construction [6] transferred exact and asymptotic saturation to balanced pencils, but its error bounds did not establish a fixed approximation gap. Here weighted singleton blocks and rational square normalization keep a constant fraction of the dimension available to test relations. All proofs of this new transfer are included.

The general phenomenon of noncomputability in quantum correlations and tracial matrix optimization is established. The source theorem is Taller–Vidick [1]; its foundations include $\text{MIP}^* = \text{RE}$ [2]. Goldbring–Hart [3] already prove noncomputability of the universal theory of the hyperfinite factor and the absence of computable general matrix-moment approximation bounds. Cleve–Culf–Taller [4] prove fixed-precision hardness for tilted XOR games and a constrained tracial Max-Cut problem. Our claim concerns the particular homogeneous determinant objective (1), with only a quadratic normalization on its variables and with (2). It is not a new general theorem about the noncomputability of tracial optimization.

2 A source game with unambiguous question positions

Throughout, a synchronous strategy means the convention used in [1, Sections 2.4 and 5.1]: a common finite-dimensional space, a maximally entangled state, and projective measurements paired with their transposes. For a constraint-variable game, this permits a joint measurement for each constraint and a separate binary measurement for each variable. The two question sets can be included in one disjoint union. We use synchronous value throughout this section; no equality with unrestricted quantum value is assumed.

Theorem 5.2 and its proof in [1] provide, for any fixed rational $0 < \varepsilon < 1/72$, a computable reduction with synchronous completeness $1 - \varepsilon$ and soundness $71/72$. Its verifier receives three parity-constrained answer bits from Alice and one bit from Bob and compares Bob’s bit with a specified Alice position. A folded query can coincide with another query. The following normalization avoids silently treating three positions as three distinct variables.

Write $\pi(x, y, i)$ for the rational joint probability of an Alice question x , a Bob question y , and the answer position $i \in \{1, 2, 3\}$ checked by the verifier. If the position is a deterministic function of (x, y) , this notation simply has support on that function. Alice's three answers must have product $\sigma_x \in \{\pm 1\}$. Let π_X be the Alice marginal.

Lemma 3 (Positional normalization). *Such a game can be transformed computably into a rationally weighted constraint-variable game with contexts of two or three distinct variables. If the old game has a synchronous strategy of error e , the new one has a synchronous strategy of error at most $e/4$. Every synchronous strategy of error e' for the new game yields one of error at most $18e'$ for the old game. Dimensions are unchanged.*

Proof. Introduce three distinct local variables $u_{x,1}, u_{x,2}, u_{x,3}$ for every Alice question, and a variable v_y for every Bob question. These two families are disjoint. Give the parity row

$$u_{x,1}u_{x,2}u_{x,3} = \sigma_x$$

weight $\pi_X(x)/2$. Give the equality row $u_{x,i}v_y = 1$ weight $\pi(x, y, i)/2$. Zero weights may be omitted; duplicate rows may be combined. A row is sampled according to these weights, and a variable uniformly within that row.

For completeness, let $P_{x,i}$ be Alice's commuting parity observables and let V_y be Bob's observables on the common space. For the new variables use $P_{x,i}$ and V_y . Alice uses the original joint measurement on parity rows; on an equality row she measures $P_{x,i}$ and returns the same bit twice. The first equality position agrees perfectly and the second has exactly the old disagreement probability. The equality rows have total weight $1/2$, giving error $e/4$. Invalid parity answers of an old strategy can first be replaced by a fixed valid answer without decreasing success.

For soundness, write $U_{x,i}$ and V_y for the new variable observables. On each parity row let Alice's commuting observables be $P_{x,i}$, with product $\sigma_x I$. On each equality row her two valid answers coincide, so they are generated by one involution $E_{x,y,i}$. Let e_{par} and e_{eq} be the conditional errors on the two types of rows; then $e' = (e_{\text{par}} + e_{\text{eq}})/2$. In normalized Hilbert–Schmidt norm,

$$e_{\text{par}} = \sum_x \frac{\pi_X(x)}{12} \sum_{i=1}^3 \|P_{x,i} - U_{x,i}\|_2^2,$$

and

$$e_{\text{eq}} = \sum_{x,y,i} \frac{\pi(x, y, i)}{8} (\|E_{x,y,i} - U_{x,i}\|_2^2 + \|E_{x,y,i} - V_y\|_2^2).$$

Decode with Alice's parity measurement $P_{x,i}$ and Bob's V_y . The squared triangle inequality gives

$$\|P_{x,i} - V_y\|_2^2 \leq 2\|P_{x,i} - U_{x,i}\|_2^2 + 2\|U_{x,i} - V_y\|_2^2.$$

Since $\sum_y \pi(x, y, i) \leq \pi_X(x)$ and $\|U - V\|_2^2 \leq 2\|U - E\|_2^2 + 2\|E - V\|_2^2$, the decoded error is at most

$$6e_{\text{par}} + 8e_{\text{eq}} \leq 16e' \leq 18e'.$$

The conservative constant 18 will suffice. $\square$

Enumerate the bounded finite sampler and aggregate probabilities of identical original questions, retaining the position actually checked; do not reveal its hidden random seed to Alice. The resulting probabilities are rational. After positional normalization, duplicating rows with a common denominator gives a uniform row list. Variables retain their global names across copies. For fixed Bob observables, identical copies have the same local optimization for Alice, so choosing a least-error measurement among the copies shows that this duplication does not change the supremal value. This may be very large but is a terminating computation. Keep

multiplicities rather than replacing them by an unweighted set. Thus the source theorem and Lemma 3 give a uniform list with

$$\text{YES error} \leq \varepsilon/4, \quad \text{NO synchronous value} \leq \sigma := 1 - \frac{1}{1296}. \quad (5)$$

The soundness follows because a new value greater than $1 - 1/1296$ would give an old value greater than $71/72$. Suprema need not be attained.

3 From game errors to averaged unitary relations

Allow nonempty contexts of sizes 1, 2, 3; this slightly generalizes the normalized source. Let there be $N \geq 1$ rows and remove unused variables. For each row, list the square relation for each variable, the signed product relation, and one positive commutator $x_j x_k x_j x_k = I$ for each pair. Complete this list to seven relations with empty words equal to I . There are

$$R = 7N$$

listed relations, each of length at most four. Occurrences in different rows remain distinct. For unitary variables $V = (V_j)$ let

$$\xi(V) = \frac{1}{R} \sum_{r=1}^R \|V_{i_1(r)} \cdots V_{i_k(r)} - \sigma_r I\|_2^2.$$

Lemma 4. *A synchronous strategy of error e yields unitary variables with $\xi \leq 34e$ in the same dimension. Conversely, unitary variables with averaged error ξ yield a synchronous strategy of error at most 50ξ in the same dimension.*

Proof. For the first direction let A_j^i be the row observables and X_j the variable involutions. Put $d_{ij} = \|A_j^i - X_j\|_2^2$. The normalized-trace expression for a maximally entangled strategy gives

$$e = \frac{1}{N} \sum_i \frac{1}{4k_i} \sum_{j \in S_i} d_{ij}, \quad \sum_{i,j} d_{ij} \leq 12Ne.$$

Unitary telescoping bounds a product residual squared by $3 \sum_j d_{ij}$. A pair commutator residual squared is at most $8(d_{ij} + d_{ik})$. Summing pairs costs at most $16 \sum_j d_{ij}$. Squares and empty words have zero error. Consequently $R\xi \leq 228Ne$, so $\xi \leq (228/7)e < 34e$.

For the reverse direction, round each unitary V_j spectrally to the closest Hermitian involution X_j . Eigenvalue by eigenvalue,

$$\|V_j - X_j\|_2^2 \leq \|V_j^2 - I\|_2^2.$$

For each row, the new product residual squared is at most four times the sum of its old product residual squared and the individual rounding errors. Write $e_{jk} = \|V_j V_k V_j V_k - I\|_2$ for the original pair residual. A new pair commutator residual squared is at most

$$3e_{jk}^2 + 12\|V_j - X_j\|_2^2 + 12\|V_k - X_k\|_2^2.$$

Each variable appears in at most two pairs. Its squared rounding error therefore has coefficient at most $4 + 2 \cdot 12 = 28$; the old product and pair residuals have coefficients at most four and three. Hence the sum of the new product and pair commutator squared errors is at most 28 times the original seven-relation squared-error sum for that row.

For a row of size three write its rounded variables as X, Y, Z . Pinch Y to the algebra commuting with X :

$$T = (Y + XYX)/2, \quad Y' = \text{sign}(T),$$

with sign $+1$ on a zero eigenvalue. Then Y' is an involution commuting with X , and

$$\|Y - Y'\|_2^2 \leq \frac{1}{2} \|XY - YX\|_2^2. \quad (6)$$

Indeed the off-diagonal part of Y in the X -eigenspaces has squared norm $\|XY - YX\|_2^2/4$. The sign-rounding error of the diagonal Hermitian contraction is no larger than this: use $(1 - |\lambda|)^2 \leq 1 - \lambda^2$ and $Y^2 = I$. The two errors are orthogonal.

Alice uses the commuting involutions $X, Y', \sigma XY'$ and their joint measurement. The sum of squared distances to X, Y, Z is at most

$$2\|XYZ - \sigma I\|_2^2 + 3\|Y - Y'\|_2^2.$$

Thus this row's game error is at most one sixth of its new product-plus-commutator error sum. For a size-two row use $X, \sigma X$, giving error $\|XY - \sigma I\|_2^2/8$; for size one use the constant σ , giving error $\|X - \sigma I\|_2^2/4$. All cases are bounded by one quarter of the new error sum. Averaging yields at most $28 \cdot 7\xi/4 = 49\xi \leq 50\xi$. Bob uses the X_j and Alice the transposed joint measurements on a maximally entangled state. These extend to a synchronous strategy on the disjoint union of question sets. $\square$

4 A homogeneous multiplication block

Put $a = 3/5$, $b = 4/5$, $c = ab = 12/25$, and define

$$H(U, V, W, Z) = \begin{pmatrix} aU & bW \\ bZ & -aV \end{pmatrix}.$$

Lemma 5. *For unitary $m \times m$ labels, set $E = UZ^* - WV^*$. Then*

$$HH^* = \begin{pmatrix} I & cE \\ cE^* & I \end{pmatrix}.$$

In particular H is unitary precisely when $(UZ^*)(VZ^*) = WZ^*$. Its squared singular values belong to $[1/25, 49/25]$, its operator norm is at most $7/5$, and

$$-\frac{1}{m} \log |\det H| \leq \frac{72}{49} \|E\|_2^2. \quad (7)$$

Proof. Direct multiplication proves the identity. Since $\|E\|_{\text{op}} \leq 2$, the eigenvalues are $1 \pm cs_j(E)$, giving the interval. Also

$$-\frac{1}{m} \log |\det H| = -\frac{1}{2m} \sum_j \log(1 - c^2 s_j(E)^2).$$

For $0 \leq u \leq 4c^2 = 576/625$, $-\log(1 - u) \leq u/(1 - 4c^2)$. Its coefficient is $c^2/[2(1 - 4c^2)] = 72/49$. $\square$

Uniform invertibility is essential here: a Hilbert–Schmidt estimate alone would not control a logarithmic determinant near zero.

For each original signed word add two padding letters pp and use fresh prefix labels of lengths $2, \dots, k-1$ for its padded length k . Introduce one common reference label z . A multiplication gate has lower-left label z , takes two successive factors as its diagonal labels, and has a fresh prefix as output, except that the final output is σz . The two input labels may coincide, but neither equals z or its own output.

For *each* original word include its own extra gate imposing $p^2 = I$, with the same label p shared across words. This deliberate repetition prevents an error in one padding gate from

being charged to every relation. Empty original words require one padded gate and their own square gate. A word of length at most four uses at most six gates including the square gate. Gates assigned to different words are disjoint occurrences.

Take 25 copies of each gate and initially one singleton block for every label. Write the resulting coefficient matrices as $A_\ell^{(0)}$. Each entry has only one label. The sums $\sum_\ell A_\ell^{(0)*} A_\ell^{(0)}$ and $\sum_\ell A_\ell^{(0)*} A_\ell^{(0)}$ are the identity block by block. For a gate this follows from $a^2 + b^2 = 1$ and the stated label separation, even when its two inputs coincide. Coefficient supports are disjoint. Let

$$q_\ell = \text{Tr}(A_\ell^{(0)*} A_\ell^{(0)}) \in \mathbb{N}.$$

These are positive integers: the 25 copies contribute 9 per diagonal occurrence and 16 per off-diagonal occurrence, while the singleton contributes 1.

5 Weighted rational normalization

Lemma 6. *The construction can be completed to rational matrices $A_\ell \in M_D(\mathbb{Q})$ and positive integers w_ℓ, t such that*

$$D = t^2 = \sum_\ell w_\ell^2, \quad \sum_\ell A_\ell A_\ell^* = \sum_\ell A_\ell^* A_\ell = I_D, \quad \text{Tr}(A_\ell^* A_h) = \delta_{\ell h} w_\ell^2, \quad (8)$$

$$\text{label } \ell \text{ has at least } w_\ell^2/2 \text{ singleton blocks}, \quad 200R \leq D \leq 8000R. \quad (9)$$

The tuple $B_\ell = tA_\ell/w_\ell$ is real rational and satisfies (2).

Proof. For each initial label choose $w_\ell = \lceil \sqrt{2q_\ell} \rceil$ and append $w_\ell^2 - q_\ell$ singleton blocks. Its squared coefficient norm becomes w_ℓ^2 and its singleton count is $1 + w_\ell^2 - q_\ell \geq w_\ell^2/2$. Moreover $2q_\ell \leq w_\ell^2 < 6q_\ell$. Let $D_0 = \sum_\ell w_\ell^2$ and $t = \lceil \sqrt{D_0} \rceil$. Append $t^2 - D_0$ entirely new labels, each appearing in one singleton only, and give each weight $w = 1$. This makes $D = t^2$ and preserves all identities. New labels obey the singleton condition as well.

If n is the number of original variables, $n \leq 3N = 3R/7$. There are at most $4R$ prefix labels, so the initial label count is $s_0 \leq n + 2 + 4R \leq 33R/7 < 5R$. The number G of gates before replication obeys $2R \leq G \leq 6R$. Thus the initial scalar dimension is

$$d_0 = \sum_\ell q_\ell = s_0 + 50G \leq 305R.$$

We have $2d_0 \leq D_0 < 6d_0$ and $D \leq 4D_0$, proving (9) with conservative constants. The construction and dense rational encoding have polynomial size in the expanded row list. Coefficient numerators and denominators have $O(\log R)$ bits.

For $B_\ell = tA_\ell/w_\ell$, (8) gives $\tau_D(B_\ell^* B_h) = \delta_{\ell h}$. Therefore, for any density Q , $\tau_D \Phi_Q(I) = 1$ and

$$\text{cap}(\Phi_Q)^{1/D} \leq \det \Phi_Q(I)^{1/D} \leq 1.$$

For $Q_0 = \text{diag}(w_\ell^2/D)$, we instead have $\Phi_{Q_0}(X) = \sum_\ell A_\ell X A_\ell^*$, a unital trace-preserving map. Its capacity is one: operator concavity of $\log$ gives $\text{Tr} \log \Phi_{Q_0}(X) \geq \text{Tr} \log X$ for $X > 0$, and $X = I$ gives the reverse bound. This is also the standard doubly stochastic capacity identity [5]. Hence $\mathcal{C}(B) = 1$. $\square$

It is the intermediate tuple A , not generally B , whose coefficient map is doubly stochastic. The final rational coefficients B_ℓ do not belong to a fixed finite alphabet. These distinctions are needed for the fixed-gap result.

6 The averaged determinant transfer

Let $\rho = R/D$, so $1/8000 \leq \rho \leq 1/200$, and write $L_A(Z) = \sum_\ell A_\ell \otimes Z_\ell$ (likewise for L_B). For an original relation r , let e_r denote its normalized Hilbert–Schmidt residual norm in the definition of ξ .

Proposition 7. *For the tuple B above, unitary variables of mean relation error ξ in dimension m give*

$$\mathcal{F}_m(B) \geq \exp(-(3600/49)\rho\xi). \quad (10)$$

Conversely, if $\mathcal{F}_m(B) \geq 1 - \delta$ with $0 < \delta < 1$, there are unitary variables in the same dimension with

$$\xi < 35\delta/\rho. \quad (11)$$

In particular the source constraint-variable game has a synchronous strategy of error at most $1750\delta/\rho$.

Proof. For the forward direction choose $z = p = I$ as matrix labels. Assign exact generator products to prefixes and I to the added free labels. Every gate except the final gate of an original relation is unitary. Set $Y_\ell = (w_\ell/\sqrt{D})V_\ell$, where V_ℓ are the unitary labels. Then $\sum \tau_m(Y_\ell^* Y_\ell) = \sum w_\ell^2/D = 1$, and the pencil equals $L_A(V)$. Each original final gate occurs 25 times. Equation (7) gives

$$-\log |\det L_A(V)|^{2/(Dm)} \leq \frac{50}{D} \frac{72}{49} \sum_{r=1}^R e_r^2 = \frac{3600}{49} \rho \xi.$$

For the reverse direction choose a maximizing feasible tuple Y . Write

$$Z_\ell = \frac{\sqrt{D}}{w_\ell} Y_\ell, \quad X = L_A(Z) = L_B(Y).$$

The Gram identity gives $\tau_{Dm}(X^* X) = 1$. A block diagonal polar unitary P satisfies

$$\Delta^2 := \|X - P\|_2^2 = 2(1 - \tau_{Dm}|X|) \leq 2(1 - \sqrt{1 - \delta}) \leq 2\delta. \quad (12)$$

This uses arithmetic–geometric mean on the singular values, and allows arbitrary nonunitary starting labels.

Let V_ℓ be polar unitaries of Z_ℓ . The singleton multiplicities in (9) give

$$\sum_\ell \frac{w_\ell^2}{D} \|Z_\ell - V_\ell\|_2^2 \leq 2\Delta^2.$$

With $H = L_A(V)$, orthogonal supports therefore imply $\|H - X\|_2 \leq \sqrt{2}\Delta$ and $\|H - P\|_2 \leq (1 + \sqrt{2})\Delta$. Lemma 5 gives $\|H\|_{\text{op}} \leq 7/5$, including all singleton blocks, so

$$\|HH^* - I\|_2^2 \leq C_0 \Delta^2, \quad C_0 = \frac{144}{25}(3 + 2\sqrt{2}).$$

For a gate g with unitary labels U, V, W, V_z , put $e_g = \|UV_z^* - WV^*\|_2$. Its normalized squared unitarity defect is $c^2 e_g^2$. Its 25 copies contribute $50c^2 e_g^2/D$ to the global squared defect. Thus

$$\frac{50c^2}{D} \sum_g e_g^2 \leq C_0 \Delta^2.$$

Right multiply all labels by V_z^* to fix the reference, and denote the gauged labels by $\tilde{V}_\ell$. Let G_r be the occurrences of gates assigned to relation r , including its own copy of the square gate $p^2 = I$. Telescoping the padded product and then removing its two padding factors gives

$$\|\tilde{V}_{i_1(r)} \cdots \tilde{V}_{i_k(r)} - \sigma_r I\|_2 \leq \sum_{g \in G_r} e_g, \quad |G_r| \leq 6.$$

The last summand can be chosen as the residual of that relation's square gate. Cauchy–Schwarz bounds the squared left-hand side by $|G_r| \sum_{g \in G_r} e_g^2$. The sets G_r are disjoint sets of occurrences, even when their labels coincide. Consequently

$$\xi \leq \frac{6}{R} \sum_g e_g^2 \leq \frac{12C_0}{50c^2} \frac{\delta}{\rho} = (18 + 12\sqrt{2}) \frac{\delta}{\rho} < 35 \frac{\delta}{\rho}.$$

Apply Lemma 4 for the game strategy. □

7 A fixed gap and nonrecursive witnesses

Proof of Theorem 1. Fix $\varepsilon = 2^{-40}$ in the source theorem and perform the positional normalization. Let

$$\beta = \frac{1 - \sigma}{3500} = \frac{1}{4536000}, \quad \gamma = \beta/8000.$$

For a YES instance, (5) and Lemma 4 yield averaged unitary errors approaching at most $34\varepsilon/4$. Taking the supremum in (10) gives

$$\mathcal{F}_\infty(B) \geq \exp\left(-\frac{3600}{49} \rho \frac{34\varepsilon}{4}\right) \geq 1 - \frac{3600}{49} \frac{1}{200} \frac{34}{4} 2^{-40}.$$

The last loss equals $153/(49 \cdot 2^{40}) < 1/145152000000 = \gamma/4$, by direct rational arithmetic. This establishes (3) without assuming an attaining strategy or perfect completeness.

For a NO instance, suppose $\mathcal{F}_\infty(B) > 1 - \beta\rho$. There would be a finite m and a feasible witness with $\mathcal{F}_m(B) > 1 - \beta\rho$. Proposition 7 would give a synchronous game strategy of error less than

$$1750\beta = (1 - \sigma)/2,$$

contradicting (5). Thus $\mathcal{F}_\infty(B) \leq 1 - \beta\rho \leq 1 - \gamma$. Lemma 6 proves the promised rationality and capacity identity. A $\gamma/8$ approximation separates the two cases, for example by comparison with $1 - 5\gamma/8$, and would decide the halting problem. □

Proof of Corollary 2. At fixed m , (1) is a compact semialgebraic optimization in the real and imaginary parts of the entries. For a rational threshold $r \geq 0$, comparison with r can be expressed using the squared determinant polynomial and r^{Dm} ; real quantifier elimination makes $\mathcal{F}_m$ uniformly computable. This is a computability statement, not an efficiency claim.

Suppose a total recursive bound $b(L)$ guaranteed a witness with additive deficit at most $\gamma/8$ for every encoded instance of length L . Evaluate all $\mathcal{F}_m$ with $1 \leq m \leq b(L)$ to error at most $\gamma/16$, and take their maximum. On a YES instance this is at least $1 - 7\gamma/16$; on a NO instance it is at most $1 - 15\gamma/16$. Comparing with $1 - 3\gamma/4$ decides halting, a contradiction. □

The absence of a uniform recursive bound does not say that every fixed tuple has a non-computable real optimum, or that every instance is difficult. At every fixed dimension the optimization remains decidable. Nor does the proof produce non-hyperlinear groups: its source completeness is strictly less than one.

8 Diagnostics and scope

The accompanying compact compiler stores scalar blocks and multiplicities rather than allocating the full dense tuple. Exact rational checks cover 70 context lists, all five encountered local label patterns, the stochastic identities, coefficient orthogonality, square normalization and singleton proportions. Forty-eight complex-matrix cases check each of the forward bound, reverse bound and unitary-to-strategy rounding, with 48 inexact relation controls. These are finite diagnostics, not a proof of undecidability.

For the six-row magic square, the compiler produces $R = 42$, 210 multiplication gates before replication, $D = 23409$, and 438 final labels. The standard four-dimensional Pauli assignment makes the determinant objective exactly one; all 512 scalar sign assignments fail the row system. A separate positional-normalization verifier covers a nonuniform rational distribution with coincident Bob queries, including exhaustive scalar cases and complex-matrix tests of both error transfers. This is a familiar illustration of the construction, not a new nonlocality result; neither verifier implements the full machine-to-game source reduction.

The proof has several explicit limits. The gap $\gamma \approx 2.76 \cdot 10^{-11}$ is a conservative theoretical constant, not a claim about achievable numerical experiments. The source-to-explicit-list expansion need not be efficient. The final coefficients are rational with controlled bit length, but are not uniformly bounded in one fixed alphabet and do not themselves form a doubly stochastic tuple. The capacity-one identity uses a displayed covariance Q_0 . The source game hardness, quantum logic background and general noncomputability of tracial approximation are credited to prior work. Targeted source searches do not certify worldwide novelty or establish an importance rating.

Assistance and assessment. Developed with substantial OpenAI Codex assistance in derivation, writing, source checking and finite computations. GPT-6 Astra at xhigh effort was requested to review the candidate construction in reused context; runtime model identity was not separately attested. Any final exact-version assessment is recorded separately and must not be inferred from an earlier paper’s score. No human peer review, certified independent context or proof-assistant verification is claimed.

References

- [1] A. Taller and T. Vidick, *Approximating the quantum value of an LCS game is RE-hard*, arXiv:2507.22444v2 (1 April 2026), especially Sections 2.4, 3.2 and 5.1, Theorem 5.2 and its proof. <https://arxiv.org/abs/2507.22444v2>.
- [2] Z. Ji, A. Natarajan, T. Vidick, J. Wright and H. Yuen, *MIP*=RE*, arXiv:2001.04383. <https://arxiv.org/abs/2001.04383>.
- [3] I. Goldbring and B. Hart, *The universal theory of the hyperfinite II_1 factor is not computable*, Bulletin of Symbolic Logic **30**(2) (2024), 181–198. <https://doi.org/10.1017/bsl.2024.7>.
- [4] R. Cleve, E. Culf and A. Taller, *XOR Games at Full Tilt: The Hardness of Binary Nonlocal Games*, arXiv:2607.06876v1 (8 July 2026), especially Section 6.4. <https://arxiv.org/abs/2607.06876v1>.
- [5] A. Garg, L. Gurvits, R. Oliveira and A. Wigderson, *Operator Scaling: Theory and Applications*, Foundations of Computational Mathematics **20** (2020), 223–290. <https://doi.org/10.1007/s10208-019-09417-z>.

- [6] L. Eriksson, *Exponential witness size and undecidable saturation for amplified determinants*, AIRR working paper, ARR-2026-438PTRJA439VSRTW, v1 (2026). <https://airr.science/papers/ARR-2026-438PTRJA439VSRTW/versions/v1/>.
- [7] W. Slofstra, *The set of quantum correlations is not closed*, Forum of Mathematics, Pi **7** (2019), e1. <https://arxiv.org/abs/1703.08618>.