Electronic Colloquium on Computational Complexity, Report No. 199 (2026)

Parallel Repetition for Entangled Games with Gap Exponent Three

Zhao Song*

August 5, 2026

Abstract

We prove that every finite two-player game G with entangled value $\omega^*(G) = 1 - \epsilon$ satisfies

$$\omega^*(G^{\otimes n}) \leq \exp(-\Omega(\frac{\epsilon^3}{\epsilon + \ell} n))$$

for every $n \geq 1$, where $\ell := \log(|A||B|)$, and A and B are the answer alphabets. Compared with Chapter 6 of the OpenAI report [Ope26], this improves the gap exponent from thirteen to three and matches the cubic gap dependence in Holenstein's general classical bound [Hol09]:

$$\omega(G^{\otimes n}) \leq \exp(-\Omega(\frac{(1 - \omega(G))^3}{1 + \ell} n)).$$

The proof replaces the randomly shifted logarithmic grid used in quantum correlated sampling by smooth soft labels. This makes the relevant label infidelity quadratic in the distance between state descriptions and avoids a Jensen loss when averaging over questions. Together with the postselection argument, these improvements yield the cubic gap dependence stated above.

1 Introduction

Parallel repetition is a basic method for amplifying the soundness of two-player games: the verifier samples independent coordinates, and the players must win every coordinate. For classical strategies, exponential decay was established by Raz [Raz98], and Holenstein [Hol09] subsequently gave a substantially simpler information-theoretic proof based on dependency breaking and correlated sampling. Here $\omega(G)$ denotes the classical value of G, the supremum of the winning probability over strategies in which the two players share randomness but no entanglement. His bound is

$$\omega(G^{\otimes n}) \leq \exp(-\Omega((1 - \omega(G))^3 n/(1 + \ell))),$$

for every $n \geq 1$, where $\ell := \log(|A||B|)$ and A, B are the answer alphabets; this remains the best known gap dependence for general classical games. Rao [Rao11] sharpened the gap dependence for projection games, whereas Raz [Raz11] showed that the strongest linear-gap bound cannot hold in full generality. The entangled setting is more delicate because conditioning on success changes a shared quantum state, while the players must reconstruct the resulting local descriptions without communication.

The general quantum analogue of Raz's theorem was explicitly identified as open by Cleve, Høyer, Toner, and Watrous [CHTW04]. Exact or exponential repetition was subsequently established under additional structure: quantum XOR games [CSUU08], entangled unique games

*magic.linuxkde@gmail.com. The author would like to thank Omri Weinstein for helpful discussions.

ISSN 1433-8092

[KRT10], projection games [DSV15], games with strictly positive Cartesian question support [CS14], and free or product-distribution games [JPY14, CS15]. Dinur, Steurer, and Vidick [DSV15] also introduced the quantum correlated-sampling primitive that later became central to general parallel-repetition arguments.

For general entangled games, Kempe and Vidick [KV11] obtained amplification by adding consistency or dummy questions, and Bavarian, Vidick, and Yuen [BVY17] proved exponential amplification after anchoring the game. Because these transformations modify the game, they do not establish a repetition theorem for the original game. Yuen [Yue16] proved polynomial decay for arbitrary unmodified games. The recent OpenAI report [Ope26] established exponential decay with explicit alphabet dependence for every finite two-player entangled game. Writing $\epsilon := 1 - \omega^*(G)$, its theorem gives

$$\omega^*(G^{\otimes n}) \leq \exp(-\Omega(\epsilon^{13} n/(\epsilon + \ell))).$$

The report does not claim that the exponent 13 is optimal. It states that this exponent arises from the quantitative loss in the standard quantum correlated-sampling lemma of Dinur, Steurer, and Vidick [DSV15, Lemma 17]. Kempe and Regev [KR10] show that no universal exponent below 2 is possible, even for constant-answer entangled unique games. Before the present work, the optimal exponent for general entangled games was therefore known only to lie between 2 and 13:

What is the optimal exponent of ϵ in general quantum parallel repetition?

1.1 Our Results

We improve Theorem 1.1 of [Ope26, Chapter 6] from ϵ^{13} to ϵ^3, which matches the classical bound of Holenstein [Hol09] quoted above.

Theorem 1.1 (Improved parallel-repetition exponent, informal version of Theorem 4.6)**.** *There is a universal constant $c_3 > 0$ such that every finite two-player entangled game G with $\omega^*(G) = 1 - \epsilon$ satisfies*

$$\omega^*(G^{\otimes n}) \leq \exp(-c_3 \epsilon^3 n/(\epsilon + \ell))$$

for every integer $n \geq 1$. Here $\ell := \log(|A||B|)$.

For fixed answer alphabets, the theorem gives a constant-factor reduction in value after $O(\epsilon^{-3})$ repetitions and exponential decay thereafter. Its cubic dependence on the gap matches Holenstein's [Hol09] general classical bound, while applying to arbitrary finite entangled games without modifying the game. The exponent 3 remains one power above the exponent-2 obstruction of Kempe and Regev [KR10], leaving open whether the bound is optimal.

Table 1: Parallel-repetition decay rates for general finite two-player games. Write $\epsilon := 1 - \omega^*(G)$ in the quantum rows and $\epsilon := 1 - \omega(G)$ in the classical rows, and $\ell := \log(|A||B|)$. Raz [Raz98] states his bound with an unspecified constant depending only on the value; the exponent 32 is implicit in his proof and was extracted by Holenstein [Hol09].

Year	Authors	Reference	Statement ID	Bound	Q/C
1998	Raz	[Raz98]	Main theorem	$\exp(-\Omega(\epsilon^{32} n/\ell))$	Classical
2009	Holenstein	[Hol09]	Theorem 2.5	$\exp(-\Omega(\epsilon^3 n/(1+\ell)))$	Classical
2026	OpenAI	[Ope26]	Theorem 1.1	$\exp(-\Omega(\epsilon^{13} n/(\epsilon+\ell)))$	Quantum
2026	This paper	This paper	Theorem 1.1	$\exp(-\Omega(\epsilon^3 n/(\epsilon+\ell)))$	Quantum

2 Technique Overview

2.1 Summary of the OpenAI Chapter 6 approach

OpenAI's proof follows the conditioning-and-sampleability framework of Yuen [Yue16], together with Holenstein's classical correlated-sampling method [Hol09] and quantum correlated sampling [Ope26, Chapter 6]. Begin with a strategy for $G^{\otimes n}$ whose probability of winning every coordinate is larger than the proposed exponential bound. A greedy conditioning argument selects a small set $D \subsetneq [n]$ of coordinates. If $p := \Pr[W_D]$ is the probability of winning the coordinates in D, $m := n - |D|$, and q is the average probability of winning a uniformly chosen coordinate outside D conditioned on W_D, then q is close to one. The information cost per remaining coordinate is $\eta := \frac{\log(1/p)+|D|\ell}{m}$, $\ell := \log(|A||B|)$. The greedy choice of D guarantees that η is small whenever the all-coordinate winning probability is too large.

Conditioning on W_D produces an ideal one-coordinate experiment, but this experiment is not immediately a legal strategy for G. Its question distribution is a posterior distribution rather than μ, its revealed history is not supplied by the referee, and its shared state depends jointly on Alice's and Bob's questions. The postselected state-alignment and history-sampleability lemma resolves these three defects. It constructs locally generated histories that are close to the ideal posterior distribution in total variation and two locally available state descriptions $|\Psi^A\rangle$ and $|\Psi^B\rangle$ satisfying $\mathbb{E}[\||\Psi^A\rangle - |\Psi^B\rangle\|^2] = O(\eta)$. The report's new sampleability estimate is proved through a purified positive-operator martingale. Its one-step entropy gap is an operator Bregman divergence in the sense of Petz, with related resolvent estimates due to Kim [Pet07, Kim14]. Classical correlated sampling synchronizes the histories. The quantum correlated-sampling lemma of Dinur, Steurer, and Vidick [DSV15] then turns the two nearby state descriptions into an actual shared state. For description distance $t := \||\Psi^A\rangle - |\Psi^B\rangle\|$, that lemma incurs vector error $O(t^{1/6})$, apart from an arbitrarily small finite-catalyst error. Averaging and applying Jensen's inequality to the mean squared bound on t therefore gives the postselection-stable rounding estimate $\omega^*(G) \geq q - O(\eta^{1/12})$.

The exponent 13 comes from the final parameter choice. If the proposed decay rate is $\exp(-\gamma n)$, greedy conditioning gives $\eta = O(\gamma(1 + \ell/\epsilon))$. To make the rounding loss $O(\eta^{1/12})$ smaller than the original gap $\epsilon := 1 - \omega^*(G)$, one needs $\eta = O(\epsilon^{12})$. Thus one sets $\gamma = \Theta(\epsilon^{13}/(\epsilon + \ell))$. The rounded single-game strategy then wins with probability strictly larger than $\omega^*(G)$, which is the desired contradiction. This proves exponential decay for the unchanged game, but the hard-bin quantum sampler and the subsequent Jensen step are quantitatively expensive.

2.2 Summary of our approach

We retain the conditioning, history-sampling, and state-alignment parts of Chapter 6. The improvement is concentrated in two places: a smoother quantum correlated-sampling primitive and a near-perfect measurement estimate. For a finite distribution of state-description pairs, write $D := \mathbb{E}[\||\psi_z\rangle - |\varphi_z\rangle\|^2]$. Instead of assigning each Schmidt coefficient [NC10] to a randomly shifted hard logarithmic bin, Lemma 4.1 assigns it a two-bin raised-cosine label. Nearby logarithmic coefficients then have quadratically close labels. A diagonal filter whose weights vary exponentially across the labels approximately restores the Schmidt coefficients, while the exact synchronization identity controls the probability that Alice and Bob keep different trials.

The construction separates into weight approximation and synchronization. If $s_\lambda := -\log(\lambda)/h$ and $|g_\lambda\rangle$ is the corresponding raised-cosine label, then the diagonal label operator T satisfies

$$\|(T - \lambda I)|g_\lambda\rangle\| \leq Ch\lambda, \qquad \lambda\mu(1 - |\langle g_\lambda|g_\mu\rangle|^2) \leq C(\lambda - \mu)^2/h^2.$$

The first estimate says that filtering approximately reconstructs the target Schmidt weight. The second makes the overlap loss quadratic in the coefficient difference. Together with the Schmidt-transport inequality, it bounds both the synchronization defect and the squared error of the synchronized output by $C(h^2 + \Delta/h^2)$ for a pair at squared distance $\Delta := \||\psi\rangle - |\varphi\rangle\|^2$. The parties implement the filters on independent finite-dimensional shared trials and retain their first local successes. A disagreement between the two first-success times is charged as an error, whereas a common first-success time produces the synchronized output. The shared resource is fixed for the entire finite family before z is known; only each party's local channel depends on its own description.

For smoothing width h, the construction gives the average infidelity bound

$$\mathbb{E}[1 - \langle\psi_z|\rho_z|\psi_z\rangle] \leq C_{\mathrm{sm}}(h^2 + D/h^2 + D) + \zeta.$$

The important feature is that the bound is linear in the squared description distance before averaging. Choosing $h := D^{1/4}$ yields average infidelity $O(D^{1/2})$ and average trace-distance error $O(D^{1/4})$. Since the postselected state-alignment lemma gives $D = O(\eta)$, Lemma 4.4 yields the generic postselection-stable rounding estimate $\omega^*(G) \geq q - B_4\eta^{1/4}$. Thus the generic rounding loss becomes $O(\eta^{1/4})$ rather than $O(\eta^{1/12})$. This estimate alone would lead to gap exponent 5. To reach exponent 3, Lemma 4.5 uses that the ideal conditioned coordinate already wins with probability q close to one. If the ideal pure state has rejection probability α and the prepared state has infidelity f, then the rejection probability after preparation is at most $\alpha + O(\sqrt{\alpha f} + f)$. Here $\alpha := 1 - q$ and $f = O(\eta^{1/2})$, so the loss is $O((1 - q)^{1/2}\eta^{1/4} + \eta^{1/2})$. This improvement comes from retaining the losing effect rather than first converting the prepared state to trace distance. If L is the losing effect, the ideal state is rejected with probability α, and the prepared state has infidelity f, Uhlmann's theorem [Uhl76, NC10] and the contraction $\sqrt{L}$ give

$$\sqrt{\mathrm{Tr}[L\rho]} \leq \sqrt{\alpha} + \sqrt{2f}, \qquad \mathrm{Tr}[L\rho] \leq \alpha + C\sqrt{\alpha f} + Cf.$$

The estimate is then averaged under the locally generated history distribution. The total-variation change and the probability of mismatching histories cost $O(\kappa)$, where $\kappa := \sqrt{3\eta/2}$. If A and R' are the resulting mean ideal rejection and mean prepared-state infidelity, respectively, then $A \leq 1 - q + O(\eta^{1/2})$ and $R' \leq O(\eta^{1/2}) + 2\zeta$. Cauchy–Schwarz bounds the averaged cross term by $\sqrt{AR'}$, producing exactly the loss in Lemma 4.5 after $\zeta \to 0$.

In the contradiction argument $1 - q = O(\epsilon)$ and it is enough to take $\eta = O(\epsilon^2)$. Finally, $\eta = O(\gamma(1 + \ell/\epsilon))$ allows $\gamma = \Theta(\epsilon^3/(\epsilon + \ell))$. The smooth sampler changes the required information-cost scale from $\eta = O(\epsilon^{12})$ to $\eta = O(\epsilon^4)$, and the near-perfect estimate relaxes it further to $\eta = O(\epsilon^2)$. Substituting this value of γ into the contradiction argument gives Theorem 4.6.

3 Preliminaries and Notation

All logarithms are natural. For an integer $n \geq 1$, write $[n] := \{1, \ldots, n\}$. For a finite set S, $|S|$ denotes its cardinality, and $D_0 \subsetneq [n]$ means that D_0 is a proper subset. Probabilities and expectations are denoted by $\Pr[\cdot]$ and $\mathbb{E}[\cdot]$; subscripts specify the underlying random variable or distribution. The symbols $O(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$ hide universal positive constants unless a dependence is stated explicitly. Symbols such as C, c, C_{sm}, B_4, B_{np}, and c_3 denote positive universal constants.

3.1 Games and parallel repetition

A finite two-player one-round game is a tuple $G := (X, Y, A, B, \mu, V)$, where X, Y are the question sets, A, B are the nonempty answer alphabets, μ is a probability distribution on $X \times Y$, and

$V : X \times Y \times A \times B \to \{0,1\}$ is the acceptance predicate. The referee samples $(x,y) \sim \mu$, sends x to Alice and y to Bob, and receives answers $a \in A$ and $b \in B$.

A POVM is a finite family of positive semidefinite operators summing to the identity. A finite-dimensional entangled strategy consists of a unit vector $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ and local POVMs $\{E_x^a : a \in A\}$ and $\{F_y^b : b \in B\}$. Thus $E_x^a, F_y^b \succeq 0$ and $\sum_a E_x^a = I_{\mathcal{H}_A}$, $\sum_b F_y^b = I_{\mathcal{H}_B}$. Its winning probability is

$$\sum_{x,y} \mu(x,y) \sum_{a,b} V(x,y,a,b) \langle \psi | E_x^a \otimes F_y^b | \psi \rangle.$$

The entangled value $\omega^*(G)$ is the supremum of this probability over all finite-dimensional tensor-product strategies. We use

$$\epsilon := 1 - \omega^*(G), \qquad \ell := \log(|A||B|)$$

for the soundness gap and the logarithmic answer-alphabet size.

The game $G^{\otimes n}$ consists of n independently sampled question pairs, and the verifier accepts only if every coordinate is won. A repeated strategy may nevertheless correlate its answers across coordinates. For $i \in [n]$, let W_i be the event that coordinate i is won, and for $S \subseteq [n]$ put $W_S := \bigcap_{i \in S} W_i$. Thus $W_{[n]}$ is the all-coordinate winning event and $\omega^*(G^{\otimes n})$ is the supremum of its probability.

3.2 States and operator norms

The space $\mathbb{C}^d$ is equipped with its Euclidean norm, denoted $\|\cdot\|$. Kets $|u\rangle$ are vectors, bras $\langle u|$ are their adjoints, and $|u\rangle\langle u|$ is the rank-one projector associated with a unit vector. A mixed state ρ is a positive semidefinite operator with $\mathrm{Tr}[\rho] = 1$. The identity on $\mathbb{C}^d$ is I_d, tensor products are denoted by $\otimes$, and M^* is the adjoint of an operator M.

For an operator M,

$$\|M\|_1 := \mathrm{Tr}[\sqrt{M^* M}], \qquad \|M\|_F := \sqrt{\mathrm{Tr}[M^* M]}$$

are the trace and Frobenius norms. We call $\|\rho - \sigma\|_1$ the trace-norm distance; the corresponding operational trace distance is one half of this quantity. For a unit vector $|\psi\rangle$ and a state ρ, $1 - \langle \psi | \rho | \psi \rangle$ is their infidelity. The bar in $\overline{P}$ denotes entrywise conjugation in the fixed maximally entangled basis. A channel is a completely positive trace-preserving map, and an isometry preserves inner products.

3.3 Tools from previous work

We use three results from Chapter 6 of the OpenAI report as black boxes [Ope26]. The first input is the quantitative greedy-conditioning lemma [Ope26, Chapter 6, Lemma 3.1].

Lemma 3.1 (Greedy conditioning)**.** *Suppose a finite-dimensional strategy for $G^{\otimes n}$ wins all coordinates with probability $\vartheta > 0$. If $0 < \delta < 1$ and $\log(1/\vartheta)/\delta < n$, then there is a proper set $D_0 \subsetneq [n]$ such that, with $m := n - |D_0|$,*

$$|D_0| \le \frac{\log(1/\vartheta)}{\delta}, \qquad p := \Pr[W_{D_0}] \ge \vartheta, \qquad q := \frac{1}{m} \sum_{i \in [n] \setminus D_0} \Pr[W_i \mid W_{D_0}] \ge 1 - \delta.$$

For the remaining two statements, fix a finite-dimensional strategy for $G^{\otimes n}$ and a proper subset $D_0 \subsetneq [n]$ with $\Pr[W_{D_0}] > 0$, and use the associated parameters p, m, q, η from Definition 4.3 below. Following Chapter 6, Section 3.2, let Q denote the ideal posterior distribution. A sample from Q is written (i, R, X_i, Y_i), where i is uniform in $[n] \setminus D_0$, R is the conditioned history, and (X_i, Y_i) are the live questions. We write (i, r, x, y) for a realization and $|\Psi_{r,x,y}\rangle$ for the corresponding ideal state.

The second input identifies local measurements on the ideal state [Ope26, Chapter 6, Lemma 3.2].

Lemma 3.2 (Ideal state simulates the remaining coordinate)**.** *For every (i, r, x, y) in the support of Q, there are local* POVMs $\{M^a_{r,x} : a \in A\}$ *and* $\{N^b_{r,y} : b \in B\}$*, determined by (r, x) and (r, y), respectively, such that*

$$\langle \Psi_{r,x,y}| M^a_{r,x} \otimes N^b_{r,y} |\Psi_{r,x,y}\rangle = \mathsf{Q}(A_i = a, B_i = b \mid R = r, X_i = x, Y_i = y).$$

Consequently, the ideal experiment has average winning probability

$$\mathbb{E}_{(i,r,x,y)\sim\mathsf{Q}}[\sum_{a,b} V(x, y, a, b)\langle \Psi_{r,x,y}| M^a_{r,x} \otimes N^b_{r,y} |\Psi_{r,x,y}\rangle] = q.$$

The third input makes the histories and state descriptions locally available [Ope26, Chapter 6, Lemma 3.3].

Lemma 3.3 (Postselected state alignment and history sampleability)**.** *Let J_A and J_B be the tuple distributions obtained when Alice and Bob generate the history from their own questions. There are locally describable states $|\Psi^A_{r,x}\rangle$ and $|\Psi^B_{r,y}\rangle$ satisfying*

$$\begin{aligned} \mathbb{E}_{(i,r,x,y)\sim\mathsf{Q}}[\| |\Psi^A_{r,x}\rangle - |\Psi_{r,x,y}\rangle \|^2] &\le 8\eta, \\ \mathbb{E}_{(i,r,x,y)\sim\mathsf{Q}}[\| |\Psi^B_{r,y}\rangle - |\Psi_{r,x,y}\rangle \|^2] &\le 8\eta. \end{aligned} \tag{1}$$

Moreover, for $\kappa := \sqrt{3\eta/2}$,

$$d_{\mathrm{TV}}(\mathsf{Q}, \mathsf{J}_A) \le \kappa, \qquad d_{\mathrm{TV}}(\mathsf{Q}, \mathsf{J}_B) \le \kappa.$$

There is finite shared randomness such that, on fresh questions $(x, y) \sim \mu$, Alice and Bob sample a common uniform $i \in [n] \setminus D_0$ and histories r_A, r_B with tuple laws $\mathsf{J}_A, \mathsf{J}_B$ and $\Pr[r_A \neq r_B] \le 4\kappa$.

Consequently,

$$\mathbb{E}_{(i,r,x,y)\sim\mathsf{Q}}[\| |\Psi^A_{r,x}\rangle - |\Psi^B_{r,y}\rangle \|^2] \le 32\eta. \tag{2}$$

4 Proof of the improved exponent

The proof has three stages. First, smooth correlated sampling converts a mean squared discrepancy between local state descriptions into a controlled state-preparation error. Second, postselection-stable rounding converts this error into a single-game strategy. Third, greedy conditioning and near-perfect rounding give the parallel-repetition contradiction. We use the trace norm $\|M\|_1 := \mathrm{Tr}[\sqrt{M^*M}]$. For unit vectors $|u\rangle$ and $|v\rangle$ we repeatedly use

$$\| |u\rangle\langle u| - |v\rangle\langle v| \|_1 \le 2\| |u\rangle - |v\rangle \|.$$

Table 2: State-preparation and rounding losses for quantum correlated sampling. Write $t := \||\psi\rangle - |\varphi\rangle\|$, $D := \mathbb{E}[t^2]$, and η for the information cost in the conditioned single-game experiment. Arbitrarily small finite-resource accuracy terms are suppressed. The 2019 result is an operational promised-distance statement and does not by itself give a uniform Chapter 6 rounding bound.

Year	Authors	Reference	Statement ID	Sampler error	Rounding loss
2015	Dinur–Steurer–Vidick	[DSV15]	Lemma 17	$O(t^{1/6})$	$O(\eta^{1/12})$
2019	Lasecki	[Las19]	Theorem 3.1.1	$O(t^{1/5})$	promised distance
2026	This paper	This paper	Lemma 4.1	$O(D^{1/4})$	$O(\eta^{1/4})$

4.1 Smooth operational correlated sampling

The rounding argument supplies a distribution of pairs of state descriptions together with an upper bound on their mean squared distance. The construction below bounds the output infidelity pointwise by a quantity linear in the squared description distance. Averaging over the distribution therefore incurs no Jensen loss, removing the square-root losses of the earlier argument.

Lemma 4.1 (Smooth operational correlated sampling). *There is a universal constant $C_{\mathrm{sm}} \geq 1$ with the following property. Let $\mathcal{Z}$ be finite, let ν be a probability distribution on $\mathcal{Z}$, and for every $z \in \mathcal{Z}$ let $|\psi_z\rangle, |\varphi_z\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ be unit vectors. Set*

$$D := \mathbb{E}_{z\sim\nu}[\||\psi_z\rangle - |\varphi_z\rangle\|^2].$$

For every $h \in (0, 1/2)$ and $\zeta > 0$ there are a finite-dimensional shared state Ω, which may depend on the finite family, ν, h, and ζ but not on the realized z, and families of local channels $\{\mathcal{A}_\psi\}$ and $\{\mathcal{B}_\varphi\}$ indexed separately by Alice's and Bob's local state descriptions. Their joint output $\rho_z := (\mathcal{A}_{\psi_z} \otimes \mathcal{B}_{\varphi_z})(\Omega)$ satisfies

$$\mathbb{E}_{z\sim\nu}[1 - \langle\psi_z|\rho_z|\psi_z\rangle] \leq C_{\mathrm{sm}}(h^2 + D/h^2 + D) + \zeta. \tag{3}$$

Consequently, optimizing h and rescaling ζ gives

$$\mathbb{E}_{z\sim\nu}[\|\rho_z - |\psi_z\rangle\langle\psi_z|\|_1] \leq C_{\mathrm{sm}} D^{1/4} + \zeta. \tag{4}$$

The channels use only finite shared randomness and finite-dimensional shared entanglement, and hence can be absorbed into a tensor-product strategy.

Proof. We give a pointwise construction and then average. It suffices to treat $0 < \zeta < 1$, since a smaller accuracy parameter only strengthens the conclusion. Fix an internal finite-resource accuracy $\xi \in (0, 1)$, to be chosen at the end. If $D \geq 1/16$, use a fixed product shared state and fixed local output channels. Since $D^{1/4} \geq 1/2$, both assertions hold after increasing C_{sm}, since an infidelity is at most one and a trace distance at most two. For the h-dependent bound, we may therefore assume $D < 1/16$ and fix an arbitrary $h \in (0, 1/2)$; this choice is admissible, including when $D = 0$. Fix one pair of descriptions and write

$$|\psi\rangle = \sum_i \lambda_i |u_i\rangle|u_i'\rangle, \qquad |\varphi\rangle = \sum_j \mu_j |v_j\rangle|v_j'\rangle$$

in the standard correlated-sampling convention. Complete the Schmidt vectors to orthonormal bases and allow zero coefficients. We define

$$\Delta := \||\psi\rangle - |\varphi\rangle\|^2, \qquad w_{ij} := |\langle u_i|v_j\rangle|^2.$$

The Schmidt transport inequality used in the proof of the quantum correlated-sampling lemma is

$$\sum_{i,j}(\lambda_i-\mu_j)^2 w_{ij} \le 2\Delta. \tag{5}$$

To justify Eq. (5), let $\rho := \mathrm{Tr}_2[|\psi\rangle\langle\psi|]$ and $\sigma := \mathrm{Tr}_2[|\varphi\rangle\langle\varphi|]$. The left-hand side is $\|\sqrt{\rho}-\sqrt{\sigma}\|_F^2$. If $F(\rho,\sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1$ is the root fidelity, the affinity–fidelity inequality $F(\rho,\sigma)^2 \le \mathrm{Tr}[\sqrt{\rho}\sqrt{\sigma}]$ [LYM+19, Eq. (4.2)] and monotonicity of fidelity under partial trace give

$$\|\sqrt{\rho}-\sqrt{\sigma}\|_F^2 \le 2(1-F(\rho,\sigma)^2) \le 2(1-|\langle\psi|\varphi\rangle|^2) \le 2\Delta.$$

The constant two is necessary: if $|\psi\rangle$ and $|\varphi\rangle$ are product states whose second-register vectors agree and whose first-register vectors have overlap $\cos\theta$, the left side equals $2\sin^2\theta$ while $\Delta = 2-2\cos\theta$.

Smooth logarithmic labels. We replace hard logarithmic bins by overlapping two-bin wave packets. For $s \in \mathbb{R}$, set $k := \lfloor s \rfloor$ and $t := s-k \in [0,1)$, and define

$$|g(s)\rangle := \cos(\pi t/2)|k\rangle + \sin(\pi t/2)|k+1\rangle.$$

For $\lambda > 0$ put $s_\lambda := -\log(\lambda)/h$ and $|g_\lambda\rangle := |g(s_\lambda)\rangle$. Put $|g_0\rangle := |\perp\rangle$, where $|\perp\rangle$ is orthogonal to all integer-labelled vectors. Since the family of descriptions is finite, there is a common integer $k_{\max} \ge 0$ such that every positive coefficient has $s_\lambda \in [0, k_{\max}]$. On the label space set

$$T := \sum_{k=0}^{k_{\max}+1} \tau_k |k\rangle\langle k|, \qquad \tau_k := e^{-kh}, \qquad T|\perp\rangle := 0,$$

and write $\mathsf{T} := T \otimes I_d$.

We first record the scalar estimates. The two T-eigenvalues on the support of g_λ are between $e^{-h}\lambda$ and $e^{h}\lambda$, and hence

$$\|(T-\lambda I)|g_\lambda\rangle\| \le Ch\lambda. \tag{6}$$

The map $s \mapsto g(s)$ is globally Lipschitz after truncating its distance at 2, so

$$\|g_\lambda - g_\mu\| \le C\min\{1, |\log\lambda - \log\mu|/h\}.$$

Together with the elementary estimates for $e^x - 1$, this gives

$$\mu|1-\langle g_\lambda|g_\mu\rangle| \le C|\lambda-\mu|/h, \tag{7}$$

$$\lambda\mu(1-|\langle g_\lambda|g_\mu\rangle|^2) \le C(\lambda-\mu)^2/h^2. \tag{8}$$

For completeness, let $x := |\log\lambda - \log\mu|$ and $M := \max\{\lambda,\mu\}$. If $x \le h$, then $|\lambda-\mu| = M(1-e^{-x}) \ge Mx/2$. Lipschitz continuity of the labels, together with $|1-\langle u|v\rangle| \le \|u-v\|$ and $1-|\langle u|v\rangle|^2 \le \|u-v\|^2$ for unit vectors, gives the two displayed bounds. If $x > h$, then $|\lambda-\mu| = M(1-e^{-x}) \ge M(1-e^{-h}) \ge Mh/2$; the trivial bounds $|1-\langle g_\lambda|g_\mu\rangle| \le 2$ and $1-|\langle g_\lambda|g_\mu\rangle|^2 \le 1$ again give the two displayed inequalities. The definitions at zero make the same bounds immediate if one coefficient vanishes.

Filter synchronization. Define orthonormal vectors and their projectors by

$$|p_i\rangle := |g_{\lambda_i}\rangle|u_i\rangle, \qquad |q_j\rangle := |g_{\mu_j}\rangle|v_j\rangle, \qquad P := \sum_i |p_i\rangle\langle p_i|, \qquad Q := \sum_j |q_j\rangle\langle q_j|.$$

Thus P and $\overline{Q}$ are legal local filters. Set $S := d\sum_{k=0}^{k_{\max}+1}\tau_k^2$. Consider the normalized shared trial with vector $S^{-1/2}$ times

$$\sum_{k=0}^{k_{\max}+1}\tau_k|k,k\rangle \otimes \sum_{r=1}^{d}|r,r\rangle.$$

Alice applies P and Bob applies $\overline{Q}$, where the bar is the standard conjugation on Bob's half of a maximally entangled state. Define the unnormalized weights associated with the events that Alice succeeds, Bob succeeds, and both succeed by

$$a := \|P\mathsf{T}\|_F^2, \qquad b := \|\mathsf{T}Q\|_F^2, \qquad c := \|Q\mathsf{T}P\|_F^2.$$

Indeed, the actual one-trial probabilities are a/S, b/S, and c/S, respectively. The coefficient of $p_i \otimes \overline{q}_j$ in the both-success branch is $\langle p_i|\mathsf{T}|q_j\rangle$. Since P, Q are projectors and T is self-adjoint, their synchronization defect has the exact form

$$\mathfrak{m} := a + b - 2c = \|P\mathsf{T} - \mathsf{T}Q\|_F^2. \tag{9}$$

We define

$$A_\psi := \sum_i \lambda_i|p_i\rangle\langle p_i|, \qquad A_\varphi := \sum_j \mu_j|q_j\rangle\langle q_j|.$$

The eigenvalue sandwich preceding Eq. (6) gives

$$\|P\mathsf{T} - A_\psi\|_F^2 + \|\mathsf{T}Q - A_\varphi\|_F^2 \le Ch^2.$$

Moreover,

$$\|A_\psi - A_\varphi\|_F^2 = \sum_{i,j}((\lambda_i - \mu_j)^2 + 2\lambda_i\mu_j(1 - |\langle g_{\lambda_i}|g_{\mu_j}\rangle|^2))w_{ij}.$$

The squared triangle inequality gives

$$\mathfrak{m} \le 3(\|P\mathsf{T} - A_\psi\|_F^2 + \|A_\psi - A_\varphi\|_F^2 + \|A_\varphi - \mathsf{T}Q\|_F^2).$$

Using Eq. (5) and Eq. (8) in Eq. (9) yields

$$\mathfrak{m} \le C(h^2 + \Delta/h^2). \tag{10}$$

Output-state comparison. It remains to compare the synchronized output with the target after the local label erasures. Because the p_i are orthonormal, Alice has a local isometry sending p_i to u_i with a fixed label register. Similarly, Bob has a local isometry sending $\overline{q}_j$ to v_j'. Under the standard vectorization convention, the resulting unnormalized synchronized vector is

$$|x\rangle = \sum_{i,j}\beta_{ij}\langle u_i|v_j\rangle|u_i\rangle|v_j'\rangle, \qquad \beta_{ij} := \langle g_{\lambda_i}|T|g_{\mu_j}\rangle,$$

and $\|x\|^2 = c$. On the other hand,

$$|\varphi\rangle = \sum_{i,j}\mu_j\langle u_i|v_j\rangle|u_i\rangle|v_j'\rangle.$$

Indeed,

$$\beta_{ij} - \mu_j = \langle g_{\lambda_i} | (T - \mu_j I) | g_{\mu_j} \rangle + \mu_j (\langle g_{\lambda_i} | g_{\mu_j} \rangle - 1).$$

Therefore Eq. (6) and Eq. (7) imply

$$|\beta_{ij} - \mu_j|^2 \le C(h^2 \mu_j^2 + (\lambda_i - \mu_j)^2 / h^2).$$

The vectors $|u_i\rangle|v'_j\rangle$ are orthonormal. Therefore

$$\| |x\rangle - |\varphi\rangle \|^2 \le C(h^2 + \Delta/h^2). \tag{11}$$

Let $H := C(h^2 + \Delta/h^2)$. If $H < 1/4$, Eq. (11) implies $|\|x\| - 1| \le \| |x\rangle - |\varphi\rangle \| \le \sqrt{H}$, so $x \ne 0$ and

$$\| |x\rangle / \|x\| - |\varphi\rangle \| \le |1 - \|x\|| + \| |x\rangle - |\varphi\rangle \| \le 2\sqrt{H}.$$

Hence normalization changes this squared distance by at most a universal factor; if it is bounded below by a constant, the trivial infidelity bound suffices. Thus the normalized synchronized state has infidelity at most $C(h^2 + \Delta/h^2)$ with $|\varphi\rangle$.

Repeated trials and finite resources. To finish the protocol, the parties use independent trials and each keeps the first trial on which its local filter succeeds. The eigenvalue sandwich preceding Eq. (6) gives $a, b \ge e^{-2h}$. In each trial the probability that at least one filter succeeds is $(a + b - c)/S$. Conditioning on the first such trial, the probability that the two first-success times differ is

$$1 - \frac{c}{a + b - c} = \frac{\mathfrak{m}}{a + b - c} \le e^{2h} \mathfrak{m}.$$

On this branch charge infidelity one. Combining this estimate with Eq. (10) and Eq. (11) shows that the pointwise output has infidelity at most

$$C(h^2 + \Delta/h^2)$$

with $|\varphi\rangle$, apart from a finite-timeout error.

All resources can be made finite directly. Since S is the trial normalization, take

$$N := \lceil e^{2h} S \log(2/\xi) \rceil.$$

By the union bound, the probability that either party aborts is at most $2(1 - e^{-2h}/S)^N \le 2\exp(-N e^{-2h}/S) \le \xi$. Take Ω to be the N-fold tensor power of the normalized trial state. On timeout, each party outputs a fixed local basis state. These prescriptions define local completely positive trace-preserving channels. The common $k_{\max}$, the trial state, and N depend only on the finite family and the global parameters, not on the realized z. The local filters and label-erasing isometries depend only on the corresponding local description.

Averaging and trace-distance conversion. Finally, write $f_\varphi := 1 - \langle \varphi | \rho | \varphi \rangle$. The triangle inequality for purified distance gives

$$1 - \langle \psi | \rho | \psi \rangle \le 2 f_\varphi + 2(1 - |\langle \psi | \varphi \rangle|^2),$$

and $1 - |\langle \psi | \varphi \rangle|^2 \le \Delta$. Hence changing the target from $|\varphi\rangle$ to $|\psi\rangle$ adds at most $C\Delta$ to the infidelity. Averaging the pointwise estimate gives $\mathbb{E}_{z \sim \nu}[1 - \langle \psi_z | \rho_z | \psi_z \rangle] \le C_{\mathrm{sm}}(h^2 + D/h^2 + D) + 2\xi$. Choosing $\xi := \zeta/2$ proves Eq. (3). To prove Eq. (4), set $\zeta_0 := \zeta^2/16$ and rerun the construction with $\xi := \zeta_0/2$.

If $0 < D < 1/16$, take $h := D^{1/4}$, which gives average infidelity $O(D^{1/2}) + \zeta_0$. If $D = 0$, choose h so that $C_{\mathrm{sm}} h^2 \le \zeta_0$; the average infidelity is then at most $2\zeta_0$, so the average trace distance is at most $2\sqrt{2\zeta_0} < \zeta$. The case $D \ge 1/16$ was handled at the outset. The pure-target bound

$$\|\rho - |\psi\rangle\langle\psi|\|_1 \le 2\sqrt{1 - \langle\psi|\rho|\psi\rangle}$$

and Jensen's inequality give Eq. (4), since $2\sqrt{\zeta_0} \le \zeta/2$; the zero-$D$ and large-D cases were handled above. Stinespring dilation turns the local channels into local unitaries with private environments, so the construction is a valid finite-dimensional tensor-product strategy. $\square$

Remark 4.2. Although Lemma 4.1 is stated for a distribution, its proof gives a pointwise infidelity bound for each pair once the common finite label range is fixed. The distributional formulation is the one needed below: the shared state and the global parameters depend on the finite family, but not on the realized questions or state descriptions.

4.2 Postselection-stable single-game rounding

Definition 4.3 (Postselection parameters)**.** *Fix a finite-dimensional strategy for the repeated game and a proper subset $D_0 \subsetneq [n]$ such that $p := \Pr[W_{D_0}] > 0$. Define*

$$m := n - |D_0|, \qquad q := \frac{1}{m} \sum_{i \in [n] \setminus D_0} \Pr[W_i \mid W_{D_0}], \qquad \eta := \frac{\log(1/p) + |D_0|\ell}{m}.$$

Postselected objects. Let Q and $|\Psi\rangle$ denote the objects from the shared setup in Subsection 3.3, and let $\mathsf{J}_A, \mathsf{J}_B, |\Psi^A\rangle, |\Psi^B\rangle$ denote the objects supplied by Lemma 3.3. Unless stated otherwise, all expectations below are taken with respect to Q. The candidate-pair consequence in Eq. (2) is

$$D := \mathbb{E}[\||\Psi^A\rangle - |\Psi^B\rangle\|^2], \qquad \text{hence } D \le 32\eta.$$

Generic acceptance probabilities. Here we make no assumption that the ideal conditioned experiment wins with probability close to one. The argument combines the ideal coordinate measurements from Lemma 3.2, the locally sampled histories from Lemma 3.3, and the trace-distance guarantee from Lemma 4.1. This gives a general $O(\eta^{1/4})$ rounding loss; the near-perfect argument below improves the dependence when q is close to one.

Lemma 4.4 (Improved postselection-stable rounding)**.** *There is a universal constant $B_4 \ge 1$ such that, in the postselected setup above, $0 \le \eta \le 1$ implies*

$$\omega^*(G) \ge q - B_4 \eta^{1/4}. \tag{12}$$

Proof. Lemma 3.2 supplies local coordinate measurements for which the ideal experiment has average winning probability q. Invoke Lemma 3.3 and apply Lemma 4.1 to its finite family of candidate pairs with $D \le 32\eta$. The shared resource and smoothing parameter are fixed from this finite family and the global parameters before the fresh questions arrive.

For a tuple in the support of Q, let $\rho_{i,r,x,y}$ be the prepared state and put

$$s(i, r, x, y) := \|\rho_{i,r,x,y} - |\Psi_{r,x,y}\rangle\langle\Psi_{r,x,y}|\|_1.$$

Eq. (4), the first estimate in Eq. (1), and Jensen's inequality give

$$\mathbb{E}_{\mathsf{Q}}[s] \le C_{\mathrm{sm}} D^{1/4} + \zeta + 2\mathbb{E}_{\mathsf{Q}}[\||\Psi^A\rangle - |\Psi\rangle\|] \le C_{\mathrm{sm}} D^{1/4} + \zeta + 4\sqrt{2\eta}.$$

We now describe the resulting legal strategy. On fresh questions $(x,y)\sim\mu$, the players use the finite shared randomness from Lemma 3.3 to choose a common coordinate i and local histories r_A, r_B. They apply the channels of Lemma 4.1 to $|\Psi^A_{r_A,x}\rangle$ and $|\Psi^B_{r_B,y}\rangle$, then apply the local coordinate POVMs from Lemma 3.2. On tuples outside the support of Q, fix arbitrary default descriptions and POVMs.

Let $w(i,r,x,y)\in[0,1]$ be the ideal branch winning probability, extended by zero outside the support of Q, and extend s by the value 2 there. Since the trace-norm distance between two states is at most 2, this gives $0\le s\le 2$ everywhere. Lemma 3.2 gives $\mathbb{E}_{\mathsf{Q}}[w]=q$. Since $d_{\mathrm{TV}}(\mathsf{Q},\mathsf{J}_A)\le\kappa$ and $\Pr[r_A\ne r_B]\le 4\kappa$, where $\kappa:=\sqrt{3\eta/2}$,

$$\mathbb{E}_{\mathsf{J}_A}[w]\ge q-\kappa,\qquad \mathbb{E}_{\mathsf{J}_A}[s]\le\mathbb{E}_{\mathsf{Q}}[s]+2\kappa.$$

On the event $r_A=r_B=:r$, a tuple in the support of Q uses the same history on both sides, and replacing the ideal state by a state at trace-norm distance s changes its winning probability by at most $s/2$. Outside the support of Q, the inequality is trivial because $w-s/2=-1$. Therefore

$$\mathrm{win}\ge\mathbb{E}_{\mathsf{J}_A}[w-s/2]-\Pr[r_A\ne r_B]\ge q-\kappa-\frac{1}{2}(\mathbb{E}_{\mathsf{Q}}[s]+2\kappa)-4\kappa=q-\frac{1}{2}\mathbb{E}_{\mathsf{Q}}[s]-6\kappa,$$

where the first step follows from the pointwise $s/2$ measurement comparison on the event $r_A=r_B$ and charging loss one when $r_A\ne r_B$, the second step follows from $\mathbb{E}_{\mathsf{J}_A}[w]\ge q-\kappa$, $\mathbb{E}_{\mathsf{J}_A}[s]\le\mathbb{E}_{\mathsf{Q}}[s]+2\kappa$, and $\Pr[r_A\ne r_B]\le 4\kappa$, the third step follows from collecting the κ terms.

Using $D\le 32\eta$, $\kappa=\sqrt{3\eta/2}$, and $\sqrt{\eta}\le\eta^{1/4}$ for $0\le\eta\le 1$, we obtain

$$\mathrm{win}\ge q-B_4\eta^{1/4}-\frac{\zeta}{2}$$

after increasing the universal constant B_4. Every positive ζ produces a finite-dimensional strategy, so the definition of $\omega^*(G)$ and the limit $\zeta\to 0$ give Eq. (12). □

Near-perfect acceptance. The preceding lemma treats an arbitrary acceptance probability. In the repetition argument the ideal conditional experiment already wins with probability close to one. Keeping squared fidelity until the final measurement gives a stronger estimate in that regime.

Lemma 4.5 (Near-perfect rounding). *There is a universal constant $B_{\mathrm{np}}\ge 1$ such that, whenever $0\le\eta\le 1$ in the postselected single-game experiment,*

$$\omega^*(G)\ge q-B_{\mathrm{np}}((1-q)^{1/2}\eta^{1/4}+\eta^{1/2}).\tag{13}$$

Proof. Let B_{np} be a universal constant to be fixed at the end, with $B_{\mathrm{np}}\ge\sqrt{512}$. If $\eta\ge 1/512$, then, since $q\le 1$,

$$q-B_{\mathrm{np}}((1-q)^{1/2}\eta^{1/4}+\eta^{1/2})\le q-B_{\mathrm{np}}\eta^{1/2}\le 0\le\omega^*(G).$$

Thus the claim is immediate in this case. Assume for now that $0<\eta<1/512$.

Apply the infidelity part of Lemma 4.1 to the candidate pairs supplied by Lemma 3.3, with $z=(i,r,x,y)\sim\mathsf{Q}$, $|\psi_z\rangle:=|\Psi^A_{r,x}\rangle$, and $|\varphi_z\rangle:=|\Psi^B_{r,y}\rangle$. Their mean squared distance is at most $D\le 32\eta$. Set $h:=(32\eta)^{1/4}\in(0,1/2)$, let ρ_z be the output state, and put

$$f_{\mathrm{cand}}(z):=1-\langle\Psi^A_{r,x}|\rho_z|\Psi^A_{r,x}\rangle.$$

Eq. (3) and $D \leq 32\eta$ give

$$\mathbb{E}_{\mathsf{Q}}[f_{\mathrm{cand}}] \leq C_{\mathrm{sm}}(h^2 + D/h^2 + D) + \zeta \leq C_0\eta^{1/2} + \zeta \tag{14}$$

for a universal constant C_0.

We next transfer the target from the locally described candidate state to the ideal conditioned state. Write

$$f(z) := 1 - \langle \Psi_{r,x,y}|\rho_z|\Psi_{r,x,y}\rangle$$

and write $\mathcal{P}$ for purified distance. The triangle inequality for $\mathcal{P}$, followed by the scalar inequality $(u+v)^2 \leq 2u^2 + 2v^2$, gives

$$\begin{aligned} f(z) &= \mathcal{P}(\rho_z, |\Psi_{r,x,y}\rangle\langle\Psi_{r,x,y}|)^2 \\ &\leq 2\mathcal{P}(\rho_z, |\Psi^A_{r,x}\rangle\langle\Psi^A_{r,x}|)^2 + 2\mathcal{P}(|\Psi^A_{r,x}\rangle, |\Psi_{r,x,y}\rangle)^2 \leq 2f_{\mathrm{cand}}(z) + 2\||\Psi^A_{r,x}\rangle - |\Psi_{r,x,y}\rangle\|^2. \end{aligned}$$

Therefore Eq. (1) and Eq. (14) imply

$$R := \mathbb{E}_{\mathsf{Q}}[f] \leq C_1\eta^{1/2} + 2\zeta \tag{15}$$

for a universal constant C_1.

Use the same legal strategy and the same default conventions outside the support of Q as in the proof of Lemma 4.4. For z in the support of Q, let L_z be the losing effect of the ideal coordinate POVMs and put

$$\alpha(z) := \langle \Psi_{r,x,y}|L_z|\Psi_{r,x,y}\rangle.$$

Extend α and f by the value 1 outside the support of Q. Then $0 \leq \alpha, f \leq 1$, and Lemma 3.2 gives $\mathbb{E}_{\mathsf{Q}}[\alpha] = 1 - q$.

We record the measurement estimate used in the hybrid. Suppose a pure state $|\Psi\rangle$ is rejected by an effect L with probability α, while a state ρ has infidelity $f := 1 - \langle\Psi|\rho|\Psi\rangle$. By Uhlmann's theorem [Uhl76, NC10], choose purifications $|\widetilde{\Psi}\rangle$ and $|\widetilde{\rho}\rangle$ satisfying $\||\widetilde{\Psi}\rangle - |\widetilde{\rho}\rangle\| \leq \sqrt{2f}$. Put $\widetilde{L} := \sqrt{L} \otimes I$. Since $0 \leq L \leq I$, we have $\|\widetilde{L}\| \leq 1$ and $\|\widetilde{L}|\widetilde{\Psi}\rangle\| = \sqrt{\alpha}$. Hence the triangle inequality gives

$$\sqrt{\mathrm{Tr}[L\rho]} = \|\widetilde{L}|\widetilde{\rho}\rangle\| \leq \|\widetilde{L}|\widetilde{\Psi}\rangle\| + \|\widetilde{L}(|\widetilde{\rho}\rangle - |\widetilde{\Psi}\rangle)\| \leq \sqrt{\alpha} + \sqrt{2f}.$$

After squaring, the rejection probability on ρ is at most $\alpha + C\sqrt{\alpha f} + Cf$ for a universal constant C.

We now apply this estimate under the actual history coupling. On $r_A = r_B =: r$, the players use the common tuple $z = (i, r, x, y)$; on $r_A \neq r_B$, charge rejection probability one. Consequently,

$$1 - \mathrm{win} \leq 4\kappa + \mathbb{E}_{\mathsf{J}_A}[\alpha + C\sqrt{\alpha f} + Cf]. \tag{16}$$

Set $A := \mathbb{E}_{\mathsf{J}_A}[\alpha]$ and $R' := \mathbb{E}_{\mathsf{J}_A}[f]$. Since $\alpha, f \in [0,1]$ and $d_{\mathrm{TV}}(\mathsf{Q}, \mathsf{J}_A) \leq \kappa$,

$$A \leq 1 - q + \kappa, \qquad R' \leq R + \kappa.$$

Put $u := 1 - q$ and $t := \eta^{1/2}$. Since $\kappa = \sqrt{3\eta/2}$ and $R \leq C_1\eta^{1/2} + 2\zeta$, there is a universal constant C_2 such that

$$A \leq u + C_2 t, \qquad R' \leq C_2 t + 2\zeta.$$

We may assume $0 < \zeta \leq 1$. Using $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$, we obtain

$$\sqrt{AR'} \leq \sqrt{(u + C_2 t)(C_2 t + 2\zeta)} \leq C_3(u^{1/2}\eta^{1/4} + \eta^{1/2} + \zeta^{1/2})$$

for a universal constant C_3. Cauchy–Schwarz in Eq. (16) now gives

$$1 - \mathrm{win} \le 4\kappa + A + C\sqrt{AR'} + CR' \le 1 - q + C_4((1-q)^{1/2}\eta^{1/4} + \eta^{1/2} + \zeta^{1/2}).$$

Every $0 < \zeta \le 1$ produces a finite-dimensional strategy. Sending $\zeta \to 0$ and using the definition of $\omega^*(G)$ therefore proves Eq. (13) for $0 < \eta < 1/512$, after choosing $B_{\mathrm{np}} \ge C_4$.

It remains to handle $\eta = 0$. In this case $D = 0$, $\kappa = 0$, and Eq. (1) forces both candidate-to-ideal mean squared errors to be zero. Apply Lemma 4.1 with arbitrary $h \in (0, 1/2)$ and $\zeta > 0$. The preceding argument then gives

$$1 - \mathrm{win} \le 1 - q + C\sqrt{(1-q)(C_{\mathrm{sm}}h^2 + 2\zeta)} + C(C_{\mathrm{sm}}h^2 + 2\zeta).$$

Sending first $\zeta \to 0$ and then $h \to 0$ yields $\omega^*(G) \ge q$, which is Eq. (13) when $\eta = 0$. □

4.3 Proof of the parallel-repetition theorem

We now combine Lemma 3.1 with the near-perfect rounding estimate in Lemma 4.5.

Theorem 4.6 (Improved parallel-repetition exponent, formal version of Theorem 1.1)**.** *There is a universal constant $c_3 > 0$ such that every finite two-player entangled game G with $\omega^*(G) = 1 - \epsilon$ satisfies*

$$\omega^*(G^{\otimes n}) \le \exp\left(-c_3 \frac{\epsilon^3}{\epsilon + \log(|A||B|)} n\right) \tag{17}$$

for every integer $n \ge 1$.

Proof. If $\epsilon = 0$, then the right-hand side of Eq. (17) is 1, so the claim follows from $\omega^*(G^{\otimes n}) \le 1$. If $\epsilon = 1$, then $\omega^*(G^{\otimes n}) \le \omega^*(G) = 0$: winning all coordinates implies winning any fixed coordinate, and marginalizing a repeated strategy gives a single-game strategy. Assume henceforth that $0 < \epsilon < 1$ and put

$$\ell := \log(|A||B|), \qquad \delta := \frac{\epsilon}{4}, \qquad K := (8B_{\mathrm{np}})^{-4}, \qquad c_3 := \frac{K}{8}, \qquad \gamma := c_3 \frac{\epsilon^3}{\epsilon + \ell}.$$

Fix $n \ge 1$ and suppose for contradiction that Eq. (17) fails. By the definition of the entangled value as a supremum, there is a finite-dimensional repeated strategy whose probability ϑ of winning every coordinate satisfies

$$\vartheta > \exp(-\gamma n).$$

Since $B_{\mathrm{np}} \ge 1$, we have $K \le 1$ and $c_3 \le 1/8$. Together with $0 < \epsilon < 1$ and $\epsilon + \ell \ge \epsilon$, this gives

$$\gamma \le c_3 \epsilon^2 \le \epsilon/8 = \delta/2,$$

where the first step uses $\epsilon + \ell \ge \epsilon$ and the definition of γ, the second step uses $c_3 \le 1/8$ and $0 < \epsilon < 1$, the third step uses $\delta = \epsilon/4$.

The preceding strict bound on ϑ therefore implies

$$\log(1/\vartheta) < \gamma n \le \delta n/2 < \delta n,$$

where the first step uses $\vartheta > \exp(-\gamma n)$ and monotonicity of log, the second step uses $\gamma \le \delta/2$, the third step uses $\delta > 0$ and $n \ge 1$.

Lemma 3.1 supplies a proper set of conditioned coordinates $D_0 \subsetneq [n]$. Its size bound gives

$$|D_0| \le \frac{\log(1/\vartheta)}{\delta} < \frac{\gamma n}{\delta} \le \frac{n}{2},$$

where the first step uses the size bound in Lemma 3.1, the second step uses $\log(1/\vartheta) < \gamma n$ and $\delta > 0$, the third step uses $\gamma \le \delta/2$.

And hence $m := n - |D_0| > n/2$. The other conclusions of the lemma give

$$m > \frac{n}{2}, \qquad |D_0| < \frac{\gamma n}{\delta}, \qquad p := \Pr[W_{D_0}] \ge \vartheta, \qquad q \ge 1 - \frac{\epsilon}{4}. \tag{18}$$

The information cost of the resulting single-coordinate experiment is

$$\eta := \frac{\log(1/p) + |D_0|\ell}{m}.$$

Eq. (18) and $\vartheta > \exp(-\gamma n)$ give

$$\log(1/p) \le \log(1/\vartheta) < \gamma n, \qquad |D_0|\ell \le \frac{\gamma n \ell}{\delta}, \qquad m > \frac{n}{2}.$$

Using $\delta = \epsilon/4$ and $(\epsilon + 4\ell)/(\epsilon + \ell) \le 4$, we obtain

$$\eta < 2\gamma(1 + 4\ell/\epsilon) = 2c_3\epsilon^2(\epsilon + 4\ell)/(\epsilon + \ell) \le 8c_3\epsilon^2 = K\epsilon^2 \le 1,$$

where the first step uses $\log(1/p) < \gamma n$, $|D_0|\ell \le \gamma n\ell/\delta$, $m > n/2$, and $\delta = \epsilon/4$, the second step uses the definition of γ, the third step uses $(\epsilon+4\ell)/(\epsilon+\ell) \le 4$, the final step uses $K \le 1$ and $0 < \epsilon < 1$.

Since $q \ge 1 - \epsilon/4$ and $\eta < K\epsilon^2$, the loss in Eq. (13) is strictly less than

$$B_{\mathrm{np}}((\epsilon/4)^{1/2}(K\epsilon^2)^{1/4} + (K\epsilon^2)^{1/2}) \le (1/16 + 1/64)\epsilon < \epsilon/8.$$

where the first step uses $K^{1/4} = (8B_{\mathrm{np}})^{-1}$, $K^{1/2} = (8B_{\mathrm{np}})^{-2}$, and $B_{\mathrm{np}} \ge 1$, the second step uses $1/16 + 1/64 = 5/64 < 1/8$ and $\epsilon > 0$.

Indeed, $K^{1/4} = (8B_{\mathrm{np}})^{-1}$ makes the first contribution $\epsilon/16$, while the second is $\epsilon/(64B_{\mathrm{np}}) \le \epsilon/64$. Lemma 4.5 therefore gives

$$\omega^*(G) > 1 - 3\epsilon/8.$$

By the definition of $\omega^*(G)$ as a supremum over finite-dimensional strategies, there is an actual single-game strategy with winning probability greater than $1 - 3\epsilon/8$. This is strictly larger than $1-\epsilon = \omega^*(G)$, a contradiction. Therefore Eq. (17) holds for the fixed n, and hence for all $n \ge 1$. □

References

[BVY17] Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Hardness amplification for entangled games via anchoring. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, pages 303–316, 2017.

[CHTW04] Richard Cleve, Peter Hoyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Conference on Computational Complexity*, pages 236–249, 2004.

[CS14] Andr'e Chailloux and Giannicola Scarpa. Parallel repetition of entangled games with exponential decay via the superposed information cost. In *Automata, Languages, and Programming*, volume 8572 of *Lecture Notes in Computer Science*, pages 296–307, 2014.

[CS15] Andr'e Chailloux and Giannicola Scarpa. Parallel repetition of free entangled games: Simplification and improvements. arXiv:1410.4397, 2015.

[CSUU08] Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, 17:282–299, 2008.

[DSV15] Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. *Computational Complexity*, 24(2):201–254, 2015.

[Hol09] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5:141–172, 2009.

[JPY14] Rahul Jain, Attila Pereszl'enyi, and Penghui Yao. A parallel repetition theorem for entangled two-player one-round games under product distributions. In *Proceedings of the 29th IEEE Conference on Computational Complexity*, pages 209–216, 2014.

[Kim14] Isaac H. Kim. Modulus of convexity for operator convex functions. *Journal of Mathematical Physics*, 55(8):082201, 2014.

[KR10] Julia Kempe and Oded Regev. No strong parallel repetition with entangled and non-signaling provers. In *Proceedings of the 25th IEEE Conference on Computational Complexity*, pages 7–15, 2010.

[KRT10] Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. *SIAM Journal on Computing*, 39(7):3207–3229, 2010.

[KV11] Julia Kempe and Thomas Vidick. Parallel repetition of entangled games. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 353–362, 2011.

[Las19] Dariusz Lasecki. Noisy embezzlement of entanglement and applications to entanglement dilution. Master's thesis, University of Waterloo, 2019.

[LYM+19] Yeong-Cherng Liang, Yu-Hao Yeh, Paulo E. M. F. Mendonça, Run Yan Teh, Margaret D. Reid, and Peter D. Drummond. Quantum fidelity measures for mixed states. *Reports on Progress in Physics*, 82(7):076001, 2019.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition, 2010.

[Ope26] OpenAI. Ten advances in mathematics and theoretical computer science. Technical report, 2026.

[Pet07] Dénes Petz. Bregman divergence as relative operator entropy. *Acta Mathematica Hungarica*, 116:127–131, 2007.

[Rao11] Anup Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998.

[Raz11] Ran Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing*, 40(3):771–777, 2011.

[Uhl76] Armin Uhlmann. The "transition probability" in the state space of a $*$-algebra. *Reports on Mathematical Physics*, 9(2):273–279, 1976.

[Yue16] Henry Yuen. A parallel repetition theorem for all entangled games. arXiv:1604.04340, 2016.

Acknowledgments

The AI tools used in preparing this draft were Gemini 3.1 Pro, Codex 5.6, and Claude Code Fable 5. After the author provided Codex with the OpenAI report, Codex proposed a plan for improving the dependence from ϵ^{13} to ϵ^9. After carefully studying the resulting ϵ^9 proof, the author independently discovered how to improve it further to ϵ^3. In addition, Gemini 3.1 Pro, Codex and Claude Code were used to improve the paper's English and grammar. All the proofs in this paper have been carefully verified by the author.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il