A Parallel Repetition Theorem for All Entangled Games*

Henry Yuen†

Massachusetts Institute of Technology, Cambridge, MA, USA
hyuen@mit.edu

Abstract

The behavior of games repeated in parallel, when played with quantumly entangled players, has received much attention in recent years. Quantum analogues of Raz's classical parallel repetition theorem have been proved for many special classes of games. However, for general entangled games no parallel repetition theorem was known.

We prove that the entangled value of a two-player game G repeated n times in parallel is at most $c_G n^{-1/4} \log n$ for a constant c_G depending on G, provided that the entangled value of G is less than 1. In particular, this gives the first proof that the entangled value of a parallel repeated game must converge to 0 for all games whose entangled value is less than 1. Central to our proof is a combination of both classical and quantum correlated sampling.

1998 ACM Subject Classification F.1.2. Modes of Computation

Keywords and phrases parallel repetition, direct product theorems, entangled games, quantum games

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.77

1 Introduction

A two-player one-round game G is played between a referee and two isolated players (who we will call Alice and Bob), who communicate only with the referee and not between themselves. The referee first samples a question pair (x, y) from some distribution μ and sends x to Alice and y to Bob. Alice and Bob respond with answers a and b respectively, and they win if $V(x, y, a, b) = 1$ for some predicate V.

The maximum winning probability of Alice and Bob in a game G is a quantity that depends on what resources they are allowed to use. If their answers are a deterministic function of their received question (and perhaps some public random string), then we call their maximum winning probability the *classical value* of G, denoted by $\mathrm{val}(G)$. However quantum mechanics allows Alice and Bob to share a resource called *entanglement*, which gives rise to correlations that cannot be reproduced with public randomness only. When Alice and Bob make use of entanglement to play a game G, we call their maximum winning probability the *entangled value* of G, denoted by $\mathrm{val}^*(G)$. For all games, the classical value is at most the entangled value. Cast in the language of games, the famous Bell's Theorem states that there exist games G where those values are different: $\mathrm{val}^*(G) > \mathrm{val}(G)$ [3].

The Parallel Repetition Question is the following natural and basic question: given a game G with value less than 1, what is the value of the game G^n, wherein Alice and Bob

* The full version of this paper can be found on the arXiv at http://arxiv.org/abs/1604.04340.
† This work was supported by Simons Foundation grant 360893 and National Science Foundation Grant 1218547.

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).
Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;
Article No. 77; pp. 77:1–77:13


Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

play n independent instances of G played *in parallel*? More formally, in the game G^n, the referee samples n independent question pairs $(x_1, y_1), \ldots, (x_n, y_n)$ from μ, and sends $(x_1, \ldots, x_n)$ to Alice, and sends $(y_1, \ldots, y_n)$ to Bob. Alice responds with answer tuple $(a_1, \ldots, a_n)$, Bob responds with $(b_1, \ldots, b_n)$, and the players win if for all coordinates $i \in [n]$, $V(x_i, y_i, a_i, b_i) = 1$.

The difficulty in relating $\mathrm{val}(G^n)$ with $\mathrm{val}(G)$ and n is that even though each of the n instances of G in G^n are independent, Alice and Bob need not play each instance independently. For example, since Alice receives $(x_1, \ldots, x_n)$ all at once, she can use some question x_j to answer the i'th game, and Bob can do something similar. Because of such strategies, for every k there are games G such that $\mathrm{val}(G^k) = \mathrm{val}(G) < 1$. This shows that the naive expectation that $\mathrm{val}(G^n) = \mathrm{val}(G)^n$ is false.

The naive expectation is not too far from the truth, however: Raz's Parallel Repetition Theorem [19] states that

$$\mathrm{val}(G^n) \leq (1 - (1 - \mathrm{val}(G))^3)^{c_G n},$$

where c_G is a constant depending on G. In particular, as n goes to infinity, the classical success probability goes to 0 exponentially fast in n (provided that $\mathrm{val}(G) < 1$). The proof is highly nontrivial, although it has been simplified and improved upon in recent years [12, 4]. Raz's Parallel Repetition Theorem has heavily influenced complexity theory, most notably in the areas of hardness of approximation [11] and communication complexity [13, 5].

One open question, which we call the Quantum Parallel Repetition Conjecture, asks whether an analogue of Raz's Parallel Repetition Theorem holds in the setting of entangled players. The Quantum Parallel Repetition Conjecture has been resolved for many special cases of games, including free games [6, 14, 7], projection games [9], XOR games [8], unique games [15], anchored games [1], and fortified games [2]. However, the general case has remained elusive. Not only do we not know of a quantum analogue of Raz's Parallel Repetition Theorem, it hasn't even been shown that if $\mathrm{val}^*(G) < 1$, then $\mathrm{val}^*(G^n)$ goes to 0 as n goes to infinity! Could quantum entanglement allow players to counteract the value-decreasing effect of parallel repetition?

In this paper we prove that for all nontrivial entangled games G (i.e. $\mathrm{val}^*(G) < 1$), the entangled value of G^n must converge to 0. This resolves a weaker version of the Quantum Parallel Repetition Conjecture for general games. Quantitatively, our result is the following:

▶ **Theorem 1** (Main Theorem). *Let G be a game involving two entangled players with $\mathrm{val}^*(G) = 1 - \varepsilon$. Then for all integer $n > 0$,*

$$\mathrm{val}^*(G^n) \leq c \cdot \frac{s_G \log n}{\varepsilon^{17} n^{1/4}}$$

where c is a universal constant and s_G is the bit-length of the players' answers in G.

This shows that the entangled value of G^n must decay at a polynomial rate with n. The full Quantum Parallel Repetition Conjecture states that the rate of decay is in fact exponential, and this remains an important open problem.

1.1 Previous work

There has been extensive work on the parallel repetition of entangled games. As stated earlier, past results have applied to various special classes of games, but there was no result that covered *all* games.

The results coming closest to the Quantum Parallel Repetition Conjecture are the work of Kempe and Vidick [16] and Bavarian, Vidick, and Yuen [1, 2]. Rather than proving parallel repetition theorems for general games, these works prove general *gap amplification* theorems, which are closely related. Instead of showing that for games G where $\mathrm{val}^*(G) < 1$ that $\mathrm{val}^*(G^n)$ goes to 0 with n, the game G is first converted to another game H where analyzing $\mathrm{val}^*(H^n)$ is much more tractable. Gap amplification is a technique used in complexity theory and cryptography to amplify the difference between two cases of a problem (usually called the *completeness* and *soundness* cases).

Kempe and Vidick showed that given an arbitrary game G, one can efficiently transform it to another game H with the following properties: if the *classical* value of G is 1 (meaning that there is a perfect deterministic strategy), then $\mathrm{val}(H^n) = 1$ (and thus $\mathrm{val}^*(H^n) = 1$). If the *entangled* value of G is less than 1, then the entangled value of H^n decays at a polynomial rate $n^{-\Omega(1)}$. In this tranformed game H, in addition to playing the game G, the referee will randomly choose to ask "consistency" questions to check that the players give the same answers on the same questions[1]. Thus [16] prove gap amplification for general games – with a caveat. Because of the random consistency checks in the game H, the "quantum completeness" is not preserved: even if $\mathrm{val}^*(G) = 1$, it is not necessarily the case that $\mathrm{val}^*(H) = 1$.

More recently, Bavarian, Vidick, and Yuen [1, 2] gave better gap amplification results for entangled games[2]. They showed that for general games G, one can apply a simple transformation to obtain another game H with the following properties:

1. If $\mathrm{val}^*(G) = 1$, then $\mathrm{val}^*(H^n) = 1$.
2. If $\mathrm{val}^*(G) < 1$, then $\mathrm{val}^*(H^n) \leq \exp(-\Omega(n))$.

Note that the transformation from G to H preserves quantum completeness, and that when $\mathrm{val}^*(G) < 1$, the entangled value of the repeated game decays *exponentially*. Like [16], the transformations of [1, 2] construct H by adding auxiliary questions to the game G. The transformation given in [1] is called *anchoring*, and the trasformation in [2] is called *fortification*. The latter transformation gives a quantum generalization of the fortification technique of [17] for *classical games*. The quantitative aspects of repeated anchored games are different from those of fortified games, but both yield general gap amplification theorems for entangled games.

The results of Bavarian, Vidick and Yuen show that, while we do not know if the Quantum Parallel Repetition Conjecture holds for all games G, we *do* know that it holds for a class of games that effectively captures the general case, in fact with exponential decay similar to Raz's theorem. Since the main application of parallel repetition in complexity theory and quantum information is gap amplification, the results of [1, 2] effectively settle the Quantum Parallel Repetition Conjecture – as far as applications are concerned.

But as a scientific question, the original Quantum Parallel Repetition Conjecture is a fundamental and basic problem about the power of entanglement in games. Prior to this work, one might have wondered whether there exists a game G such that $\mathrm{val}^*(G) < 1$, but there is some constant δ such that for infinitely many n there is a nefarious entangled strategy for G^n with success probability at least δ? Here we prove that this cannot happen.

1.2 Proof overview

Theorem 1 is proved via reduction: if $\mathrm{val}^*(G^n)$ is too large, then from an optimal entangled strategy for G^n we can construct an entangled strategy for the single-shot game G that wins

[1] This transformation is to due to Feige and Kilian [10], who proved a similar result for classical games.
[2] They also obtain general gap amplification results for games with more than two players.

with probability strictly greater than $\mathrm{val}^*(G)$, which would be a contradiction.

In more detail, suppose that $\mathrm{val}^*(G) = 1 - \varepsilon$. If the success probability of the players in G^n is dramatically larger than our target bound (which in our case is $\sim n^{-O(1)}$), then we can identify a set of coordinates $C \subseteq [n]$ that is not too large, but has the property that for a uniformly random coordinate $i \in [n] - C$,

$$\Pr(\text{Win game } i \mid \text{Win games in } C) > 1 - \varepsilon/2 \tag{1}$$

where here the probability is both over the randomness of the questions in G^n, the randomness of the players' entangled strategy, and the randomly chosen index i. Thus it would be advantageous if Alice and Bob could play the single-shot game G by "embedding" it in a randomly chosen ith coordinate of G^n, and playing G^n *conditioned* on the event that the games indexed by C have been won. If they could do this, then by (1), the probability they win the ith coordinate of G^n, and hence the original game G, is at least $1 - \varepsilon/2 > \mathrm{val}^*(G)$, which would be a contradiction.

If the players are classical (i.e. use deterministic strategies), this embedding is performed in the following way. Alice and Bob are first given questions (X_i, Y_i) for the i'th game. Based on their received question, Alice and Bob jointly sample a *dependency-breaking variable* R. The essential features of this dependency-breaking variable are:

1. **Usefulness:**[3] $\mathsf{P}_{A_iB_i|RX_iY_iW_C} = \mathsf{P}_{A_i|RX_iW_C} \cdot \mathsf{P}_{B_i|RY_iW_C}$
2. **Sampleability:** $\mathsf{P}_{R|X_iY_iW} \approx \mathsf{P}_{R|X_iW_C} \approx \mathsf{P}_{R|Y_iW_C}$

where "$\approx$" means closeness in statistical distance. Here, W_C denotes the event that the players win all the games in C. $\mathsf{P}_{A_iB_i|RX_iY_iW_C}$ denotes the probability distribution of Alice's and Bob's answers in the ith coordinate when playing G^n, conditioned on the dependency-breaking variable R, their received questions for the ith game (X_i, Y_i), *and* the event W_C. The "Usefulness property" states that, the players' answers in the ith round are independent of each other, conditioned on R, their own questions, and W_C. Thus, given R distributed according to $\mathsf{P}_{R|X_iY_iW_C}$, Alice can sample A_i on her own, because she possesses R and X_i, and similarly Bob can sample B_i on his own, because he possesses knowledge of R and Y_i. By (1), the probability that $V(X_i, Y_i, A_i, B_i) = 1$ will be strictly greater than $\mathrm{val}^*(G)$, wherein we would arrive at a contradiction.

As the name suggests, the "sampleability property" implies that Alice and Bob can (approximately) jointly sample the variable R. Even though the distribution $\mathsf{P}_{R|X_iY_iW_C}$ may depend on both players' questions, the sampleability property shows R, up to some error, only depends on X_i or Y_i, but not both. Using the *correlated sampling procedure* of [12], Alice and Bob can jointly sample R from $\mathsf{P}_{R|X_iY_iW}$ with high probability.

At a high level, the proof of our quantum parallel repetition theorem is similar. However instead of sampling a dependency-breaking variable R, the players will need to sample a *dependency-breaking state*. It is an entangled state $|\Psi_{x_iy_i}\rangle$ that depends on both Alice's and Bob's questions (x_i, y_i), and satisfies similar Usefulness and Sampleability properties:

1. **Usefulness:** The distribution of measurement outcomes by making local measurements on $|\Psi_{X_iY_i}\rangle$ is equal to $\mathsf{P}_{A_iB_i|X_iY_iW_C}$.
2. **Sampleability:** There exist states $|\Phi_{X_i}\rangle$ and $|\Gamma_{Y_i}\rangle$ such that $|\Psi_{X_iY_i}\rangle \approx |\Phi_{X_i}\rangle \approx |\Gamma_{Y_i}\rangle$

where "$\approx$" means closeness in ℓ_2 distance, and the statements hold on average over X_iY_i.

[3] We will let P denote the probability distribution that describes the joint distribution of the random variables relevant in an execution of the strategy for G^n, including the players' questions $X_1, \ldots, X_n, Y_1, \ldots, Y_n$, the players' answers $A_1, \ldots, A_n, B_1, \ldots, B_n$, and the dependency-breaking variable R.

The Usefulness property states that if on input (x_i, y_i), Alice and Bob were to share the entangled state $|\Psi_{x_i y_i}\rangle$, then they could make local measurements to obtain outcomes distributed according to $\mathsf{P}_{A_i B_i | X_i Y_i W_C}$, which would mean that their success probability would be $\Pr(\text{Win } i \mid \text{Win } C)$, which is greater than $\text{val}^*(G)$, an impossibility.

The Sampleability property implies that on input (x_i, y_i) Alice and Bob are actually able to approximately prepare the state $|\Psi_{x_i y_i}\rangle$. This is because of the *quantum correlated sampling procedure* of Dinur, Steurer, and Vidick, who used it to prove a parallel repetition theorem for entangled projection games [9]. It is entirely analogous to Holenstein's correlated sampling procedure: Alice has a description of a state $|\Phi_{X_i}\rangle$ that's close to $|\Psi_{X_i Y_i}\rangle$, and Bob has a description of a state $|\Gamma_{Y_i}\rangle$ that is also close to $|\Psi_{X_i Y_i}\rangle$. Via local transformations on preshared quantum entanglement, Alice and Bob can generate an approximation of $|\Psi_{X_i Y_i}\rangle$. Combined with the Usefulness property, Alice and Bob are then able to win the ith game with too high probability.

It is not difficult to define states that satisfy the Usefulness property. Consider an execution of the entangled strategy for G^n. In the beginning, the players share some entangled state $|\psi\rangle$, and upon obtaining questions $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$, the players apply local measurements depending on these questions to $|\psi\rangle$ to obtain answer tuples $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$. One can define an ensemble of states $\{|\Psi_{x_i, y_i}\rangle\}$ that are, roughly speaking, derived from the post-measurement state of the players *conditioned* on the players having won all the games in C (that is, conditioned on the event W_C), *and* having received a specific question pair (x_i, y_i) in the i'th coordinate. Such an ensemble of states would satisfy the Usefulness property.

However, the primary challenge is achieving Sampleability property, that is, to show the states $|\Psi_{x_i, y_i}\rangle$ only depend on one player's question, but not both. One major obstacle to proving the Sampleability property is the following: in the players' strategy for G^n, Bob (say) may elect to "print" his entire vector of questions $(y_1, \ldots, y_n)$ into the entangled state $|\psi\rangle$. He can do this by applying a local unitary operation controlled on his questions on some ancilla qubits in $|\psi\rangle$. We cannot say he does not do this, because the shared entangled state $|\psi\rangle$ and the players' measurements are completely arbitrary. But this implies that we cannot hope to prove that the post-measurement state is independent of y_i, conditioned on x_i.

Despite such barriers, we are able to define the $|\Psi_{x_i, y_i}\rangle$ in such a way that removes such adversarial dependencies on the players' questions. Assuming (for contradiction) that the players' probability of success is at least $n^{-O(1)}$, then we are able to prove that these states satisfy the Sampleability property. We build upon many previous works: we use the information theoretic framework of [6, 14], carefully combined with the operator analysis techniques from [9]. The definition of the dependency-breaking states $|\Psi_{x_i, y_i}\rangle$ includes the classical dependency-breaking variables of [12] used to prove Raz's parallel repetition theorem. Our final constructed strategy for the single-shot game G uses both classical and quantum correlated sampling procedures.

2 Preliminaries

2.1 Probability distributions

We largely adopt the notational conventions from [12] for probability distributions. We let capital letters denote random variables and lower case letters denote specific samples. We will use subscripted sets to denote tuples, e.g., $X_{[n]} := (X_1, \ldots, X_n)$, $x_{[n]} = (x_1, \ldots, x_n)$, and if $C \subset [n]$ is some subset then X_C will denote the sub-tuple of $X_{[n]}$ indexed by C. We use P_X to denote the probability distribution of random variable X, and $\mathsf{P}_X(x)$ to

denote the probability that $X = x$ for some value x. For multiple random variables, e.g., X, Y, Z, $\mathsf{P}_{XYZ}(x, y, z)$ denotes their joint distribution with respect to some probability space understood from context.

We use $\mathsf{P}_{Y|X=x}(y)$ to denote the conditional distribution $\mathsf{P}_{YX}(y, x)/\mathsf{P}_X(x)$, which is defined when $\mathsf{P}_X(x) > 0$. When conditioning on many variables, we usually use the shorthand $\mathsf{P}_{X|y,z}$ to denote the distribution $\mathsf{P}_{X|Y=y,Z=z}$. For example, we write $\mathsf{P}_{V|\omega_{-i},x_i,y_i}$ to denote $\mathsf{P}_{V|\Omega_{-i}=\omega_{-i},X_i=x_i,Y_i=y_i}$. For an event W we let $\mathsf{P}_{XY|W}$ denote the distribution conditioned on W. We use the notation $\mathbb{E}_X f(x)$ and $\mathbb{E}_{\mathsf{P}_X} f(x)$ to denote the expectation $\sum_x \mathsf{P}_X(x) f(x)$.

Let P_{X_0} be a distribution of $\mathcal{X}$, and for every x in the support of P_{X_0}, let $\mathsf{P}_{Y|X_1=x}$ be a conditional distribution defined over $\mathcal{Y}$. We define the distribution $\mathsf{P}_{X_0}\mathsf{P}_{Y|X_1}$ over $\mathcal{X} \times \mathcal{Y}$ as

$$(\mathsf{P}_{X_0}\mathsf{P}_{Y|X_1})(x, y) := \mathsf{P}_{X_0}(x) \cdot \mathsf{P}_{Y|X_1=x}(y).$$

Additionally, we write $\mathsf{P}_{X_0Z}\mathsf{P}_{Y|X_1}$ to denote the distribution $(\mathsf{P}_{X_0Z}\mathsf{P}_{Y|X_1})(x, z, y) := \mathsf{P}_{X_0Z}(x, z) \cdot \mathsf{P}_{Y|X_1=x}(y)$.

For two random variables X_0 and X_1 over the same set $\mathcal{X}$, we use

$$\|\mathsf{P}_{X_0} - \mathsf{P}_{X_1}\| := \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathsf{P}_{X_0}(x) - \mathsf{P}_{X_1}(x)|,$$

to denote the total variation distance between P_{X_0} and P_{X_1}.

2.2 Quantum information theory

For comprehensive references on quantum information we refer the reader to [18, 21].

For a vector $|\psi\rangle$, we use $\||\psi\rangle\|$ to denote its Euclidean length. For a matrix A, we will use $\|A\|_1$ to denote its *trace norm* $\mathrm{Tr}(\sqrt{AA^\dagger})$, and $\|A\|_F$ to denote its *Frobenius norm* $\sqrt{\mathrm{Tr}(AA^\dagger)}$. A density matrix is a positive semidefinite matrix with trace 1. The *fidelity* between two density matrices ρ and σ is defined as $F(\rho, \sigma) = \|\sqrt{\rho}\sqrt{\sigma}\|_1$. For Hermitian matrices A, B we write $A \preceq B$ to indicate that $A - B$ is positive semidefinite. We use $\mathbb{I}$ to denote the identity matrix. A *positive operator valued measurement* (POVM) with outcome set $\mathcal{A}$ is a set of positive semidefinite matrices $\{E^a\}$ labeled by $a \in \mathcal{A}$ that sum to the identity.

We will use the convention that, when $|\psi\rangle$ is a pure state, ψ refers to the rank-1 density matrix $|\psi\rangle\langle\psi|$. We use subscripts to denote system labels; so ρ_{AB} will denote the density matrix on the systems A and B. A *classical-quantum* state ρ_{XE} is classical on X and quantum on E if it can be written as $\rho_{XE} = \sum_x p(x)|x\rangle\langle x|_X \otimes \rho_{E|X=x}$ for some probability measure $p(\cdot)$. The state $\rho_{E|X=x}$ is by definition the E part of the state ρ_{XE}, conditioned on the classical register $X = x$. We write $\rho_{XE|X=x}$ to denote the state $|x\rangle\langle x|_X \otimes \rho_{E|X=x}$. We often write expressions such as $\rho_{E|x}$ as shorthand for $\rho_{E|X=x}$ when it is clear from context which registers are being conditioned on. This will be useful when there are many classical variables to be conditioned on.

2.3 Classical and quantum correlated sampling

Correlated sampling is a key component of Holenstein's proof of the classical parallel repetition theorem.

▶ **Lemma 2** (Classical correlated sampling [12]). *Let* P *and* Q *be two probability distributions over a universe* $\mathcal{U}$ *such that* $\|\mathsf{P} - \mathsf{Q}\|_1 \leq \varepsilon < 1$*. Then there exists a zero communication two-player protocol using shared randomness where the first player outputs an element* $p \in \mathcal{U}$ *distributed according to* P*, the second player samples an element* $q \in \mathcal{U}$ *distributed according to* Q*, and with probability at least* $1 - O(\varepsilon)$*, the two elements are identical (i.e.* $p = q$*).*

We call the protocol in the Lemma above the *classical correlated sampling procedure*. The next lemma is the quantum extension of the correlated sampling lemma, proved by [9] in order to obtain a parallel repetition theorem for entangled projection games, a class of two-player games. Their lemma is a robust version of the quantum state embezzlement procedure of [20].

▶ **Lemma 3** (Quantum correlated sampling [9]). *Let d be an integer. Then there exists an integer d' and a collection of unitaries V_ψ, W_ψ acting on $\mathbb{C}^{dd'}$ for every state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$, such that the following holds: for any two states $|\varphi\rangle, |\theta\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$,*

$$\|\overline{V}_\varphi \otimes W_\theta |E_{dd'}\rangle - |\varphi\rangle|E_{d'}\rangle\| \leq O(\| \, |\varphi\rangle - |\theta\rangle \, \|^{1/6})$$

where $|E_d\rangle \propto \sum_{j=1}^d \frac{1}{\sqrt{j}}|j\rangle|j\rangle$ is the d-dimensional embezzlement state.

We shall call the protocol in the Lemma above the *quantum correlated sampling procedure*.

3 Proof of the Main Theorem

Let G be a two-player one-round game with question distribution μ and referee predicate $V(x, y, a, b)$. Let $\mathcal{A}$ and $\mathcal{B}$ denote the alphabets of Alice's and Bob's answers, respectively. Let $\mathrm{val}^*(G) = 1 - \varepsilon$.

Consider an optimal entangled strategy for G^n, which consists of a shared entangled state $|\psi\rangle^{E_A E_B} \in \mathbb{C}^d \otimes \mathbb{C}^d$ and measurement POVMs for Alice and Bob, $\{A^{a_{[n]}}_{x_{[n]}}\}$ and $\{B^{b_{[n]}}_{y_{[n]}}\}$ respectively. We will assume that $|\psi\rangle$ is symmetric; i.e., $|\psi\rangle = \sum_i \sqrt{\lambda_i}|v_i\rangle|v_i\rangle$ for some orthonormal basis $\{|v_i\rangle\}$. This is without loss of generality, as we can always rotate (say) Bob's basis vectors to match Alice's basis vectors, and fold the unitary rotation into Bob's measurements. For $i \in [n]$, let W_i denote the event that the players win coordinate i using this optimal strategy. Let $W = W_1 \wedge \cdots \wedge W_n$ denote the event that the players win all coordinates. For a set $C \subseteq [n]$, let $W_C = \wedge_{i \in C} W_i$.

▶ **Proposition 4.** *Suppose that $\log 1/\Pr(W) \leq \varepsilon n/16 - \log 4/\varepsilon$. Then there exists a set $C \subseteq [n]$ of size at most $t = \frac{8}{\varepsilon}(\log 4/\varepsilon + \log 1/\Pr(W))$ such that*

$$\Pr_{i \notin C}(W_i|W_C) \geq 1 - \varepsilon/2.$$

where i is chosen uniformly from $[n] - C$.

Proof. Set $\delta = \varepsilon/8$. Let $W_{>1-\delta}$ denote the event that the players won more than $(1-\delta)n$ rounds. To show existence of such a set C, we will show that $\mathbb{E}_C \Pr(\neg W_i|W_C) \leq \varepsilon/2$, where C is a (multi)set of t independently chosen indices in $[n]$. This implies that there exists a particular set C such that $\Pr(\neg W_i|W_C) \leq \varepsilon/2$, which concludes the claim.

First we write, for a fixed C, $\Pr(\neg W_i|W_C) = \Pr(\neg W_i|W_C, W_{>1-\delta}) \Pr(W_{>1-\delta}|W_C) + \Pr(\neg W_i|W_C, \neg W_{>1-\delta}) \Pr(\neg W_{>1-\delta}|W_C)$. Observe that $\Pr(\neg W_i|W_C \wedge W_{>1-\delta})$ is the probability that, conditioned on winning all rounds in C, the randomly selected coordinate $i \in [n]-C$ happens to be one of the (at most) δn lost rounds. This is at most $\delta n/(n-t) \leq \varepsilon/4$, where we use our assumption on t from the Proposition statement. Now observe that

$$\mathop{\mathbb{E}}_C \Pr(\neg W_{>1-\delta}|W_C) \leq \mathop{\mathbb{E}}_C \frac{\Pr(W_C|\neg W_{>1-\delta})}{\Pr(W_C)} \leq \frac{1}{\Pr(W)}(1-\delta)^t \leq \varepsilon/4$$

where in the second line we used the fact that $\Pr(W_C) \geq \Pr(W)$. ◀

For the rest of the proof we will fix a set C given by Proposition 4.

3.1 Dependency-breaking variables

We introduce the random variables that play an important role in the proof of Theorem 1. Let $C \subseteq [n]$ be as given by Proposition 4. We fix $C = \{m+1, m+2, \ldots, n\}$, where $m = n - |C|$, as this will easily be seen to hold without loss of generality. Let $(X_{[n]}, Y_{[n]})$ be distributed according to $\mu_{[n]}$ and $(A_{[n]}, B_{[n]})$ be defined from $X_{[n]}$ and $Y_{[n]}$ as follows:

$$\mathsf{P}_{A_{[n]}B_{[n]}|x_{[n]},y_{[n]}}(a_{[n]}, b_{[n]}) = \langle\psi| A^{a_{[n]}}_{x_{[n]}} \otimes B^{b_{[n]}}_{y_{[n]}} |\psi\rangle.$$

Let (X_C, Y_C) and $\mathrm{Z} = (A_C, B_C)$ be random variables that denote the players' questions and answers respectively associated with the coordinates indexed by C.

We use the random variables Ω and R that are crucially used in Holenstein's proof of Raz's parallel repetition theorem. Let $D_1, \ldots, D_m$ be independent and uniformly distributed in $\{Alice, Bob\}$. Let $M_1, \ldots, M_m$ be independent random variables defined in the following way: for each $i \in [m]$,

$$M_i = \begin{cases} X_i & \text{if } D_i = Alice \\ Y_i & \text{if } D_i = Bob \end{cases}$$

Now for $i \in [m]$, we define $\Omega_i := (D_i, M_i)$. We say that Ω_i *fixes Alice's input* if $D_i = Alice$, and otherwise Ω_i fixes Bob's input. We write Ω to denote the random variable $(\Omega_1, \ldots, \Omega_m, X_C, Y_C)$, where $X_C Y_C$ are Alice and Bob's questions in the coordinates indexed by C. For $i \in [m]$ we write Ω_{-i} to denote the random variable Ω with Ω_i omitted.

▶ **Proposition 5.** *Conditioned on Ω, $X_{[n]}$ and $Y_{[n]}$ are independent.*

Finally, we will define a *dependency-breaking variable* $R := (\Omega, A_C, B_C)$, where A_C and B_C are the players' answers in the coordinates indexed by C. For $i \notin C$, we let $R_{-i} := (\Omega_{-i}, A_C, B_C)$. R_i will refer to Ω_i. We will use lowercase letters to denote instantiations of these random variables: e.g., r_{-i}, x_i, and y_i refer to specific values of R_{-i}, X_i, and Y_i.

Throughout our proofs, all expectations are implicitly over the measure defined by P. For example, the expectation $\mathbb{E}_{\Omega_{-i}Z|x_i,y_i}$ indicates $\sum_{\omega_{-i},a_C,b_C} \mathsf{P}_{\Omega_{-i}A_CB_C|x_i,y_i}(\omega_{-i}, a_C, b_C)$. Given an event such as W (winning all the coordinates) or W_C (winning all the coordinates in C), $\mathsf{P}(W)$ and $\mathsf{P}(W_C)$ will mean the probability of these events with respect to the distribution P.

The following Lemma expresses the idea that, because W_C is an event that occurs with not-too-small probability, conditioning on it cannot skew the distribution of variables corresponding to an average coordinate by too much. This Lemma follows in a straightforward manner from the [12].

▶ **Lemma 6.** *The following statements hold on, average over i chosen uniformly in $[m]$:*
1. $\mathbb{E}_i \left\| \mathsf{P}_{R_i X_i Y_i | W_C} - \mathsf{P}_{R_i X_i Y_i} \right\|_1 \leq O(\sqrt{\delta})$
2. $\mathbb{E}_i \left\| \mathsf{P}_{X_i Y_i R_{-i} | W_C} - \mathsf{P}_{X_i Y_i} \cdot \mathsf{P}_{R_{-i}|X_i W_C} \right\|_1 \leq O(\sqrt{\delta})$
3. $\mathbb{E}_i \left\| \mathsf{P}_{X_i Y_i R_{-i} | W_C} - \mathsf{P}_{X_i Y_i} \cdot \mathsf{P}_{R_{-i}|Y_i W_C} \right\|_1 \leq O(\sqrt{\delta})$

where $\delta := \frac{1}{m}\left(\log 1/\mathsf{P}(W_C) + |C| \log |\mathcal{A}||\mathcal{B}|\right)$.

3.2 Two key Lemmas, and proof of the Main Theorem

For every $i \in [n] - C$, we will construct a collection of bipartite states $\{|\Psi_{r_{-i},x_i,y_i}\rangle\} \subseteq \mathbb{C}^d \otimes \mathbb{C}^d$, which we call dependency-breaking states, that are indexed by the dependency-breaking variable r_{-i} defined above, and questions (x_i, y_i). The following lemmas state the important properties of this collection of states:

▶ **Lemma 7** (Usefulness Lemma). *For all r_{-i}, x_i, y_i, there exist POVMs $\{\widehat{A}^{a_i}_{r_{-i},x_i}\}$ and $\{\widehat{B}^{b_i}_{r_{-i},y_i}\}$ acting on $\mathbb{C}^d$ such that*

$$\mathsf{P}_{A_iB_i|r_{-i},x_i,y_i}(a_i,b_i) = \mathrm{Tr}\left(\widehat{A}^{a_i}_{r_{-i},x_i} \otimes \widehat{B}^{b_i}_{r_{-i},y_i} \Psi_{r_{-i},x_i,y_i}\right).$$

▶ **Lemma 8** (Sampleability Lemma). *There exists an integer $d' \geq d$ such that for every i, r_{-i}, x_i, y_i, there exist local unitaries $U_{r_{-i},x_i}, V_{r_{-i},y_i}$ acting on $\mathbb{C}^{d'}$ such that*

$$\mathop{\mathbb{E}}_{i} \mathop{\mathbb{E}}_{X_iY_i} \left[\mathop{\mathbb{E}}_{R_{-i}|x_i,y_i,W_C} \left\| U_{r_{-i},x_i} \otimes V_{r_{-i},y_i} |E_{dd'}\rangle - |\Psi_{r_{-i},x_i,y_i}\rangle |E_{d'}\rangle \right\| \right] \leq O((\delta^{1/4}/\mathsf{P}(W_C))^{1/12})$$

where $|E_{dd'}\rangle$ and $|E_{d'}\rangle$ are dd' and d'-dimensional embezzlement states, respectively, and δ is defined to be $\frac{1}{m}(\log 1/\mathsf{P}(W_C) + |C| \log |\mathcal{A}||\mathcal{B}|)$.

Lemma 7 shows that the states $|\Psi_{r_{-i},x_i,y_i}\rangle$ are *useful* to have; they allow Alice and Bob to produce answers in the i'th coordinate whose statistics are consistent with the dependency-breaking variable r_{-i} and their inputs (x_i, y_i). Lemma 8 shows that these states are *locally generatable* by Alice and Bob, when given joint access to preshared entanglement, the dependency-breaking variable r_{-i} and their own inputs x_i and y_i respectively.

Using these two Lemmas we can prove the Main Theorem.

Proof of the Main Theorem. Consider the following strategy for the game G. Alice and Bob share beforehand the embezzlement state $|E_{dd'}\rangle$ of dimension dd' given by Lemma 8, and they also have access to shared randomness. Given inputs (x_i, y_i) distributed according to $\mathsf{P}_{X_iY_i} = \mu$:

1. Alice and Bob jointly sample a uniformly random $i \in [n] - C$.
2. Alice and Bob jointly, approximately sample R_{-i} from $\mathsf{P}_{R_{-i}|x_i,y_i,W_C}$ using the classical correlated sampling procedure.
3. Alice applies U_{r_{-i},x_i} to her side of $|E_{dd'}\rangle$
4. Bob applies V_{r_{-i},y_i} to his side of $|E_{dd'}\rangle$
5. Alice measures her side of the entanglement using $\{\widehat{A}^{a_i}_{r_{-i},x_i}\}$ and outputs the outcome a_i
6. Bob measures his side of the entanglement using $\{\widehat{B}^{b_i}_{r_{-i},y_i}\}$ and outputs the outcome b_i

We now analyze the success probability of this strategy. We will use $\widetilde{\mathsf{P}}$ to denote the distribution of variables in the probability space associated with an execution of this strategy. For example, we will write $\widetilde{\mathsf{P}}_{R_{-i}|X_iY_i}$ to denote the distribution of R_{-i} conditioned on X_iY_i that is sampled in Step 1. From Lemma 6 we have that on average over i, $\mathsf{P}_{X_iY_iR_{-i}|W_C} \approx \mathsf{P}_{X_iY_i} \cdot \mathsf{P}_{R_{-i}|X_iW_C} \approx \mathsf{P}_{X_iY_i} \cdot \mathsf{P}_{R_{-i}|Y_iW_C}$, where "$\approx$" means closeness in statistical distance. By invoking the classical correlated sampling procedure of Lemma 2, we get

$$\mathop{\mathbb{E}}_{i} \|\mathsf{P}_{X_iY_i} \cdot \widetilde{\mathsf{P}}_{R_{-i}|X_iY_i} - \mathsf{P}_{X_iY_iR_{-i}|W_C}\|_1 \leq O(\sqrt{\delta}).$$

After Step 3, Alice and Bob will possess a state $|\Lambda_{r_{-i},x_i,y_i}\rangle$ such that

$$\mathop{\mathbb{E}}_{i} \mathop{\mathbb{E}}_{X_iY_i} \left[\mathop{\mathbb{E}}_{R_{-i}|x_i,y_i,W_C} \|\Lambda_{r_{-i},x_i,y_i} - \Psi_{r_{-i},x_i,y_i}\|_1 \right] \leq \eta$$

where $\eta = O((\delta^{1/4}/\mathsf{P}(W_C))^{1/12})$. Consider the measurement process in Steps 4 and 5. Let $\widetilde{\mathsf{P}}_{A_iB_i|r_{-i},x_i,y_i}$ denote the distribution of measurement outcomes in this strategy, conditioned

on their inputs and a sampled value of r_{-i}. By Lemma 7 and the fact that the trace norm is nonincreasing under quantum operations, we have that

$$\mathop{\mathbb{E}}_{i} \mathop{\mathbb{E}}_{X_iY_i} \left[\mathop{\mathbb{E}}_{R_{-i}|x_i,y_i,W_C} \|\widetilde{\mathsf{P}}_{A_iB_i|x_i,y_i,r_{-i}} - \mathsf{P}_{A_iB_i|x_i,y_i,r_{-i}}\|_1 \right] \leq \eta$$

or equivalently

$$\mathop{\mathbb{E}}_{i} \|\mathsf{P}_{X_iY_i} \cdot \widetilde{\mathsf{P}}_{R_{-i}|X_iY_iW_C} \cdot \widetilde{\mathsf{P}}_{A_iB_i|x_i,y_i,r_{-i}} - \mathsf{P}_{X_iY_i} \cdot \mathsf{P}_{R_{-i}|X_iY_iW_C} \cdot \mathsf{P}_{A_iB_iR_{-i}|X_iY_iW_C}\|_1 \leq \eta.$$

By Lemma 6 we have $\mathbb{E}_i \|\mathsf{P}_{X_iY_i|W_C} - \mathsf{P}_{X_iY_i}\| \leq \sqrt{\delta}$. By triangle inequality and that $\widetilde{\mathsf{P}}_{X_iY_i} = \mathsf{P}_{X_iY_i}$, we have

$$\mathop{\mathbb{E}}_{i} \|\widetilde{\mathsf{P}}_{X_iY_iR_{-i}A_iB_i} - \mathsf{P}_{X_iY_iR_{-i}A_iB_i|W_C}\|_1 \leq O(\eta).$$

Note that $\widetilde{\mathsf{P}}_{X_iY_iR_{-i}A_iB_i}$ represents the probability distribution of all the variables present in the strategy above. Let W_i denote the probability the players win the ith coordinate. Thus we get

$$\mathop{\mathbb{E}}_{i} |\widetilde{\mathsf{P}}(W_i) - \mathsf{P}(W_i|W_C)| \leq O(\eta). \tag{2}$$

Assume that

$$\mathsf{P}(W) \geq \frac{cs \log n}{\varepsilon^{17} n^{1/4}}$$

where $c > 0$ is a universal constant, and s is the bit-length of the players' answers. Since $\mathsf{P}(W_C) \geq \mathsf{P}(W)$, and using our bound on $|C|$ (from Proposition 4) and our bound on δ (from Lemma 6), this implies that the right hand side of (2) is at most $\varepsilon/4$ (for an appropriate choice of c). This implies that

$$\mathop{\mathbb{E}}_{i} \widetilde{\mathsf{P}}(W_i) \geq \mathop{\mathbb{E}}_{i} \mathsf{P}(W_i|W_C) - \varepsilon/4 \geq 1 - \varepsilon/2 - \varepsilon/4 > \mathrm{val}^*(G)$$

where in the second line we used the bound from Proposition 4. However, this implies that there exists an i such that $\widetilde{\mathsf{P}}(W_i) > \mathrm{val}^*(G)$, which is a contradiction. Therefore $\mathsf{P}(W) \leq \frac{cs \log n}{\varepsilon^{17} n^{1/4}}$. ◀

Now we turn to defining the states and operators promised in the two key lemmas above, as well as giving an intuition for them.

3.3 Quantum states and operators

In this subsection we define the states $|\Psi_{r_{-i},x_i,y_i}\rangle$ and measurement operators $\{\widehat{A}^{a_i}_{r_{-i},x_i}\}$ and $\{\widehat{B}^{b_i}_{r_{-i},y_i}\}$. Recall that the dependency-breaking variable R consists of the set of fixed questions $\Omega = (X_C, Y_C, \Omega_1, \ldots, \Omega_m)$ and fixed answers $Z = (A_C, B_C)$ for the coordinates in C.

Coarse-grained measurements. We first *coarsen* the measurement POVMs $\{A^{a_{[n]}}_{x_{[n]}}\}$ and $\{B^{b_{[n]}}_{y_{[n]}}\}$ that constitute Alice and Bob's strategy in G^n to construct a set of *intermediate measurements*, which essentially produce answers for the games in set C, conditioned on a setting of Ω.

Fix i, ω, a_C, b_C, x_i, y_i. Define $A^{a_C}_{\omega_{-i},x_i} = \sum_{a_{[n]}|a_C} \mathbb{E}_{X_{[n]}|\omega_{-i},x_i} A^{a_{[n]}}_{x_{[n]}}$, and $B^{b_C}_{\omega_{-i},y_i} = \sum_{b_{[n]}|b_C} \mathbb{E}_{Y_{[n]}|\omega_{-i},y_i} B^{b_{[n]}}_{y_{[n]}}$ where $a_{[n]}|a_C$ (resp. $b_{[n]}|b_C$) indicates summing over all tuples $a_{[n]}$ consistent with the suffix a_C (resp. $b_{[n]}$ consistent with suffix b_C) and recall that $\mathbb{E}_{X_{[n]}|\omega_{-i},x_i}$ is shorthand for $\sum_{x_{[n]}} \mathsf{P}_{X_{[n]}|\Omega_{-i}=\omega_{-i},X_i=x_i}(x_{[n]})$. We also define $A^{a_C}_{\omega} = \mathbb{E}_{X_{[n]}|\omega} A^{a_C}_{x_{[n]}}$ and $B^{b_C}_{\omega} = \mathbb{E}_{Y_{[n]}|\omega} B^{b_C}_{y_{[n]}}$.

Let ρ denote the reduced density matrix of $|\psi\rangle$ on Alice's side. Since we have assumed that $|\psi\rangle$ is symmetric, ρ is also the reduced density matrix on Bob's side. For all i, ω, x_i, y_i, a_C, b_C, let U_{ω_{-i},x_i,a_C}, U_{ω,a_C}, V_{ω_{-i},y_i,b_C}, and V_{ω,b_C} be unitaries such that

$$U_{\omega_{-i},x_i,a_C}(A^{a_C}_{\omega_{-i},x_i})^{1/2}\sqrt{\rho} \qquad V_{\omega_{-i},y_i,b_C}(B^{b_C}_{\omega_{-i},y_i})^{1/2}\sqrt{\rho}$$
$$U_{\omega,a_C}(A^{a_C}_{\omega})^{1/2}\sqrt{\rho} \qquad V_{\omega,b_C}(B^{b_C}_{\omega})^{1/2}\sqrt{\rho}$$

are positive semidefinite. Such unitaries can be found via singular value decompositions. For notational convenience, let

$$S_{\omega_{-i},x_i,a_C} = U_{\omega_{-i},x_i,a_C}(A^{a_C}_{\omega_{-i},x_i})^{1/2} \qquad T_{\omega_{-i},y_i,b_C} = V_{\omega_{-i},y_i,b_C}(B^{b_C}_{\omega_{-i},y_i})^{1/2}$$
$$S_{\omega,a_C} = U_{\omega,a_C}(A^{a_C}_{\omega})^{1/2} \qquad T_{\omega,b_C} = V_{\omega,b_C}(B^{b_C}_{\omega})^{1/2}$$

Fine-grained measurements. Now we can define the *fine-grained measurements* that Alice and Bob can apply to obtain answers for the i'th game. Define

$$\widehat{A}^{a_i}_{r_{-i},x_i} = S^{-1}_{\omega_{-i},x_i,a_C} A^{a_C,a_i}_{\omega_{-i},x_i} S^{-1}_{\omega_{-i},x_i,a_C} \qquad \widehat{B}^{b_i}_{r_{-i},y_i} = T^{-1}_{\omega_{-i},y_i,b_C} B^{b_C,b_i}_{\omega_{-i},y_i} T^{-1}_{\omega_{-i},y_i,b_C}$$

where

$$A^{a_C,a_i}_{\omega_{-i},x_i} = \sum_{a_{[n]}|a_C,a_i} \mathbb{E}_{X_{[n]}|\omega_{-i},x_i} A^{a_{[n]}}_{x_{[n]}} \qquad B^{b_C,b_i}_{\omega_{-i},y_i} = \sum_{b_{[n]}|b_C,b_i} \mathbb{E}_{Y_{[n]}|\omega_{-i},y_i} B^{b_{[n]}}_{y_{[n]}}$$

and $a_{[n]}|a_C,a_i$ (resp. $b_{[n]}|b_C,b_i$) denotes summing over all $a_{[n]}$ consistent with a_C and a_i (resp. all $b_{[n]}$ consistent with b_C and b_i). It is easy to verify that the sets $\{\widehat{A}^{a_i}_{r_{-i},x_i}\}_{a_i\in\mathcal{A}}$ and $\{\widehat{B}^{b_i}_{r_{-i},y_i}\}_{b_i\in\mathcal{B}}$ form POVMs. Here, for a square matrix A, A^{-1} denotes its generalized inverse.

States. Now we are ready to define the states. Fix i, $r_{-i} = (\omega_{-i}, a_C, b_C)$, and x_i, y_i. Then let

$$|\Psi_{r_{-i},x_i,y_i}\rangle = \frac{S_{\omega_{-i},a_C,x_i} \otimes T_{\omega_{-i},b_C,y_i}|\psi\rangle}{\left\|S_{\omega_{-i},a_C,x_i} \otimes T_{\omega_{-i},b_C,y_i}|\psi\rangle\right\|}.$$

Observe that the normalization $\left\|S_{\omega_{-i},a_C,x_i} \otimes T_{\omega_{-i},b_C,y_i}|\psi\rangle\right\|^2$ is equal to $\mathsf{P}_{A_C B_C|\omega_{-i},x_i,y_i}(a_C, b_C)$.

3.4 Proof of Usefulness Lemma (Lemma 7)

This Lemma follows from a simple calculation: for every $x_i, y_i, a_i, b_i, r_{-i}$:

$$\begin{aligned}
&\mathrm{Tr}\left(\widehat{A}^{a_i}_{r_{-i},x_i} \otimes \widehat{B}^{b_i}_{r_{-i},y_i} \Psi_{r_{-i},x_i,y_i}\right) \\
&= \frac{1}{\left\| S_{\omega_{-i},a_C,x_i} \otimes T_{\omega_{-i},b_C,y_i} |\psi\rangle \right\|^2} \mathrm{Tr}\left(A^{a_C,a_i}_{\omega_{-i},x_i} \otimes B^{b_C,b_i}_{\omega_{-i},y_i} |\psi\rangle\langle\psi|\right) \\
&= \frac{1}{\mathsf{P}_{A_C B_C|\omega_{-i},x_i,y_i}(a_C,b_C)} \sum_{a_{[n]}|a_C,a_i} \sum_{b_{[n]}|b_C,b_i} \mathop{\mathbb{E}}_{X_{[n]}Y_{[n]}|\omega_{-i},x_i,y_i} \mathrm{Tr}\left(A^{a_{[n]}}_{x_{[n]}} \otimes B^{b_{[n]}}_{y_{[n]}} |\psi\rangle\langle\psi|\right) \\
&= \frac{\mathsf{P}_{A_i B_i A_C B_C|\omega_{-i},x_i,y_i}(a_i,b_i,a_C,b_C)}{\mathsf{P}_{A_C B_C|\omega_{-i},x_i,y_i}(a_C,b_C)} \\
&= \mathsf{P}_{A_i B_i|r_{-i},x_i,y_i}(a_i,b_i).
\end{aligned}$$

In the second equality we used that conditioned on Ω, $X_{[n]}$ and $Y_{[n]}$ are independent, so therefore $\mathbb{E}_{X_{[n]}|\omega_{-i},x_i} \mathbb{E}_{Y_{[n]}|\omega_{-i},y_i} = \mathbb{E}_{X_{[n]}Y_{[n]}|\omega_{-i},x_i,y_i}$. In the last equality we used that $r_{-i} = (\omega_{-i}, a_C, b_C)$. This concludes the Usefulness Lemma.

3.5 Proof of the Sampleability Lemma (Lemma 8)

Overview. Here we give some intuition. We first analyze an ensemble of states $\{|\Gamma_{x_i,x_C,a_C}\rangle\}$ (for now we omit mention of the dependency-breaking variable R for simplicity). These are indexed by Alice's questions in the i'th coordinate, her questions in the C coordinates, as well as her answers in the C coordinates. The state $|\Gamma_{x_i,x_C,a_C}\rangle$ roughly represents the state of the players where only Alice has applied her measurements – Bob hasn't done anything yet.

Fix a y_i, x_C, a_C. For average x_i, x_i' that are independently sampled from the marginal distribution $\mathsf{P}_{X_i|Y_i=y_i}$, we will show that $\||\Gamma_{x_i,x_C,a_C}\rangle - |\Gamma_{x_i',x_C,a_C}\rangle\| \sim \frac{1}{n}$. To handle issues such as Alice "printing" her input onto the state $|\psi\rangle$ (as discussed in the introduction), the definition of $|\Gamma_{x_i,x_C,a_C}\rangle$ requires local unitaries that "undo" such overt actions of Alice and Bob – this is accomplished by the unitaries U and V defined in Section 3.3.

Then, we consider what happens when we apply Bob's measurement to both states $|\Gamma_{x_i,x_C,a_C}\rangle$ and $|\Gamma_{x_i',x_C,a_C}\rangle$, and condition on obtaining answers b_C for the C coordinates. His measurement will depend on the questions y_i and y_C. The post-measurement states will be precisely $|\Psi_{x_i,y_i,x_C,y_C,a_C,b_C}\rangle$ and $|\Psi_{x_i',y_i,x_C,y_C,a_C,b_C}\rangle$. The distance between these states will be, roughly speaking, the distance between $|\Gamma_{x_i,x_C,a_C}\rangle$ and $|\Gamma_{x_i',x_C,a_C}\rangle$ divided by the probability of Bob obtaining outcome b_C conditioned on Alice obtaining a_C. If we average this distance over all choices of x_C, y_C, a_C, b_C that imply the event W_C, we get that the average distance between $|\Psi_{x_i,y_i,x_C,y_C,a_C,b_C}\rangle$ and $|\Psi_{x_i',y_i,x_C,y_C,a_C,b_C}\rangle$ is approximately $\frac{1}{n\mathsf{P}(W_C)}$. If $\mathsf{P}(W)$ is much greater than $1/n$, then this distance is small. We then invoke quantum correlated sampling (Lemma 3), and that proves the Sampleability Lemma.

Because of space constraints, we omit the proof and refer the reader to the full version at `http://arxiv.org/abs/1604.04340`.

Acknowledgments. The author thanks both the Institute of Mathematical Sciences at the National University of Singapore, and the Weizmann Institute of Science for hospitable stays during which this research was conducted. The author also thanks Corinna Li, Mohammad Bavarian, Govind Ramnarayan, and anonymous referees for helpful feedback and discussions.

References

1 Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Anchoring games for parallel repetition. *arXiv preprint arXiv:1509.07466*, 2015.

2 Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Parallel repetition via fortification: analytic view and the quantum case. *arXiv preprint arXiv:1603.05349*, 2016.

3 John S Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1(3), 1964.

4 Mark Braverman and Ankit Garg. Small value parallel repetition for general games. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*, 2015.

5 Mark Braverman, Anup Rao, Omri Weinstein, and Amir Yehudayoff. Direct products in communication complexity. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 746–755. IEEE, 2013.

6 André Chailloux and Giannicola Scarpa. Parallel repetition of entangled games with exponential decay via the superposed information cost. In *Automata, Languages, and Programming*, pages 296–307. Springer, 2014.

7 Kai-Min Chung, Xiaodi Wu, and Henry Yuen. Parallel repetition for entangled k-player games via fast quantum search. In *the 30th Conference on Computational Complexity (CCC)*, pages 512–536, 2015.

8 Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum xor proof systems. *Computational Complexity*, 17(2):282–299, 2008.

9 Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. In *the 29th Conference on Computational Complexity (CCC)*, pages 197–208, 2014.

10 Uriel Feige and Joe Kilian. Two-prover protocols – low error at affordable rates. *SIAM Journal on Computing*, 30(1):324–346, 2000.

11 Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4), 2001.

12 Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(8):141–172, 2009. doi:10.4086/toc.2009.v005a008.

13 Rahul Jain. New strong direct product results in communication complexity. *Electronic Colloquium on Computational Complexity (ECCC)*, 18(24):2, 2011.

14 Rahul Jain, Attila Pereszlényi, and Penghui Yao. A parallel repetition theorem for entangled two-player one-round games under product distributions. In *Proceedings of 29th Conference on Computational Complexity (CCC)*, pages 209–216, 2014.

15 Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. In *Proceedings of Foundations of Computer Science (FOCS)*, 2008.

16 Julia Kempe and Thomas Vidick. Parallel repetition of entangled games. In *Proceedings of the forty-third annual ACM symposium on Theory of computing (STOC)*, pages 353–362, 2011.

17 Dana Moshkovitz. Parallel repetition from fortification. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 414–423. IEEE, 2014.

18 Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.

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

20 Wim van Dam and Patrick Hayden. Universal entanglement transformations without communication. *Physical Review A*, 67(6):060302, 2003.

21 Mark M Wilde. *Quantum information theory*. Cambridge University Press, 2013.