

Multiplayer parallel repetition for expanding games

Irit Dinur ^{*†} Prahladh Harsha ^{*‡} Rakesh Venkat [‡] Henry Yuen [§]

June 26, 2017

Abstract

We investigate the value of parallel repetition of one-round games with any number of players $k \geq 2$. It has been an open question whether an analogue of Raz’s Parallel Repetition Theorem holds for games with more than two players, i.e., whether the value of the repeated game decays exponentially with the number of repetitions. Verbitsky has shown, via a reduction to the density Hales-Jewett theorem, that the value of the repeated game must approach zero, as the number of repetitions increases. However, the rate of decay obtained in this way is extremely slow, and it is an open question whether the true rate is exponential as is the case for all two-player games.

Exponential decay bounds are known for several special cases of multi-player games, e.g., free games and anchored games. In this work, we identify a certain expansion property of the base game and show all games with this property satisfy an exponential decay parallel repetition bound. Free games and anchored games satisfy this expansion property, and thus our parallel repetition theorem reproduces all earlier exponential-decay bounds for multiplayer games. More generally, our parallel repetition bound applies to all multiplayer games that are *connected* in a certain sense.

We also describe a very simple game, called the GHZ game, that does *not* satisfy this connectivity property, and for which we do not know an exponential decay bound. We suspect that progress on bounding the value of this the parallel repetition of the GHZ game will lead to further progress on the general question.

1 Introduction and Results

We consider multi-player one-round games, and their parallel repetition. In a k -player game G , a referee chooses a k -tuple of questions (x^1, \dots, x^k) from some question distribution μ , and sends x^t to player t . Each player t gives an answer a^t that only depends on their question (i.e., they cannot communicate with each other). The referee evaluates the players’ questions and answers according to some predicate $V((x^1, \dots, x^k), (a^1, \dots, a^k))$, and the players win if this predicate evaluates

*Research of the first and second author supported in part by an ISF-UGC grant number 1399/14.

†Weizmann Institute of Science. email: irit.dinur@weizmann.ac.il

‡Tata Institute of Fundamental Research. email: prahladh@tifr.res.in, rakesh09@gmail.com

§UC Berkeley. email: hyuen@cs.berkeley.edu

to 1. The *value* of a game G , denoted by $\text{val}(G)$, is the players' maximum success probability over all possible strategies they may use.

Here is a very natural operation on games, called *parallel repetition*: starting with a k -player game G , we can construct a new k -player game $G^{\otimes n}$, called the n -fold parallel repetition of G . In $G^{\otimes n}$, the referee will select n independent question tuples (x_i^1, \dots, x_i^t) from μ for each *coordinate* $i = 1, \dots, n$, and send (x_1^t, \dots, x_n^t) to each player t . Each player has to respond with n answers, and they win this repeated game if their answers and questions for each coordinate i satisfy the original game predicate V . We call G the *base game* of the parallel repeated game $G^{\otimes n}$.

The central question we consider is how $\text{val}(G^{\otimes n})$ depends on the base game G and the number of repetitions n . When G is a two-player game, the behavior of $\text{val}(G^{\otimes n})$ has been extensively studied, especially due to its applications to probabilistically checkable proofs and hardness of approximation. The central result in this area is Raz's Parallel Repetition Theorem [12] (coupled with subsequent improvements due to Holenstein [9]), which states the following:

Theorem 1 (Two-player parallel repetition). *Let G be a one-round two-player game with $\text{val}(G) \leq 1 - \varepsilon$ for some $\varepsilon \in (0, 1)$. Then for all $n \geq 0$,*

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c\varepsilon^3 n}{\log |\mathcal{A}|}\right)$$

where $|\mathcal{A}|$ is the answer alphabet size of G and $c > 0$ is a universal constant.

In other words, for *nontrivial* two-player games G (i.e., games whose value is less than 1), $\text{val}(G^{\otimes n})$ decays exponentially fast in n .

What about parallel repetition for games involving more than two players? It remains an intriguing open question whether Raz's Theorem can be extended to the multiplayer case. An early result of Verbitsky [14] shows that for multiplayer games G with $\text{val}(G) < 1$, the value of the repeated game $G^{\otimes n}$ must decay to 0 as n goes to infinity. However, the bound on the rate of decay is extremely weak: his result only shows that $\text{val}(G^{\otimes n})$ is bounded by a function that is inversely proportional to the inverse Ackermann function of n [11]! This poor rate of decay comes from its black-box usage of the density Hales-Jewett theorem from extremal combinatorics.

So far, Verbitsky's theorem is still the only result that gives a general parallel repetition bound for all multiplayer games. Exponential-decay parallel repetition bounds (à la Raz) for multiplayer games have been proven when there are additional assumptions on the game; for example, it has long been a folklore result that multiplayer *free* games satisfy an exponential-decay parallel repetition theorem [5].¹ Recently, Bavarian, Vidick and Yuen [1] studied a variant of parallel repetition (called "anchoring") where the base game G is first modified to an equivalent game \tilde{G} before being repeated in parallel, producing $\tilde{G}^{\otimes n}$. They proved that the value of $\tilde{G}^{\otimes n}$ is exponentially small in n when $\text{val}(G) < 1$, and otherwise $\text{val}(\tilde{G}^{\otimes n}) = 1$.² Buhrman et al. [4] show that the *non-signalling* value (which upper bounds $\text{val}(G)$) of a multiplayer game decays exponentially under parallel

¹A free game is one where each players' question is independent of all the other players'.

²In fact, Bavarian, Vidick and Yuen were motivated by the question of parallel repetition for *quantum* players; they showed that so-called "anchored" games satisfy quantum parallel repetition theorems.

repetition if the game has *full support*, that is, all possible question tuples to the players occur with non-zero probability. However, this does not necessarily imply a non-trivial decay of the value in general for such games, since there are games G with non-signalling value 1, but $\text{val}(G) < 1$.

We observe that the class of multiplayer games for which we have exponential-decay parallel repetition bounds (or, for that matter, any rate of decay better than inverse Ackermann!) all share a particular feature in common: when viewed as hypergraphs, the games all possess a certain “well-connectedness” property. For example, consider any two question tuples x, \hat{x} in the support of the question distribution μ of a free game. The question tuple $x = (x^1, \dots, x^k)$ can be “locally morphed” to $\hat{x} = (\hat{x}^1, \dots, \hat{x}^k)$ via a sequence of question tuples $(\hat{x}^1, \dots, \hat{x}^j, x^{j+1}, \dots, x^k)$ for $j = 1, \dots, k$, each of which remain in the support of μ . Furthermore, the anchoring transformations of [1] can be understood as improving the connectivity properties of the base game before repetition. In this paper, we formalize this well-connectedness property as a type of *expansion* of the base game, and show that any connected multiplayer game has exponential-decay parallel repetition bounds. We associate with every base game G , a related graph H_G (see Definition 2) and show that if H_G is connected, then the value of the repeated game, $\text{val}(G^{\otimes n})$, goes down exponentially in n , more precisely, for sufficiently large n ,

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c\varepsilon^5\lambda^2n}{\log|\mathcal{A}|}\right),$$

where λ is the spectral gap of the Laplacian of the graph H_G and c is some universal constant (see Theorem 6 for an exact statement of the result). Thus, if the graph H_G is connected (i.e., $\lambda > 0$), we have an exponential decay in n . In the case of games G , wherein the associated graph H_G is not only connected but also expanding (i.e., λ is a constant), as is the case with free games, and anchored games, the rate of exponential decay is a function of alphabet size $|\mathcal{A}|$ of the base game G as in Raz’s theorem. The class of *full-support* games investigated in Buhrman et al. [4] have a connected H_G , and hence our result shows an exponential decay bound on the value of such games under parallel repetition.

Why care about games with more than two players? The notion of a game is an extremely basic notion, and it’s use is pervasive in communication complexity, probabilistically checkable proofs (PCPs), etc. Whereas two-player games are already quite powerful and give us a lot, many problems are inherently higher-dimensional, i.e., would more naturally be cast as games with more than two players. The reason this is not commonly done is because we don’t know how to analyze these creatures. For example, constraint-satisfaction-problems with arity k are naturally cast as a k -player game. They can be reduced to a two-player game in the same way that a hypergraph can be converted to a graph, but this reduction in dimensionality might be lossy. Indeed, it is empirically true that PCPs with 3 or more queries are much more powerful than 2-query PCPs, but what is the reason for this?

Furthermore, this sudden jump in difficulty in going from two-player problems to three or more players is encountered also when studying multiparty communication complexity, and seemingly because of the same technique limitations. While direct sum and direct product theorems

are known for two-party communication complexity, nothing is known for the multiparty setting (in the so-called *number-on-forehead* model), and in fact making progress on this is connected to hard problems in circuit complexity.

We feel that the study of games with three or more players is a very important component in understanding such questions.

1.1 Notation

We establish some notational conventions, before stating our results formally.

For a k -player game G , we will let \mathcal{X}^t denote the question alphabet for player t , and $\mathcal{X} = \mathcal{X}^1 \times \mathcal{X}^2 \times \dots \times \mathcal{X}^k$ is the question alphabet for all the players together, underlying the question distribution μ . We will let \mathcal{A}^t denote the answer alphabet for player t , and $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^k$ to denote the answer alphabet for all the players together.

We will use superscripts to denote the players, and subscripts to denote the coordinate in parallel repetition. For example, x_i^t denotes the question received by player t in coordinate i . A single variable x can denote questions to the players in some coordinate clear from context, or a single coordinate game. We use x^{-t} to denote the questions to all but the t -th player in a single coordinate. When talking about multiple coordinates, we will use subscripts: x_{-i} denotes the questions to players in all but the i 'th coordinate, and x_{-i}^t denotes the all questions to player t in the repeated game except for the i 'th coordinate. To denote the question to player t in all coordinates, we use $x_{[n]}^t$.

We largely adopt the notational conventions from [9] for probability distributions. We let capital letters denote random variables and lower case letters denote specific samples. We use P_X to denote the probability distribution of random variable X , and $P_X(x)$ to denote the probability that $X = x$ for some value x . For multiple random variables, e.g., X, Y, Z , $P_{XYZ}(x, y, z)$ denotes their joint distribution with respect to some probability space understood from context.

We use $P_{Y|X=x}(y)$ to denote the conditional distribution $P_{YX}(y, x)/P_X(x)$, which is defined when $P_X(x) > 0$. When conditioning on many variables, we usually use the shorthand $P_{X|y,z}$ to denote the distribution $P_{X|Y=y, Z=z}$. For an event W we let $P_{XY|W}$ denote the distribution conditioned on W . We use the notation $\mathbb{E}_X f(x)$ and $\mathbb{E}_{P_X} f(x)$ to denote the expectation $\sum_x P_X(x) f(x)$.

Let P_{X_0} be a distribution over \mathcal{X} and $P_{X_1, Y}$ a joint distribution over $\mathcal{X} \times \mathcal{Y}$. Suppose for every x in the support of P_{X_0} , the conditional distribution $P_{Y|X_1=x}$ defined over \mathcal{Y} is well-defined. We then define the distribution $P_{X_0} P_{Y|X_1}$ over $\mathcal{X} \times \mathcal{Y}$ as

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

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

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

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

1.2 Our results

To define our class of connected and expanding games, we need the following notion of the $(k-1)$ -connection graph of a game G . This graph, denoted H_G , has a vertex for every k -tuple of questions, and two k -tuples are connected by an edge if they agree on $(k-1)$ coordinates. A game is $(k-1)$ -connected iff H_G is connected.

To further define our notion of expansion for a k -player game we need to take the weights of G into account when defining H_G . For this it is instructive to think of an intermediate bipartite graph $B_G = (\mathcal{X}', \mathcal{X}, E)$ as follows. The right hand vertices is simply \mathcal{X} , the set of all k -tuples of questions, and we endow these vertices with weights as given by G . The left hand vertices consists of all punctured k -tuples, which are k -tuples of questions where exactly one of the entries is replaced by a special \star symbol. Connect each k -tuple of questions to all of the k ways to make it into a punctured k -tuple. Now, consider the distribution on punctured tuples obtained by selecting a random k -tuple from \mathcal{X} according to the game distribution, and then puncturing it in a random location. The graph H_G is defined by selecting a random punctured tuple according to this distribution, and then selecting independently two k -tuples conditioned on this puncturing. Note that each completion is distributed exactly according to the original game distribution.

We now move to a completely explicit description consistent with the above. In what follows, $P_X(x)$ denotes the probability of question tuple x under the question distribution μ , $P_{X^t}(x^t)$ denotes the marginal probability of player t 's question, and $P_{X^t|X^{-t}=x^{-t}}(x^t)$ denotes the same probability, conditioned on the other players having received x^{-t} .

Definition 2 ($(k-1)$ -connection graph of G). *Let $G = (\mu, V)$ be a k -player game with question set $\mathcal{X} = \mathcal{X}^1 \times \dots \times \mathcal{X}^k$. The $(k-1)$ -connection graph of G is the weighted graph $H_G = (V_H, \rho)$ with vertex set $V_H = \mathcal{X}$ and weight function $\rho : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$, defined as follows: for every $x, x' \in \mathcal{X}$,*

$$\rho(x, x') = \begin{cases} \frac{1}{k} P_X(x) \left[\sum_{t \in [k]} P_{X^t|X^{-t}}(x^t) \right] & \text{if } x = x', \\ \frac{1}{k} P_{X^{-t}}(x^{-t}) \cdot P_{X^t|X^{-t}}(x^t) \cdot P_{X^t|X^{-t}}(x'^t) & \text{if } \exists t \text{ s.t. } x^{-t} = x'^{-t}, x^t \neq x'^t, \\ 0 & \text{otherwise.} \end{cases}$$

The weight function $\rho(x, x')$ can be viewed as the probability of generating the pair (x, x') according to the following random process: first, $x \in \mathcal{X}$ is sampled from the distribution P_X . Then, a coordinate $t \in [k]$ is chosen uniformly at random, and x' is sampled from the conditional distribution $P_{X|X^{-t}}$ (that is, the distribution μ conditioned on x^{-t}).

Observe that ρ is symmetric, i.e., $\rho(x, x') = \rho(x', x)$. Furthermore, note that the weight on any given vertex is exactly:

$$\begin{aligned} \rho(x, \cdot) &= \sum_{x'} \rho(x, x') = \\ &= P_X(x) \cdot \frac{1}{k} \sum_{t \in [k]} P_{X^t|X^{-t}}(x^t) + P_X(x) \cdot \frac{1}{k} \sum_{t \in [k]} \sum_{x'^t \neq x^t} P_{X^t|X^{-t}}(x'^t) \\ &= P_X(x). \end{aligned}$$

Therefore $\rho(\cdot, \cdot)$ is a probability distribution over $\mathcal{X} \times \mathcal{X}$.

Remark 3. Henceforth, when we talk about graph properties such as diameter, connectedness or expansion of H_G , we will do so only with respect to the vertices having non-zero weight.

We now recall the definition of a graph with a weight function ρ being a spectral expander:

Definition 4 (Normalized Laplacian). *Let H be a weighted graph where $\rho(u, v) \leq 1$ is the weight between vertices u and v . The normalized Laplacian $L_H \in \mathbb{R}^{|V| \times |V|}$ of H is defined to be*

$$(L_H)_{u,v} = \begin{cases} 1 - \frac{\rho(u,v)}{\rho(v)} & \text{if } u = v \text{ and } \rho(v) \neq 0 \\ -\frac{\rho(u,v)}{\sqrt{\rho(u)\rho(v)}} & \text{if } \rho(u), \rho(v) \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

where $\rho(u) = \sum_v \rho(u, v)$ and $\rho(v) = \sum_u \rho(u, v)$.

It is well-known that the second smallest eigenvalue of H is given by the following variational formula: for all $r \in \mathbb{N}$,

$$\lambda(H) = \inf_g \frac{\sum_{u,v} \rho(u,v) \|g(u) - g(v)\|^2}{\sum_u \rho(u) \|g(u) - \bar{g}\|^2} \quad (1)$$

where the infimum is over all vector-valued functions $g : V(H) \rightarrow \mathbb{R}^r$ defined on the vertices of H , and \bar{g} is a vector in \mathbb{R}^r where for each $i \in [r]$, $\bar{g}_i = \sum_u \rho(u) g(u)_i$.

Definition 5 (Expander graph). *Let $\lambda \in (0, 1)$. A graph H is a λ -expander if $\lambda(H) \geq \lambda$.*

Our main result is an exponential-decay parallel repetition bound for multiplayer games whose $(k-1)$ -connection graph is expanding:

Theorem 6 (Main theorem). *Let $\epsilon, \lambda \in (0, 1)$. Let G be a k -player game with $\text{val}(G) \leq 1 - \epsilon$. If the $(k-1)$ -connection graph H_G is a λ -expander, then we have, for all $n \geq \frac{\log 4/\epsilon}{\epsilon^5 \lambda^2}$:*

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c\epsilon^5 \lambda^2 n}{\log |\mathcal{A}|}\right)$$

where $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^k$ is the answer alphabet in G , and c is a universal constant.

By applying our main theorem to free games and anchored games, we recover existing exponential-decay parallel repetition results for multiplayer games [5, 1]. We also get an exponential-decay lower bound for *connected games* – games whose $(k-1)$ -connection graph is connected. We record these consequences in the following corollary:

Corollary 7. *Let G be a k -player game with $\text{val}(G) \leq 1 - \epsilon$, and let $n \geq \frac{\log 4/\epsilon}{\epsilon^5 \lambda^2}$. If G is:*

1. Free, i.e., $\mu(x) = \mu^1(x^1) \times \dots \times \mu^k(x^k)$, then

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c\epsilon^5 n}{k^2 \log |\mathcal{A}|}\right).$$

2. α -Anchored (see Definition 14, and [1] for more details), then

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c \alpha^{2k} \varepsilon^5 n}{64 k^2 \log |\mathcal{A}|}\right).$$

3. Connected, i.e., the $(k - 1)$ -connection graph is connected, then

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c \rho_{\min}^2 \varepsilon^5 n}{\log |\mathcal{A}|}\right)$$

where $\rho_{\min} = \min_{u,v:\rho(u,v)>0} \rho(u,v)$. In particular, if the game G is such that μ is the uniform distribution over some set $S \subseteq \mathcal{X}$, then $\rho_{\min} \geq (k|S|^2)^{-1}$.

where c is a universal constant.

The proof of Corollary 7 can be found in Appendix A.

Observe that our proof of exponential decay for games whose corresponding $(k - 1)$ -connection graph is connected proves a rate of exponential decay that is dependent on the size of the base game G . It is conceivable that this rate of decay can be further improved to depend only on the alphabet size $|\mathcal{A}|$ of the base game and be independent of the size of the base game (as is the case in Raz’s theorem for 2 player games). For games whose corresponding $(k - 1)$ -connection graph is expanding (as is the case with free games and anchoring games), we obtain a rate of exponential decay which is a function of only the base game’s alphabet size.

Remark 8. For simplicity, we state Theorem 6 assuming the base game has a connected $(k - 1)$ -connection graph. It is easy to check (from the proof of Theorem 6) that it also extends to games that are disjoint union of games each of which has a connected $(k - 1)$ -connection graph. By disjoint union we mean that each question occurs only in one of the components. For $k = 2$, this captures all possible games since every game is a union of disjoint games whose $(k - 1 = 2 - 1 = 1)$ -connection graphs are connected). We note that for 2-player games, there are alternate proof techniques [12, 9] using correlated sampling which prove even better rate of exponential decay (our proof of Theorem 6 does not use correlated sampling). In contrast, for $k > 2$, there are many games that are not captured by our theorem. We will see below an example of such a $k = 3$ -player game called the GHZ game.

A comment about fortified games. Bavarian, Vidick and Yuen also proved a parallel repetition bound for a special class of multiplayer games *fortified games* [2] (a class of games introduced by Moshkovitz [10]). However, we do not consider this a “true” exponential-decay parallel repetition bound, because it does not establish a decay bound of the form $\text{val}(G^{\otimes n}) \leq \exp(-\beta n)$ for some constant β that depends on the game G , but is independent of n . Instead, it proves a decay bound that is exponential only for a small number of repetitions (depending on the base game). After this small number of repetitions, there are no guarantees about any further value decay (other than that promised by Verbitsky’s theorem). Because we are interested in the asymptotic behavior of an n -repeated multiplayer game as n goes to infinity, we do not consider the parallel repetition of fortified games here.

A disconnected three-player game. It may seem that, given Corollary 7, we have established a general exponential-decay parallel repetition bound for *all* multiplayer games, albeit with some slightly annoying dependency on a quantity related to the minimum probability of any question from μ . Unfortunately, this is far from the case.

Here is a simple three-player game called the *GHZ game* whose parallel repetition resists analysis; the best decay bound we have comes from Verbitsky’s theorem [14]. The GHZ game is a three-player game³ where the referee samples a question triple (x, y, z) uniformly at random from $\{(1, 0, 0), (0, 1, 0), (0, 0, 1), (1, 1, 1)\}$, and sends each bit of the triple to the corresponding player. The players respond with bits a, b, c respectively, and they win iff $x \wedge y \wedge z = a \oplus b \oplus c$. It is easy to see that $\text{val}(\text{GHZ}) = 3/4$ (achieved by the strategy where all players always output “0”). However, the best general bound we have on $\text{val}(\text{GHZ}^{\otimes n})$ is the weak inverse-Ackermann decay given by Verbitsky’s theorem.

Our main theorem does not apply because the $(k - 1)$ -connection graph H_{GHZ} of the GHZ game is actually *disconnected*; no two question triples are connected via a single coordinate change. One necessary criterion for the $(k - 1)$ -connection graph to be connected is that, after fixing any subset of $(k - 1)$ players’ questions, the remaining player’s question is yet undetermined. On the other hand, the players’ questions in the GHZ game satisfy a linear relation (i.e. $x \oplus y \oplus z = 1$), and thus fixing two players’ questions also fixes the third.

We believe that the strong correlations present in the GHZ question distribution represent the “hardest instance” of the multiplayer parallel repetition problem. Existing techniques from the two-player case (which we leverage in this paper) appear to be incapable of analyzing games with question distributions with such strong correlations. Thus we explicitly raise the open question of proving an exponential-decay parallel repetition bound for the GHZ game:

Conjecture 9 (GHZ parallel repetition). There exists a constant $\beta > 0$ such that for all n ,

$$\text{val}(\text{GHZ}^{\otimes n}) \leq \exp(-\beta n).$$

Finally, we remark that this challenge of handling strongly correlated question distributions is reminiscent of the challenge of proving *direct sum* theorems for multiparty communication complexity in the *Number-on-Forehead* (NOF) model. There, each player sees every players’ inputs but their own, so fixing $(k - 1)$ out of k players’ inputs will fix the remaining player’s inputs. Proving direct sum results in NOF communication complexity has resisted progress for reasons that appear to be related to the multiplayer parallel repetition problem.

³The GHZ game comes from the study of non-locality in quantum physics; when the players use classical strategies, their maximum success probability is $\text{val}(G) = 3/4$, but using quantum entanglement, the GHZ can be won with certainty [7].

2 Proof of Theorem 6

2.1 Proof outline

We first give a brief overview of the information-theoretic approach to proving two-player parallel repetition as in [12, 9], and explain the technical barrier to extending the proof to three or players. We then will describe how we circumvent this technical barrier.

Essentially all known proofs of parallel repetition proceed via reduction, showing how a “too good” strategy for the repeated game G^n can be “rounded” into a strategy for G with success probability strictly greater than $\text{val}(G)$, yielding a contradiction.

Let \mathcal{S}^n be a strategy for G^n that has a high success probability. Either by induction or via a probabilistic argument one can identify a set of coordinates S and an index i such that $\Pr(\text{Players win round } i | W_S) > \text{val}(G) + \delta$, where W is the event that the players’ answers satisfy the predicate V in all instances of G indexed by S . Given a pair of questions (x, y) in G the strategy \mathcal{S} embeds them in the i -th coordinate of a n -tuple of questions

$$x_{[n]}y_{[n]} = \begin{pmatrix} x_1, x_2, \dots, x_{i-1}, & x & , x_{i+1}, \dots, x_n \\ y_1, y_2, \dots, y_{i-1}, & y & , y_{i+1}, \dots, y_n \end{pmatrix}$$

that is distributed according to $P_{X_{[n]}Y_{[n]}|X_i=x, Y_i=y, W}$. The players then simulate \mathcal{S}^n on $x_{[n]}$ and $y_{[n]}$ respectively to obtain answers (a_1, \dots, a_n) and (b_1, \dots, b_n) , and return (a_i, b_i) as their answers in G . The single-shot strategy \mathcal{S} succeeds with probability precisely $\Pr(\text{Win } i | W_S)$ in G , yielding the desired contradiction.

As \mathcal{S}^n need not be a product strategy, conditioning on W_S may introduce correlations that make $P_{X_{[n]}Y_{[n]}|X_i=x, Y_i=y, W_S}$ impossible to sample exactly. A key insight in Raz’ proof of parallel repetition is that it is still possible for the players to *approximately* sample from this distribution. For this, we introduce a *dependency-breaking variable* R with the following properties:

- (a) Given $r \sim P_R$ the players can locally sample $x_{[n]}$ and $y_{[n]}$ according to $P_{X_{[n]}Y_{[n]}|X_i=x, Y_i=y, W_S, r}$
- (b) The players can jointly sample from P_R using shared randomness.

In [9] R is defined so that a sample r fixes at least one of $\{x_{i'}, y_{i'}\}$ for each $i' \neq i$. It can then be shown that conditioned on x , R is nearly (though not exactly) independent of y , and vice-versa. In other words,

$$P_{R|X_i=x, W_S} \approx P_{R|X_i=x, Y_i=y, W_S} \approx P_{R|Y_i=y, W_S} \tag{2}$$

where “ \approx ” denotes closeness in statistical distance. Eq. (2) suffices to guarantee that the players can *approximately* sample the same r from $P_{R|X_i=x, Y_i=y, W_S}$ with high probability, achieving point (b) above. This sampling is accomplished through a technique called *correlated sampling*.

This argument relies heavily on the assumption that there are only two players who employ a deterministic strategy. With more than two players, it is not known how to design an appropriate

dependency-breaking variable R that satisfies *both* items (a) and (b) above: in order to be jointly sampleable, R needs to fix as few inputs as possible; in particular, no single player should require knowledge of the other player's questions to sample R . On the other hand, in order to allow players to locally sample their inputs conditioned on R , the variable needs to fix as many inputs as possible. These two requirements turn out to be in direct conflict as soon as there are more than two players, and a straightforward generalization of the two-player version of the dependency-breaking variable cannot be "correlatedly sampled" by all players, unless every player has knowledge of the question received by some other player.

We avoid this roadblock by proving in Section 2.4 that if the $(k-1)$ -connected graph H_G is connected, then the players can avoid correlated sampling altogether. In fact, they can sample an appropriate dependency-breaking variable from a *global* distribution that does not depend on any player's question.

2.2 Following Raz-Holenstein

Fix a k -player game $G = (\mu, V)$, with answer alphabet $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_k$ and $\text{val}(G) = 1 - \varepsilon$. Consider the n -fold parallel repetition $G^{\otimes n}$ and consider an optimal strategy $\{f^t : (\mathcal{X}^t)^{\otimes n} \rightarrow (\mathcal{A}^t)^{\otimes n}\}_{t \in [k]}$ for the k players.

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 $S \subseteq [n]$, let $W_S = \bigwedge_{i \in S} W_i$. In the following, all probabilities are with respect to this optimal strategy.

Proposition 10. *Let $\varepsilon > 0$. Suppose that $\log 1/\Pr(W) \leq \varepsilon n/16 - \log 4/\varepsilon$. Then there exists a set $S \subseteq [n]$ of size at most $t = \frac{8}{\varepsilon} (\log 4/\varepsilon + \log 1/\Pr(W))$ such that*

$$\Pr_{i \notin S}(\neg W_i | W_S) \leq \varepsilon/2$$

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

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 S , we will show that $\mathbb{E}_S \Pr(\neg W_i | W_S) \leq \varepsilon/2$, where S is a (multi)set of t independently chosen indices in $[n]$. This implies that there exists a particular set S such that $\Pr(\neg W_i | W_S) \leq \varepsilon/2$, which concludes the claim.

First we write, for a fixed S ,

$$\begin{aligned} \Pr(\neg W_i | W_S) &= \Pr(\neg W_i | W_S, W_{>1-\delta}) \Pr(W_{>1-\delta} | W_S) \\ &\quad + \Pr(\neg W_i | W_S, \neg W_{>1-\delta}) \Pr(\neg W_{>1-\delta} | W_S). \end{aligned}$$

Observe that $\Pr(\neg W_i | W_S \wedge W_{>1-\delta})$ is the probability that, conditioned on winning all rounds in S , the randomly selected coordinate $i \in [n] - S$ happens to be one of the (at most) δn lost rounds. This is at most $\delta n / (n - t) \leq \varepsilon/4$. Now observe that

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

where for the second inequality we used the fact that $P(W_S) \geq P(W)$. \square

For the remainder of this proof we will fix a set S as given by Proposition 10. By renaming coordinates, we will assume without loss of generality that S is the last t coordinates of $[n]$. We will let $m = n - |S|$. We will refer to the games indexed by set S as the S -games.

2.3 Dependency-breaking variables

We define the k -player analogue of the dependency-breaking variable R that is used so crucially in information-theoretic proofs of parallel repetition [12, 9, 3]. R will consist of a variable Ω , which fixes the questions for the S -games, and at least $(k - 1)$ -of- k questions in every other coordinate, and a variable $Z = (A_S)$, which fixes the answers of S -games. More formally, $\Omega = (\Omega_1, \dots, \Omega_m, X_S)$, where X_S are fixed questions for the S -games. Each $\Omega_i = (D_i, M_i)$, for $i \in \bar{S}$, where D_i is a uniformly random value in $[k]$, and

$$M_i = X_i^{-t} \quad \text{if } D_i = t$$

In other words, D_i specifies which player's question to omit; the other $(k - 1)$ players are fixed.

For $i \notin S$, we let Ω_{-i} denote Ω with Ω_i omitted. We let $R_{-i} := (\Omega_{-i}, A_S)$. 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^t refer to specific values of R_{-i}, X_i^t respectively.

Claim 11. *Conditioned on R , $\{X_{[n]}^t\}_{t \in [k]}$ are independent.*

In the following, P_I denotes the distribution of a uniformly random $i \in [m]$, and “ $P \approx_\delta Q$ ” indicates that the probability distributions P and Q are δ -close in statistical distance. We will fix

$$\delta = \frac{1}{m} \left(\log \frac{1}{P(W_S)} + |S| \log |\mathcal{A}| \right).$$

The next lemma states that for an average i , if we sample questions x_i, \hat{x}_i from the joint probability distribution $\rho(x_i, \hat{x}_i)$, the distributions of the corresponding dependency-breaking variables will be close.

Lemma 12.

$$\frac{1}{m} \sum_i \sum_{x_i, \hat{x}_i \in \mathcal{X}} \rho(x_i, \hat{x}_i) \left\| P_{R_{-i}|x_i, W_S} - P_{R_{-i}|\hat{x}_i, W_S} \right\|_1 \leq O(\sqrt{\delta})$$

where $\rho(\cdot, \cdot)$ is the weight function of the $(k - 1)$ -connection graph H_G .

Proof. First, we establish the following: for all $t \in [k]$, we have

$$\mathbb{E}_i \sum_{x_i^{-t}, x_i^t, \hat{x}_i^t} P_{X_i^{-t}}(x_i^{-t}) P_{X_i^t|x_i^{-t}}(x_i^t) \cdot P_{\hat{X}_i^t|x_i^{-t}}(\hat{x}_i^t) \left\| P_{R_{-i}|x_i, W_S} - P_{R_{-i}|\hat{x}_i, W_S} \right\|_1 \leq O(\sqrt{\delta}) \quad (3)$$

where we use the shorthand $x_i := x_i^{-t} x_i^t$ and $\hat{x}_i := x_i^{-t} \hat{x}_i^t$. This follows from the same arguments found in [9, 3]; for each player t , we can treat the other $(k - 1)$ players as one “meta player”, and apply the two-player analysis to obtain (3).

Observe that when $x_i^t \neq \hat{x}_i^t$, we have

$$\mathbb{P}_{X_i^{-t}}(x_i^{-t}) \cdot \mathbb{P}_{X_i^t|x_i^{-t}}(x_i^t) \cdot \mathbb{P}_{\hat{X}_i^t|x_i^{-t}}(\hat{x}_i^t) = k\rho(x_i, \hat{x}_i).$$

On the other hand, when $x_i^t = \hat{x}_i^t$, $x_i = \hat{x}_i$ so therefore $\|\mathbb{P}_{R_{-i}|x_i, W_S} - \mathbb{P}_{R_{-i}|\hat{x}_i, W_S}\|_1 = 0$. Furthermore, for x_i and \hat{x}_i that differ in more than 1 coordinate, we have $\rho(x_i, \hat{x}_i) = 0$, and for every x_i, \hat{x}_i such that $\rho(x_i, \hat{x}_i) \neq 0$, there exists a unique $t \in [k]$ such that $x_i^t \neq \hat{x}_i^t$. Thus we can bound for every i :

$$\begin{aligned} & \sum_{x_i, \hat{x}_i \in \mathcal{X}} \rho(x_i, \hat{x}_i) \|\mathbb{P}_{R_{-i}|x_i, W_S} - \mathbb{P}_{R_{-i}|\hat{x}_i, W_S}\|_1 \\ &= \frac{1}{k} \sum_{t \in [k]} \sum_{x_i^{-t}, \hat{x}_i^{-t}} \mathbb{P}_{X_i^{-t}}(x_i^{-t}) \cdot \mathbb{P}_{X_i^t|x_i^{-t}}(x_i^t) \cdot \mathbb{P}_{\hat{X}_i^t|x_i^{-t}}(\hat{x}_i^t) \|\mathbb{P}_{R_{-i}|x_i, W_S} - \mathbb{P}_{R_{-i}|\hat{x}_i, W_S}\|_1. \end{aligned}$$

Averaging over i and using (3), we obtain the statement of the lemma. \square

2.4 Avoiding correlated sampling using expansion

At this point, ideally, every player would like to sample from $R_{-i}|x_i, W_S$. Lemma 12 establishes that $R_{-i}|x_i, W_S$ is close to $R_{-i}|x_i^{-t}, W_S$ for each $t \in [k]$. None of the players alone has knowledge of x_i^{-t} , however. We will show now that nevertheless, there is a *global* distribution known to all the players, from which the players can approximately sample $R_{-i}|x_i, W_S$.

Lemma 13. *For all $i \in [m]$ there exists a distribution $\tilde{\mathbb{P}}_{R_{-i}}$ over R_{-i} such that*

$$\frac{1}{m} \sum_i \sum_x \rho(x) \|\mathbb{P}_{R_{-i}|x, W_S} - \tilde{\mathbb{P}}_{R_{-i}}\|_1 \leq O\left(\frac{\delta^{1/4}}{\sqrt{\lambda}}\right).$$

Proof. For each i , define the vector-valued function $g_i : \mathcal{X} \rightarrow \mathbb{R}^{R_{-i}}$ as follows: for all $x \in \mathcal{X}$,⁴

$$g_i(x) = \sqrt{\mathbb{P}_{R_{-i}|x, W_S}}$$

where $\sqrt{\mathbb{P}_{R_{-i}|x, W_S}}$ denotes the entry-wise square root of the probability distribution $\mathbb{P}_{R_{-i}|x, W_S}$, viewed as a vector. In other words, the entries of $g_i(x)$ are indexed by different values r_{-i} of the random variable R_{-i} . Thus, g_i is a unit vector in the ℓ_2 norm.

For any i and any $x, \hat{x} \in \mathcal{X}$, the quantity $\|g_i(x) - g_i(\hat{x})\|^2$ is simply the square of the *Hellinger distance* between $\mathbb{P}_{R_{-i}|x, W_S}$ and $\mathbb{P}_{R_{-i}|\hat{x}, W_S}$, which can be related to their statistical distance by

$$\|g_i(x) - g_i(\hat{x})\|^2 \leq \|\mathbb{P}_{R_{-i}|x, W_S} - \mathbb{P}_{R_{-i}|\hat{x}, W_S}\|_1.$$

By Lemma 12, we can average the above inequality over all i and choosing x, \hat{x} according to the probability distribution $\rho(x, \hat{x})$, we get

$$\frac{1}{m} \sum_i \sum_{x, \hat{x}} \rho(x, \hat{x}) \|g_i(x) - g_i(\hat{x})\|^2 \leq \mathbb{E}_i \sum_{x, \hat{x}} \rho(x, \hat{x}) \|\mathbb{P}_{R_{-i}|x, W_S} - \mathbb{P}_{R_{-i}|\hat{x}, W_S}\|_1 \leq O(\sqrt{\delta})$$

⁴Here, when we write x , we are implicitly mean x_i ; we drop the subscript i for notational convenience.

But now we can leverage Equation (1). For every i , define the vector $\bar{g}_i = \sum_x P_X(x) g_i(x)$. This is not necessarily a unit vector, but we have the relation

$$\frac{1}{m} \sum_i \sum_x \rho(x) \|g_i(x) - \bar{g}_i\|^2 \leq \frac{1}{\lambda m} \sum_i \sum_{x, \hat{x}} \rho(x, \hat{x}) \|g_i(x) - g_i(\hat{x})\|^2 \leq O\left(\frac{\sqrt{\delta}}{\lambda}\right).$$

If $O(\sqrt{\delta}/\lambda)$ is small, then this implies that on average, the vectors $g_i(x)$ are all close to a fixed state \bar{g}_i . Since $g_i(x)$ are all unit vectors, this implies that \bar{g}_i is close to a unit vector. By increasing the error by a constant factor, we can assume that \bar{g}_i is in fact a unit vector. Thus we can construct the probability distribution

$$\tilde{P}_{R_{-i}}(r_{-i}) = \bar{g}_i(r_{-i})^2.$$

Using that the statistical distance is at most (up to constant factors) the square root of the Hellinger distance, we get that

$$\frac{1}{m} \sum_i \sum_x \rho(x) \|P_{R_{-i}|x, W_S} - \tilde{P}_{R_{-i}}\|_1 \leq O(\delta^{1/4} \lambda^{-1/2}).$$

□

2.5 Finishing the proof

Let $\{f^t\}$ be an optimal strategy for the game $G^{\otimes n}$. If $P(W) \leq \frac{4}{\varepsilon} 2^{-\varepsilon n/16}$, then we are done. Otherwise, suppose $\log 1/P(W) \leq \varepsilon n/16 - \log 4/\varepsilon$. Let the subset S be as given by Proposition 10, and assume the coordinates are numbered so that S is the last $|S|$ coordinates of $[n]$. For all $i \in [m]$, let $\tilde{P}_{R_{-i}}$ be as given by Lemma 13. Consider the following single-shot strategy by the players, where x is drawn from μ and x^t is given to player t :

1. Using shared randomness, the players sample an $i \in [m]$ uniformly at random, and sample r_{-i} from $\tilde{P}_{R_{-i}}$. Each player t then sets x_i^t to be their “true” question x^t they received from the referee.
2. Using private randomness, each player t samples x_{-i}^t from $P_{X_{-i}^t|x_i^t, r_{-i}}$. That is, each player samples questions for the n coordinates that come from the repeated game, conditioned on their own true input x_i^t and the dependency-breaking variable r_{-i} .
3. Player t outputs the i 'th component of the answer vector $f^t(x_{[n]}^t)$.

Lemma 13 implies that after the first step, the sample r_{-i} each player possesses will be, up to statistical error $O(\delta^{1/4}/\sqrt{\lambda})$, distributed according to $P_{R_{-i}|x, W_S}$ (on average over i and x). Then, by Claim 11, the joint distribution of the random variables $\{X_{[n]}^t\}$ that the players have sampled is

$$P_{X_{[n]}^1|x_i^1, r_{-i}} \times \cdots \times P_{X_{[n]}^k|x_i^k, r_{-i}} = P_{X_{[n]}|x_i, r_{-i}}.$$

Thus, conditioned on r_{-i} and x_i , the distribution of their answers a_i will be distributed according to $P_{A_i|x_i, r_{-i}}$. Averaging over i , x_i , and r_{-i} , we get that their answers are $O(\delta^{1/4}/\sqrt{\lambda})$ -close to being distributed according to

$$P_I \cdot P_{X_i} \cdot P_{R_{-i}|X_i, W_S} \cdot P_{A_i|X_i, R_{-i}}$$

where P_I stands for the uniform distribution over $i \in [m]$. We also have that, on average i , $P_{X_i|W_S}$ is $O(\sqrt{\delta})$ -close in statistical distance to P_{X_i} . Thus their answers are $O(\delta^{1/4}/\sqrt{\lambda}) + O(\sqrt{\delta})$ close to being distributed as

$$P_I \cdot P_{A_i|W_S}.$$

Thus by Proposition 10, the probability that the players win G is at least

$$1 - \varepsilon/2 - \left(O(\delta^{1/4}/\sqrt{\lambda}) + O(\sqrt{\delta}) \right).$$

If $O(\delta^{1/4}/\sqrt{\lambda}) + O(\sqrt{\delta}) < \varepsilon/2$, then we would contradict the fact that $\text{val}(G) = 1 - \varepsilon$. This implies that we must have $\delta = \Omega(\varepsilon^4 \lambda^2)$. If we let $P(W) = 2^{-\gamma n}$, then we can write

$$\delta \leq \frac{16}{\varepsilon} \left[\frac{1}{n} \log \frac{4}{\varepsilon} + 2 \log |\mathcal{A}| \gamma \right]$$

where we plugged in the bound on $|S| \leq n/2$ from Proposition 10. This implies the lower bound

$$\gamma \geq \Omega \left(\frac{\varepsilon^5 \lambda^2}{\log |\mathcal{A}|} \right) \quad (4)$$

when $n \geq \frac{\log 4/\varepsilon}{\varepsilon^5 \lambda^2}$, proving the theorem.

References

- [1] Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Anchoring games for parallel repetition. (manuscript), 2015. [arXiv:1509.07466](#).
- [2] Mohammad Bavarian, Thomas Vidick, and Henry Yuen. Parallel repetition via fortification: analytic view and the quantum case. (manuscript), 2016. [arXiv:1603.05349](#).
- [3] Mark Braverman and Ankit Garg. Small value parallel repetition for general games. In *Proc. 47th ACM Symp. on Theory of Computing (STOC)*, pages 335–340, 2015. [doi:10.1145/2746539.2746565](#).
- [4] Harry Buhrman, Serge Fehr, and Christian Schaffner. On the Parallel Repetition of Multi-Player Games: The No-Signaling Case. In *Proc. 9th Conf. on the Theory of Quantum Computation, Communication and Cryptography (TQC 2014)*, pages 24–35, 2014. [doi:10.4230/LIPIcs.TQC.2014.24](#).
- [5] Kai-Min Chung, Xiaodi Wu, and Henry Yuen. Parallel repetition for entangled k -player games via fast quantum search. In *Proc. 30th Comput. Complexity Conf.*, pages 512–536, 2015. [arXiv:1501.00033](#), [doi:10.4230/LIPIcs.CCC.2015.512](#).
- [6] Persi Diaconis and Daniel Stroock. Geometric bounds for eigenvalues of Markov chains. *Ann. Appl. Probab.*, 1(1):36–61, 1991. [doi:10.1214/aoap/1177005980](#).
- [7] Daniel M. Greenberger, Michael A. Horne, Abner Shimony, and Anton Zeilinger. Bells theorem without inequalities. *Am. J. Phys.*, 58(12):1131–1143, 1990. [doi:10.1119/1.16243](#).
- [8] Venkatesan Guruswami. Rapidly mixing Markov chains: a comparison of techniques. (survey), 2016. [arXiv:1603.01512](#).

- [9] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory Comput.*, 5(1):141–172, 2009. (Preliminary version in *39th STOC*, 2007). [arXiv:cs/0607139](https://arxiv.org/abs/cs/0607139), [doi:10.4086/toc.2009.v005a008](https://doi.org/10.4086/toc.2009.v005a008).
- [10] Dana Moshkovitz. Parallel repetition from fortification. In *Proc. 55th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 414–423, 2014. [doi:10.1109/FOCS.2014.51](https://doi.org/10.1109/FOCS.2014.51).
- [11] D. H. J. Polymath. A new proof of the density Hales-Jewett theorem. *Ann. of Math.*, 175:1283–1327, 2012. [arXiv:0910.3926](https://arxiv.org/abs/0910.3926), [doi:10.4007/annals.2012.175.3.6](https://doi.org/10.4007/annals.2012.175.3.6).
- [12] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, June 1998. (Preliminary version in *27th STOC*, 1995). [doi:10.1137/S0097539795280895](https://doi.org/10.1137/S0097539795280895).
- [13] Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combin. Probab. Comput.*, 1(4):351–370, 1992. (Preliminary version in *1st LATIN*, 1992). [doi:10.1017/S0963548300000390](https://doi.org/10.1017/S0963548300000390).
- [14] Oleg Verbitsky. Towards the parallel repetition conjecture. *Theoret. Comput. Sci.*, 157(2):277–282, 1996. (Preliminary version in *9th Structure in Complexity Theory Conference*, 1994). [doi:10.1016/0304-3975\(95\)00165-4](https://doi.org/10.1016/0304-3975(95)00165-4).

A Proof of Corollary 7

For each type of game, we compute a lower bound on the second-smallest eigenvalue of the corresponding $(k - 1)$ -connection graph. Applying Theorem 6 then yields the statements of the corollary.

A.1 Free games

For simplicity, assume that $\mu(x)$ is the uniform distribution over $[d]^k$, where $d = |\mathcal{X}^1| = \dots = |\mathcal{X}^k|$.⁵ Then the $(k - 1)$ -connection graph is a weighted version of the d -ary, k -dimensional hypercube (with self loops). Indeed, the corresponding weight function ρ behaves as follows: for $x, x' \in [d]^k$, we have $\rho(x, x) = d^{-(k+1)}$, and $\rho(x, x') = d^{-(k+1)}/k$ when x and x' differ in exactly one coordinate, and is 0 otherwise. If we compute the normalized Laplacian L_H , we get that

$$(L_H)_{u,v} = \begin{cases} 1 - \frac{1}{d} & \text{if } u = v \text{ and } \rho(v) \neq 0 \\ -\frac{1}{kd} & \text{if } \rho(u), \rho(v) \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

This is the normalized Laplacian corresponding to the Cayley graph over the Abelian group $(\mathbb{Z}/d\mathbb{Z})^k$ with (weighted) generators $\{g \in (\mathbb{Z}/d\mathbb{Z})^k : |g| \leq 1\}$ where $|g|$ is the number of non-zero components of g . If $g = (0, 0, \dots, 0)$, then its weight is d^{-1} , and if $|g| = 1$, then its weight is $(kd)^{-1}$. The spectrum of Cayley graphs is well understood; we have that the smallest non-zero eigenvalue of L_H is therefore $\lambda(H) = \frac{1}{k}$. Thus H_G is a $1/k$ -expander.

⁵Indeed, by letting d be large enough, we can approximate μ arbitrarily well through discretization and identifying $[d]$ with \mathcal{X}^t for $t = 1, \dots, k$ in a many-to-one-fashion. Our bounds will not depend on d , so d can be taken to be arbitrarily large.

A.2 Anchored games

Here, we prove a lower bound on the second eigenvalue of the $(k - 1)$ -connection graph of an *anchored game*, and show that it is at least $8k/\alpha^k$. Plugging in this bound into Theorem 6 gives us

$$\text{val}(G^{\otimes n}) \leq \exp\left(-\frac{c \alpha^{2k} \varepsilon^5 n}{64 k^2 \log |\mathcal{A}|}\right).$$

This asymptotically matches the bounds obtained in [1] in terms of the dependence on α and k .

Let us first recall the definition of an anchored game.

Definition 14 (α -anchored games [1]). *Given a k -prover game G , and a parameter $\alpha < 1$ we define the α anchored game G_\perp as follows: the referee chooses a question tuple (x^1, \dots, x^k) , according to G , and independently, for every $t \in [k]$, replaces x^t by the anchoring symbol \perp with probability α to get the tuple (x'^1, \dots, x'^k) . The new domain is thus $\mathcal{X}^1 \times \mathcal{X}^2 \dots \mathcal{X}^k$, where $X^i = \mathcal{X} \cup \{\perp\}$. If any of the x^i 's are \perp , the verifier accepts trivially, otherwise the verifier accepts according to the predicate of the game G .*

For convenience, we will denote the α -anchored game itself by G in this section, and its $(k - 1)$ -connection graph by H_G . We will show the following lemma.

Lemma 15. ⁶ $\lambda(H_G) \geq \alpha^k/8k$, when $\alpha < 1/2$.

In order to prove Lemma 15, we need to make a couple of observations. First, note that the 1-connection graph H_G 's vertices can be partitioned into disjoint sets V_0, V_1, \dots, V_k , where V_i has vertices of all question-tuples with exactly i bottom symbols. Thus, V_0 has vertices corresponding to the original question tuples, and $V_k = \{(\perp, \perp, \dots, \perp)\}$. While V_0 has edges between its own vertices (corresponding to edges in the 1-connection graph of the un-anchored game), all other edges in H_G go between V_i and V_{i+1} .

We will lower bound $\lambda(H_G)$ using the notion of *congestion* in the graph. This technique was first introduced by Diaconis and Strook [6], and improved by Sinclair [13]. The below form can be found in the survey [8, Section 4].

Let us view H_G as an undirected graph⁷, with weight function ρ on the edges. Since $\rho(x, y) = \rho(y, x)$ by our definition, this is well-defined. A set of canonical paths in H_G is a set \mathcal{P} of simple paths, one between every ordered pair (x, y) in H_G . The *path congestion parameter* of this set of canonical paths is defined as follows:

$$\zeta(\mathcal{P}) \triangleq \max_{e \in E(H_G)} \frac{1}{\rho(e)} \sum_{p_{xy} \ni e} \rho(x)\rho(y)|p_{xy}|$$

Above, p_{xy} denotes the path from x to y in \mathcal{P} , and $|p_{xy}|$ is its length. Intuitively, the numerator in the above equation defines the ‘load’ on the edge (x, y) , while $\rho(x, y)$ can be interpreted as its

⁶Although the proof of the lemma can be easily seen to show a bound dependent only on α and k for all $\alpha < 1$, the anchored game definition in [1] sets α to be a constant $< 1/2$. We only state this case, for clarity of exposition and comparison to their result.

⁷On the other hand, if viewed as a directed Markov chain, the transition probability $\Pr[y | x]$ for moving from x to y is exactly $\rho(x, y)/\rho(x)$. The stationary distribution on every vertex is $\rho(x)$.

capacity. Thus, one would naturally expect that if we could find a set of canonical paths with low congestion parameter, the graph must be expanding in some sense. This is formalized in the following theorem:

Theorem 16 ([13], see also [8, Theorem 4.3]). *For any set of canonical paths \mathcal{P} ,*

$$\lambda(H_G) \geq \frac{1}{\zeta(\mathcal{P})}$$

We will prove Lemma 15, by choosing a good set of canonical paths in H_G .

Proof of Lemma 15. Consider two vertices x, y in H_G . Let $\Delta(x, y) = \{i_1, \dots, i_s\} \subseteq [k]$ be the set of (player) indices where the tuples differ, with $i_1 \leq i_2 \leq \dots \leq i_s$. We will define the canonical path from x to y to be the one obtained by flipping each of x^{i_1}, \dots, x^{i_s} to \perp in order, and then flip these to y^{i_1}, \dots, y^{i_s} , but in the reverse order $i_s \rightarrow \dots \rightarrow i_1$. Each flip corresponds to moving along an edge in H_G . Call the set of these canonical paths \mathcal{P} . The path from x to y in \mathcal{P} is exactly the reverse of the path from y to x .

We will upper bound the congestion through any edge $e = \{u, v\}$ caused by \mathcal{P} . If $u, v \in V_0$, then no path in \mathcal{P} passes through this edge, and hence the congestion on e is 0. Suppose that $u \in V_l$, and $v \in V_{l+1}$ for some $l < k$.

We need to find which vertices x would use a canonical path that passes from u to v to reach another vertex. To identify this set, define $B_v \triangleq \{i \in [k] : v_i = \perp\}$, and similarly $B_u \triangleq \{i \in [k] : u_i = \perp\}$. Clearly $|B_v| = l + 1$, $|B_u| = l$, and $B_u \subseteq B_v$. Let us write u as $u = (\perp^l, z_u)$, where the indices are appropriately ordered (with z_u in \bar{B}_u).

For $0 \leq r \leq l$, a vertex $w \in V_r$ will be said to be in the r -th *shadow* of u (denoted by $S_r(u)$), if:

- (a) $w|_{\bar{B}_u} = z_u$, and
- (b) If $B_u = \{j_1, \dots, j_l\}$, with $j_1 \leq \dots \leq j_l$, then $w_{j_q} \neq \perp$ for every $q > l - r$.

The following Claim is easy to verify:

Claim 17. $\rho(S_r(u)) = \Pr_{x \sim \rho}[x^{\bar{B}_u} = z_u] \times \alpha^r (1 - \alpha)^{l-r}$

Proof. Any vertex in $S_r(u)$ can be seen to be generated by the verifier in the following way: pick a random question in the original (un-anchored) game conditioned on $x^{\bar{B}_u} = z_u$, then flip j_1, \dots, j_r to \perp (happens with probability α^r), and leave the others unflipped (happens with probability $(1 - \alpha)^{k-r}$). The probability of not flipping \bar{B}_u (i.e. $(1 - \alpha)^{k-l}$) is accounted for in the distribution ρ of the anchored game. This yields the measure of the set $S_r(u)$ as being the expression given above. \square

Any path in \mathcal{P} that passes through u will necessarily either originate or end in one of its shadows. The length of any canonical path as defined above is at most $2k$. Hence, the load through the edge (u, v) can be upper bounded as follows (denoting $\Pr_{x \sim \rho}[x^{\bar{B}_u} = z_u]$ by $\Pr[z_u]$ for clarity):

$$\begin{aligned}
\sum_{p_{xy} \ni e} \rho(x)\rho(y)|p_{xy}| &\leq 2k \sum_{p_{xy} \ni e} \rho(x)\rho(y) \\
&\leq 4k \sum_{r=0}^l \rho(S_r(x)) \\
&= 4k \sum_{r=0}^l \Pr_{x \sim \rho} [x^{B_u} = z_u] \times \alpha^r (1 - \alpha)^{l-r} \\
&= 4k(1 - \alpha)^l \Pr[z_u] \sum_{r=0}^l \left(\frac{\alpha}{1 - \alpha} \right)^r \\
&\leq 4k(l + 1)(1 - \alpha)^l \Pr[z_u] && \dots \text{ since } \alpha < 1/2 \\
&\leq 8k \Pr[z_u]
\end{aligned}$$

The capacity of edge (u, v) is $\rho(u, v) = \Pr[z_u] \times \alpha^l$. Thus, the congestion along the edge is bounded by

$$\zeta(e) \leq \frac{8k \Pr[z_u]}{\Pr[z_u] \times \alpha^l} = \frac{8k}{\alpha^l}$$

Hence, the maximum congestion is bounded by $\zeta(\mathcal{P}) \leq \frac{8k}{\alpha^k}$, which yields the lower bound $\lambda(H_G) > \frac{\alpha^k}{8k}$, by invoking Theorem 16. \square

A.3 Connected games

This follows from the observation that $\lambda(H) \geq \rho_{min}$ when the graph H is connected. The “in particular” statement follows from the definition of the weight function ρ of the $(k - 1)$ -connection graph: $P_X(x)$ is simply $1/|S|$, and $P_{X^t|x^t}(x^{t'})$ is also at least $1/|S|$.