

Strong Parallel Repetition for Unique Games on Small Set Expanders

Dana Moshkovitz*

March 15, 2021

Abstract

We show that NP-hardness of approximating Boolean unique games on small set expanders can be amplified to the full Unique Games Conjecture on small set expanders. The latter conjecture is known to imply hardness results for problems like BALANCED-SEPARATOR, MINIMUM-LINEAR-REARRANGEMENT and SMALL-SET-EXPANSION that are not known under the Unique Games Conjecture.

Our result follows from: (1) A transformation that fortifies any unique game on a small set expander. (2) An improved parallel repetition theorem for fortified games on small set expanders. Our approach avoids the known limitation on strong parallel repetition of unique games by enlarging the alphabet considerably before repetition. The limitation does not hold for unique games with a sufficiently large alphabet. Prior to this work, parallel repetition from fortification was only known for projection games, and seemed hopeless for unique games. Perhaps surprisingly, the idea for fortification of unique games comes from a reduction of Raghavendra and Steurer from SMALL-SET-EXPANSION to UNIQUE GAMES.

*danama@cs.utexas.edu. Department of Computer Science, UT Austin. This material is based upon work supported by the National Science Foundation under grants number 1648712 and 1705028.

1 Introduction

1.1 Unique Games on Small Set Expanders

The Unique Games Conjecture (UGC) [19] implies optimal inapproximability results that we do not know how to prove otherwise, e.g., for problems like MAX-CUT [20] and VERTEX-COVER [25], and, more broadly, for constraint satisfaction problems on constant-sized alphabets [32]. The conjecture postulates the NP-hardness of the following gap problem:

Definition 1.1 (Unique Games). The input is a regular graph $G = (V, E)$, an alphabet Σ and permutations $\pi_e : \Sigma \rightarrow \Sigma$ for all $e \in E$. A labeling to G is a function $\sigma : V \rightarrow \Sigma$. We say that a labeling satisfies an edge $e = (u, v) \in E$ if $\pi_e(\sigma(u)) = \sigma(v)$. In the $1 - \delta$ vs. ε gap problem the goal is to distinguish between the following two cases:

- *Completeness*: There is a labeling that satisfies $1 - \delta$ fraction of edges.
- *Soundness*: Every labeling satisfies at most ε fraction of the edges.

The fraction of edges that a labeling satisfies is called its value. The maximum value over all labelings for a game is called the value of the game, $\text{value}(\mathcal{G})$. The size of a game $\text{size}(\mathcal{G})$ is the size $|V| + |E|$ of the graph underlying the game

Conjecture 1.2 (Unique Games Conjecture [19]). *For any $0 < \varepsilon, \delta < 1$, there exists $k \geq 1$, such that the $1 - \delta$ vs. ε gap problem of unique games on alphabet of size k is NP-hard.*

The name “unique games” refers to the following two prover game: a verifier picks a uniform edge $e = (u, v) \in E$; sends u to one prover, and sends v to the other prover; each prover returns a label to the vertex it got; the verifier checks the constraint on the edge. Note that given the answer of one prover, there is exactly one satisfying answer for the other prover. (One subtlety is that the two provers should employ the same strategy, i.e., have identical answers for identical questions; otherwise the two prover game corresponds to a bipartite G .)

Despite its many applications, the Unique Games Conjecture does not seem strong enough to prove strong inapproximability results for problems like BALANCED-SEPARATOR and SMALL-SET-EXPANSION. The Unique Games on Small Set Expanders Conjecture [33, 34] is a stronger version of the Unique Games Conjecture that implies strong inapproximability results for problems like BALANCED-SEPARATOR and MINIMUM-LINEAR-ARRANGEMENT [34]. It is equivalent to the hardness of approximating SMALL-SET-EXPANSION [33, 34]. The conjecture requires the constraint graph of the unique game to be a small set expander (SSE), and rules out small clusters of vertices that are largely disconnected from the rest of the graph. Note that we cannot require the constraint graph to be an expander, since unique games on expanders can be solved efficiently [1]. Small set expansion is the next best thing.

Definition 1.3 (Non-expansion). Let $G = (V, E)$ be a D -regular graph. The non-expansion of a set $S \subseteq V$ is the fraction of edges inside S among all edges with at least one endpoint in S , i.e., $\bar{\Phi}(S) = \frac{|E(S, S)|}{D|S|}$.

In an expander, like in a random graph, the non-expansion of a set is equal to its fraction, up to a small additive error term. In a “small set expander” like the noisy hypercube (vertices correspond to binary strings, edges correspond to pairs of strings with a low Hamming distance between them) a set of fraction μ can have large non-expansion μ^ϵ , for $\epsilon > 0$ related to the Hamming distance on edges. In particular, sets of constant fraction $1/e$ can have large non-expansion $1 - \epsilon$, however, small sets of fraction $\mu \ll e^{-1/\epsilon}$ mostly expand.

Small set expansion follows from (and is equivalent to, for a strong definition of small set expansion) an algebraic property called *hypercontractivity*:

Definition 1.4 (Hypercontractivity). We say that $G = (V, E)$ is (p, q) -hypercontractive if for every $f, g : V \rightarrow \mathbb{R}$,

$$\mathbf{E}_{(u,v) \in E} [f(u)f(v)] \leq \|f\|_p \|g\|_{q'},$$

where $q' = \frac{q}{q-1}$.

For example, the adjacency matrix of the noisy hypercube is (p, q) -hypercontractive for $\sqrt{\frac{p-1}{q-1}} \geq \rho$, where $0 < \rho < 1$ corresponds to the average agreement between adjacent strings of the hypercube. (p, q) -hypercontractivity implies that for all $A \subseteq V$ we have $\overline{\Phi}(A) \leq \Pr[A]^{1/p-1/q}$. In this work when we refer to small set expansion, we refer to (p, q) -hypercontractivity for constants $q > p > 0$. Hypercontractivity is convenient to work with since it tensorizes: if it holds for G , it also holds for products of G .

Unique games on small set expanders are known to be hard for basic semidefinite programming algorithms [26], and existing hardness results for unique games are on graphs that are closely related to small set expanders [22, 10, 9, 4, 21, 23] (See Section 1.6 for discussion). Raghavendra and Steurer [33] originally defined the Unique Games on Small Set Expanders Conjecture. They did so with a weak definition of small set expansion, only considering sufficiently small sets. With Tulsiani [34] they proved it implies a stronger conjecture characterizing the full expansion profile of the graph and implying hypercontractivity [30, 41].

Conjecture 1.5 (Unique Games Conjecture on Small Set Expanders). *For any $\varepsilon, \delta > 0$ there is an alphabet Σ such that the $1 - \delta$ vs. ε gap unique games problem on graphs that are small set expanders is NP-hard.*

1.2 Soundness Amplification for Unique Games

The simplest case of unique games is when the alphabet Σ is Boolean. Then each constraint is either equality or no-equality of two bits. If all the constraints are no-equalities, we get MAX-CUT. Goemans and Williamson [15] gave an efficient approximation algorithm for the $1 - \delta$ vs. $1 - \Theta(\sqrt{\delta})$ gap problem of Boolean unique games where there is a specific constant in the $\Theta(\cdot)$ and $\delta > 0$ is small. Assuming the Unique Games Conjecture, there is an NP-hardness result for Boolean unique games that matches the performance of the Goemans-Williamson algorithm [20]. The best NP-hardness we can currently prove (without assumptions) falls short of this, but the case of Boolean alphabet is where the community has come the furthest towards proving hardness of unique games for completeness $1 - \delta$ for small $\delta > 0$: there is NP-hardness for a very narrow gap $1 - \delta$ vs. $1 - 2\delta$ [17, 23], and there is an approach to prove NP-hardness for a much wider gap $1 - \delta$ vs. $1 - C \cdot \delta$ for any constant $C \geq 1$ and sufficiently small $\delta > 0$ [24, 12]. A natural question is:

Can hardness of Boolean unique games be lifted to the full Unique Games Conjecture?

Parallel repetition is an operation that, for any $k \geq 1$, maps a unique game \mathcal{G} over alphabet Σ to a unique game $\mathcal{G}^{\otimes k}$ over larger alphabet Σ^k , decreasing the value exponentially in k [36, 18, 35, 11, 6].

Definition 1.6 (Parallel repetition). Let \mathcal{G} be a game on a constraint graph $G = (V, E)$ with alphabet Σ and constraints π_e on the edges. Let $k \geq 1$. The k -repeated game $\mathcal{G}^{\otimes k}$ is on the constraint graph $G^{\otimes k} = (V^k, E^{\otimes k})$ and on alphabet Σ^k . There is an edge between (u_1, \dots, u_k) and (v_1, \dots, v_k) if $e_i = (u_i, v_i) \in E$ for all $1 \leq i \leq k$. The constraint on the edge is that all the k edges are satisfied.

In the language of two prover games, in the repeated game the verifier picks k tests $e_1 = (u_1, v_1), \dots, e_k = (v_k, u_k)$ of the original game, sends one prover (u_1, \dots, u_k) , and sends the

other prover (v_1, \dots, v_k) . Each prover responds with labels for all the vertices it got, and the verifier checks that all k tests are satisfied.

Ideally, for any constant $\varepsilon > 0$, starting with a Boolean unique games gap of $1 - \delta$ vs. $1 - \frac{\ln(1/\varepsilon)}{\varepsilon} \cdot \delta$ for a very small $\delta > 0$, one could pick $k \approx \frac{\varepsilon}{\delta}$, and decrease the soundness to

$$\left(1 - \frac{\ln(1/\varepsilon)}{\varepsilon} \cdot \delta\right)^{\frac{\varepsilon}{\delta}} \leq e^{-\frac{\ln(1/\varepsilon)\delta}{\varepsilon} \cdot \frac{\varepsilon}{\delta}} = \varepsilon,$$

while keeping the completeness $(1 - \delta)^{\frac{\varepsilon}{\delta}} \geq 1 - \delta \cdot \frac{\varepsilon}{\delta} = 1 - \varepsilon$. Unfortunately, the parameters that parallel repetition is known to achieve, given in the following theorem, fall short of this ideal:

Theorem 1.7 (Parallel repetition theorem [36, 18, 35, 11]). *Let $\Delta > 0$ and $k \geq 1$ where $\Delta \ll 1/\sqrt{k}$. Suppose that \mathcal{G} is a unique game of $\text{value}(\mathcal{G}) \leq 1 - \Delta$. Then, $\text{value}(\mathcal{G}^{\otimes k}) \leq 1 - \Omega(\sqrt{k}\Delta)$.*

Note that even assuming the strongest possible hardness for Boolean unique games, $\Delta = \Theta(\sqrt{\delta})$, Theorem 1.7 with its dependence on \sqrt{k} cannot improve the gap. This led researchers to ask whether there is *strong parallel repetition*:

Definition 1.8 (Strong parallel repetition). Let $\Delta > 0$, $k \geq 1$, where $\Delta \ll 1/k$. Strong parallel repetition is an operation that maps unique games of value at most $1 - \Delta$ to unique games of value at most $1 - \Omega(k\Delta)$.

Strong parallel repetition is known for unique games whose underlying graphs are expanders [38], as well as for a generalization of expanders where only $O(1)$ eigenvalues are close to 1 [40] (like unique games on expanders, these are known to be efficiently solvable [33]; in fact, small set expanders can have a large number $\exp((\log n)^{\Omega(1)})$ of eigenvalues close to 1 [2]).

In contrast, [37, 3] showed that parallel repetition has parameters similar to those of Theorem 1.7 for *any* potentially hard unique game. Let $\text{sdp.value}(\mathcal{G})$ denote the value of a basic semidefinite program for \mathcal{G} . Semidefinite programs consider a relaxed notion of labelings and satisfaction and therefore satisfy $\text{sdp.value}(\mathcal{G}) \geq \text{value}(\mathcal{G})$. Moreover, $\text{sdp.value}(\mathcal{G})$ is (approximately) efficiently computable. Hence, if the $1 - \delta$ vs. $1 - \Delta$ gap problem of unique games is NP-hard, we expect hard instances \mathcal{G} with $\text{value}(\mathcal{G}) \leq 1 - \Delta$ to have $\text{sdp.value}(\mathcal{G}) \geq 1 - \delta$. The limitation is as follows:

Theorem 1.9 (Limitation on strong parallel repetition of unique games [3]). *Let $\delta > 0$, $k \geq 1$. Suppose that \mathcal{G} is a unique game with $\text{sdp.value}(\mathcal{G}) \geq 1 - \delta$. Then, $\text{value}(\mathcal{G}^{\otimes k}) \geq 1 - O(\sqrt{k\delta \log(|\Sigma|/\delta)})$. Moreover, for the special case of Boolean unique games ($|\Sigma| = 2$), $\text{value}(\mathcal{G}^k) \geq 1 - 4\sqrt{k\delta}$.*

Therefore, as long as the alphabet is small, $|\Sigma| \ll 2^{1/\sqrt{\delta}}$, parallel repetition cannot amplify even a $1 - \delta$ vs. $1 - \Theta(\sqrt{\delta})$ gap for unique games, and for Boolean alphabet parallel repetition cannot improve over the parameters in Theorem 1.7.

We wish to emphasize that if one believes that unique games on small set expanders could be hard, then the limitation implies that *parallel repetition of small set expanders behaves completely differently from parallel repetition of expanders*. Strong parallel repetition of the latter is known, whereas strong parallel repetition of the former cannot be true in general.

1.3 The Main Theorem

Our main contribution is soundness amplification for Boolean unique games on small set expanders. A corollary of our main theorem is that NP-hardness for Boolean unique games on small set expanders with gap $1 - \delta$ vs. $1 - C\delta$ for any $C \geq 1$ and sufficiently small $\delta > 0$ implies the Unique Games Conjecture on small set expanders and its applications.

Theorem 1.10 (Main). *Let $q > p > 0$ be constants. For any $\varepsilon > 0$, sufficiently small $\delta > 0$ and $\Delta > \frac{\ln(1/\varepsilon)}{\varepsilon} \cdot \delta$, any Boolean unique game \mathcal{G} on a (p, q) -hypercontractive graph can be efficiently transformed into a unique game \mathcal{G}^* on $((1-2\delta)q+2\delta, q)$ -hypercontractive graph, over an alphabet of size $\exp(\tilde{O}(1/\delta))O(1/\varepsilon)$, and whose size is $\text{size}(\mathcal{G})^{\exp(\tilde{O}(1/\delta))O(\log 1/\varepsilon)}$, such that:*

- *If $\text{value}(\mathcal{G}) \geq 1 - \delta$, then $\text{value}(\mathcal{G}^*) \geq 1 - O(\varepsilon)$.*
- *If $\text{value}(\mathcal{G}) \leq 1 - \Delta$, then $\text{value}(\mathcal{G}^*) \leq O(\varepsilon)$.*

We note that the small set expansion of the graph that underlies \mathcal{G}^* is weaker than the small set expansion of the graph that underlies \mathcal{G} , yet for constants p, q and ε , we still get a small set expander. Moreover, given a unique games theorem as implied by Theorem 1.10, one could get better small set expansion parameters by applying the transformation of Raghavendra-Steurer-Tulsiani [34].

To prove Theorem 1.10 we follow the general approach of “parallel repetition from fortification” [14, 28]:

1. *Fortification:* Transform the Boolean unique game to a different game that satisfies a desired property (“fortified”), and evades the limitation of [37, 3].
2. *Repetition:* Show a strong parallel repetition theorem for unique games that satisfy the property achieved in the previous step.

This approach was successful for *projection games* [28]. In projection games the constraints on the edges are functions, not necessarily permutations. Appropriately, projection games are usually defined on bipartite graphs, and different alphabets and associated with the two parts. It is known how to amplify soundness of projection games [36, 29, 11, 28], and this is used for the optimal inapproximability results known today (e.g., for many constraint satisfaction problems [5, 16, 29, 8] or for SET-COVER [13, 27, 11]).

Fortified games were defined in [28] as follows:

Definition 1.11 (fortified game). Let $\varepsilon > 0$. We say that a projection game \mathcal{G} on a graph $G = (V, E)$ is *fortified*, if for every $S \subseteq V$ with $|S| \geq \varepsilon|V|$, the sub-game induced¹ on S has value at most $\text{value}(\mathcal{G}) + \varepsilon$.

Fortification allows for an incredibly simple analysis of parallel repetition with near ideal decay $\text{value}(\mathcal{G})^k$, and, perhaps surprisingly, it is easy to fortify projection games! (See Sub-section 1.4.4 below). However, the fortification of [28] inherently produces projection games that are not unique games (in fact, this is how it evades the limitation of [36, 3]!), and therefore cannot be used for amplification of unique games.

For reasons that will be clarified in Sub-section 1.5.1, in this work we consider a slightly stronger property than fortification that we call *non-expansion fortification*:

Definition 1.12 (non-expansion fortified value). For a parameter $\varepsilon > 0$ and a unique game \mathcal{G} , the non-expansion fortified value of \mathcal{G} is

$$\text{value}_\varepsilon(\mathcal{G}) = \max_{S \subseteq V: \bar{\Phi}(S) \geq \varepsilon} \text{value}(\mathcal{G}|_S).$$

We say that the game is *non-expansion fortified* if $\text{value}_\varepsilon(\mathcal{G})$ is upper bounded using $\text{value}(\mathcal{G})$. Note that sets of size $\varepsilon|V|$ typically have non-expansion at least ε . Non-expansion fortification applies to smaller sets that have ε non-expansion. In a (p, q) -hypercontractive graph, there are no sets of fraction smaller than ε^c for $1/c = 1/p - 1/q$ with non-expansion ε , and hence non-expansion fortification follows from fortification against sets of fraction ε^c .

In this work we show how to transform any unique game on a small set expander to a non-expansion fortified unique game on a small set expander and over a large alphabet. We then

¹Namely, the sub-game obtained by conditioning on the verifier’s edge landing inside S .

show a strong parallel repetition theorem that – assuming the initial game is non-expansion fortified – shows that the repeated game has low value. Next we give more details about each of these contributions.

1.4 Fortification of Unique Games on Small Set Expanders

We show non-expansion fortification for unique games on small set expanders:

Theorem 1.13 (Fortification). *Let $q > p > 0$ be constants. For any $0 < \varepsilon < 1$, for any sufficiently small $\delta > 0$, and sufficiently large $\Delta = \Theta(\delta)$, any Boolean unique game \mathcal{G} on a (p, q) -hypercontractive graph can be efficiently transformed into a unique game \mathcal{G}' on a $((1 - 2\delta)q + 2\delta, q)$ -hypercontractive graph, over an alphabet of size $\exp(\tilde{O}(1/\delta))$ and whose size is $\text{size}(\mathcal{G})^{\exp(\tilde{O}(1/\delta))}$ such that:*

- If $\text{value}(\mathcal{G}) \geq 1 - \delta$, then $\text{value}(\mathcal{G}') \geq 1 - O(\delta)$.
- If $\text{value}(\mathcal{G}) \leq 1 - \Delta$, then $\text{value}_{\varepsilon^2}(\mathcal{G}') \leq 1 - O(\Delta)$.

Next we explain the main difficulty in fortification and the main ideas that go into solving it. Fortification holds for a repetition of a game. The intuition is that for a product, even restriction to a fairly small set S of vertices leaves a single typical coordinate sufficiently close to uniform. The issue is that proximity to uniform has to be $\Theta(\delta)$, and this requires roughly $\Theta(1/\delta^2)$ repetitions. This number of repetitions is prohibitive for unique games, since the δ completeness error accumulates across repetitions. (This is also the reason why strong parallel repetition fails for unique games.)

The inspiration for fortification comes from a reduction of Raghavendra and Steurer [33] from SMALL-SET-EXPANSION to UNIQUE GAMES, however we find it natural to describe fortification without explicitly mentioning small set expansion. We draw the analogy to small set expansion in Sub-section 1.4.3.

Fortification is done in two steps: In the first step (“take it or leave it”) we decrease the completeness error. This allows us to make repetitions in the second step (“multiple rounds”). The decrease in the completeness error in the first step comes at the cost of introducing a global constraint on the strategy of the provers. The second step removes the constraint by ensuring it holds. The first step keeps the alphabet Boolean; it is one test that averages the outcomes of many tests. The second step increases the alphabet.

Fortification has a parameter $l \approx 1/\delta$. The first step uses δ -noise to keep the graph underlying the take-it-or-leave-it game a small set expander, even if worse than the initial small set expander. It decreases the completeness error to about $2^{-l}\delta$. In the second step the alphabet increases to about 2^l . We omit some of the details in the description below. The full description of fortification can be found in Section 3.

1. *Take-it-or-leave-it:* The verifier picks uniformly edges $e_1 = (u_1, v_1), \dots, e_l = (u_l, v_l) \in E$ and labels $\sigma_1, \dots, \sigma_l \in \Sigma$. It picks labels $\sigma'_1, \dots, \sigma'_l \in \Sigma$ by setting $\sigma'_i = \pi_{e_i}(\sigma_i)$ with probability $1 - \delta$ and picking σ'_i uniformly at random with probability δ . The first prover gets $(u_1, \sigma_1), \dots, (u_l, \sigma_l)$, and the second prover gets $(v_1, \sigma'_1), \dots, (v_l, \sigma'_l)$. Each of those should be interpreted as proposed labels to the vertices. Each prover needs to decide whether it takes the proposal (as a bundle) or leaves it. The goal of the provers is to make the same decision (take it or leave it) on the correlated proposals they got. The global constraint is that each prover must take $\approx 2^{-l}$ fraction of all possible proposals and leave the rest.
2. *Multiple-rounds:* The verifier picks about 2^l correlated pairs of take-it-or-leave-it proposals. The verifier randomly shuffles the proposals for each prover. Each prover must pick about one proposal among the proposals it got. The verifier checks that the provers picked matching proposals. (We remark that the actual multiple-rounds transformation is more

complicated than this because of possible fluctuations in the number of taken proposals within a sample. The final game has several independent repetitions of shuffling, and the players get to choose which to play).

Note that both steps maintain uniqueness of the game. We elaborate on their analysis next.

1.4.1 Analyzing The “Take It Or Leave It” Game

The analysis of the take-it-or-leave-it game is one of the main contributions of the current work. Given \mathcal{G} with $\text{value}(\mathcal{G}) \geq 1 - \delta$ in the completeness case, the intended strategy is to pick suggestions that are *mostly* consistent with a single labeling of value at least $1 - \delta$ (note that within $\approx 1/\delta$ suggestions, one expects to encounter rejecting edges that create inconsistency between suggestions given to different players). By choosing a random cutoff of sufficiently small magnitude, the fraction of taken suggestions is sufficiently small, say about $2^{0.001l} \cdot 2^{-l}$, and the probability of agreement between the players once one of them decides to take the suggestion is $1 - O(\delta)$.

For the soundness, we should show that any set of taken suggestions of fraction $2^{0.001l} \cdot 2^{-l}$, on which the players agree with probability at least $1 - \Delta$ once one of them decides to take the suggestion, can be used to construct a labeling for \mathcal{G} with value at least $1 - O(\Delta)$. The intuition is that the set of taken suggestions is a small set that does not expand in the graph that underlies the take-it-or-leave-it-game. That graph has two components: $G^{\otimes l}$, which is a small set expander since G is a small set expander, and a component that corresponds to labelings in $\mathcal{G}^{\otimes l}$. Any small set that does not expand must be due to the second component, and thereby imply a good labeling for $\mathcal{G}^{\otimes l}$ and hence for \mathcal{G} .

The actual analysis is subtle because of two important issues: (i) Different tuples in V^l may appear with different numbers of labels in the set of taken suggestions, and this induces a non-uniform distribution over V . (ii) The set of taken suggestions is quite large $2^{0.001l} \cdot 2^{-l}$, which implies that a constant fraction of V may get two labels.

Therefore, the labeling we construct is for all vertices but, perhaps, a subset $V_{err} \subseteq V$ of possibly constant fraction of the vertices. We show that there exists a probability distribution \mathcal{D} over the edges E , in which edge gets a probability within a constant factor of $1/|E|$, such that with probability $1 - O(\Delta)$ over the choice of an edge from \mathcal{D} , either the edge lands completely outside V_{err} and its labeling is accepted, or the edge lands completely inside V_{err} . Because of G 's small set expansion, V_{err} must be of fraction at most $O(\Delta)$. Thus, the value of \mathcal{G} when edges are picked according to \mathcal{D} and the labeling is as constructed is at least $1 - O(\Delta)$. This implies the same for the value of \mathcal{G} when edges are picked uniformly.

To construct the labeling along with V_{err} and \mathcal{D} we follow the following plan:

1. Show that there exists a number of labels L , such that the subset of tuples in V^l with approximately (up to a constant factor) L different labels taken is non-expanding in $G^{\otimes l}$.
2. Deduce from the small set expansion of $G^{\otimes l}$ that this subset of V^l is in fact of fraction $1 - O(\Delta)$, and by the upper bound on the number of taken suggestions, $L \leq O(2^{0.001l})$.
3. Pick vertices $v_1, \dots, v_l \in V$ and labels $\sigma_1, \dots, \sigma_l \in \Sigma$ according to the uniform distribution over taken suggestions with approximately L different labels. Pick uniformly $i_0 \in [l]$. Fix $v_1, \dots, v_{i_0-1}, v_{i_0+1}, \dots, v_l$ and $\sigma_1, \dots, \sigma_{i_0-1}, \sigma_{i_0+1}, \dots, \sigma_l$. The distribution over $(v_{i_0}, \sigma_{i_0}) \in V \times \Sigma$ gives rise to:
 - A labeling $V \rightarrow \Sigma$ (by the chain rule, the entropy in σ_{i_0} conditioned on v_{i_0} is small, and hence there is typically one likely label σ_{i_0} per vertex v_{i_0}).
 - A subset $V_{err} \subseteq V$ (that contains the few vertices v_{i_0} with two likely labels σ_{i_0}).
 - A distribution \mathcal{D} over E (which corresponds to the non-uniform distribution over vertices, since some tuples in V^l have more taken labels than others (by a constant factor)).

1.4.2 Analyzing Multiple Rounds

Since the graph that underlies the take-it-or-leave-it game is a small set expander, so is the graph that underlies the shuffling game. Therefore non-expansion fortification follows from fortification. We use the product structure to show that the game is fortified. Our main observation is that we can apply the information-theoretic ideas in the analysis of parallel repetition [36, 18] to do exactly that! For brevity, we quote those ideas as lemmas in the preliminaries (with the proofs appearing in [36, 18, 35] as well as many other sources on parallel repetition). Our analysis is highly condensed by referring to the lemmas.

As we remarked when we introduced multiple-rounds, the actual transformation is more complicated than just repetition and shuffling, and so is the analysis. The verifier generates several rounds of the shuffling game, and asks the players to choose one of them. Interestingly, to analyze this game we rely on the fortification of the shuffling game.

1.4.3 Comparison To Previous Work on Unique Games and Small Set Expansion

In analogy to the work of Raghavendra and Steurer [33] the fortification operation in this paper can be thought of as consisting of two reductions:

$$\text{UNIQUE-GAME-ON-SSE} \rightarrow \text{SMALL-SET-EXPANSION} \rightarrow \text{UNIQUE-GAME-ON-SSE},$$

where take-it-or-leave-it corresponds to the first reduction, and multiple-rounds corresponds to the second reduction. The goal of our reductions (fortification) and their setup (gap, expansion) are different from that of the existing reductions between unique games and small set expansion [33, 34], and therefore our reductions and their analysis are different from existing reductions:

1. There is an extremely simple reduction from UNIQUE-GAME-ON-SSE to SMALL-SET-EXPANSION (see [39]), however the reduction works in the case of *large completeness-soundness gap* (like in the full Unique Games Conjecture), whereas we need to work in the case of *narrow gap* (like in Boolean Unique Games). Our take-it-or-leave-it reduction is therefore more complex, and its analysis is far more subtle.
2. Raghavendra and Steurer [33] showed a reduction from SMALL-SET-EXPANSION to UNIQUE GAMES, and together with Tulsiani [34], they showed a reduction from SMALL-SET-EXPANSION to UNIQUE-GAME-ON-SSE. The reduction of Raghavendra-Steurer can be thought of as a noisy version of our multiple-rounds. Since we are to argue fortification, our analysis is completely different from the RS analysis. We cannot use the reduction of RST for fortification, and we do not need it thanks to starting with a small set expander.

1.4.4 Comparison To Previous Work on Fortification

The fortification operation in [28] is (a derandomization of) the following reduction: The verifier picks uniformly at random $(u, v) \in E$ and vertices $u_1, \dots, u_{l-1}, v_1, \dots, v_{l-1} \in V$. The verifier shuffles u, u_1, \dots, u_{l-1} and sends them to one prover; it shuffles v, v_1, \dots, v_{l-1} and sends them to the other prover. Each prover responds with labels to all the vertices it got. The verifier checks that the edge (u, v) is satisfied. Note that this game is not unique. It does have the same value as the original game, and it introduces a product structure on the questions of each of the provers. It is quite different from the fortification operation in this paper.

1.5 Strong Parallel Repetition of Non-Expansion Fortified Unique Games

In this paper we also give an improved analysis of strong parallel repetition for non-expansion fortified games:

Theorem 1.14 (Strong parallel repetition for non-expansion fortified games). *Let $\varepsilon > 0$ be a sufficiently small constant. Let \mathcal{G} be a game with $\text{value}_{\varepsilon^2}(\mathcal{G}) \leq 1 - \Delta$ and $k = \lceil \ln(1/\varepsilon)/\Delta \rceil$, then $\text{value}(\mathcal{G}^{\otimes k}) < 3\varepsilon$.*

Importantly, the number k of repetitions is about $1/\Delta$, and not $1/\Delta^2$, i.e., the parallel repetition theorem is strong. Also importantly, the non-expansion parameter depends only on the desired value of the repeated game ε , and not on the alphabet Σ of the game (We do not try to optimize the parameter for non-expansion and keep it ε^2 , whereas we only need it to be about $\frac{\varepsilon}{3 \ln(1/\varepsilon)}$). In contrast, the existing analysis of parallel repetition for fortified games had a dependence on the alphabet (See Sub-section 1.5.1 for comparison).

To understand the idea of the analysis, it is instructive to consider the case of $k = 2$ repetitions (also called “rounds”). Assume on way of contradiction that there is a labeling to \mathcal{G}^2 with value much larger than $\text{value}(\mathcal{G})^2$. By the chain rule, the probability of winning the second round *conditioned* on winning the first round is much larger than $\text{value}(\mathcal{G})$. We will show that this cannot happen.

Fix the questions of the first round, so the remaining round is in one-to-one correspondence with the game \mathcal{G} , and the labeling in \mathcal{G}^2 implies a labeling to \mathcal{G} . Partition the questions in the second round $V = V_1 \cup \dots \cup V_M$ according to the label implied for the first round. Edges between different V_i 's do not win the first round of \mathcal{G}^2 . Hence, in the second round, when the game is conditioned on winning the first round, we really play \mathcal{G} conditioned on the edge landing within V_i for some i . Since the provers win \mathcal{G}^2 with noticeable probability, many of the edges land in V_i for some i . By non-expansion fortification, each one of the V_i 's that has many edges within it must have value $\approx \text{value}(\mathcal{G})$. V_i 's that do not have many edges within them have low probability of coming up.

1.5.1 Comparison To Previous Work on Strong Parallel Repetition for Fortified Games

The paper [28] gives a strong parallel repetition for fortified games, but requires fortification against sets of fraction $|\Sigma|^{-k} \varepsilon$ in order to argue that the k -repeated game has value ε . The dependence on the alphabet is prohibitive in our setting, since fortification increases the alphabet² Σ to more than $2^{1/\delta}$ (and necessarily so!). Our analysis of parallel repetition uses the stronger premise on non-expansion fortification as opposed to fortification to eliminate the dependence on $|\Sigma|$.

1.6 An Open Problem Towards The Unique Games Conjecture

It is unknown whether unique games on small set expanders could be hard or not, and it is possible that unique games could be hard, while unique games on small set expanders would be easy. However it seems that hard unique games could have some small set expansion like property: the proof of the 2-to-2 Games Theorem [22, 10, 9, 4, 21, 23] constructs hard instances that are not small set expanders, however *the value of the game restricted to small sets that do not expand is bounded*. We therefore believe that generalizing our main theorem as follows might be crucial to the proof of the Unique Games Conjecture:

Conjecture 1.15 (Soundness amplification for unique games on quasi small set expanders). *For any $\varepsilon > 0$, for sufficiently small $\delta > 0$ and for sufficiently large $\Delta = \Theta(\delta)$, there exists a transformation that maps unique games \mathcal{G} to unique games \mathcal{G}^* such that:*

- *If $\text{value}(\mathcal{G}) \geq 1 - \delta$, then $\text{value}(\mathcal{G}^*) \geq 1 - O(\varepsilon)$.*

²For context, in projection games there are two alphabets: large and small. The fortification of [28] increases the large alphabet but not the small alphabet, and the dependence on the alphabet size in the parallel repetition theorem of [28] is in the small alphabet.

- If $\text{value}_{1-\Delta}(\mathcal{G}) \leq 1 - \Delta$, then $\text{value}(\mathcal{G}^*) \leq O(\varepsilon)$.

The conjecture does not require that the graph that underlies \mathcal{G} is a small set expander, but it does assume that \mathcal{G} restricted to sets with $1 - \Delta$ non-expansion has bounded value, even if the sets have a much smaller fraction than $1 - \Delta$. Note that $\text{value}(\mathcal{G}) \leq 1 - 2\Delta$ implies that for every $S \subseteq V$ of fraction $|S|/|V| \geq 1 - \Delta$, it holds $\text{value}(\mathcal{G}|_S) \leq 1 - \Delta$.

2 Preliminaries

2.1 Small Set Expansion

The following lemma that shows that a product of small set expanders is a small set expander is key to our analysis:

Lemma 2.1 (Product of small set expanders). [30, 41] If $G = (V_G, E_G)$ and $H = (V_H, E_H)$ are (p, q) -hypercontractive, then $G \otimes H = (V_G \times V_H, E_G \otimes E_H)$ is (p, q) -hypercontractive.

The noisy hypercube is a small set expander, see, e.g., [31].

Lemma 2.2 (Noisy hypercube). The δ -noisy hypercube, namely the weighted graph on $\{0, 1\}^n$ where a random edge is picked by picking uniformly $x \in \{0, 1\}^n$ and letting $y_i = x_i$ with probability $1 - \delta$, and letting y_i be random with probability δ , leading to an edge (x, y) , is (p, q) -hypercontractive for $\sqrt{\frac{p-1}{q-1}} \geq 1 - \delta$.

2.2 Information-Theoretic Analysis of Parallel Repetition

We will use the ideas of the information-theoretic analysis of parallel repetition, which we summarize in the next few lemmas: (Below $|\cdot|$ for random variables refers to total variation distance between distributions.)

The first lemma upper bounds the effect of restricting a product distribution to a sufficiently large subset, when examining a single random coordinate:

Lemma 2.3. [36] Let X_1, \dots, X_t be independent random variables. Let W be an event with $\Pr[W] \geq 2^{-w}$. Pick $i_0 \in [t]$ uniformly at random. Then

$$\mathbf{E}_{i_0} [|X_{i_0} - X_{i_0}|W|] \leq \sqrt{\frac{w}{t}}.$$

This lemma does not immediately imply an analysis of repetitions of a two prover game. The reason is that X_{i_0} and $X^{-i_0} = X_1, \dots, X_{i_0-1}, X_{i_0+1}, \dots, X_t$ may be highly dependent given W . For example, assume that X_1, \dots, X_t are random bits, and let W be the event that $\bigoplus_i X_i = 0$. Then $X^{-i_0}|(X_{i_0}, W)$ satisfies $\bigoplus_{i \neq i_0} X_i = X_{i_0}$. Suppose that $(X_1, Y_1), \dots, (X_t, Y_t)$ are questions to the provers in $\mathcal{G}^{\otimes t}$. To show that a labeling that achieves high value in $\mathcal{G}^{\otimes t}$ conditioned on W induces a labeling of high value for \mathcal{G} , a prover that gets X_{i_0} must be able to decide on X^{-i_0} conditioned on W (without any knowledge of Y_{i_0}), and a prover that gets Y_{i_0} must be able to decide on Y^{-i_0} conditioned on W , where $Y^{-i_0} = Y_1, \dots, Y_{i_0-1}, Y_{i_0+1}, \dots, Y_t$ (without any knowledge of X_{i_0}), in such a way that $(X_i, Y_i) \in E$ for all $i \neq i_0$. This seemingly impossible task is the crux of the difficulty in analyzing parallel repetition.

The solution of [36, 18] centers around the definition of a *correlation breaking variable* R : For $i \in [t]$ let $R_i = X_i$ with probability $1/2$, and $R_i = Y_i$ with probability $1/2$ independently at random. Denote $R^{-i_0} = R_1, \dots, R_{i_0-1}, R_{i_0+1}, \dots, R_t$. Note that X^{-i_0} and Y^{-i_0} are independent conditioned on R^{-i_0} and on W .

1. Prove (See Lemma 2.4 below) that $R^{-i_0}|(X_{i_0}, W)$ and $R^{-i_0}|(Y_{i_0}, W)$ are typically close.

2. Use *correlated sampling* (See Lemma 2.5 below) so the provers can agree on a value for R^{-i_0} with high probability.
3. Let the prover holding X_{i_0} pick X^{-i_0} conditioned on R^{-i_0}, W and let the prover holding Y_{i_0} pick Y^{-i_0} conditioned on R^{-i_0}, W .

The next lemma builds on Lemma 2.3. It uses that $(X_1, Y_1), \dots, (X_t, Y_t)$ remain independent when conditioning on R .

Lemma 2.4. [36] *Let $(X_1, Y_1), \dots, (X_t, Y_t)$ be t independent pairs of random variables, where X_1, \dots, X_t and Y_1, \dots, Y_t are over the same sample space. Let W be an event in this sample space, where $\Pr[W] \geq 2^{-w}$. Pick $i_0 \in [t]$ uniformly. For $i \in [t]$ let $R_i = X_i$ with probability $1/2$, and $R_i = Y_i$ with probability $1/2$ independently at random. Denote $R^{-i_0} = R_1, \dots, R_{i_0-1}, R_{i_0+1}, \dots, R_t$. Then*

$$\mathbf{E}_{i_0} [|R^{-i_0}|(X_{i_0}, W) - R^{-i_0}|(Y_{i_0}, W)|] \leq O\left(\sqrt{\frac{w}{t}}\right).$$

Next we give the short definition and analysis of correlated sampling, which was suggested in [18] in this context:

Lemma 2.5 (Correlated sampling). *Let D_1, D_2 be distributions over a finite sample space Ω with $|D_1 - D_2| \leq \varepsilon$. Suppose two non-communicating provers with access to shared randomness. Then the first prover can sample x_1 from D_1 and the second prover can sample x_2 from D_2 , such that $x_1 = x_2$ with probability at least $1 - 2\varepsilon$.*

Proof. Pick uniformly at random pairs $(a_i, p_i) \in \Omega \times [0, 1]$ for $i = 1, 2, 3, \dots$. Each prover picks the first a_i where p_i is at most the probability of a_i according to the prover's distribution. The probability that the provers do not pick the same i is $2|D_1 - D_2| \leq 2\varepsilon$. \square

3 Non-Expansion Fortification for Unique Games

Fortification is done in two steps. In the first, the “take-it-or-leave-it” transformation decreases the completeness error significantly by adding a global constraint on the labeling; the alphabet remains Boolean. In the second, “multiple-rounds” of the game are performed and the alphabet size increases. The low completeness error achieved in the first step allows the many rounds of the second step. The second step gets rid of the global constraint (by verifying it directly), and the product structure it introduces fortifies the game.

3.1 Take It Or Leave It

Let $\Sigma = \{0, 1\}$. Let $\delta > 0$ be sufficiently small. Pick $l = \lceil 10 \log(1/\delta)/\delta \rceil$. Given a unique game \mathcal{G} on a small set expander $G = (V, E)$ with constraints $\pi_e : \Sigma \rightarrow \Sigma$, we define a unique game that we call *take-it-or-leave-it*. To make the description vivid, we use the language of two prover games, with the understanding that the two provers employ the same strategy, i.e., have identical answers for identical questions. In the take-it-or-leave-it game the verifier sends each of the provers a suggestion, and the prover needs to decide whether to take the suggestion or leave it. The suggestions the prover sends the two provers are correlated, and the provers win if they make the same decision on both suggestions. Importantly, the provers must take a noticeable fraction of the suggestions and must leave most of the suggestions. The game is as follows. The verifier picks uniformly at random edges $e_1 = (u_1, v_1), \dots, e_l = (u_l, v_l) \in E$, as well as l labels $\sigma_1, \dots, \sigma_l \in \Sigma$. For every $1 \leq i \leq l$, the verifier sets $\sigma'_i = \pi_{e_i}(\sigma_i)$ with probability $1 - \delta$ and picks σ'_i uniformly at random with the remaining probability. The verifier sends $(u_1, \sigma_1), \dots, (u_l, \sigma_l)$ to one prover. These are suggested labels to the l vertices. Similarly,

the verifier sends $(v_1, \sigma'_1), \dots, (v_l, \sigma'_l)$ to the other prover. Each prover should either take or leave the verifier's suggestion (as a bundle). The verifier checks that the two provers made the same decision, i.e., both took their suggestions or both left their suggestions. Note that the take-it-or-leave-it game is unique.

The size of the game is $\text{size}(\mathcal{G})^{\tilde{O}(1/\delta)}$. The alphabet of the game is Boolean. By Lemma 2.1, if G is (p, q) -hypercontractive, then so is $G^{\otimes l}$. We will use this property repeatedly in the soundness analysis below. The graph underlying the take-it-or-leave-it game is the product of $G^{\otimes l}$ (for the vertices) and a graph (for the labels) that is $((1 - 2\delta)q + 2\delta, q)$ -hypercontractive (like the δ -noisy hypercube; see lemma 2.2). For sufficiently small $\delta > 0$ we know that $p \leq (1 - 2\delta)q + 2\delta$. By the monotonicity of the p -norm, we know that $G^{\otimes l}$ is also $((1 - 2\delta)q + 2\delta, q)$ -hypercontractive. Thus, by Lemma 2.1, the graph underlying the take-it-or-leave-it game is $((1 - 2\delta)q + 2\delta, q)$ -hypercontractive.

A key observation is that the completeness error is about $2^{-l}\delta$ by having the provers take only suggestions that are mostly consistent with a single high value labeling to \mathcal{G} :

Lemma 3.1 (Completeness). *If $\text{value}(\mathcal{G}) \geq 1 - \delta$, then there is a strategy for the provers that accepts $2^{H(0.001)^{l-o(l)}/2^l} \leq \gamma \leq 2^{H(0.01)^l}/2^l$ fraction of the verifier's suggestions, and the probability that one prover accepts while the other rejects is $O(\gamma\delta)$.*

Proof. We describe a strategy for the take-it-or-leave-it game. The provers agree on a strategy for \mathcal{G} that succeeds with probability at least $1 - \delta$. Then they pick a random $\alpha \in [0.001, 0.01]$, and accept if the verifier's suggestions differ from the fixed strategy on at most αl of the l labels they got. Let Diff_1 be the difference for the first prover, and let Diff_2 be the difference of the second prover. There are two sources for difference between Diff_1 and Diff_2 : edges e_i that reject the fixed strategy, and noise in σ'_i . By a Chernoff bound, $|\text{Diff}_1 - \text{Diff}_2| \leq 3\delta l$ except with probability $4e^{-4\delta l/3} \leq \delta$. The probability that the provers' responses are different, conditioned on one of them accepting, is bounded by the probability that the threshold αl fell between Diff_1 and Diff_2 . When $|\text{Diff}_1 - \text{Diff}_2| \leq 3\delta l$, this probability is at most $3\delta l / (0.01 - 0.001)l = O(\delta)$. Moreover, the probability that any one of the prover accepts is the fractional volume of the ball of radius αl around the fixed strategy, and the bound on γ follows. \square

Thanks to the small set expansion of $G^{\otimes l}$ we can also prove soundness, namely that every set of taken suggestions of fraction roughly 2^{-l} on which the provers agree with high probability in the take-it-or-leave-it game leads to a labeling that achieves high value in \mathcal{G} . In the analysis below we follow the plan articulated in the introduction (Section 1.4.1).

Lemma 3.2 (Soundness). *Let $0 < p < q$ be constants. Let $\Delta, \eta > 0$ be sufficiently small constants. If \mathcal{G} is a unique game on a (p, q) -hypercontractive graph with $\text{value}(\mathcal{G}) < 1 - \Delta$, then for every strategy that takes $2^{-(1-\eta)l}$ fraction of the verifier's suggestions, the probability of agreement is at most $1 - O(\Delta)$.*

Proof. Assume on way of contradiction that for a set of taken suggestions of fraction $2^{-(1-\eta)l}$ the probability of agreement between the provers is at least $1 - \Delta'$. For every $\underline{v} = (v_1, \dots, v_l) \in V^l$ denote by $l(\underline{v})$ the number of taken labels for \underline{v} , i.e., the number of $(\sigma_1, \dots, \sigma_l) \in \Sigma^l$ such that the suggestion $(v_1, \sigma_1), \dots, (v_l, \sigma_l)$ is taken. Partition V^l according to the number of taken labels for each tuple: Set $t_0 = 0$. Pick uniformly at random thresholds $0 < t_1 \leq 1 < t_2 \leq 2 \dots < t_l \leq l$. Let S_i contain the tuples \underline{v} such that $2^{t_{i-1}} \leq l(\underline{v}) < 2^{t_i}$. Let S_0 contain the tuples in V^l that are never taken.

Let \mathcal{T} be the distribution induced on V^l by picking a random taken suggestion and considering the tuple in V^l . Pick a random $\underline{u} \in V^l$ according to \mathcal{T} and a uniform choice of a neighboring \underline{v} . Next we will bound by $O(\Delta')$ the probability that $\underline{u} \in S_i$ and $\underline{v} \in S_{i'}$ for $i \neq i'$. Without loss of generality, we focus on the case that $i' < i$, since the distribution of \underline{v} is $O(\Delta')$ -close to \mathcal{T} . The probability that $\underline{u} \in S_i$ and $\underline{v} \in S_{i'}$ for $i' < i$ is at most $\log(l(\underline{u})) - \log(l(\underline{v})) = \log(l(\underline{u})/l(\underline{v}))$.

Only $l(\underline{v})$ of the $l(\underline{u})$ suggestions associated with \underline{u} can lead to agreement in the take-it-or-leave-it game. Write $l(\underline{v})/l(\underline{u}) = 1 - \varepsilon_{\underline{v},\underline{u}}$. Since the probability of agreement is at least $1 - \Delta'$, we have that $\mathbf{E}[l(\underline{v})/l(\underline{u})] \geq 1 - \Delta'$, and thus $\mathbf{E}[\varepsilon_{\underline{v},\underline{u}}] \leq \Delta'$. By a Markov inequality, except with probability $O(\Delta')$, it holds that $\varepsilon_{\underline{v},\underline{u}} \leq 1/2$. In this case,

$$\log(l(\underline{u})/l(\underline{v})) = \log\left(\frac{1}{1 - \varepsilon_{\underline{v},\underline{u}}}\right) \leq \log(1 + O(\varepsilon_{\underline{v},\underline{u}})) \leq O(\varepsilon_{\underline{v},\underline{u}}).$$

Hence, except with probability $O(\Delta')$, we have $\log(l(\underline{u})/l(\underline{v})) \leq O(\varepsilon_{\underline{v},\underline{u}})$. Therefore, one can bound the probability that $\underline{u} \in S_i$ and $\underline{v} \in S_{i'}$ for $i' < i$ by $O(\Delta')$. Hence, the probability that $\underline{u} \in S_i$ and $\underline{v} \in S_{i'}$ for $i \neq i'$ is $O(\Delta')$. It follows that there exists i such that $\overline{\Phi}(S_i) \leq 1 - O(\Delta')$.

By the small set expansion of $G^{\otimes l}$, we know that $|S_i| \geq (1 - O(\Delta'))|V^l|$. Let \mathcal{T}_i be the probability distribution induced on V^l by picking a tuple from \mathcal{T} and conditioning on the tuple landing in S_i . By the definition of S_i , the probability of each tuple in S_i according to \mathcal{T}_i is the same up to a constant factor. Pick (v_1, \dots, v_l) from \mathcal{T}_i and pick a uniform taken labeling for the vertices $(\sigma_1, \dots, \sigma_l) \in \Sigma$. By the upper bound on the number of taken suggestions,

$$H(\sigma_1, \dots, \sigma_l | v_1, \dots, v_l) \leq O(\eta l).$$

By the chain rule,

$$H(\sigma_1, \dots, \sigma_l | v_1, \dots, v_l) = H(\sigma_1 | v_1, \dots, v_l) + H(\sigma_2 | \sigma_1, v_1, \dots, v_l) + \dots$$

Pick uniformly at random $i_0 \in [l]$, then the following is implied:

$$\mathbf{E}_{i_0} [H(\sigma_{i_0} | v_1, \dots, v_l, \sigma_1, \dots, \sigma_{i_0-1}, \sigma_{i_0+1}, \dots, \sigma_l)] \leq O(\eta).$$

Fix $i_0 \in [l]$, vertices $v_1, \dots, v_{i_0-1}, v_{i_0+1}, \dots, v_l \in V$, labels $\sigma_1, \dots, \sigma_{i_0-1}, \sigma_{i_0+1}, \dots, \sigma_l \in \Sigma$, as well as vertices $u_1, \dots, u_{i_0-1}, u_{i_0+1}, \dots, u_l \in V$ with $(v_i, u_i) \in E$ for all i , such that:

- $H(\sigma_{i_0} | v_{i_0}) \leq O(\eta)$.
- The probability of rejection in take-it-or-leave-it conditioned on the fixing is $O(\Delta')$.

(Note that there exists a fixing that satisfies both items by a Markov inequality and a union bound). Let \mathcal{D} be the distribution over (v_{i_0}, u_{i_0}) . Note that each edge has \mathcal{D} probability that is within a constant factor of $1/|E|$.

Let $V_{err} \subseteq V$ contain all fixings of the vertex v_{i_0} such that $H(\sigma_{i_0} | v_{i_0}) > 1/2$, so $|V_{err}| \leq O(\eta)|V|$. For all vertices outside V_{err} there is one label that is more likely than the other, and we consider this labeling. For the vertices inside V_{err} that are two labels that have a constant probability each. By the small set expansion of G , conditioned on $v_{i_0} \in V_{err}$, the probability that $u_{i_0} \in V_{err}$ is $O(\eta)^{1/p-1/q}$. When $u_{i_0} \notin V_{err}$, the probability of accepting in take-it-or-leave-it is at most a constant (since one of v_{i_0} 's labels occurs with constant probability but is inconsistent with u_{i_0} 's label). It follows that $\eta \leq O(\Delta')$. Therefore, with probability at least $1 - O(\Delta')$, we have $v_{i_0}, u_{i_0} \notin V_{err}$ and the labeling of v_{i_0}, u_{i_0} is accepted in \mathcal{G} . The same must be true for the uniform distribution over edges rather than \mathcal{D} , since the probabilities are within a constant from one another. \square

3.2 Multiple Rounds

The previous step puts a global constraint on the strategy: about 2^{-l} fraction of the suggestions are taken. To get rid of the global constraint we generate about 2^l different pairs of correlated suggestions, and check that matching $O(1)$ of them are taken. Importantly, we shuffle each prover's suggestions before giving them to the prover. This ensures that all the rounds have the same set of taken suggestions, and this set must be of the right fraction. If we didn't shuffle,

the provers could win by fixing i_0 and always taking the i_0 'th suggestion (in such a strategy no round has a set of taken suggestions that is of the right fraction: the i_0 'th round has a set of taken suggestions of fraction 1, whereas the other rounds have a set of fraction 0).

One subtlety is that in the completeness strategy of take-it-or-leave-it the fraction of taken suggestions varies. Therefore the number taken suggestions within the verifier's sample may vary. Hence, we first consider the shuffling game described above, but only analyze it on sub-games with $O(1)$ taken suggestions. In the final game the verifier generates many instances of the shuffling game and lets the provers decide which to play.

3.2.1 Shuffling

For a parameter $0 < \delta < 1$, let \mathcal{TL}_δ denote the take-it-or-leave-it game. Let $G[\mathcal{TL}_\delta] = (V[\mathcal{TL}_\delta], E[\mathcal{TL}_\delta])$ be its graph. Let γ be an upper bound on the fraction of taken suggestions in the completeness strategy of take-it-or-leave-it (Lemma 3.1). The multiple rounds game \mathcal{M}_δ is defined as follows: Set $t = \lceil 100/\gamma \rceil = \exp(\Theta(1/\delta))$. Pick $e_1 = (v_1, u_1), \dots, e_t = (v_t, u_t) \in E[\mathcal{TL}_\delta]$ independently at random. Send one prover a random permutation of v_1, \dots, v_t and send the other prover a random permutation of u_1, \dots, u_t . Each prover should take between 1 and 1000 of the t suggestions it got and leave the rest. The provers win if they took corresponding suggestions.

The alphabet of the multiple rounds game, which we denote $\Sigma[\mathcal{M}_\delta]$, is of size $t^{1000} = \exp(\Theta(1/\delta))$. The size of the game is $\text{size}(\mathcal{TL}_\delta)^{\exp(\Theta(1/\delta))}$. Denote the graph that underlies the multiple rounds game \mathcal{M}_δ by $G[\mathcal{M}_\delta]$. This graph is a small set expander with the same parameters as $G[\mathcal{TL}_\delta]$ by Lemma 2.1 (clearly the random permutations do not hurt small set expansion).

Next we prove completeness, but only for a certain sub-game, and soundness for all sufficiently large sub-games.

Lemma 3.3 (Completeness of a sub-game). *Assume that there is a strategy for \mathcal{TL}_δ such that the probability that one prover takes its suggestion and the other does not is $O(\gamma\delta)$. Let*

$$S = \{(v_1, \dots, v_t) \in V[\mathcal{M}_\delta] \mid \mathcal{TL}_\delta \text{ strategy takes between 1 and 1000 of } v_1, \dots, v_t\}.$$

Then there is a strategy for \mathcal{M}_δ , such that if one prover's question is restricted to S , then the verifier accepts with probability $1 - O(\delta)$.

Proof. The provers follow the strategy for \mathcal{TL}_δ for each of the suggestions they get. By a union bound, except with probability $O(\delta)$, every taken suggestion of one prover yields a taken suggestion of the other prover. In particular, if one prover's question falls in S , so is the other prover's question, and the verifier accepts. \square

Lemma 3.4. *Let $\eta \geq 0$. If the strategy for \mathcal{TL}_δ takes $\gamma^{1+\eta}$ fraction of its suggestions, then*

$$S = \{(v_1, \dots, v_t) \in V[\mathcal{M}_\delta] \mid \text{between 1 and 1000 of } v_1, \dots, v_t \text{ are taken}\}$$

is of fraction at least $\min\{\gamma^\eta, 0.99\}$.

Proof. The probability that none of the t suggestions is taken is $(1 - \gamma^{1+\eta})^t \leq 1 - 100\gamma^\eta$. The expected number of taken suggestions is $\gamma^{1+\eta}t \leq \gamma t = 100$, and thus by the Chernoff bound, the probability that more than 1000 suggestions are taken is at most 0.01. \square

In the next lemma we use the ideas of the information-theoretic analysis of parallel repetition (quoted as lemmas in the preliminaries) in order to prove that the game is fortified:

Lemma 3.5 (Fortified soundness). *Let $0 < \Delta < 1$ be sufficiently large $\Delta = \Theta(\delta)$. Suppose that there are $S, T \subseteq V[\mathcal{M}_\delta]$ of fraction 2^{-s} for $s \leq \delta^2 t$, such that the value of \mathcal{M}_δ when one prover's question is restricted to S and the other prover's question is restricted to T is at least $1 - \Delta$, then there must exist a strategy for \mathcal{TL}_δ where $\Theta(\gamma)$ fraction of the suggestions are taken, and the probability that the provers do not win the game is $O(\gamma\Delta)$.*

Proof. Let W be the event that $(u_1, \dots, u_t) \in S$ and $(v_1, \dots, v_t) \in T$. Pick $i_0 \in [t]$ uniformly at random, and let us (approximately) embed \mathcal{TL}_δ in the i_0 'th coordinate of \mathcal{M}_δ (before the random permutations) conditioned on W . By embedding we mean that one can obtain a strategy for \mathcal{TL}_δ as follows: we start with a prover that gets u_{i_0} and a prover that gets v_{i_0} . The first prover can complete u_{i_0} to $(u_1, \dots, u_t) \in S$, and the second prover can complete v_{i_0} to $(v_1, \dots, v_t) \in T$, so overall they generate a test of the \mathcal{M}_δ verifier that is $O(\delta)$ -close to uniform conditioned on landing in S and T . Once we have an approximate embedding we are done: The answer for (u_1, \dots, u_t) induces a take-it-or-leave-it decision for u_{i_0} , and the answer for (v_1, \dots, v_t) induces a take-it-or-leave-it decision for v_{i_0} . If the provers win for $(u_1, \dots, u_t), (v_1, \dots, v_t)$ in \mathcal{M}_δ , then they also win for u_{i_0} and v_{i_0} in \mathcal{TL}_δ . By the random permutation, $\Theta(\gamma)$ fraction of the proposals are taken by the strategy we defined for u_{i_0}, v_{i_0} .

The embedding is defined as follows: By Lemma 2.3,

$$\mathbf{E}_{i_0} [|e_{i_0}|W - e_{i_0}|] \leq \sqrt{\frac{s}{t}} \leq \delta.$$

By Lemma 2.4, for R that satisfies $R_i = u_i$ with probability $1/2$ and $R_i = v_i$ with probability $1/2$ for all $i \neq i_0$,

$$\mathbf{E}_{i_0} [|R|(u_{i_0}, W) - R|(v_{i_0}, W)|] \leq O\left(\sqrt{\frac{s}{t}}\right) \leq O(\delta).$$

Fix $i_0 \in [t]$ such that $|e_{i_0}|W - e_{i_0}| \leq O(\delta)$ and $|R|(u_{i_0}, W) - R|(v_{i_0}, W)| \leq O(\delta)$. The provers, one given u_{i_0} and one given v_{i_0} , can use correlated sampling (Lemma 2.5) and pick the same $R = r$ with probability at least $1 - O(\delta)$. Conditioned on $R = r$ and W , u_1, \dots, u_t is independent of v_1, \dots, v_t , so the prover holding u_{i_0} can complete all of $(u_1, \dots, u_t) \in S$ and the prover holding v_{i_0} can complete all of $(v_1, \dots, v_t) \in T$, such that $(u_1, \dots, u_t), (v_1, \dots, v_t)$ is nearly a uniform test of \mathcal{M}_δ conditioned on W . \square

3.2.2 Varying Fractions of Taken Suggestions

To handle varying fractions of taken suggestions we generate many rounds of the game \mathcal{M}_δ . We show that within all those rounds, with high probability, there exists a round in which the number of taken suggestions is between 1 and 1000. We let the provers choose that round and only play it. One way the provers can take advantage of their choice of a round is to pick a subset S of the questions of \mathcal{M}_δ and only choose rounds that contain questions from S . Fortunately, \mathcal{M}_δ is fortified, so the value of such a sub-game is bounded.

Let $\delta > 0$. Assume that the number of taken suggestions in \mathcal{TL}_δ is between $\gamma^{1+\eta}$ and γ , where $0 < \eta < 1$. We define the final multiple rounds game \mathcal{M}'_δ as follows: Set $k = \lceil 100 \log(1/\delta) \gamma^{-\eta} \rceil = \exp(\tilde{\Theta}(1/\delta))$. Pick uniformly at random edges $e_1 = (u_1, v_1), \dots, e_k = (u_k, v_k) \in E[\mathcal{M}_\delta]$. Send u_1, \dots, u_k to one prover, and send v_1, \dots, v_k to another prover. Each prover should respond with a round $i \in [k]$, as well as with a label $\sigma \in \Sigma[\mathcal{M}_\delta]$. Suppose that the first prover responds with i, σ and the second prover responds with i', σ' . The verifier accepts if $i = i'$ and σ, σ' satisfy the unique test of e_i .

Note that this is a unique game. The alphabet of the game, which we denote by $\Sigma[\mathcal{M}'_\delta]$, is of size $\exp(\tilde{\Theta}(1/\delta))$. The size of the game is $\text{size}(\mathcal{TL}_\delta)^{\exp(\tilde{\Theta}(1/\delta))}$. The graph underlying the game is $G[\mathcal{M}_\delta]^{\otimes k}$, and it is a small set expander with the same parameters as $G[\mathcal{M}_\delta]$.

Lemma 3.6 (Completeness). *Assume that there is a strategy for \mathcal{TL}_δ where between $\gamma^{1+\eta}$ and γ fraction of the suggestions are taken, and whenever one prover takes its suggestion, the other prover takes its suggestion too with probability $1 - O(\delta)$. Then there is a strategy for \mathcal{M}'_δ that the verifier accepts with probability $1 - O(\delta)$.*

Proof. Let $S \subseteq V[\mathcal{M}_\delta]$ be the set defined in Lemma 3.3 for the strategy for \mathcal{TL}_δ . The provers pick the first $i \in [k]$ such that their i 'th vertex is in S , if such exists. They then follow the strategy from Lemma 3.3 to label that vertex. By Lemma 3.4, for every $i \in [k]$ there is probability at least $\min\{\gamma^\eta, 0.99\}$ that u_i falls into S . By a Chernoff bound, the probability that for all $i \in [k]$ we have $u_i \notin S$ is at most $O(\delta)$. By Lemma 3.3, with probability at least $1 - O(\delta)$ both provers pick the same i and the verifier accepts their labels for the i 'th round. \square

Lemma 3.7 (Fortified soundness). *Assume that $0 < \Delta < 1$ is sufficiently large $\Delta = \Theta(\delta)$. Suppose that a set $S \subseteq V[\mathcal{M}'_\delta]$ of fraction at least $(2k/\delta) \cdot 2^{-s}$ for $s \leq 100\delta^2/\gamma$ is such that the value of \mathcal{M}'_δ restricted to S is at least $1 - \Delta$. Then there exists a strategy for \mathcal{TL}_δ where $\Theta(\gamma)$ fraction of the suggestions are taken, and whenever one prover takes its suggestion, the other prover takes its suggestion with probability $1 - O(\Delta)$.*

Proof. Consider the strategy that achieves value at least $1 - \Delta$ for \mathcal{M}'_δ restricted to S . Let $i_0 \in [k]$ be such that the provers choose i_0 with probability at least $1/2k$ (note that there must exist such i_0). Sample $e_1 = (u_1, v_1), \dots, e_{i_0-1} = (u_{i_0-1}, v_{i_0-1}), e_{i_0+1} = (u_{i_0+1}, v_{i_0+1}), \dots, e_k = (u_k, v_k) \in E[\mathcal{M}_\delta]$ conditioned on $(u_1, \dots, u_k), (v_1, \dots, v_k) \in S$ and i_0 being chosen by the provers. Let

$$S' = \{u_{i_0} \in V[\mathcal{M}_\delta] \mid (u_1, \dots, u_k) \in S \wedge i_0 \text{ chosen by prover}\}.$$

$$T' = \{v_{i_0} \in V[\mathcal{M}_\delta] \mid (v_1, \dots, v_k) \in S \wedge i_0 \text{ chosen by prover}\}.$$

The strategy for \mathcal{M}'_δ induces a strategy that achieves expected value $1 - \Delta$ for \mathcal{M}_δ restricted to S', T' (the expectation is over the choice of $e_1, \dots, e_{i_0-1}, e_{i_0+1}, \dots, e_k$). Moreover, S', T' are of fraction at least 2^{-s} for any fixing of $e_1, \dots, e_{i_0-1}, e_{i_0+1}, \dots, e_k$ that has probability at least $\delta \cdot 1/|E[\mathcal{M}_\delta]|^{k-1}$. By a union bound, the probability of *all other* fixings of $e_1, \dots, e_{i_0-1}, e_{i_0+1}, \dots, e_k$ combined is at most δ . Hence, there exist S', T' of fraction at least 2^{-s} and a strategy for \mathcal{M}_δ restricted to S', T' that achieves value at least $1 - O(\Delta)$. By the fortified soundness of \mathcal{M}_δ (Lemma 3.5), there is a strategy for \mathcal{TL}_δ where $\Theta(\gamma)$ fraction of the suggestions are taken, and whenever one prover takes its suggestion, the other prover takes its suggestion too with probability $1 - O(\Delta)$. \square

3.3 Fortification Summary

In this sub-section we compute the parameters of fortification and summarize the proof of Theorem 1.13. Given a Boolean unique game \mathcal{G} over a (p, q) -hypercontractive graph and a completeness parameter $0 < \delta < 1$ we consider \mathcal{M}'_δ . The underlying graph is $((1 - 2\delta)q + 2\delta, q)$ -hypercontractive. The alphabet is of size $\exp(\Theta(1/\delta))$. The size is $\text{size}(\mathcal{G})^{\exp(\Theta(1/\delta))}$. The completeness in Theorem 1.13 follows from the completeness of multiple-rounds (Lemma 3.6) and the completeness of take-it-or-leave-it (Lemma 3.1). To get non-expansion fortification note that since the underlying graph is $((1 - 2\delta)q + 2\delta, q)$ -hypercontractive, there exists $c(\delta) = \Theta_q(1/\delta)$ such that fortification against sets of fraction $\varepsilon^{c(\delta)}$ implies non-expansion fortification against sets of fraction ε^2 . Hence, non-expansion fortified soundness in Theorem 1.13 follows from the fortified soundness of multiple-rounds (Lemma 3.7), and the soundness of take-it-or-leave-it (Lemma 3.2).

4 Strong Parallel Repetition For Non-Expansion Fortified Games

In this section we prove a strong parallel repetition theorem for games that are non-expansion fortified. This theorem applies to projection games, and not just for unique games.

For convenience, we reproduce the definition of non-expansion fortification here. It guarantees bounded value for sub-games induced on sets that do not expand well. Recall that $\bar{\Phi}(S) = \frac{E(S,S)}{|S|D}$.

Definition 4.1 (non-expansion fortified). For a parameter $\varepsilon > 0$ and a unique game \mathcal{G} , the non-expansion fortified value of \mathcal{G} is

$$\text{value}_\varepsilon(\mathcal{G}) = \max_{S \subseteq V: \bar{\Phi}(S) \geq \varepsilon} \text{value}(\mathcal{G}|_S).$$

We say that the game is non-expansion fortified if $\text{value}_\varepsilon(\mathcal{G}) \leq 1 - \Delta$.

Our parallel repetition theorem guarantees near ideal decay assuming that the game is non-expansion fortified (we remark that we do not try to optimize the parameter for non-expansion fortification and keep it ε^2 , whereas we only need it to be about $\frac{\varepsilon}{3 \ln(1/\varepsilon)}$):

Theorem 4.2 (Strong parallel repetition for non-expansion fortified games). *Let $\varepsilon > 0$ be a sufficiently small constant. Let \mathcal{G} be a game with $\text{value}_{\varepsilon^2}(\mathcal{G}) \leq 1 - \Delta$ and $k = \lceil \ln(1/\varepsilon)/\Delta \rceil$, then $\text{value}(\mathcal{G}^{\otimes k}) < 3\varepsilon$.*

Let $\sigma^{*k} : V^k \rightarrow \Sigma^k$ be a labeling for $\mathcal{G}^{\otimes k}$. Pick independently and uniformly at random edges $e_1, \dots, e_k \in E$. In the next lemma we show that for every $1 \leq i < k$, as long as the probability that the first i edges are satisfied is at least ε , the probability that the $(i+1)$ 'th edge is satisfied, conditioned on the first i edges being satisfied, is (approximately) at most $1 - \Delta$. The theorem follows by applying the lemma iteratively.

Lemma 4.3 (Single round). *Let \mathcal{G} be a game with $\text{value}_{\varepsilon^2}(\mathcal{G}) \leq 1 - \Delta$. Let $\sigma^{*k} : V^k \rightarrow \Sigma^k$ be a labeling for $\mathcal{G}^{\otimes k}$. Pick uniformly constraints e_1, \dots, e_k from $\mathcal{G}^{\otimes k}$. Then for every $1 \leq i < k$ where*

$$P_i := \Pr_{e_1, \dots, e_i} [\bigwedge_{j=1}^i \pi_{e_j}] \geq \varepsilon,$$

we have

$$P_{i+1} := \Pr_{e_1, \dots, e_{i+1}} [\bigwedge_{j=1}^{i+1} \pi_{e_j}] < P_i \cdot (1 - \Delta) + 2\varepsilon^2 \Delta.$$

Proof. Let $1 \leq i < k$ where the probability that the first i constraints are satisfied is $P_i \geq \varepsilon$, and assume on way of contradiction that $P_{i+1} > P_i \cdot (1 - \Delta)$.

First we fix the coordinates outside the $(i+1)$ 'th to uniformly chosen edges $e_1 = (u_1, v_1), \dots, e_i = (u_i, v_i), e_{i+2} = (u_{i+2}, v_{i+2}), \dots, e_k = (u_k, v_k) \in E$. Then $\mathcal{G}^{\otimes k}$ restricted to the $(i+1)$ 'th coordinate is in one-to-one correspondence with \mathcal{G} . We use the correspondence to define two partitions of the vertices in V :

1. For every $\sigma_1, \dots, \sigma_i \in \Sigma$ there is a part $S_{\sigma_1, \dots, \sigma_i}$ that contains all those $u \in V$ such that $\sigma^{*k}(u_1, \dots, u_i, u, u_{i+2}, \dots, u_k) = (\sigma_1, \dots, \sigma_i, \dots)$.
2. For every $\sigma_1, \dots, \sigma_i \in \Sigma$ there is a part $T_{\sigma_1, \dots, \sigma_i}$ that contains all those $v \in V$ such that $\sigma^{*k}(v_1, \dots, v_i, v, v_{i+2}, \dots, v_k) = (\sigma_1, \dots, \sigma_i, \dots)$.

Each edge $(u, v) \in E$ where $u \in S_{\sigma_1, \dots, \sigma_i}$ and $v \in T_{\sigma'_1, \dots, \sigma'_i}$, where $(\sigma'_1, \dots, \sigma'_i) \neq (\pi_{e_1}(\sigma_1), \dots, \pi_{e_i}(\sigma_i))$, corresponds to an edge of $E^{\otimes k}$ that is not satisfied by σ^{*k} . Hence,

$$P_i \cdot |E| = \mathbf{E}_{e_1, \dots, e_i, e_{i+2}, \dots, e_k} \left[\sum_{\sigma_1, \dots, \sigma_i} \left| E(S_{\sigma_1, \dots, \sigma_i}, T_{\pi_{e_1}(\sigma_1), \dots, \pi_{e_i}(\sigma_i)}) \right| \right] \geq \varepsilon |E|. \quad (1)$$

With probability at least $P_i - 2\varepsilon^2$, we have (Denoting by D the degree of the graph underlying \mathcal{G})

$$\left| E(S_{\sigma_1, \dots, \sigma_i}, T_{\pi_{e_1}(\sigma_1), \dots, \pi_{e_i}(\sigma_i)}) \right| \geq \varepsilon^2 \cdot \left| S_{\sigma_1, \dots, \sigma_i} \cup T_{\pi_{e_1}(\sigma_1), \dots, \pi_{e_i}(\sigma_i)} \right| \cdot D. \quad (2)$$

Since otherwise we can strictly upper bound the average $P_i |E|$ from (1) by

$$(P_i - 2\varepsilon^2) |E| + \varepsilon^2 \cdot \mathbf{E}_{e_1, \dots, e_i, e_{i+2}, \dots, e_k} \left[\sum_{\sigma_1, \dots, \sigma_i} \left(|S_{\sigma_1, \dots, \sigma_i}| + |T_{\pi_{e_1}(\sigma_1), \dots, \pi_{e_i}(\sigma_i)}| \right) \cdot D \right] = P_i |E|,$$

which leads to a contradiction.

By non-expansion fortification, for every $\sigma_1, \dots, \sigma_i$ such that Inequality (2) holds, the fraction of satisfied edges within $S_{\sigma_1, \dots, \sigma_i} \cup T_{\pi_{e_1}(\sigma_1), \dots, \pi_{e_i}(\sigma_i)}$ is at most $1 - \Delta$. Therefore, $P_{i+1} \leq (P_i - 2\varepsilon^2)(1 - \Delta) + 2\varepsilon^2 = P_i(1 - \Delta) + 2\varepsilon^2\Delta$. \square

Finally, we prove our parallel repetition theorem by applying our one-round lemma:

Proof. (of Theorem 4.2 from Lemma 4.3) Let $\sigma^{*k} : V^k \rightarrow \Sigma^k$ be a labeling for $\mathcal{G}^{\otimes k}$. We upper bound the probability P_i that the first i edges among the k are satisfied by σ^{*k} by induction on i . Specifically, we prove:

$$P_i \leq (1 - \Delta)^i + \varepsilon + 2\varepsilon^2\Delta(i - 1).$$

For $i = 1$, $P_1 = \text{value}(\mathcal{G}) \leq 1 - \Delta$. Assume that the claim holds for i , and let us bound P_{i+1} . If $P_i \leq \varepsilon$ we are done, since $P_{i+1} \leq P_i \leq \varepsilon$. Otherwise, by Lemma 4.3,

$$\begin{aligned} P_{i+1} &< P_i(1 - \Delta) + 2\varepsilon^2\Delta \\ &\leq ((1 - \Delta)^i + \varepsilon + 2\varepsilon^2\Delta(i - 1))(1 - \Delta) + 2\varepsilon^2\Delta \\ &\leq (1 - \Delta)^{i+1} + \varepsilon + 2\varepsilon^2\Delta i. \end{aligned}$$

The inductive claim follows. Note that $P_k \leq (1 - \Delta)^k + \varepsilon + 2\varepsilon^2\Delta(k - 1) \leq \varepsilon + \varepsilon + O(\varepsilon^2 \log(1/\varepsilon)) \leq 3\varepsilon$ for sufficiently small $\varepsilon > 0$. \square

Acknowledgement

The author is grateful to Govind Ramnarayan and Vignesh Manoharan for their collaboration during part of the work on this paper.

References

- [1] S. Arora, S. A. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *Proc. 40th ACM Symp. on Theory of Computing*, pages 21–28, 2008.
- [2] B. Barak, P. Gopalan, J. Håstad, R. Meka, P. Raghavendra, and D. Steurer. Making the long code shorter. In *Proc. 53rd IEEE Symp. on Foundations of Computer Science*, pages 370–379, 2012.
- [3] B. Barak, M. Hardt, I. Haviv, A. Rao, O. Regev, and D. Steurer. Rounding parallel repetitions of unique games. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 374–383, 2008.
- [4] B. Barak, J. A. Kelner, and D. Steurer. Rounding sum-of-squares relaxations. In *Proc. 46th ACM Symp. on Theory of Computing*, pages 31–40, 2014.

- [5] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915, 1998.
- [6] M. Braverman and A. Garg. Small value parallel repetition for general games. In *Proc. 47th ACM Symp. on Theory of Computing*, pages 335–340, 2015.
- [7] M. Braverman and D. Minzer. Optimal tiling of the euclidean space using symmetric bodies. Technical Report 2011.04071, ArXiv, 2020.
- [8] S. O. Chan. Approximation resistance from pairwise independent subgroups. In *Proc. 45th ACM Symp. on Theory of Computing*, pages 447–456, 2013.
- [9] I. Dinur, S. Khot, G. Kindler, D. Minzer, and S. Safra. On non-optimally expanding sets in grassmann graphs. In *Proc. 50th ACM Symp. on Theory of Computing*, 2018.
- [10] I. Dinur, S. Khot, G. Kindler, D. Minzer, and S. Safra. Towards a proof of the 2-to-1 games conjecture? In *Proc. 50th ACM Symp. on Theory of Computing*, pages 376–389, 2018.
- [11] I. Dinur and D. Steurer. Analytical approach to parallel repetition. In *Proc. 46th ACM Symp. on Theory of Computing*, 2014.
- [12] R. Eldan and D. Moshkovitz. Reduction from non-unique games to boolean unique games. Technical Report TR20-093, ECCG, 2020.
- [13] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [14] U. Feige and J. Kilian. Two-prover protocols - low error at affordable rates. *SIAM Journal on Computing*, 30(1):324–346, 2000.
- [15] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.
- [16] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [17] J. Håstad, S. Huang, R. Manokaran, R. O’Donnell, and J. Wright. Improved NP-inapproximability for 2-variable linear equations. *Theory of Computing*, 13(19):1–51, 2017.
- [18] T. Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.
- [19] S. Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 767–775, 2002.
- [20] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for MAX-CUT and other two-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [21] S. Khot, D. Minzer, D. Moshkovitz, and S. Safra. Small set expansion in the johnson graph. Technical Report TR18-078, ECCG, 2018.
- [22] S. Khot, D. Minzer, and S. Safra. On independent sets, 2-to-2 games, and grassmann graphs. In *Proc. 49th ACM Symp. on Theory of Computing*, pages 576–589, 2017.
- [23] S. Khot, D. Minzer, and S. Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. In *Proc. 59th IEEE Symp. on Foundations of Computer Science*, 2018.
- [24] S. Khot and D. Moshkovitz. Candidate hard unique game. In *Proc. 48th ACM Symp. on Theory of Computing*, pages 63–76, 2016.
- [25] S. Khot and O. Regev. Vertex cover might be hard to approximate to within 2-epsilon. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.

- [26] S. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into l_1 . In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 53–62, 2005.
- [27] D. Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$ -approximating set-cover. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX 2012*, volume 7408, pages 276–287, 2012.
- [28] D. Moshkovitz. Parallel repetition from fortification. In *Proc. 55th IEEE Symp. on Foundations of Computer Science*, 2014.
- [29] D. Moshkovitz and R. Raz. Two query PCP with sub-constant error. *Journal of the ACM*, 57(5), 2010.
- [30] C. Nair. Equivalent formulations of hypercontractivity using information measures. In *International Zurich Seminar on Information and Communication*, 2014.
- [31] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [32] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proc. 40th ACM Symp. on Theory of Computing*, pages 245–254, 2008.
- [33] P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. In *Proc. 42nd ACM Symp. on Theory of Computing*, pages 755–764, 2010.
- [34] P. Raghavendra, D. Steurer, and M. Tulsiani. Reductions between expansion problems. In *Proceedings of the 27th Conference on Computational Complexity*, pages 64–73, 2012.
- [35] A. Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.
- [36] R. Raz. A parallel repetition theorem. In *SIAM Journal on Computing*, volume 27, pages 763–803, 1998.
- [37] R. Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing*, 40(3):771–777, 2011.
- [38] R. Raz and R. Rosen. A strong parallel repetition theorem for projection games on expanders. In *IEEE Conference on Computational Complexity*, pages 247–257, 2012.
- [39] D. Steurer. *On the Complexity of Unique Games and Graph Expansion*. PhD thesis, Princeton University, 2010.
- [40] M. Tulsiani, J. Wright, and Y. Zhou. Optimal strong parallel repetition for projection games on low threshold rank graphs. In *ICALP*, pages 1003–1014, 2014.
- [41] Y. Zhao. *Generations and Applications of Hypercontractivity and Small-Set Expansion*. PhD thesis, Carnegie Mellon University, 2020.