The Parallel Repetition of Non-signaling Games

Counterexamples and Dichotomy

Justin Holmgren

Princeton University

Princeton, NJ, USA

justin.holmgren@princeton.edu

Lisa Yang

MIT

Cambridge, MA, USA

lisayang@mit.edu

ABSTRACT

Non-signaling games are an important object of study in the theory of computation, for their role both in quantum information and in (classical) cryptography. In this work, we study the behavior of these games under parallel repetition.

We show that, unlike the situation both for classical games and for two-player non-signaling games, there are k -player non-signaling games (for $k \geq 3$) whose values do not tend to 0 with sufficient parallel repetition. In fact, parallel repetition sometimes does not decrease their value whatsoever.

We show that in general, every game's non-signaling value under parallel repetition is either lower bounded by a positive constant or decreases exponentially with the number of repetitions. Furthermore, exponential decrease occurs if and only if the game's sub-non-signaling value (Lancien and Winter, CJTCS '16) is less than 1.

CCS CONCEPTS

• Theory of computation → Interactive proof systems; Complexity theory and logic.

KEYWORDS

parallel repetition, non-signaling, multi-player games

ACM Reference Format:

Justin Holmgren and Lisa Yang. 2019. The Parallel Repetition of Non-signaling Games: Counterexamples and Dichotomy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC '19)*, June 23–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3313276.3316367>

1 INTRODUCTION

A multi-player game $\mathcal{G}$ consists of an interaction between a referee and k players $P_1, \dots, P_k$. The referee samples k questions $q_1, \dots, q_k$ from some joint distribution π and sends q_i to P_i . The players then, without communicating amongst themselves, respond with answers $a_1, \dots, a_k$, and are judged to *win* or *lose* the game according to a predicate $W(q_1, \dots, q_k, a_1, \dots, a_k)$. The (classical) value of the game, denoted $v(\mathcal{G})$, is the maximum probability with

which players can win. The study of multi-player games has been profoundly impactful in diverse areas of theoretical computer science, with foundational applications in complexity theory, hardness of approximation, and cryptography (see e.g., [BGKW88, BFL91, BFLS91, Kil92, FGL⁺96, ALM⁺98, Hås01]).

Parallel repetition is a natural and important operation on any game, originally motivated by its potential application to the soundness amplification of interactive proof systems [FRS88, FL92]. In the n -fold parallel repetition of $\mathcal{G}$, denoted $\mathcal{G}^n$, the referee now samples n independent sets of questions $\{(q_1^{(i)}, \dots, q_k^{(i)})\}_{i=1}^n$ from π , and sends *all* n questions $\{q_j^{(i)}\}_{i=1}^n$ to the j^{th} player at once for each $1 \leq j \leq k$. The players are then required to win all n copies of the game; that is, the j^{th} player must produce $\{a_j^{(i)}\}_{i=1}^n$ so that $W(q_1^{(i)}, \dots, q_k^{(i)}, a_1^{(i)}, \dots, a_k^{(i)}) = 1$ for every $1 \leq i \leq n$. Despite the simplicity of this transformation, it is surprisingly tricky to analyze. It was initially claimed that $v(\mathcal{G}^n)$ is always equal to $v(\mathcal{G})^n$ [FRS88]. This was quickly disproved, however, by Fortnow [For89] and Feige [Fei91], who constructed a two-player game $\mathcal{G}$ satisfying $v(\mathcal{G}^2) = v(\mathcal{G}) < 1$. However, Raz's celebrated parallel repetition theorem proves that for any non-trivial two-player game, the value of $\mathcal{G}^n$ decreases exponentially with n [Raz98]. That is, if $v(\mathcal{G}) < 1$, then $v(\mathcal{G}^n) \leq \bar{v}^n$ for some $\bar{v} < 1$ that depends on $\mathcal{G}$ (but may be significantly larger than $v(\mathcal{G})$).

An important variant on the aforementioned classical value of a game is its non-signaling value v_{ns} , originally defined with the intent of modeling players that may share quantum entanglement. This is the players' maximum winning probability with a relaxed non-communication restriction. Specifically, they may communicate freely with one another to generate any distribution of answers $(a_1, \dots, a_k)$ with a single caveat: for each $S \subseteq [k]$, the *marginal distribution* a_S must be a function only of q_S . While the distinction is perhaps unintuitive, a game's non-signaling value and classical value often differ. The usual example of this is the (two-player) CHSH game [CHSH69], whose classical and non-signaling values are respectively $3/4$ and 1 .¹

Specific non-signaling games have also recently found applications in cryptography. In particular, multi-prover interactive proofs (MIPs) with soundness against non-signaling provers [IKM09, KRR13, KRR14] are indispensable in obtaining non-interactive delegation of computation from standard assumptions [KRR13, KRR14, BHK17, BKK⁺18]. In this work, we focus on the following question:

How does parallel repetition affect a game's non-signaling value?

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
STOC '19, June 23–26, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6705-9/19/06...\$15.00
<https://doi.org/10.1145/3313276.3316367>

¹An arbitrarily large gap between the two values can be achieved by parallel repeating the CHSH game.

Prior Work on Non-Signaling Parallel Repetition. In a simplification and generalization of Raz’s work, Holenstein showed that if a two-player game $\mathcal{G}$ has non-signaling value $v < 1$, then the non-signaling value of $\mathcal{G}^n$ is bounded by $\bar{v}^n$ for some $\bar{v} < 1$ that depends only on v [Hol09].

The multi-player case was considered by Buhrman, Fehr, and Schaffner, who proved a similar result for *complete support games* – games in which every combination of queries is asked with positive probability [BFS14]. Their result comes with the caveat that $\bar{v}$ depends not only on v , but also on the minimum probability with which a query is asked. Arnon-Friedman, Renner, and Vidick give an alternative proof of the same result, and they further note that any game $\mathcal{G}$ can be modified (by adding dummy queries) to have complete support without significantly affecting its non-signaling value [ARV16].

At this point, we find it prudent to point out that games without complete support are quite important in “practical” applications. In particular the non-signaling MIPs of [KRR13, KRR14] and of follow-up works heavily rely on games whose query distributions have incomplete support. These MIPs can be viewed as a family of games indexed by binary strings x , where for some time- T decidable language $\mathcal{L}$, the non-signaling value of game x indicates whether $x \in \mathcal{L}$. In the MIPs of [KRR13, KRR14] and follow-ups, even the *support* of the query distribution has size that is superpolynomial in T . Thus, the “dummy query” approach of [ARV16] (or any approach that only tweaks the query distribution) *inherently* assigns $T^{-\omega(1)}$ probability to some query. As a result, the aforementioned bounds of [BFS14, ARV16] do not achieve even a constant soundness error unless the number of repetitions is super-polynomial in T . This corresponds to a trivial MIP, as a verifier can check for himself whether $x \in \mathcal{L}$ in time T with no error.

The results of [BFS14, ARV16] were generalized by Lancien and Winter [LW16], who defined a further relaxation of a game’s non-signaling value, called its *sub-non-signaling value* (v_{ns}). They show that this value does behave nicely under parallel repetition: for any k -player game $\mathcal{G}$ with $v_{\text{ns}}(\mathcal{G}) < 1$, it holds for all n that $v_{\text{ns}}(\mathcal{G}^n) \leq \bar{v}^n$ for a value $\bar{v} < 1$ that depends only on v and k .

Still, the question of how parallel repetition can affect a general game’s non-signaling value has remained open.

1.1 Our Results

Perhaps our most surprising result is the following counterexample to parallel repetition. In multi-player classical games [Ver96], two-player entangled games [Yue16], and two-player non-signaling games [Hol09], it is known that the value of any non-trivial repeated game decreases to 0 after sufficiently many parallel repetitions. For multi-player non-signaling games, we show that this does not hold in general.

Theorem 1 (Parallel Repetition Counterexample). *For every $k \geq 3$, there exists a k -player game $\mathcal{G}$ such that $v_{\text{ns}}(\mathcal{G}) < 1$ and $v_{\text{ns}}(\mathcal{G}^n) \geq \Omega(1)$.*

We also prove that if the non-signaling value of the repeated game decays to 0, then the decay must be exponential. This establishes a converse to the main result of [LW16], and is not at all obvious. In fact, there are several contexts (e.g., classical multi-player games and entangled two-player games) where it is known

that the value of a game decays to 0 with sufficient parallel repetition, but it is not known (and is a major open question) whether this decay is always exponential.

Theorem 2 (Parallel Repetition Dichotomy). *For every game $\mathcal{G}$, either $v_{\text{ns}}(\mathcal{G}^n) \geq \Omega(1)$ or $v_{\text{ns}}(\mathcal{G}^n) \leq \exp(-\Omega(n))$. The former occurs when $v_{\text{ns}}(\mathcal{G}) = 1$ and the latter occurs when $v_{\text{ns}}(\mathcal{G}) < 1$.*

Thus our counterexample is general: for *any* game $\mathcal{G}$ with sub-non-signaling value 1, it holds that $v_{\text{ns}}(\mathcal{G}^n) \geq \Omega(1)$. Finally, we show that if a game has sufficiently small *non-signaling* value, then the non-signaling value of the repeated game decays exponentially.

Theorem 3 (Parallel Repetition Magic Value). *For every k , there exists a constant $\alpha_k > 0$ such that if $v_{\text{ns}}(\mathcal{G}) < \alpha_k$, then $v_{\text{ns}}(\mathcal{G}^n) \leq \exp(-\Omega(n))$.*

Theorems 1 to 3 each directly follow from our main lemma, which is simple both to state and to use, and may be of independent interest.

Lemma 1 (Non-signaling Value Lower Bound). *For every k , there exists a constant $\alpha_k \geq 2^{-O(k^2)}$ such that for any k -player game $\mathcal{G}$ with $v_{\text{ns}}(\mathcal{G}) = 1$, it holds that $v_{\text{ns}}(\mathcal{G}) \geq \alpha_k$.*

Beyond our general results, we also show that there are *specific* games whose non-signaling values do not decrease *at all* under parallel repetition, subsuming Theorem 1.

Theorem 4 (Strong Parallel Repetition Counterexample). *For every $k \geq 3$, there exists a k -player game $\mathcal{G}$ such that $v_{\text{ns}}(\mathcal{G}) < 1$ and $v_{\text{ns}}(\mathcal{G}^n) = v_{\text{ns}}(\mathcal{G})$ for any $n \geq 1$.*

Table 1: Known bounds on the worst-case (slowest) decay for $v(\mathcal{G}^n)$ or $v_{\text{ns}}(\mathcal{G}^n)$ for a game $\mathcal{G}$ with $v(\mathcal{G}) = 1 - \epsilon$ or $v_{\text{ns}}(\mathcal{G}) = 1 - \epsilon$ respectively. $\mathcal{A}$ denotes the set of possible player answers in $\mathcal{G}$. Ackermann⁻¹ denotes the inverse Ackermann function.

	Two-player games	Multi-player games
Classical	$e^{-\Omega\left(\frac{\epsilon^3 n}{\log \mathcal{A} }\right)}$	$O\left(\frac{1}{\text{Ackermann}^{-1}(n)}\right)$
Non-Signaling	$e^{-\Omega(\epsilon^2 n)}$	$\geq \Omega(1)$ [This Work]

1.2 Our Techniques

We give an overview of the proof of our main technical lemma, Lemma 1. We prove Lemma 1 without referencing sub-non-signaling strategies by (1) introducing a condition that is equivalent to a game $\mathcal{G}$ having sub-non-signaling value 1 and (2) proving implications of this condition on the non-signaling value of $\mathcal{G}$.

Suppose that $\mathcal{G}$ is a k -player game in which the joint distribution of questions asked is π (over a set $Q = Q_1 \times \dots \times Q_k$), and the i^{th} player’s answer is expected to lie in a set $\mathcal{A}_i$. Let $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k$. We will often say “the support of $\mathcal{G}$ ” to refer to the support of its question distribution π . We say “the S -marginal support of $\mathcal{G}$ ” to refer to the support of the marginal distribution π_S .

An honest-referee non-signaling strategy for $\mathcal{G}$ is a function $\tilde{P}$ that maps $q = (q_1, \dots, q_k)$ in the support of $\mathcal{G}$ to a distribution on

$\mathcal{A}$, satisfying the following non-signaling “consistency” condition: for any q, q' in the support of $\mathcal{G}$, if $S \subseteq [k]$ is a set of players for which $q_S = q'_S$, then it must also hold that $\tilde{P}(q)_S = \tilde{P}(q')_S$. This is in contrast to a (full) non-signaling strategy, which must be defined on the entirety of $\mathcal{Q}$ and must satisfy the non-signaling condition for any $q, q' \in \mathcal{Q}$.

It turns out (Proposition 3.6) that a game $\mathcal{G}$ has sub-non-signaling value 1 iff there is an honest-referee non-signaling strategy $\tilde{P}$ that wins $\mathcal{G}$ with probability 1. To prove Lemma 1, we construct a full non-signaling strategy P that has significant “agreement” with $\tilde{P}$ for any honest-referee non-signaling strategy $\tilde{P}$.

A natural first attempt at constructing P might be to somehow extend $\tilde{P}$ into a full non-signaling strategy. However, this approach is doomed to fail in general. A specific counterexample is given by the “anti-correlation” games, first defined in [LW16, ARV16]. We include a formal description of these games for completeness (and with slightly greater generality) in Definition 4.2. These games have honest-referee non-signaling value 1 but non-signaling value less than 1.

Instead, we show that it is always possible to extend $\alpha_k \cdot \tilde{P}$ for some constant $0 < \alpha_k < 1$. Specifically, we construct a non-signaling strategy P such that for every q in the support of $\mathcal{G}$ and every $a \in \mathcal{A}$, it holds that $\Pr_{A \leftarrow P(q)}[A = a] \geq \alpha_k \cdot \Pr_{A \leftarrow \tilde{P}(q)}[A = a]$. This implies that P ’s winning probability for $\mathcal{G}$ is at least α_k times that of $\tilde{P}$.

1.2.1 Three-Player Games. For simplicity of exposition, we first consider three-player games.

The following somewhat easier task serves both as a building block and as a warm-up in our construction of P . For every $S \subseteq \{1, 2, 3\}$ with $|S| \leq 2$ and every q in $\mathcal{Q}_S$ (in particular, q may not be in the S -marginal support of $\mathcal{G}$), we construct a distribution $\mu_S(q)$ on $\mathcal{A}_S$ such that $\{\mu_S(q)\}_{S,q}$ satisfies the following consistency conditions.

Condition 1 (Agreement with $\tilde{P}$). *For every q in the support of $\mathcal{G}$ and every $S \subseteq \{1, 2, 3\}$ with $|S| \leq 2$, the marginal distribution $\tilde{P}(q)_S$ is equal to $\mu_S(q_S)$.*

Condition 2 (Non-Signaling). *The distribution $\mu_S(q)_T$ depends only on T and q_T , and not on S or the rest of q . Formally: for any $S, S' \subseteq [3]$ with $|S|, |S'| \leq 2$, any $T \subseteq S \cap S'$, any $q \in \mathcal{Q}_S$ and $q' \in \mathcal{Q}_{S'}$, if $q_T = q'_T$ then $\mu_S(q)_T = \mu_{S'}(q')_T$.*

We define $\mu_S(q)$ in one of two different ways, depending on whether q is in the S -marginal support of $\mathcal{G}$.

If Yes: Let $\bar{q}$ be in the support of $\mathcal{G}$ such that $\bar{q}_S = q$. We then define $\mu_S(q) \stackrel{\text{def}}{=} \tilde{P}(\bar{q})_S$. This is well-defined (i.e. does not depend on the choice of $\bar{q}$) because $\tilde{P}$ is honest-referee non-signaling.

This “Yes” case behavior is sufficient to ensure that Condition 1 holds.

If No: This means that $|S| = 2$. Say that $S = \{i, j\}$ for $i < j$. We define $\mu_S(q)$ to be the distribution on (a_i, a_j) obtained by independently sampling $a_i \leftarrow \mu_{\{i\}}(q_i)$ and $a_j \leftarrow \mu_{\{j\}}(q_j)$.

Condition 2 holds because if $(S, q) \neq (S', q')$ but $q_T = q'_T$, then $|T| = 1$. Both the “Yes” case and the “No” case define $\mu_S(q)$ so that any 1-marginal $\mu_S(q)_T$ depends only on q_T . In the “Yes” case this

follows from the fact that $\tilde{P}$ is honest-referee non-signaling, and in the “No” case this follows from reduction to the “Yes” case.

We now construct the (full) non-signaling strategy P in terms of the distributions $\{\mu_S(q)\}$. We again have two cases, this time corresponding to whether q is in the support of $\mathcal{G}$.

If Yes: We define $P(q)$ to be the distribution on (a_1, a_2, a_3) obtained as follows. With probability $1/3$ sample $a \leftarrow \tilde{P}(q)$. Otherwise (with probability $2/3$), independently sample $a_1 \leftarrow \tilde{P}(q)_1$, $a_2 \leftarrow \tilde{P}(q)_2$, and $a_3 \leftarrow \tilde{P}(q)_3$. This “Yes” case behavior is sufficient to ensure that P extends $\frac{1}{3} \cdot \tilde{P}$.

If No: We define $P(q)$ to be the distribution on (a_1, a_2, a_3) obtained as follows. Pick a uniformly random $i \leftarrow \{1, 2, 3\}$, define $S = \{1, 2, 3\} \setminus \{i\}$, and independently sample $a_S \leftarrow \mu_S(q_S)$ and $a_i \leftarrow \mu_{\{i\}}(q_i)$.

It remains to verify that P is non-signaling. For this, we must check that for all $q, q' \in \mathcal{Q}$ and all $S \subseteq \{1, 2, 3\}$ such that $q_S = q'_S$, it holds that

$$P(q)_S = P(q')_S. \quad (1)$$

If $|S| = 0$ then Eq. (1) is vacuous, and if $|S| = 3$ then Eq. (1) is trivial. We claim that without loss of generality, we can focus on $|S| = 2$ as the remaining case. Indeed, suppose that we have established Eq. (1) for all $|S| = 2$ (and all q, q' for which $q_S = q'_S$). If $q_T = q'_T$ for $|T| = 1$, then we can construct $\hat{q} \in \mathcal{Q}$ and sets $S, S' \supset T$ (with $|S| = |S'| = 2$) such that $q_S = \hat{q}_S$ and $q_{S'} = q'_{S'}$. Thus,

$$\begin{aligned} P(q)_T &= (P(q)_S)_T = (P(\hat{q})_S)_T = P(\hat{q})_T \\ &= (P(\hat{q})_{S'})_T = (P(q')_{S'})_T = P(q')_T. \end{aligned} \quad (2)$$

Now without loss of generality suppose that $S = \{1, 2\}$ (the other cases follow analogously due to the symmetry of P). For any $q \in \mathcal{Q}$, if q is in the support of $\mathcal{G}$, then

$$P(q) = \frac{1}{3} \cdot \tilde{P}(q)_{\{1,2,3\}} + \frac{2}{3} \cdot \tilde{P}(q)_{\{1\}} \times \tilde{P}(q)_{\{2\}} \times \tilde{P}(q)_{\{3\}}$$

and so

$$P(q)_S = \frac{1}{3} \cdot \tilde{P}(q)_{\{1,2\}} + \frac{2}{3} \cdot \tilde{P}(q)_{\{1\}} \times \tilde{P}(q)_{\{2\}}. \quad (3)$$

If q is not in the support of $\mathcal{G}$, then

$$P(q) = \frac{1}{3} \cdot \begin{pmatrix} \mu_{\{1,2\}}(q_{\{1,2\}}) \times \mu_{\{3\}}(q_{\{3\}}) \\ + \mu_{\{1,3\}}(q_{\{1,3\}}) \times \mu_{\{2\}}(q_{\{2\}}) \\ + \mu_{\{2,3\}}(q_{\{2,3\}}) \times \mu_{\{1\}}(q_{\{1\}}) \end{pmatrix}$$

and so

$$P(q)_S = \frac{1}{3} \cdot \mu_{\{1,2\}}(q_{\{1,2\}}) + \frac{2}{3} \cdot \mu_{\{1\}}(q_1) \times \mu_{\{2\}}(q_2). \quad (4)$$

We now show that Eq. (1) holds for all $q, q' \in \mathcal{Q}$ with $q_S = q'_S$ for $|S| = 2$.

- If both q and q' are in the support of $\mathcal{G}$, then both $P(q)_S$ and $P(q')_S$ are both given by Eq. (3), and Eq. (1) follows respectively from the fact that $\tilde{P}$ is honest-referee non-signaling.
- If neither q nor q' are in the support of $\mathcal{G}$, then $P(q)_S$ and $P(q')_S$ are given by Eq. (4), and Eq. (1) follows from Condition 2.
- If exactly one of q and q' is in the support of $\mathcal{G}$, then Eq. (1) follows from Condition 1.

1.2.2 Games With More Than Three Players. We now return to the case that $\mathcal{G}$ is a general k -player game. In this case, we construct the non-signaling strategy P from $\tilde{P}$ iteratively, as follows. We will maintain the invariant that at the i^{th} step of our construction, we have two components:

- A collection of distributions $\mu^{(i)} = \{\mu_S^{(i)}\}_{|S| \leq i}$, where for each $S \subseteq [k]$ and $q \in Q_S$, $\mu_S^{(i)}(q)$ is a distribution on $\mathcal{A}_S$.
- An honest-referee non-signaling strategy $\tilde{P}^{(i)}$.

We will ensure that these components satisfy the following consistency conditions.

Condition 3 (Non-Signaling). For any S and $q \in Q_S$, the distribution $\mu_S^{(i)}(q)_T$ depends only on T and q_T , and not on S or the rest of q . Formally: for any $S, S' \subseteq [k]$ with $|S|, |S'| \leq i$, any $T \subseteq S \cap S'$, any $q \in Q_S$ and $q' \in Q_{S'}$, if $q_T = q'_T$ then $\mu_S^{(i)}(q)_T = \mu_{S'}^{(i)}(q')_T$.

Condition 4 (Agreement). For every q in the support of $\mathcal{G}$ and every $S \subseteq [k]$ with $|S| \leq i$, the marginal distribution $\tilde{P}^{(i)}(q)_S$ is equal to $\mu_S^{(i)}(q_S)$.

Condition 5. $\tilde{P}^{(i)}$ is an extension of $\alpha_i \cdot \tilde{P}$ for some $\alpha_i > 0$ – i.e., for each q in the support of $\mathcal{G}$ and each $a \in \mathcal{A}$, we have $\Pr_{A \leftarrow \tilde{P}^{(i)}(q)}[A = a] \geq \alpha_i \cdot \Pr_{A \leftarrow \tilde{P}(q)}[A = a]$.

To begin, we define $\mu_{\{i\}}^{(1)}(q) \stackrel{\text{def}}{=} \tilde{P}(q')_{\{i\}}$ for any q' that satisfies $q'_i = q$, and we define $\tilde{P}^{(1)} = \tilde{P}$. After k steps, the resulting $\{\mu_S^{(k)}\}_{|S| \leq k}$ constitutes our desired non-signaling strategy P .

Symmetric Marginal Compositions of $\tilde{P}^{(i)}$. We will only ever construct $\tilde{P}^{(i+1)}$ and $\mu^{(i+1)}$ by applying a very specific type of transformation, which we call a symmetric marginal composition (SMC), to $\tilde{P}^{(i)}$ and $\mu^{(i)}$. This type of transformation generalizes what we did for three-player games in Section 1.2.1. There, P ($\tilde{P}^{(3)}$ in our current terminology) was constructed as a combination of two transformations applied to $\tilde{P}^{(2)}$ (which happened to be equal to $\tilde{P}$).

- The first transformation, which we denote by [3], generally takes as input some $\tilde{P}^{(i)}$ and produces $\tilde{P}^{(i+1)}$ that on input q , samples and outputs the 3-marginal $\tilde{P}^{(i)}(q)_{\{1,2,3\}}$ – which of course is all of $\tilde{P}^{(i)}(q)$. Thus [3] is the identity transformation.
- The second transformation, which we denote by [1, 1, 1], generally takes as input $\tilde{P}^{(i)}$ and produces $\tilde{P}^{(i+1)}$ that on input q , samples answers from each of the 1-marginals of $\tilde{P}^{(i)}$ – that is, $\tilde{P}^{(i+1)}(q)$ independently samples $a_j \leftarrow \tilde{P}^{(i)}(q)_j$ for each $j \in \{1, 2, 3\}$ and outputs (a_1, a_2, a_3) .

In Section 1.2.1, we in fact defined $\tilde{P}^{(3)}$ to select the first transformation of $\tilde{P}^{(2)}$ with probability 1/3, and to otherwise (with probability 2/3) select the second transformation of $\tilde{P}^{(2)}$. We call this a “mixed” transformation (in contrast to [3] and [1, 1, 1], which we call “pure” transformations), and denote it by the linear combination $\frac{1}{3} \cdot [3] + \frac{2}{3} \cdot [1, 1, 1]$.

In general, we represent a pure transformation by $\lambda = [\lambda_1, \dots, \lambda_m]$ for any $\lambda_1 \geq \dots \geq \lambda_m \geq 1$ with $\sum_j \lambda_j = k$. This transformation acts on a k -player honest-referee non-signaling strategy $\tilde{P}^{(i)}$ to produce another k -player honest-referee strategy $\tilde{P}^{(i+1)}$. This $\tilde{P}^{(i+1)}$, on any input q in the support of $\mathcal{G}$, acts as follows.

- (1) Sample m pairwise disjoint subsets $S_1, \dots, S_m \subseteq [k]$ with $|S_j| = \lambda_j$ for each $j \in [m]$.
- (2) Independently sample $a_j \leftarrow \tilde{P}^{(i)}(q)_{S_j}$ for each $j \in [m]$.
- (3) Output the unique $a^* \in \mathcal{A}$ obtained by composing each of the a_j 's. That is, the a^* for which $a_{S_j}^* = a_j$ for all $j \in [m]$.

Finally, we say that a linear combination $\mathbf{u} = \sum \mathbf{u}_\lambda \cdot \lambda$ (for $\mathbf{u}_\lambda \in \mathbb{R}_{\geq 0}$) is an SMC if it is a probability distribution over pure transformations – i.e., if $\sum \mathbf{u}_\lambda = 1$. It turns out that any SMC preserves honest-referee non-signaling. In general, we will define $\tilde{P}^{(i+1)}$ by applying an SMC $\mathbf{u}^{(i)}$ to $\tilde{P}^{(i)}$. For the moment, we leave $\mathbf{u}^{(i)}$ unspecified.

Symmetric Marginal Compositions of $\mu^{(i)}$. We now give a similar parameterization of our construction of $\mu^{(i+1)}$. First, for all S and all q in the S -marginal support of $\mathcal{G}$, $\mu_S^{(i+1)}(q)$ is determined by our choice of $\tilde{P}^{(i+1)}$ and Condition 4. As for defining $\mu_S^{(i+1)}(q)$ when q is not in the S -marginal support of $\mathcal{G}$, we first simplify matters in two ways.

- First, we note that it suffices to define $\mu_S^{(i+1)}(q)$ for $|S| = i + 1$. Indeed, if these distributions satisfy the analogue of Condition 3, then for $|T| \leq i$, we can define $\mu_T^{(i+1)}(q)$ as $(\mu_S^{(i+1)}(q'))_T$ for any $S \supseteq T$ with $|S| = i + 1$, and any $q' \in Q_S$ such that $q'_S = q$.
- Second, we simplify Condition 3. Recall this requires that for all S, S' , all $T \subseteq S \cap S'$, and all $q \in Q_S, q' \in Q_{S'}$ such that $q_T = q'_T$, it holds that $(\mu_S^{(i+1)}(q))_T = (\mu_{S'}^{(i+1)}(q'))_T$. In fact this holds as long as it holds when $|T| = i$. Indeed, for smaller T we can deduce this via a chain of equalities similar to Eq. (2).

Now fix $|S| = i + 1$, and fix q not in the S -marginal support of $\mathcal{G}$. We will define $\mu_S^{(i+1)}(q)$ in terms of $\mu^{(i)}$ via a similar class of transformations to those that we used to define $\tilde{P}^{(i+1)}$. Specifically, suppose that $i \geq \lambda_1 \geq \dots \geq \lambda_m \geq 1$ with $\sum_j \lambda_j = i + 1$. We then view $\lambda = [\lambda_1, \dots, \lambda_m]$ as acting on $\mu^{(i)}$ to produce $\mu_S^{(i+1)}(q)$ as follows. Specifically, $\mu_S^{(i+1)}(q)$ is the distribution obtained from the following sampling process.

- (1) Sample m pairwise disjoint subsets $S_1, \dots, S_m \subseteq S$ with $|S_j| = \lambda_j$ for each $j \in [m]$.
- (2) Independently sample $a_j \leftarrow \mu_{S_j}^{(i)}(q_{S_j})$ for each $j \in [m]$.
- (3) Output the unique $a^* \in \mathcal{A}_S$ obtained by composing each of the a_j 's. That is, the a^* for which $a_{S_j}^* = a_j$ for all $j \in [m]$.

In general, we will consider a mixed SMC of $\mu^{(i)}$, given by a vector $\mathbf{v}^{(i)} = \sum \mathbf{v}_\lambda^{(i)} \cdot \lambda$. We will use the same vector $\mathbf{v}^{(i)}$ for every S and q for which $|S| = i + 1$ and q is not in the S -marginal support of $\mathcal{G}$.

Thus, we have parameterized our candidate construction of $\tilde{P}^{(i+1)}$ and $\mu^{(i+1)}$ by two SMCs $\mathbf{u}^{(i)}$ and $\mathbf{v}^{(i)}$. In this construction, $\tilde{P}^{(i+1)}$ and $\mu^{(i+1)}$ satisfy the “Agreement” property automatically. There are three remaining properties that must hold.

- $\tilde{P}^{(i+1)}$ must extend a constant multiple of $\tilde{P}^{(i)}$. This corresponds to $\mathbf{u}^{(i)}$ containing $[k]$ with positive probability.
- $\tilde{P}^{(i+1)}$ must be non-signaling. As noted earlier, this follows from the fact that all SMCs preserve non-signaling.

- $\mu^{(i+1)}$ must be non-signaling, despite the fact that $\mu_S^{(i+1)}(q)$ is defined via $\mathbf{u}^{(i)}$ when q is in the S -marginal support of $\mathcal{G}$, and via $\mathbf{v}^{(i)}$ otherwise.

The first and last conditions impose constraints on $\mathbf{u}^{(i)}$ and $\mathbf{v}^{(i)}$ that we show are always satisfiable.

2 PRELIMINARIES

2.1 Notation and Terminology

A probability distribution on a finite set Ω is a function $P : \Omega \rightarrow \mathbb{R}_{\geq 0}$ satisfying $\sum_{\omega \in \Omega} P(\omega) = 1$. We write $\Delta(\Omega)$ to denote the set of all probability distributions on Ω . A sub-probability distribution on Ω is a function $P : \Omega \rightarrow \mathbb{R}_{\geq 0}$ satisfying $\sum_{\omega \in \Omega} P(\omega) \leq 1$. We write $\Delta^{\leq 1}(\Omega)$ to denote the set of all sub-probability distributions on Ω .

- For any (sub-)probability distribution P , let $\text{Supp}(P)$ denote the set of all $\omega \in \Omega$ for which $P(\omega) > 0$.
- For any (sub-)probability distribution P and $n \in \mathbb{N}$, let P^n denote the product (sub-)probability distribution $P^n : \Omega^n \rightarrow \mathbb{R}_{\geq 0}$ defined by $P^n(\omega^{(1)}, \dots, \omega^{(n)}) = \prod_i P(\omega^{(i)})$.
- We will denote the total variational distance between two (sub-)probability distributions P and Q by $d_{\text{TV}}(P, Q)$, which is defined as half of their ℓ_1 -distance:

$$d_{\text{TV}}(X, Y) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{z \in \text{Supp}(X) \cup \text{Supp}(Y)} |\Pr[X = z] - \Pr[Y = z]|.$$

We write $P \approx_{\epsilon} Q$ if $d_{\text{TV}}(P, Q) \leq \epsilon$.

- We now consider the case when Ω is a product of sets $\Omega_1 \times \dots \times \Omega_k$ and S is a subset of $[k]$.
  - The restriction of Ω to S is $\Omega_S = \prod_{i \in S} \Omega_i$. For any $\omega \in \Omega$, the restriction of ω to S is $\omega_S = (\omega_i)_{i \in S}$.
  - For any (sub-)probability distribution P , the marginal (sub-)probability distribution $P_S : \Omega_S \rightarrow \mathbb{R}_{\geq 0}$ is defined by $P_S(\omega_S) = \sum_{\omega' \in \Omega : \omega'_S = \omega_S} P(\omega')$.

2.2 Multiplayer Games

Definition 2.1. A k -player game is a tuple $(Q, \mathcal{A}, \pi, W)$, where $Q = Q_1 \times \dots \times Q_k$ and $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k$ are finite sets, $\pi : Q \rightarrow \mathbb{R}_{\geq 0}$ is a probability distribution, and $W : Q \times \mathcal{A} \rightarrow [0, 1]$ is a “winning probability” function.

Definition 2.2. Given a k -player game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ where $Q = Q_1 \times \dots \times Q_k$ and $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_k$, its n -fold parallel repetition is defined as the k -player game $\mathcal{G}^n \stackrel{\text{def}}{=} (Q', \mathcal{A}', \pi', W')$ where $Q' \stackrel{\text{def}}{=} Q_1^n \times \dots \times Q_k^n$, $\mathcal{A}' \stackrel{\text{def}}{=} \mathcal{A}_1^n \times \dots \times \mathcal{A}_k^n$, $\pi'(q) \stackrel{\text{def}}{=} \prod_{i=1}^n \pi(q^{(i)})$, and $W'(q, a) \stackrel{\text{def}}{=} \prod_{i=1}^n W(q^{(i)}, a^{(i)})$.

In the above we write elements $q \in Q'$ as $(\{q_1^{(i)}\}_{i \in [n]}, \dots, \{q_k^{(i)}\}_{i \in [n]})$, we write q_j to denote $(q_j^{(1)}, \dots, q_j^{(n)})$, and we write $q^{(i)}$ to denote $(q_1^{(i)}, \dots, q_k^{(i)})$. Our notation for components of elements of $\mathcal{A}'$ is analogous.

Definition 2.3. A strategy for a game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ is a function $P : Q \rightarrow \Delta(\mathcal{A})$.

Definition 2.4. A strategy P for a k -player game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ is local if there are functions $\{P_i : Q_i \rightarrow \mathcal{A}_i\}_{i \in [k]}$ such that for

every q , $P(q)$ is the distribution outputting $(P_1(q_1), \dots, P_k(q_k))$ with probability 1.

Definition 2.5. A strategy P for a k -player game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ is non-signaling if for every subset $S \subseteq [k]$, there exists a function $\text{Sim}_S : Q_S \rightarrow \Delta(\mathcal{A}_S)$ where for all $q \in Q$, it holds that $P(q)_S = \text{Sim}_S(q_S)$.

Definition 2.6 (Values of a Game). The value of a game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ with respect to a strategy P for $\mathcal{G}$, denoted by $v[P](\mathcal{G})$, is the expected value of $W(Q, A)$ in the probability space obtained by sampling $Q \leftarrow \pi$ and $A \leftarrow P(Q)$. The classical value of $\mathcal{G}$, denoted by $v(\mathcal{G})$ is the maximum value of $\mathcal{G}$ with respect to any local strategy. The non-signaling value of $\mathcal{G}$, denoted by $v_{\text{ns}}(\mathcal{G})$, is the maximum value of $\mathcal{G}$ with respect to any non-signaling strategy.

3 GENERALIZATIONS OF NON-SIGNALING STRATEGIES

[LW16] introduces the following generalization of non-signaling strategies, called sub-non-signaling strategies, where $P(q)$ is only required to be a sub-probability distribution.

Definition 3.1 ([LW16]). A sub-non-signaling strategy P for a k -player game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ is a function $P : Q \rightarrow \Delta^{\leq 1}(\mathcal{A})$ satisfying the following property: for every $S \subseteq [k]$, there exists a function $\text{Sim}_S : Q_S \rightarrow \Delta(\mathcal{A}_S)$ such that for every $q \in Q$ and $a \in \mathcal{A}_S$, it holds that $P(q)_S(a) \leq \text{Sim}_S(q_S)(a)$.

Definition 3.2. The sub-non-signaling value of $\mathcal{G}$, denoted by $v_{\text{sns}}(\mathcal{G})$, is the maximum value of

$$\sum_{q \in Q} \pi(q) \sum_{a \in \mathcal{A}} P(q)(a) W(q, a)$$

attained by any sub-non-signaling strategy P .

In our work, it will be useful to consider the following generalization of non-signaling strategies where the non-signaling condition is only required to hold over queries in $\text{Supp}(\pi)$. Such strategies appear to be non-signaling to an honest referee asking queries according to π .

Definition 3.3. An honest-referee non-signaling strategy P for a k -player game $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ is a function $P : \text{Supp}(\pi) \rightarrow \Delta(\mathcal{A})$ satisfying the following property: for every $S \subseteq [k]$, there exists a function $\text{Sim}_S : \text{Supp}(\pi_S) \rightarrow \Delta(\mathcal{A}_S)$ such that for every $q \in \text{Supp}(\pi)$, it holds that $P(q)_S = \text{Sim}_S(q_S)$.

Remark 3.4. We can view any non-signaling strategy P as a complete set of marginals $\{P_S : Q_S \rightarrow \Delta(\mathcal{A}_S)\}_{S \subseteq [k]}$ for any $|S| \subseteq [k]$

and $q \in Q_S$, $P_S(q) \stackrel{\text{def}}{=} P(q)_S$ for any $q' \in Q$ such that $q'_S = q$. This is well-defined (i.e., does not depend on the choice of q') because P is non-signaling, so we have $P_S = \text{Sim}_S$. Similarly, we can view any honest-referee non-signaling strategy P as a complete set of marginals $\{P_S : \text{Supp}(\pi_S) \rightarrow \Delta(\mathcal{A}_S)\}_{S \subseteq [k]}$ where $P_S = \text{Sim}_S$.

The following proposition is analogous to the fact that $v_{\text{ns}}(\mathcal{G}^n) \geq v_{\text{ns}}(\mathcal{G})^n$ and implies that if $v_{\text{hr-ns}}(\mathcal{G}) = 1$, then also $v_{\text{hr-ns}}(\mathcal{G}^n) = 1$ for every n .

Proposition 3.5. For any game $\mathcal{G}$ and any number of repetitions n , it holds that $v_{\text{hr-ns}}(\mathcal{G}^n) \geq v_{\text{hr-ns}}(\mathcal{G})^n$.

PROOF. Let $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ be a k -player game, and let $P : \text{Supp}(\pi) \rightarrow \Delta(\mathcal{A})$ be an honest-referee non-signaling strategy for $\mathcal{G}$. Let $\mathcal{G}^n = (Q^n, \mathcal{A}^n, \pi^n, W^n)$ denote the n -fold parallel repetition of $\mathcal{G}$.

Define

$$P^n : \text{Supp}(\pi^n) \rightarrow \Delta(\mathcal{A}^n)$$

$$P^n(q^{(1)}, \dots, q^{(n)})(a^{(1)}, \dots, a^{(n)}) = \prod_{i=1}^n P(q_1^{(i)}, \dots, q_k^{(i)})(a_1^{(i)}, \dots, a_k^{(i)}),$$

which is well-defined because $(q^{(1)}, \dots, q^{(n)}) \in \text{Supp}(\pi^n)$ iff for all $i \in [n]$, $(q_1^{(i)}, \dots, q_k^{(i)}) \in \text{Supp}(\pi)$. It is easy to verify that P^n is an honest-referee non-signaling strategy for $\mathcal{G}^n$ since for any $(q^{(1)}, \dots, q^{(n)}) \in \text{Supp}(\pi^n)$, we can define $\text{Sim}_S(q_S^{(1)}, \dots, q_S^{(n)}) = \prod_i \text{Sim}_S(q_S^{(i)})$. Also, it is clear that $v[P^n](\mathcal{G}^n) = v[P](\mathcal{G})^n$. $\square$

For several of our corollaries, it will be useful to relate the sub-non-signaling value and honest-referee non-signaling value of a game $\mathcal{G}$. For this, we have the following proposition.

Proposition 3.6. *For any game $\mathcal{G}$, $v_{\text{sns}}(\mathcal{G}) = 1$ iff $v_{\text{hr-ns}}(\mathcal{G}) = 1$.*

PROOF. Let $\mathcal{G} = (Q, \mathcal{A}, \pi, W)$ be any k -player game.

Claim 3.7. *If $v_{\text{sns}}(\mathcal{G}) = 1$ then $v_{\text{hr-ns}}(\mathcal{G}) = 1$.*

PROOF. Suppose that $P : Q \rightarrow \Delta^{\leq 1}(\mathcal{A})$ is a sub-non-signaling strategy for $\mathcal{G}$ such that $v[P](\mathcal{G}) = 1$. Since $v[P](\mathcal{G}) = 1$, it must hold for all $q \in \text{Supp}(\pi)$ that $P(q)$ is a probability distribution, i.e. $\sum_{a \in \mathcal{A}} P(q)(a) = 1$. Therefore for each $S \subseteq [k]$, $P(q)_S$ is a probability distribution.

Since P is sub-non-signaling, there exists $\text{Sim}_S : Q_S \rightarrow \Delta(\mathcal{A}_S)$ such that for all $q \in Q$ and all $a \in \mathcal{A}_S$, it holds that $P(q)_S(a) \leq \text{Sim}_S(q_S)(a)$. Since $P(q)_S$ is a probability distribution, $\text{Sim}_S(q_S)$ must also be a probability distribution. Therefore, for any $q \in \text{Supp}(\pi)$ and any $S \subseteq [k]$, it holds that $P(q)_S = \text{Sim}_S(q_S)$ so P is in fact honest-referee non-signaling. $\square$

Claim 3.8. *If $v_{\text{hr-ns}}(\mathcal{G}) = 1$ then $v_{\text{sns}}(\mathcal{G}) = 1$.*

PROOF. Suppose that $\tilde{P} : \text{Supp}(\pi) \rightarrow \Delta(\mathcal{A})$ is an honest-referee non-signaling strategy with $v[\tilde{P}](\mathcal{G}) = 1$. Extend $\tilde{P}$ to $P : Q \rightarrow \Delta^{\leq 1}(\mathcal{A})$ by defining $P(q)$ to be the zero sub-probability distribution for all $q \notin \text{Supp}(\pi)$. Clearly $v[P](\mathcal{G}) = v[\tilde{P}](\mathcal{G}) = 1$.

It remains to establish that P is sub-non-signaling: For each $S \subseteq [k]$, define Sim_S so that for any $q \in \text{Supp}(\pi)$, $\text{Sim}_S(q_S) = \tilde{P}(q)_S$ – this is possible because $\tilde{P}$ is honest-referee non-signaling. For $q' \notin \text{Supp}(\pi_S)$, define $\text{Sim}_S(q')$ arbitrarily. Now clearly for any $q \in Q$ and any $a \in \mathcal{A}_S$, it holds that $P(q)_S(a) \leq \text{Sim}_S(q_S)(a)$. $\square$

Proposition 3.6 follows directly from Claims 3.7 and 3.8. $\square$

4 MAIN LEMMA AND IMPLICATIONS

We prove the theorems outlined in Section 1.1 assuming our main lemma, whose proof is deferred to the full version. The main theorems are stated here with greater generality in the context of honest-referee non-signaling strategies.

Lemma 1 (Non-signaling Value Lower Bound). *For every k , there exists a constant $\alpha_k > 0$ such that for any k -player game $\mathcal{G}$, it holds that $v_{\text{ns}}(\mathcal{G}) \geq \alpha_k \cdot v_{\text{hr-ns}}(\mathcal{G})$.*

From here on, we write α_k to refer to the maximum value for which Lemma 1 holds. We have not attempted to optimize α_k , but a loose accounting of our proof shows that $\alpha_k \geq 2^{-O(k^2)}$.

To show that there are counterexamples to non-signaling parallel repetition, we show that there exist games whose sub-non-signaling values are 1 but whose non-signaling values are less than 1.

Proposition 4.1. *For every $k \geq 3$, there exists a k -player game $\mathcal{G}$ for which $v_{\text{hr-ns}}(\mathcal{G}) = 1$ (i.e. $v_{\text{sns}}(\mathcal{G}) = 1$) and $v_{\text{ns}}(\mathcal{G}) < 1$.*

PROOF. One example of such a game is the “anti-correlation game”. The three-player version of this game was described in [ARV16, LW16] as an example of a game whose non-signaling value is less than 1, but whose sub-non-signaling value is 1.

Definition 4.2. *The k -player anti-correlation game $\tilde{\mathcal{G}}_k$ is the game $(\tilde{Q}_k, \tilde{\mathcal{A}}_k, \tilde{\pi}_k, \tilde{W}_k)$ where $\tilde{Q}_k = \{0, 1\}^k$, $\tilde{\mathcal{A}}_k = [k-1]^k$, $\tilde{\pi}_k$ is the uniform distribution on k -bit strings of Hamming weight $k-1$, and*

$$\tilde{W}_k(q, a) = \begin{cases} 1 & \text{if } (i \neq j \wedge q_i = q_j = 1) \implies a_i \neq a_j \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mathcal{G}$ denote the k -player anti-correlation game $\tilde{\mathcal{G}}_k$. For $1 \leq \ell \leq k$, let $h_\ell \in \{0, 1\}^k$ denote the string of Hamming weight $k-1$ with 0 in the ℓ th coordinate.

Define $P : \text{Supp}(\tilde{\pi}_k) \rightarrow \Delta(\tilde{\mathcal{A}}_k)$ as follows: $P(h_\ell)$ chooses a random element of $[k-1]$ for a_ℓ and a random permutation of $[k-1]$ for $(a_i)_{i \neq \ell}$. This is an honest-referee non-signaling strategy: for any $h_\ell, h_m \in \text{Supp}(\tilde{\pi}_k)$, the marginals $P(h_\ell)_S$ and $P(h_m)_S$, where $S = [k] \setminus \{\ell, m\}$, are both equivalent to choosing a $(k-2)$ -sized subset of $[k-1]$ and a random permutation of these elements. For distinct $i, j \neq \ell$, we have $a_i \neq a_j$ so $v_{\text{hr-ns}}[P](\tilde{\mathcal{G}}_k) = 1$. Thus $v_{\text{hr-ns}}(\tilde{\mathcal{G}}_k) = 1$ (by Proposition 3.6, this is equivalent to $v_{\text{sns}}(\tilde{\mathcal{G}}_k) = 1$).

For any non-signaling $P' : \tilde{Q}_k \rightarrow \Delta(\tilde{\mathcal{A}}_k)$ and any $h_\ell \in \text{Supp}(\tilde{\pi}_k)$, $P'(h_\ell)_{[k] \setminus \{\ell\}} = P'(1^k)_{[k] \setminus \{\ell\}}$ by non-signaling. Thus

$$\Pr_{\ell, P'(h_\ell)} [\tilde{W}_k(h_\ell, P'(h_\ell)) = 1] = \Pr_{\ell, P'(1^k)} [\tilde{W}_k(h_\ell, P'(1^k)) = 1].$$

This equals the probability that $P'(1^k)_{[k] \setminus \{\ell\}}$ is a permutation of $[k-1]$. By the Pigeonhole principle, every element in $\text{Supp}(P'(1^k))$ has an element of $[k-1]$ appearing in at least two indices so this probability is at most $2/k$. $\square$

Theorem 1 (Parallel Repetition Counterexample). *For every $k \geq 3$, there exists a k -player game $\mathcal{G}$ such that $v_{\text{ns}}(\mathcal{G}) < 1$ and $v_{\text{ns}}(\mathcal{G}^n) \geq \alpha_k$ for all $n \geq 1$, where α_k is the constant in Lemma 1.*

PROOF. This follows from Proposition 4.1 and Lemma 1: for the anti-correlation game $\tilde{\mathcal{G}}_k$, it holds that $v_{\text{hr-ns}}(\tilde{\mathcal{G}}_k) = 1 \implies v_{\text{hr-ns}}(\tilde{\mathcal{G}}_k^n) = 1 \implies v_{\text{ns}}(\tilde{\mathcal{G}}_k^n) \geq \alpha_k$. $\square$

Theorem 2 (Parallel Repetition Dichotomy). *For every game $\mathcal{G}$, either $v_{\text{ns}}(\mathcal{G}^n) \geq \Omega(1)$ or $v_{\text{ns}}(\mathcal{G}^n) \leq \exp(-\Omega(n))$. The former occurs when $v_{\text{sns}}(\mathcal{G}) = 1$ and the latter occurs when $v_{\text{sns}}(\mathcal{G}) < 1$.*

PROOF. If $v_{\text{sns}}(\mathcal{G}) < 1$, then by [LW16, Theorem 4], we have $v_{\text{ns}}(\mathcal{G}^n) \leq v_{\text{sns}}(\mathcal{G}^n) \leq e^{-\Omega(n)}$.

If $v_{\text{sns}}(\mathcal{G}) = 1$, then by Proposition 3.6 $v_{\text{hr-ns}}(\mathcal{G}) = 1$. Then by Proposition 3.5, $v_{\text{hr-ns}}(\mathcal{G}^n) \geq v_{\text{hr-ns}}(\mathcal{G})^n = 1$, so $v_{\text{hr-ns}}(\mathcal{G}^n) = 1$ for all $n \geq 1$. By Lemma 1, this implies that $v_{\text{ns}}(\mathcal{G}^n) \geq \alpha_k = \Omega(1)$. $\square$

Theorem 3 (Parallel Repetition Magic Value). *For every k -player game $\mathcal{G}$, if $v_{\text{ns}}(\mathcal{G}) < \alpha_k$, then $v_{\text{ns}}(\mathcal{G}^n) \leq \exp(-\Omega(n))$ where α_k is the constant in Lemma 1.*

PROOF. If $v_{\text{ns}}(\mathcal{G}) < \alpha_k$, then by Lemma 1 and Proposition 3.6, we have $v_{\text{sns}}(\mathcal{G}) < 1$. Then by [LW16], it holds that $v_{\text{ns}}(\mathcal{G}^n) \leq v_{\text{sns}}(\mathcal{G}^n) \leq \exp(-\Omega(n))$. $\square$

Upper Bounding the Magic Value. Given Theorem 3, it is natural to ask whether α_k is optimal: For each k , what is the largest value β_k such that for every k -player game $\mathcal{G}$, $v_{\text{ns}}(\mathcal{G}) < \beta_k$ implies $v_{\text{ns}}(\mathcal{G}^n) \leq \exp(-\Omega(n))$? Theorem 1 shows that $\beta_k < 1$, and Theorem 3 shows that $\beta_k \geq \alpha_k$, which our proof of Lemma 1 shows to be at least 2^{-k^2} . This is a large gap, but Theorem 4 shows that $\beta_k \leq \exp(-\Omega(k))$. The proof is in Section 5: the main tool is a “distributed” form of repetition that does exponentially decrease the non-signaling value of a game, but linearly increases the number of players.

Theorem 4 (Strong Parallel Repetition Counterexample). *There is a sequence of games $\{\mathcal{G}_k\}_{k \geq 3}$, with $\mathcal{G}_k$ a k -player game, such that $v_{\text{ns}}(\mathcal{G}_k) \leq \exp(-\Omega(k))$ and $v_{\text{ns}}(\mathcal{G}_k^n) \geq \alpha_k$ for all $n \geq 1$.*

5 STRONG PARALLEL REPETITION COUNTEREXAMPLE

Theorem 4. *For every $k \geq 3$, there exists a k -player game $\mathcal{G}$ such that $v_{\text{ns}}(\mathcal{G}^n) = (2/3)^{\lfloor k/3 \rfloor}$ for any $n \geq 1$.*

We prove the theorem in the case $k = 3$ in Section 5.1, and defer the remaining cases to the full version. In this case, the game is the three-player anti-correlation game $\tilde{\mathcal{G}}_3 = (\{0, 1\}^3, \{0, 1\}^3, \tilde{\pi}_3, \tilde{W}_3)$ defined in Definition 4.2.

5.1 A Non-signaling Strategy for $\tilde{\mathcal{G}}_3^n$ with Value $2/3$

Proposition 4.1 shows that the non-signaling value of $\tilde{\mathcal{G}}_3$ (and hence of $\tilde{\mathcal{G}}_3^n$) is at most $2/3$. Now we show that the hardness of this game is *completely non-amplifying*. That is, we construct a non-signaling strategy for $\tilde{\mathcal{G}}_3^n$ that achieves value $2/3$.

Proposition 5.1. *For every $n \geq 1$, $v_{\text{ns}}(\tilde{\mathcal{G}}_3^n) = 2/3$.*

PROOF. We give a non-signaling strategy P for $\tilde{\mathcal{G}}_3^n$ and show that $v[P](\tilde{\mathcal{G}}_3^n) = 2/3$.

Construction 5.2. *Given $q = (q^{(1)}, \dots, q^{(n)})$, P samples answers $a = (a^{(1)}, \dots, a^{(n)})$ as follows.*

- If no $q^{(i)}$ is equal to 111, then:
  - (1) Sample $b \leftarrow \text{Ber}(1/3)$, i.e. $b = 1$ with probability $2/3$ and $b = 0$ otherwise.
  - (2) Sample each $a^{(i)}$ independently and uniformly at random, subject to the constraint that if $q_j^{(i)} = q_k^{(i)} = 1$ for some $j \neq k$, then $a_j^{(i)} \oplus a_k^{(i)} = b$.
- If some $q^{(i)}$ is equal to 111, then:
  - (1) Sample $t \leftarrow \{1, 2, 3\}$ uniformly at random.

- (2) Sample each $a^{(i)}$ independently and uniformly at random, subject to the constraint that if $q_j^{(i)} = q_k^{(i)} = 1$ for some $j \neq k$, then:
  - If $j = t$ or $k = t$, then $a_j^{(i)} \neq a_k^{(i)}$.
  - Otherwise, $a_j^{(i)} = a_k^{(i)}$.

Loosely speaking, in the first case, P randomly decides with probability $1/3$ to lose all instances of $\tilde{\mathcal{G}}_3$. In the second case, P randomly chooses a designated player t to disagree with all other players receiving 1. It may seem strange that in the first case, P artificially chooses to lose all the games. However, this is necessary for the existence of a consistent answer distribution when all players receive 1 queries.

We claim that the value of $\tilde{\mathcal{G}}_3^n$ with respect to P is $2/3$. This is solely determined by P 's behavior in the first case, because for honestly generated queries, no $q^{(i)}$ is ever equal to 111. In the first case, with probability $2/3$ (whenever $b = 1$), the answers $a_j^{(i)}$ and $a_k^{(i)}$ corresponding to the “1” queries in the i^{th} game satisfy $a_j^{(i)} \neq a_k^{(i)}$ for every i , so P wins $\tilde{\mathcal{G}}_3^n$ with probability $2/3$.

It remains to verify that P is non-signaling, i.e. that for all sets $S \subseteq \{1, 2, 3\}$, the distribution $P(q)_S$ depends only on q_S . This is trivially true when $|S| = 0$ or $|S| = 3$. The remaining cases are $|S| = 1$ and $|S| = 2$.

These cases are easier to verify when keeping in mind the structure of P : based on q , P probabilistically chooses a set of constraints on $a^{(1)}, \dots, a^{(n)}$. Each constraint specifies the equality or inequality of different components of each $a^{(i)}$. P then independently chooses $a^{(i)}$ satisfying the constraints. Thus, to demonstrate that the distribution of a_S depends only on q_S , it suffices to show that the distribution of the constraints on a_S depends only on q_S .

Case 1: $|S| = 1$. For any q , we claim that the distribution $P(q)_S$ is uniformly random on $\{0, 1\}^n$, and thus depends only on q_S (in fact, on nothing) as required.

This holds because all constraints chosen by P satisfy

- **Symmetry:** The constraints only enforce equality or inequality of specific bits of a . Thus, when a is chosen uniformly at random to satisfy these constraints, each individual bit of a is equally likely to be 0 or 1.
- **Independence:** Each constraint only relates the bits of a single $a^{(i)}$. Thus, $a_S^{(1)}, \dots, a_S^{(n)}$ are independent as random variables.

Case 2: $|S| = 2$. For any q , we claim that $(a_S^{(1)}, \dots, a_S^{(n)}) = P(q)_S$ is distributed as follows. For concreteness say that $S = \{j, k\}$. For any q , we have:

- With probability $2/3$, the constraints generated by P on $a_S^{(i)}$ are that $a_S^{(i)} \in \{01, 10\}$ for all i for which $q_S^{(i)} = 11$. In particular, P generates these constraints if $b = 1$ (when no $q^{(i)}$ is 111), and when $t \in S$ (when some $q^{(i)}$ is 111).
- Otherwise the constraints generated by P on $a_S^{(i)}$ are that $a_S^{(i)} \in \{00, 11\}$ for all i for which $q_S^{(i)} = 11$.

We note that P may also generate constraints on $a^{(i)}$ beyond those explicitly mentioned above, specifically when $q_j^{(i)} = 1$ for some

$j \notin S$. However, inspection of P reveals that these constraints do not affect the distribution of $a_S^{(i)}$. For example, suppose that $S = \{1, 2\}$, $q^{(i)} = 111$, and $t = 2$. Then the constraints generated by P require not only that $a_1^{(i)} \neq a_2^{(i)}$, but also that $a_1^{(i)} = a_3^{(i)}$ and $a_2^{(i)} \neq a_3^{(i)}$. In this case, the latter two constraints are *redundant*: whenever $a_1^{(i)} \neq a_2^{(i)}$, they are satisfiable for a unique choice of $a_3^{(i)}$. Thus, the redundant constraints do not affect the distribution of $a_S^{(i)}$. $\square$

ACKNOWLEDGMENTS

Research conducted while Justin Holmgren was a graduate student at MIT, supported in part by the NSF MACS project – CNS-1413920, DARPA IBM – W911NF-15-C-0236, and the SIMONS Investigator Award Agreement Dated 6-5-12, and Lisa Yang was an undergraduate student at MIT.

We thank Ron Rothblum for his generous help in reading and commenting on an early version of this manuscript.

REFERENCES

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [ARV16] Rotem Arnon-Friedman, Renato Renner, and Thomas Vidick. Non-signaling parallel repetition using de Finetti reductions. *IEEE Trans. Information Theory*, 62(3):1440–1457, 2016.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *STOC*, pages 21–31. ACM, 1991.
- [BFS14] Harry Buhrman, Serge Fehr, and Christian Schaffner. On the parallel repetition of multi-player games: The no-signaling case. In *TQC*, volume 27 of *LIPICs*, pages 24–35. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.
- [BGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *STOC*, pages 113–131. ACM, 1988.
- [BHK17] Zvika Brakerski, Justin Holmgren, and Yael Tauman Kalai. Non-interactive delegation and batch NP verification from standard computational assumptions. In *STOC*, pages 474–482. ACM, 2017.
- [BKK⁺18] Saikrishna Badrinarayanan, Yael Tauman Kalai, Dakshita Khurana, Amit Sahai, and Daniel Wichs. Non-interactive delegation for low-space non-deterministic computation. In *STOC*, page to appear. ACM, 2018.
- [CHSH69] John F Clauser, Michael A Horne, Abner Shimony, and Richard A Holt. Proposed experiment to test local hidden-variable theories. *Physical review letters*, 23(15):880, 1969.
- [Fei91] Uriel Feige. On the success probability of the two provers in one-round proof systems. In *Structure in Complexity Theory Conference*, pages 116–123. IEEE Computer Society, 1991.
- [FGL⁺96] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, 1996.
- [FL92] Uriel Feige and László Lovász. Two-prover one-round proof systems: Their power and their problems (extended abstract). In *STOC*, pages 733–744. ACM, 1992.
- [For89] Lance Jeremy Fortnow. *Complexity-theoretic aspects of interactive proof systems*. PhD thesis, 1989.
- [FRS88] Lance Fortnow, John Rempel, and Michael Sipser. On the power of multi-prover interactive protocols. In *Structure in Complexity Theory Conference*, pages 156–161. IEEE Computer Society, 1988.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [Hol09] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.
- [IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *IEEE Conference on Computational Complexity*, pages 217–228. IEEE Computer Society, 2009.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *STOC*, pages 723–732. ACM, 1992.
- [KRR13] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. Delegation for bounded space. In *STOC*, pages 565–574. ACM, 2013.
- [KRR14] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to delegate computations: the power of no-signaling proofs. In *STOC*, pages 485–494. ACM, 2014.
- [LW16] Cécilia Lancien and Andreas Winter. Parallel repetition and concentration for (sub-)no-signalling games via a flexible constrained de Finetti reduction. *Chicago J. Theor. Comput. Sci.*, 2016, 2016.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [Ver96] Oleg Verbitsky. Towards the parallel repetition conjecture. *Theor. Comput. Sci.*, 157(2):277–282, 1996.
- [Yue16] Henry Yuen. A parallel repetition theorem for all entangled games. In *ICALP*, volume 55 of *LIPICs*, pages 77:1–77:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.