(A Counterexample to) Parallel Repetition for Non-Signaling Multi-Player Games

Justin Holmgren*
holmgren@mit.edu

Lisa Yang
lisayang@mit.edu

Abstract

We give a three-player game whose non-signaling value is constant (2/3) under any number of parallel repetitions. This is the first known setting where parallel repetition completely fails to reduce the maximum winning probability of computationally unbounded players.

We also show that the best known results on non-signaling parallel repetition apply to a relatively limited class of games. In particular, these games cannot yield log-prover MIPs for languages beyond PSPACE.

1 Introduction

A *multi-player game* $\mathcal{G}$ consists of an interaction between a referee and k players $P_1, \ldots, P_k$. The referee samples k questions $q_1, \ldots, q_k$ from some joint distribution π and sends q_i to P_i. The players respond with answers $a_1, \ldots, a_k$, and are judged to *win* or *lose* the game according to a predicate $W(q_1, \ldots, q_k, a_1, \ldots, a_k)$. The *value* of the game, denoted $v(\mathcal{G})$, is the maximum probability with which players can win. The study of multi-player games is a rich and active research topic in diverse areas of theoretical computer science, with foundational applications to complexity theory, hardness of approximation, and cryptography (see e.g., [BGKW88, BFL91, BFLS91, Kil92, FGL+96, ALM+98, Hås01].

In a classical game, the players are restricted to *local* strategies: each answer a_i is a (without loss of generality, deterministic) function of q_i. While natural, the assumption that players act locally is in some cases overly optimistic. One example is in quantum information theory, where shared entanglement can allow provers to implement certain non-local strategies [Bel64]. A similar phenomenon arises in cryptography [BMW98, DLN+01, KRR14, DHRW16] when considering mappings from independently-keyed ciphertexts $(\mathsf{Enc}_{K_1}(q_1), \ldots, \mathsf{Enc}_{K_k}(q_k))$ to $(\mathsf{Enc}_{K_1}(a_1), \ldots, \mathsf{Enc}_{K_k}(a_k))$ which are computable in polynomial-time. In both settings, the precise set of attainable strategies is difficult to characterize, but is bounded by the set of *non-signaling*[1] strategies where $(a_1, \ldots, a_k)$ can jointly depend on all of $(q_1, \ldots, q_k)$, but for any subset $S \subseteq [k]$, the *distribution* of a_S must depend only on q_S.

Parallel repetition is a natural idea for reducing the soundness error of proof systems (i.e., the players' maximum winning probability on false statements) by having the k players simultaneously play λ independent copies of the game. The referee now samples λ independent sets of questions $\{(q_1^{(i)}, \ldots, q_k^{(i)})\}_{i=1}^{\lambda}$ from π, and sends *all* λ queries $\{q_j^{(i)}\}_{i=1}^{\lambda}$ to the j^{th} player at once. The players are then required to win all λ copies of the game; that is, the j^{th} player must produce $\{a_j^{(i)}\}_{i=1}^{\lambda}$ so that $W(q_1^{(i)}, \ldots, q_k^{(i)}, a_1^{(i)}, \ldots, a_k^{(i)}) = 1$ for every i.

Parallel repetition was first studied in the context of *classical* games where provers are allowed to use local strategies. It was initially conjectured that $v(\mathcal{G}^\lambda) = v(\mathcal{G})^\lambda$ [FRS88]. This was quickly disproved, however, by Fortnow [For89] and Feige [Fei91], who constructed a two-player game $\mathcal{G}$ satisfying $v(\mathcal{G}^2) = v(\mathcal{G}) < 1$. On

*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

[1] or at least *approximately* non-signaling

the other hand, Raz's celebrated parallel repetition theorem establishes that for two-player games, the value of $\mathcal{G}^\lambda$ is small for *large* λ. In particular, if $v(\mathcal{G}) < 1$, then $v(\mathcal{G}^\lambda) \leq \bar{v}^\lambda$ for some $\bar{v} < 1$ which depends on $\mathcal{G}$ (but may be significantly larger than $v(\mathcal{G})$).

This result was simplified and extended to non-signaling two-player games by Holenstein [Hol09]. But such games are relatively uninteresting: in a proof system, two non-signaling provers can only prove statements in PSPACE [IKM09, Ito10] which is the same as a single prover given many rounds of interaction. In contrast, polynomially many non-signaling provers can prove statements in EXP [KRR14]. Even so, the question of parallel repetition for multi-player non-signaling games remained wide open, with a few positive results for special cases [BFS14, FRV16, LW16]. To the best of our knowledge, most researchers believed that a general analogue of Raz's result should also hold in this setting.

We show this is not true. Our main contribution is to exhibit a three-player game $\mathcal{G}$ such that the value of $\mathcal{G}^\lambda$ with respect to non-signaling strategies is $2/3$ for all λ. Our result is contrasted with known parallel repetition results in Table 1.

	Two-player games	Multi-player games
Classical	$\exp\left(-\Omega\left(\frac{\epsilon^3\lambda}{\log\lvert\mathcal{A}\rvert}\right)\right)$ [Raz98, Hol09]	$\frac{1}{\Omega(\mathsf{Ackermann}^{-1}(\lambda))}$ [Ver96]
Non-Signaling	$\exp\left(-\Omega\left(\epsilon^2\lambda\right)\right)$ [Hol09]	$\geq \frac{2}{3}$ [**This Work**]

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

Although our counterexample precludes a parallel repetition theorem for all non-signaling games, one may ask whether known parallel repetition results (which hold only for a restricted set of games) suffice for applications (e.g., constructing MIPs). In Section 4, we argue that this is not the case.

1.1 Related Work

The question of how parallel repetition of a k-player game $\mathcal{G}$ (for $k > 2$) affects its non-signaling value was first studied by Buhrman, Fehr, and Schaffner [BFS14]. They considered games $(\mathcal{Q}, \mathcal{A}, \pi, W)$ with *complete support*, where $\pi(q) > 0$ for all $q \in \mathcal{Q}$ (this is not without loss of generality because $\mathcal{Q}$ has the form $\mathcal{Q}_1 \times \cdots \times \mathcal{Q}_k$). For such games, they show that the non-signaling value of $\mathcal{G}^\lambda$ is exponentially small in λ, with a rate of exponential decay that depends on $\mathcal{G}$ (in particular, on $\min_{q\in\mathcal{Q}} \pi(q)$).

Arnon-Friedman, Renner, and Vidick give an alternative proof of this fact, and observe that one can always add uniformly distributed dummy queries to *any* game $\mathcal{G}$ to obtain a closely related game $\tilde{\mathcal{G}}$ with complete support, so that the non-signaling value of $\tilde{\mathcal{G}}^\lambda$ is exponentially small in λ. We note that the bound obtained in this way can be as large as $e^{-\lambda/|\mathcal{Q}|}$, where $\mathcal{Q}$ is the query alphabet for $\mathcal{G}$; in particular, $\mathsf{polylog}(|\mathcal{Q}|)$ repetitions may not even achieve constant soundness error. Therefore, this result is not generally applicable towards reducing the soundness error of a MIP, which may associate an n-bit input to a game with a query alphabet of size $2^{\mathsf{poly}(n)}$.

2 Preliminaries

2.1 Notation

For any finite set $\Sigma = \Sigma_1 \times \ldots \times \Sigma_k$ and subset $S \subseteq [k]$, we denote the restriction of Σ to the coordinates in S by $\Sigma_S = \prod_{i\in S} \Sigma_i$. For any string $\sigma \in \Sigma$, we denote the restriction of σ to the coordinates in S by $\sigma_S = (\sigma_i)_{i\in S}$.

For a probability mass function $P : \Sigma_1 \times \cdots \times \Sigma_k \to \mathbb{R}$ and any subset $S \subseteq [k]$, the marginal probability mass functions $P_S : \left(\prod_{i\in S} \Sigma_i\right) \to \mathbb{R}$ is defined by $P_S(\sigma) = \sum_{\sigma' \in \Sigma : \sigma'_S = \sigma} P(\sigma')$.

We will denote the total variational distance between two random variables X and Y by $d_{\mathsf{TV}}(X,Y)$, which is defined as half of the ℓ_1-distance $\|\cdot\|_1$ between their probability laws (i.e. if $\mathrm{Supp}(X)$ and $\mathrm{Supp}(Y)$ are the supports of X and Y, then $d_{\mathsf{TV}}(X,Y) = \frac{1}{2}\sum_{z\in\mathrm{Supp}(X)\cup\mathrm{Supp}(Y)} |\Pr[X=z] - \Pr[Y=z]|$). We write $X \approx_\epsilon Y$ if $d_{\mathsf{TV}}(X,Y) \leq \epsilon$.

2.2 Multi-player Games

Definition 2.1 (Multi-player Games)**.** *A k-*player game *is a tuple $(\mathcal{Q},\mathcal{A},\pi,W)$, where $\mathcal{Q} = \mathcal{Q}_1\times\cdots\times\mathcal{Q}_k$ and $\mathcal{A} = \mathcal{A}_1\times\cdots\times\mathcal{A}_k$ are finite sets, $\pi : \mathcal{Q}\to\mathbb{R}_{\geq 0}$ is a probability mass function, and $W : \mathcal{Q}\times\mathcal{A}\to[0,1]$ is a "winning probability" function.*

Remark 2.2. *In this work, when we represent a game as a string (particularly in Section 4), we explicitly list*

- *Every element of $\mathcal{Q}_i$ and $\mathcal{A}_i$ for each $i\in[k]$,*
- *$\pi(q)$ for each $q\in\mathcal{Q}$,*
- *$W(q,a)$ for each $q\in\mathcal{Q}$ and each $a\in\mathcal{A}$.*

Definition 2.3 (Repeated Games)**.** *Given a game $\mathcal{G} = (\mathcal{Q},\mathcal{A},\pi,W)$ where $\mathcal{Q} = \mathcal{Q}_1\times\cdots\times\mathcal{Q}_k$ and $\mathcal{A} = \mathcal{A}_1\times\cdots\times\mathcal{A}_k$, its λ-fold parallel repetition is defined as $\mathcal{G}^\lambda = (\mathcal{Q}',\mathcal{A}',\pi',W')$ where $\mathcal{Q}' = \mathcal{Q}_1^\lambda\times\cdots\times\mathcal{Q}_k^\lambda$, $\mathcal{A}' = \mathcal{A}_1^\lambda\times\cdots\times\mathcal{A}_k^\lambda$, $\pi'(q) = \prod_{i=1}^\lambda \pi(q^{(i)})$, and $W'(q,a) = \prod_{i=1}^\lambda W(q^{(i)},a^{(i)})$.*

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

Definition 2.4 (Strategies)**.** *A* strategy *for a game $\mathcal{G} = (\mathcal{Q},\mathcal{A},\pi,W)$ is a collection of distributions $\{P(q)\}_{q\in\mathcal{Q}}$ over $\mathcal{A}$.*

Definition 2.5 (Local Strategies)**.** *A strategy P for a k-player game $\mathcal{G} = (\mathcal{Q},\mathcal{A},\pi,W)$ is* local *if there are functions $P_1,\ldots,P_k$ such that $P(q)\equiv(P_1(q_1),\ldots,P_k(q_k))$.*

Definition 2.6 (Non-signaling Strategies)**.** *A strategy P for a k-player game $\mathcal{G} = (\mathcal{Q},\mathcal{A},\pi,W)$ is* non-signaling *if for every subset $S\subseteq[k]$, there is a (possibly inefficient) probabilistic algorithm Sim_S such that for all $q\in\mathcal{Q}$, it holds that $P(q)_S\equiv\mathsf{Sim}_S(q_S)$. If instead there exists $\epsilon>0$ such that for all $q\in\mathcal{Q}$, it holds that $P(q)_S\approx_\epsilon\mathsf{Sim}_S(q_S)$, then P is ϵ-*non-signaling.

Definition 2.7. *The* value *of a game $\mathcal{G}$ with respect to a strategy P for $\mathcal{G}$, denoted by $v[P](\mathcal{G})$, is the expected value of $W(Q,A)$ where Q is distributed according to π and A is distributed according to $P(q)$ conditioned on $Q=q$. 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 $v_{\mathsf{ns}}(\mathcal{G})$, is the maximum value of $\mathcal{G}$ with respect to any non-signaling strategy.*

2.3 Multi-prover Interactive Proofs

Definition 2.8. *A $k(\cdot)$-*prover multi-prover interactive proof system (MIP) *is a p.p.t. interactive Turing machine[2] V with the following syntax:*

1. *On input $x\in\{0,1\}^n$, V sends a message $q = (q_1,\ldots,q_{k(n)})$.*
2. *Upon receiving a response $a = (a_1,\ldots,a_{k(n)})$, V outputs either 0 or 1.*

Definition 2.9. *For a MIP V, we define for every n-bit string x its* associated game *$\mathcal{G}_x = (\mathcal{Q},\mathcal{A},\pi,W)$, where*

[2] A formal definition of an interactive Turing machine is given in [Gol01]

- $\mathcal{Q}$ *and* $\mathcal{A}$ *are both* $\left(\{0,1\}^{T(n)}\right)^{k(n)}$*, i.e. the set of all* $k(n)$*-tuples of* $T(n)$*-bit strings, where* $T(\cdot)$ *is a bound on the running time of* V*.*
- $\pi(q)$ *is the probability that* V *sends the message* q *on input* x*.*
- $W(q,a)$ *is the probability that* V *outputs* 1*, conditioned on sending* q *and receiving* a*.*

Fact 2.10. *For any MIP* V*, the mapping from* x *to* $\mathcal{G}_x$ *is computable in polynomial space.*

Definition 2.11. *A MIP is said to recognize a language* $\mathcal{L}$ *with*

- Completeness $c(\cdot)$ *if whenever* $x \in \mathcal{L}$*,* $v(\mathcal{G}_x) \geq c(|x|)$*.*
- Soundness error $s(\cdot)$ *if whenever* $x \notin \mathcal{L}$*,* $v(\mathcal{G}_x) \leq s(|x|)$*.*
- Non-signaling soundness error $s(\cdot)$ *if whenever* $x \notin \mathcal{L}$*,* $v_{\mathsf{ns}}(\mathcal{G}_x) \leq s(|x|)$*.*

Unless otherwise specified, we consider MIPs with completeness 1 *and (non-signaling) soundness error* $1 - 1/\mathsf{poly}(|x|)$*.*

3 The Counterexample

Theorem 3.1. *There is a 3-player game* $\mathcal{G}$ *such that* $v_{\mathsf{ns}}(\mathcal{G}^\lambda) = 2/3$ *for all* $\lambda \geq 1$*.*

The game is $\mathcal{G} = (\{0,1\}^3, \{0,1\}^3, \pi, W)$, where π is the uniform distribution on strings of Hamming weight 2, namely $\{011, 101, 110\}$. If $\{i : q_i = 1\} = \{j, k\}$, then

$$W(q,a) \stackrel{\mathsf{def}}{=} \begin{cases} 1 & \text{if } a_j \neq a_k \\ 0 & \text{otherwise.} \end{cases}$$

This is called the *anti-correlation game* [LW16, FRV16] because informally, it asks a random pair of players to output distinct values. These previous works identified this as an example of a game whose non-signaling value is $2/3$, but whose *sub-non-signaling* value (a value which is known to decrease under parallel repetition) is 1.

3.1 Upper Bounding $v_{\mathsf{ns}}(\mathcal{G}^\lambda)$

We first upper-bound the non-signaling value $v_{\mathsf{ns}}(\mathcal{G})$ (and hence also $v_{\mathsf{ns}}(\mathcal{G}^\lambda)$ for all $\lambda \geq 1$). This upper bound is folklore, but we include it for completeness.

Claim 3.2. $v_{\mathsf{ns}}(\mathcal{G}) \leq \frac{2}{3}$*.*

Proof. Let P be any non-signaling strategy for $\mathcal{G}$. We have

$$v[P](\mathcal{G}) = \frac{1}{3} \cdot (\Pr[A_1 \neq A_2 | Q = 110] + \Pr[A_1 \neq A_3 | Q = 101] + \Pr[A_2 \neq A_3 | Q = 011])$$

in the probability space where Q is uniformly distributed on $\{110, 101, 011\}$, and the distribution of $A = (A_1, A_2, A_3)$ conditioned on $Q = q$ is $P(q)$. Because P is non-signaling, this is equal to

$$\frac{1}{3} \cdot (\Pr[A_1 \neq A_2] + \Pr[A_1 \neq A_3] + \Pr[A_2 \neq A_3])$$

in the probability space where A is distributed as $P(111)$. But in any probability space where A_1, A_2, and A_3 are binary-valued random variables, we have

$$\Pr[A_1 \neq A_2] + \Pr[A_1 \neq A_3] + \Pr[A_2 \neq A_3] \leq 2,$$

because the pigeonhole principle rules out all three events occurring simultaneously in any outcome. □

3.2 Lower Bounding $v_{\mathsf{ns}}(\mathcal{G}^\lambda)$

Claim 3.3. *For all $\lambda \geq 1$, $v_{\mathsf{ns}}(\mathcal{G}^\lambda) \geq \frac{2}{3}$.*

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

Construction 3.4. *Given $q = (q^{(1)}, \ldots, q^{(\lambda)})$, P samples answers $a = (a^{(1)}, \ldots, a^{(\lambda)})$ as follows.*

- *If no $q^{(i)}$ is equal to* 111*, then:*
  1. *Sample $b \leftarrow \mathrm{Ber}(1/3)$, 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 $\mathcal{G}$. 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. The requirement that consistent answer distributions exist for queries which are never asked is a strange but integral[3] part of the definition of non-signaling strategies.

We claim that the value of $\mathcal{G}^\lambda$ 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 $\mathcal{G}^\lambda$ 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 $a^{(1)}, \ldots, a^{(\lambda)}$. 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 $\{0, 1\}^\lambda$, 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, $a_S^{(1)}, \ldots, a_S^{(\lambda)}$ are independent as random variables.

[3] If a player could refuse to answer on queries which are never asked, then this itself would signal information about the other players' queries. Thus, the players must be able to answer on all queries.

Case 2: $|S| = 2$. For any q, we claim that $(a_S^{(1)}, \ldots, a_S^{(\lambda)}) = P(q)_S$ 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 $a_S^{(i)}$. □

4 Multi-player Games with Known Parallel Repetition Bounds

Despite our specific counterexample in Section 3, there *are* classes of multi-player games (most notably, two-player games [Hol09] and games with complete support [BFS14, FRV16]) whose values decrease exponentially under parallel repetition. These results are generalized and subsumed by [LW16], who define a "proxy value" $v_{\mathsf{sns}}(\mathcal{G})$ for every k-player game $\mathcal{G}$ such that (i) $v_{\mathsf{ns}}(\mathcal{G}) \leq v_{\mathsf{sns}}(\mathcal{G})$ and (ii) if $v_{\mathsf{sns}}(\mathcal{G}) = 1 - \delta$, then $v_{\mathsf{sns}}(\mathcal{G}^\lambda) \leq \left(1 - \frac{\delta^2}{O(2^{2k})}\right)^\lambda$. Whenever $v_{\mathsf{sns}}(\mathcal{G}) < 1$, one thus obtains $v_{\mathsf{ns}}(\mathcal{G}^\lambda) \leq v_{\mathsf{sns}}(\mathcal{G}^\lambda) \leq \exp(-\Omega(\lambda))$.

Specifically, [LW16] studies the following family of strategies.

Definition 4.1 (Sub-non-signaling Strategies [LW16])**.** *A* sub-non-signaling *strategy P for a k-player game $\mathcal{G} = (\mathcal{Q}, \mathcal{A}, \pi, W)$ is a collection of non-negative* densities $\{P(q)\}_{q \in \mathcal{Q}}$ *over $\mathcal{A}$ such that for every $S \subseteq [k]$, there exists a (possibly inefficient) probabilistic algorithm Sim_S such that for every $q \in \mathcal{Q}$ and $a_S \in \mathcal{A}_S$, $P(q)_S(a_S) \leq \mathsf{Sim}_S(q_S)(a_S)$.*

Definition 4.2. *The* sub-non-signaling value $v_{\mathsf{sns}}(\mathcal{G})$ *is the maximum of $\sum_{q \in \mathcal{Q}} \pi(q) \sum_{a \in \mathcal{A}} P(q)(a) W(q, a)$ with respect to any sub-non-signaling strategy P.*

Are these existing positive results sufficient for applications? In particular, do they help us build k-prover MIPs for hard languages with small non-signaling soundness error?

We give a partial negative answer in Theorem 4.3. We give a relatively efficient algorithm that distinguishes for any k-player game $\mathcal{G}$ whether $v(\mathcal{G}) = 1$ or $v_{\mathsf{sns}}(\mathcal{G}) \leq 1 - \delta$.

Theorem 4.3. *There is an algorithm that given a k-player game $\mathcal{G}$ and $\delta > 0$, distinguishes in space $\mathsf{poly}(\log |\mathcal{G}|, 2^k/\delta)$ whether $v(\mathcal{G}) = 1$ or $v_{\mathsf{sns}}(\mathcal{G}) \leq 1 - \delta$.*

The following corollary is an immediate application. Informally, this states that any log-prover MIP obtaining non-signaling soundness error lower than $1 - 1/\mathsf{poly}(n)$ via the parallel repetition result of [LW16] is limited to languages in PSPACE.

Corollary 4.4. *If a language $\mathcal{L}$ has a* log*-prover MIP with completeness* 1 *and sub-non-signaling soundness error $1 - 1/\mathsf{poly}(n)$, then $\mathcal{L} \in$ PSPACE.*

Proof. This follows directly from combining Fact 2.10 and Theorem 4.3. □

On the other hand, as far as we know even *three*-prover MIPs with constant non-signaling soundness error may exist for all of EXP. This contrast may be interpreted as evidence that known parallel repetition results apply to a relatively limited class of games.

Towards proving Theorem 4.3, we first prove an analogous statement Lemma 4.7 for the value $v_{\mathsf{hv\text{-}\epsilon\text{-}ns}}(\mathcal{G})$ against *honest-verifier ϵ-non-signaling* strategies. Then we show in Lemma 4.10 that with an appropriate choice of ϵ, $v_{\mathsf{hv\text{-}\epsilon\text{-}ns}}(\mathcal{G})$ is not much larger than $v_{\mathsf{sns}}(\mathcal{G})$.

Definition 4.5 (Honest-verifier ϵ-non-signaling Strategies)**.** *For any $\epsilon \geq 0$, a* honest-verifier ϵ-non-signaling *strategy P for a k-player game $\mathcal{G} = (\mathcal{Q}, \mathcal{A}, \pi, W)$ is a strategy where for every $S \subseteq [k]$, there exists a (possibly inefficient) probabilistic algorithm Sim_S such that $(\mathcal{Q}, P(\mathcal{Q})_S) \approx_\epsilon (\mathcal{Q}, \mathsf{Sim}_S(\mathcal{Q}_S))$ or equivalently $\sum_{q\in\mathcal{Q}} \pi(q) \cdot d_{\mathsf{TV}}(P(q)_S, \mathsf{Sim}_S(q_S)) \leq \epsilon$.*

Remark 4.6. *For all $\epsilon \geq 0$, every ϵ-non-signaling strategy is also an honest-verifier ϵ-non-signaling strategy.*

By constructing two-player simulations of k-player games, we show the following.

Lemma 4.7. *There is an algorithm that given a k-player game $\mathcal{G}$ and $\epsilon, \delta > 0$, distinguishes in space $\mathsf{poly}(\log|\mathcal{G}|, 2^k/\delta\epsilon)$ whether $v(\mathcal{G}) = 1$ or $v_{\mathsf{hv\text{-}\epsilon\text{-}ns}}(\mathcal{G}) \leq 1 - \delta$.*

Proof. For a k-player game $\mathcal{G} = (\mathcal{Q}, \mathcal{A}, \pi, W)$, consider the two-player game $\hat{\mathcal{G}} = (\hat{\mathcal{Q}}, \hat{\mathcal{A}}, \hat{\pi}, \hat{W})$ where

- $\hat{\mathcal{Q}} = \{(q, (S, q')) : q \in \mathcal{Q}, S \subseteq [k], q' \in \mathcal{Q}_S\}$,
- $\hat{\mathcal{A}} = \{(a, a') : a \in \mathcal{A}, a' \in \mathcal{A}_S \text{ for some } S \subseteq [k]\}$,
- $\hat{\pi}(q, (S, q')) = \begin{cases} 2^{-k} \cdot \pi(q) & \text{if } q' = q_S \\ 0 & \text{otherwise,} \end{cases}$
- $\hat{W}((q, (S, q')), (a, a')) = W(q, a) \wedge (a' = a_S)$.

Intuitively, in $\hat{\mathcal{G}}$ the first player is given all k queries from $\mathcal{G}$, and the second player is given a random subset of these queries. The second player is used to check that the first player's answer distribution is a honest-verifier ϵ-non-signaling strategy.

Claim 4.8. *If $v(\mathcal{G}) = 1$, then $v(\hat{\mathcal{G}}) = 1$.*

Proof. Let $P = (P_1, \ldots, P_k)$ be a local strategy for $\mathcal{G}$ such that $v[P](\mathcal{G}) = 1$. Then define a local strategy $\hat{P} = (\hat{P}_1, \hat{P}_2)$ for $\hat{\mathcal{G}}$ as follows: $\hat{P}_1(q_1, \ldots, q_k) = (P_1(q_1), \ldots, P_k(q_k))$ and $\hat{P}_2(S, q') = (P_i(q'_i))_{i\in S}$. □

Claim 4.9. $v_{\mathsf{ns}}(\hat{\mathcal{G}}) \leq \max(v_{\mathsf{hv\text{-}\epsilon\text{-}ns}}(\mathcal{G}), 1 - \epsilon/2^k)$.

Proof. For any non-signaling strategy $\hat{P} = \{\hat{P}(q, (S, q'))\}_{q\in\mathcal{Q}, S\subseteq[k], q'\in\mathcal{Q}_S}$ for $\hat{\mathcal{G}}$, there exist distributions $\{\hat{P}_1(q)\}_{q\in\mathcal{Q}}$ and $\{\hat{P}_2(S, q')\}_{S\subseteq[k], q'\in\mathcal{Q}_S}$ such that

$$\forall q \in \mathcal{Q}, \hat{P}(q, (S, q'))_{\{1\}} \equiv \hat{P}_1(q)$$
$$\forall S \subseteq [k] \text{ and } q' \in \mathcal{Q}_S, \hat{P}(q, (S, q'))_{\{2\}} \equiv \hat{P}_2(S, q').$$

If $\{\hat{P}_1(q)\}_{q\in\mathcal{Q}}$ is a honest-verifier ϵ-non-signaling strategy for $\mathcal{G}$, then the probability that $\hat{P}$ wins $\hat{\mathcal{G}}$ is at most $v_{\mathsf{hv\text{-}\epsilon\text{-}ns}}(\mathcal{G})$.

Otherwise for some $S^* \subseteq [k]$, we have $\sum_{q\in\mathcal{Q}} \pi(q) d_{\mathsf{TV}}(\hat{P}_1(q)_{S^*}, \hat{P}_2(S^*, q_{S^*})) > \epsilon$. Then let $\epsilon_q = d_{\mathsf{TV}}(\hat{P}_1(q)_{S^*}, \hat{P}_2(S^*, q_{S^*}))$. Note that for any jointly distributed random variables (X, Y), the probability that $X = Y$ is at most $1 - d_{\mathsf{TV}}(X, Y)$. Therefore, the probability that $A_{S^*} = A'$, when (A, A') is distributed according to $\hat{P}(Q, (S^*, Q_{S^*}))$ and Q is distributed according to π, is at most $\sum_{q\in\mathcal{Q}} \pi(q)(1 - \epsilon_q) < 1 - \epsilon$. Therefore, the probability that $\hat{P}$ wins $\hat{\mathcal{G}}$ is at most

$$\Pr[S \neq S^*] + \Pr[S = S^*] \Pr[A_{S^*} = A' | S = S^*] \leq 1 - \frac{1}{2^k} + \frac{1}{2^k}(1 - \epsilon) = 1 - \frac{\epsilon}{2^k}.$$

Combining these two cases, we obtain $v_{\mathsf{ns}}(\hat{\mathcal{G}}) \leq \max(v_{\mathsf{hv\text{-}\epsilon\text{-}ns}}(\mathcal{G}), 1 - \epsilon/2^k)$. □

Finally, we use an algorithm of Ito [Ito10, Theorem 2] which additively ϵ'-approximates the non-signaling value of any two-player game $\mathcal{G}$ in space $\mathsf{poly}(\log|\mathcal{G}|, 1/\epsilon')$. This algorithms decides whether $v(\hat{\mathcal{G}}) = 1$ or $v_{\mathsf{ns}}(\hat{\mathcal{G}}) \leq \max(1-\delta, 1-\epsilon/2^k))$ in space $\mathsf{poly}(\log|\mathcal{G}|, 2^k/\delta\epsilon)$, hence deciding whether $v(\mathcal{G}) = 1$ or $v_{\mathsf{hv}\text{-}\epsilon\text{-}\mathsf{ns}}(\mathcal{G}) \leq 1-\delta$. □

To finish the proof of Theorem 4.3, we relate $v_{\mathsf{hv}\text{-}\epsilon\text{-}\mathsf{ns}}(\mathcal{G})$ to $v_{\mathsf{sns}}(\mathcal{G})$.

Lemma 4.10. $v_{\mathsf{hv}\text{-}\epsilon\text{-}\mathsf{ns}}(\mathcal{G}) \leq v_{\mathsf{sns}}(\mathcal{G}) + 2\epsilon \cdot 2^k$.

Proof. We show that for any honest-verifier ϵ-non-signaling strategy $P = \{P(q)\}_{q\in\mathcal{Q}}$ for $\mathcal{G} = (\mathcal{Q}, \mathcal{A}, \pi, W)$, there is a sub-non-signaling strategy $P' = \{P'(q)\}_{q\in\mathcal{Q}}$ with $v[P'](\mathcal{G}) \geq v[P](\mathcal{G}) - 2\epsilon \cdot 2^k$. Let $\{\mathsf{Sim}_S\}_{S\subseteq[k]}$ be the simulators for P as in Definition 4.5. For any q, we view $P(q)$ and $\mathsf{Sim}_S(q_S)$ as probability mass functions $P(q) : \mathcal{A} \to \mathbb{R}_{\geq 0}$ and $\mathsf{Sim}_S(q_S) : \mathcal{A}_S \to \mathbb{R}_{\geq 0}$.

We define P' for each $q \in \mathcal{Q}, a \in \mathcal{A}$ as follows: if there exists $S \subset [k]$ such that $P(q)_S(a_S) - \mathsf{Sim}_S(q_S)(a_S) > 0$, then

$$P'(q)(a) = P(q)(a) - \max_{S\subset[k]} \left(\frac{P(q)(a)}{P(q)_S(a_S)} \left(P(q)_S(a_S) - \mathsf{Sim}_S(q_S)(a_S)\right)\right).$$

Otherwise, $P'(q)(a) = P(q)(a)$. Note that $P'(q)(a) \geq 0$ for all $q \in \mathcal{Q}, a \in \mathcal{A}$.

We verify that $\{P'(q)\}_{q\in\mathcal{Q}}$ satisfies the sub-non-signaling constraints. Note that for any $S \subsetneq [k], q \in \mathcal{Q}$, and $a \in \mathcal{A}$,

$$P'(q)(a) \leq P(q)(a) - \left(\frac{P(q)(a)}{P(q)_S(a_S)} \left(P(q)_S(a_S) - \mathsf{Sim}_S(q_S)(a_S)\right)\right) = \frac{P(q)(a)}{P(q)_S(a_S)} \mathsf{Sim}_S(q_S)(a_S).$$

Thus for any $S \subsetneq [k], q \in \mathcal{Q}$, and $a_S \in \mathcal{A}_S$, we have

$$P'(q)_S(a_S) = \sum_{a'\in\mathcal{A}: a'_S = a_S} P'(q)(a') \leq \frac{\mathsf{Sim}_S(q_S)(a_S)}{P(q)_S(a_S)} \cdot \sum_{a'\in\mathcal{A}: a'_S = a_S} P(q)(a') = \mathsf{Sim}_S(q_S)(a_S).$$

Note that for any $q \in \mathcal{Q}$,

$$\begin{aligned}
\sum_{a\in\mathcal{A}} |P(q)(a) - P'(q)(a)| &\leq \sum_{a\in\mathcal{A}} \max_{S\subset[k]} \left(\frac{P(q)(a)}{P(q)_S(a_S)} \Big| P(q)_S(a_S) - \mathsf{Sim}_S(q_S)(a_S)\Big|\right) \\
&\leq \sum_{S\subset[k]} \sum_{a\in\mathcal{A}} \frac{P(q)(a)}{P(q)_S(a_S)} \Big| P(q)_S(a_S) - \mathsf{Sim}_S(q_S)(a_S)\Big| \\
&= \sum_{S\subset[k]} \sum_{a'\in\mathcal{A}_S} \Big| P(q)_S(a') - \mathsf{Sim}_S(q_S)(a')\Big| \\
&= 2 \sum_{S\subset[k]} d_{\mathsf{TV}}(P(q)_S, \mathsf{Sim}_S(q_S)).
\end{aligned}$$

Therefore, the difference $\Big|v[P](\mathcal{G}) - v[P'](\mathcal{G})\Big|$ is at most

$$\sum_{q\in\mathcal{Q}} \pi(q) \sum_{a\in\mathcal{A}} \Big| P(q)(a) - P'(q)(a)\Big| \leq 2 \sum_{q\in\mathcal{Q}} \pi(q) \sum_{S\subset[k]} d_{\mathsf{TV}}(P(q)_S, \mathsf{Sim}_S(q_S)),$$

which is bounded by $2 \cdot 2^k \cdot \epsilon$ because $\{P(q)\}_{q\in\mathcal{Q}}$ is honest-verifier ϵ-non-signaling. □

Proof of Theorem 4.3. If $v_{\mathsf{sns}}(\mathcal{G}) \leq 1-\delta$, then for $\epsilon = \frac{\delta}{4\cdot 2^k}$, $v_{\mathsf{hv}\text{-}\epsilon\text{-}\mathsf{ns}}(\mathcal{G}) \leq 1 - (\delta - 2\cdot 2^k \cdot \epsilon) = 1 - \delta/2$. Then by Lemma 4.7, there is an algorithm that decides in space $\mathsf{poly}(\log|\mathcal{G}|, 2^k/\delta)$ whether $v_{\mathsf{ns}}(\mathcal{G}) = 1$ or $v_{\mathsf{sns}}(\mathcal{G}) \leq 1-\delta$. □

Let $\mathsf{MIP}_{\mathsf{ns}}(k)$ and $\mathsf{MIP}_{\mathsf{sns}}(k)$ denote the classes of languages with k-prover non-signaling MIPs and k-prover sub-non-signaling MIPs, respectively. In Table 2, we summarize what is known about $\mathsf{MIP}_{\mathsf{ns}}(k)$ and $\mathsf{MIP}_{\mathsf{sns}}(k)$. Note that $\mathsf{MIP}_{\mathsf{sns}}(k) \subseteq \mathsf{MIP}_{\mathsf{ns}}(k)$ since the set of sub-non-signaling strategies strictly contains the set of non-signaling strategies. It is known that $\mathsf{MIP}_{\mathsf{ns}}(2) = \mathsf{PSPACE}$ [IKM09, Ito10] and $\mathsf{MIP}_{\mathsf{ns}}(\mathsf{poly}) = \mathsf{EXP}$ [KRR14]. If $\mathsf{MIP}_{\mathsf{ns}}(\mathsf{log})$ strictly contains PSPACE, then Corollary 4.4 shows that non-signaling MIPs using only games with known parallel repetition bounds are strictly weaker than non-signaling MIPs in general. We note that it even remains open to resolve whether $\mathsf{MIP}_{\mathsf{ns}}(3)$ strictly contains PSPACE.

Number of Provers k	$\mathsf{MIP}_{\mathsf{ns}}(k)$	$\mathsf{MIP}_{\mathsf{sns}}(k)$
2	$= \mathsf{PSPACE}$ [IKM09, Ito10]	$= \mathsf{PSPACE}$
log	$\subseteq \mathsf{EXP}$	$\subseteq \mathsf{PSPACE}$ [**This Work**]
poly	$= \mathsf{EXP}$ [KRR14]	$\subseteq \mathsf{EXP}$

Table 2: Known results on $\mathsf{MIP}_{\mathsf{ns}}(k)$ and $\mathsf{MIP}_{\mathsf{sns}}(k)$.

Acknowledgments

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.

[Bel64] J. S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.

[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.

[BMW98] Ingrid Biehl, Bernd Meyer, and Susanne Wetzel. Ensuring the integrity of agent-based computations by short proofs. In *Mobile Agents*, volume 1477 of *Lecture Notes in Computer Science*, pages 183–194. Springer, 1998.

[DHRW16] Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In *CRYPTO (3)*, volume 9816 of *Lecture Notes in Computer Science*, pages 93–122. Springer, 2016.

[DLN+01] Cynthia Dwork, Michael Langberg, Moni Naor, Kobbi Nissim, and Omer Reingold. Succinct proofs for np and spooky interactions. Unpublished manuscript, 2001.

[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.

[For89] Lance Jeremy Fortnow. *Complexity-theoretic aspects of interactive proof systems*. PhD thesis, 1989.

[FRS88] Lance Fortnow, John Rompel, and Michael Sipser. On the power of multi-prover interactive protocols. In *Structure in Complexity Theory Conference*, pages 156–161. IEEE Computer Society, 1988.

[FRV16] 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.

[Gol01] Oded Goldreich. *The Foundations of Cryptography - Volume 1, Basic Techniques*. Cambridge University Press, 2001.

[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.

[Ito10] Tsuyoshi Ito. Polynomial-space approximation of no-signaling provers. In *ICALP (1)*, volume 6198 of *Lecture Notes in Computer Science*, pages 140–151. Springer, 2010.

[Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *STOC*, pages 723–732. ACM, 1992.

[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.