# Algorithms, Bounds, and Strategies for Entangled XOR Games

**Adam Bene Watts**[1]
MIT Center for Theoretical Physics, 77 Massachusetts Ave, 6-304, Cambridge, MA, USA
abenewat@mit.edu

**Aram W. Harrow**[2]
MIT Center for Theoretical Physics, 77 Massachusetts Ave, 6-304, Cambridge, MA, USA
aram@mit.edu

**Gurtej Kanwar**[3]
MIT Center for Theoretical Physics, 77 Massachusetts Ave, 6-304, Cambridge, MA, USA
gurtej@mit.edu

**Anand Natarajan**[4]
California Institute of Technology, 1200 E. California Blvd, Pasadena, CA, USA
anandn@caltech.edu

## Abstract

Entangled games are a quantum analog of constraint satisfaction problems and have had important applications to quantum complexity theory, quantum cryptography, and the foundations of quantum mechanics. Given a game, the basic computational problem is to compute its *entangled value*: the supremum success probability attainable by a quantum strategy. We study the complexity of computing the (commuting-operator) entangled value $\omega^*$ of entangled XOR games with any number of players. Based on a duality theory for systems of operator equations, we introduce necessary and sufficient criteria for an XOR game to have $\omega^* = 1$, and use these criteria to derive the following results:

1. An algorithm for symmetric games that decides in polynomial time whether $\omega^* = 1$ or $\omega^* < 1$, a task that was not previously known to be decidable, together with a simple tensor-product strategy that achieves value 1 in the former case. The only previous candidate algorithm for this problem was the Navascués-Pironio-Acín (also known as noncommutative Sum of Squares or ncSoS) hierarchy, but no convergence bounds were known.
2. A family of games with three players and with $\omega^* < 1$, where it takes doubly exponential time for the ncSoS algorithm to witness this. By contrast, our algorithm runs in polynomial time.
3. Existence of an unsatisfiable phase for random (non-symmetric) XOR games. We show that there exists a constant $C_k^{\text{unsat}}$ depending only on the number $k$ of players, such that a random $k$-XOR game over an alphabet of size $n$ has $\omega^* < 1$ with high probability when the number of clauses is above $C_k^{\text{unsat}} n$.
4. A lower bound of $\Omega(n \log(n)/\log\log(n))$ on the number of levels in the ncSoS hierarchy required to detect unsatisfiability for most random 3-XOR games. This is in contrast with the classical case where the $(3n)^{\text{th}}$ level of the sum-of-squares hierarchy is equivalent to brute-force enumeration of all possible solutions.

**2012 ACM Subject Classification** Theory of computation → Quantum complexity theory

[1] ABW was supported by NSF grant CCF-1729369.
[2] AWH was funded by NSF grants CCF-1452616, CCF-1729369, ARO contract W911NF-17-1-0433 and the MIT-IBM Watson AI Lab under the project *Machine Learning in Hilbert space*.
[3] GK was supported by DOE grant DE-SC0011090.
[4] AN was supported by NSF grant CCF-1452616.


10th Innovations in Theoretical Computer Science (ITCS 2019).
Editor: Avrim Blum; Article No. 10; pp. 10:1–10:18

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

**Keywords and phrases** Nonlocal games, XOR Games, Pseudotelepathy games, Multipartite entanglement

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2019.10

**Related Version** A full version of the paper is available at `https://arxiv.org/abs/1801.00821`.

**Acknowledgements** The authors would like to thank Jöp Briet for helpful discussions, Oded Regev for pointing out the behavior of parallel-repeated GHZ games, anonymous reviewers for helpful comments on an earlier draft of this work, as well as George Wang, who contributed to this project as part of the RSI summer program at MIT. Parts of this work were completed while AWH and AN were hosted at the Institut Henri Poincaré in Paris, as part of the special semester on Analysis in Quantum Information Theory (Fall 2017), supported by NSF Grant DMS-1700168. The hospitality of the IHP is gratefully acknowledged.

## 1 Introduction

Constraint satisfaction problems (CSPs) are a fundamental object of study in theoretical computer science. In quantum information theory, there are two natural analogues of CSPs, which both play important roles: local Hamiltonians and (our focus) non-local games. Non-local games originate from Bell's pioneering 1964 paper, which showed how to test for quantum entanglement in a device with which we can interact only via classical inputs and outputs. In modern language, the tests developed by Bell are games: a referee presents two or more players with classical questions drawn from some distribution and demands answers from them. Each combination of question and answers receives some score and the players cooperate (but may not communicate) to maximize their expected score. These games are interesting because often the players can win the game with a higher probability if they share an entangled quantum state, so a high average score can certify the presence of quantum entanglement. Such tests are not only of scientific interest, but have had wide application to proof systems [8, 19], quantum key distribution [1, 13, 29], delegated computation [26], and randomness generation [10], among others.

To be able to use a nonlocal game as a test for entanglement, it is essential to be able to approximately compute two quantities: the best possible expected score when the players share either classical correlations or entangled states, respectively called the "classical" and "quantum" (or "entangled") values of the game, and denoted $\omega$ and $\omega^*$. Classically, our understanding of the complexity of computing $\omega$ rests on the intimate connection between games and CSPs. Indeed, there are several natural to ways to map a CSP into a game. Perhaps the most commonly used is the "clause-variable game," in which a CSP of any arity $k$ is mapped to a two-player game, where one player is asked for the assignment to a clause of the CSP, and the other for the assignment to a single variable. However, there is another natural yet perhaps less-studied reduction that maps a $k$-ary CSP to a $k$-player nonlocal game, which moreover is symmetric under exchange of the players. In this reduction, given a CSP with a $k$-ary predicate, the referee of the game chooses uniformly at random a single clause, consisting of a $k$-tuple of variables and set of accepted assignments. The referee will then ask each of the $k$ players for the value of one of the $k$ variables in the clause, and accept if and only if the returned values constitute an accepted assignment to the clause. Classically, a simple convexity argument shows that the players can always stick to *deterministic* strategies, where each question is assigned a fixed answer, and for odd $k$,

it is easy to show that there is a close relation between $\omega$ and the CSP value: if the CSP has value 1 (i.e. all clauses are satisfiable), $\omega = 1$, and if its value is at most $1 - \delta$, then $\omega \leq 1 - \delta/k$. Hence, thanks to various dichotomy theorems, we have a good understanding of the difficulty of computing $\omega$ for symmetric games[5]: in some cases, we know a P algorithm, and for most others, we know it is NP-complete. In particular, thanks to [18], this is known even for games where the referee's acceptance depends only the XOR of single-bit answers from the players. Such games are known as XOR games.

The hardness of computing quantum value $\omega^*$ is not as well understood, both in terms of upper and lower bounds. We know striking examples of quantum "advantage" (i.e. cases where the quantum value of a game is higher than the classical value), such as a Magic Square game, a game arising from an unsatisfiable CSP which nevertheless has an entangled strategy that succeeds with certainty, and thus quantum value $\omega^* = 1$. This advantage is also the main obstacle to our understanding, in that the set of entangled strategies is very rich: the "assignment" to each variable is no longer a value from a discrete set, but a linear operator over a Hilbert space of potentially unbounded dimension. Indeed, if infinite-dimensional entanglement is allowed, then depending on how one implements the requirement of non-communication between the players, one can obtain two different notions of entangled value – the tensor product value $\omega^*_{TP}$ and the commuting operator value $\omega^*_{CO}$ – which are not known to be equal.

As a result of the difficulties of unbounded-dimensional entanglement, we can say very little in terms of upper bounds on the complexity of computing either version of $\omega^*$, and in fact, it is not known whether even a constant-factor (additive) approximation to either is Turing-computable. For general games, the best we can say is they are recursively enumerable: for $\omega^*_{TP}$, there is a straightforward brute-force search over all strategies that in the limit of infinite time converges from below, and for $\omega^*_{CO}$, there is an algorithm, called the NPA or ncSoS hierarchy [22, 11], that in the limit of infinite time converges from above, but with no bound on the speed of convergence for either algorithm. On the hardness side, what we know is based on exploiting the CSP-game connection outlined above, but technically this has proved significantly more challenging than in the classical case. For instance, it was shown by Vidick that in the worst case, computing a constant-factor approximation to $\omega^*_{TP}$ for 3-player XOR games is NP-hard [30], matching the classical hardness of [18], but this required redoing the soundness analysis of a PCP construction in the presence of entanglement. For general (non-XOR) games and tighter approximations we have super-classical hardness results [20, 21, 14]. Moreover, families of games with a "clause-variable" structure have been found for which deciding whether $\omega^* = 1$ is uncomputable [27]. At the same time, we know that for certain families of games, $\omega^*$ is easy to compute. Perhaps the best understood case is two-player XOR-games, for which Tsirelson showed that a simple semidefinite program (the lowest level of the ncSoS) exactly computes $\omega^*_{CO} = \omega^*_{TP}$, in contrast to the classical case where $\omega$ for such games is NP-hard to approximate. A second family of games where results are known is XOR games with a maximum of two questions per player, but any number of players. Here there is a classification of all correlations achievable by quantum players, as well as a description of the measurement strategy players use to achieve these correlations. Interestingly, we arrive at the same measurement strategy later in this work through independent techniques.

[5] Classically, there are simple reductions from the general case to the symmetric case, but as we discuss below, these fail to preserve completeness in the presence of entanglement.

From the preceding results, XOR games emerge as a natural class of games to understand on the road to a full "dichotomy theorem" for quantum games. Classically, XOR games are also convenient to analyze because of their linear structure: a $k$-player XOR game represents a CSP whose clauses are linear equations over the finite field $\mathbb{F}_2$, each containing $k$ variables. As a result, classically XOR games are always easy in the "perfect completeness" regime: we can determine whether an XOR game is perfectly satisfiable in polynomial time using Gaussian elimination over $\mathbb{F}_2$, even though distinguishing $\geq 1-\varepsilon$ satisfiability from $\leq \frac{1}{2}+\varepsilon$ satisfiability is NP-complete. This linear structure also makes it easy to reason about the classical value of random instances of XOR games using linear algebra. However, this simple linear structure does not capture entangled strategies and neither the Gaussian elimination algorithm for the perfect completeness regime, nor the classical analysis of random instances generalizes easily to the quantum case. Indeed, the undecidability result of [27] applies to the perfect completeness regime for games based on systems of linear equations, though these systems are not over $\mathbb{F}_2$ and the games are in the two-player "clause-variable" format. Is the perfect completeness regime for quantum XOR games easy, as in the classical case, or hard, as suggested by Slofstra's results? And what can we say about random instances?

In this work, we make progress on these questions for the subclass of *symmetric* XOR games: those for which the game remains invariant under any permutation of the players. This class of games includes those arising from CSPs via the reduction described above, as well as the hard instances of [30]. Our main results are captured by the following theorem.

▶ **Theorem 1** (Theorems 14 and 15 in the body). *A symmetric $k$-player XOR game has entangled value $\omega^*_{CO} = \omega^*_{TP} = 1$ if and only if an associated system of linear Diophantine equations has no solution. This condition can be checked in polynomial time, and whenever it is satisfied, the perfect tensor-product strategy can be found in polynomial time. When it is not satisfied, a succinct description of an ncSoS dual certificate that $\omega^* < 1$ can be found in polynomial time (even though the certificate may be exponentially long).*

We achieve these results by viewing an XOR game as a "non-commuting" generalization of linear systems of equations, in which the expectation of differences between products of operator-valued variables and plus or minus the identity operator are constrained to be zero. We develop a "duality theory" for these systems of operator equations, where the dual certificates of infeasibility correspond to a special class of ncSoS proofs which we call "refutations." For symmetric games, we show that a dual certificate exists if and only if a certain system of linear Diophantine equations has a solution (which we call a "PREF"). An important feature of our algorithm is that while it is inspired by ncSoS, its performance can be significantly superior: it can detect in polynomial time the existence of an exponentially long ncSoS dual certificate. Indeed, we show (in Theorem 23) a concrete family of games where our algorithm can detect that $\omega^* < 1$ in time which scales polynomially in the game size $n$, whereas ncSoS takes doubly exponential time. We believe this result may be interesting in its own right to those who study the Sum of Squares algorithm, and hope that our techniques inspire further efficient algorithms that "simulate" high levels of SoS. Additionally, by further considering the dual of the system of Diophantine equations we construct, we are able to extract a simple finite-dimensional (and hence tensor product) strategy (which we call "MERP") that achieves $\omega^*_{TP} = \omega^*_{CO} = 1$ whenever the system of equations has a solution. A diagram illustrating the dualities we use is given in Figure 1.

Our notion of refutation is similar to the "substitution method" in the prior work of [9], used there to analyze clause-variable style games (there called Binary Constraint System Games) in the perfect-completeness regime. However, the connection we show between refutations and linear Diophantine equations, which is the heart of our efficient algorithm for

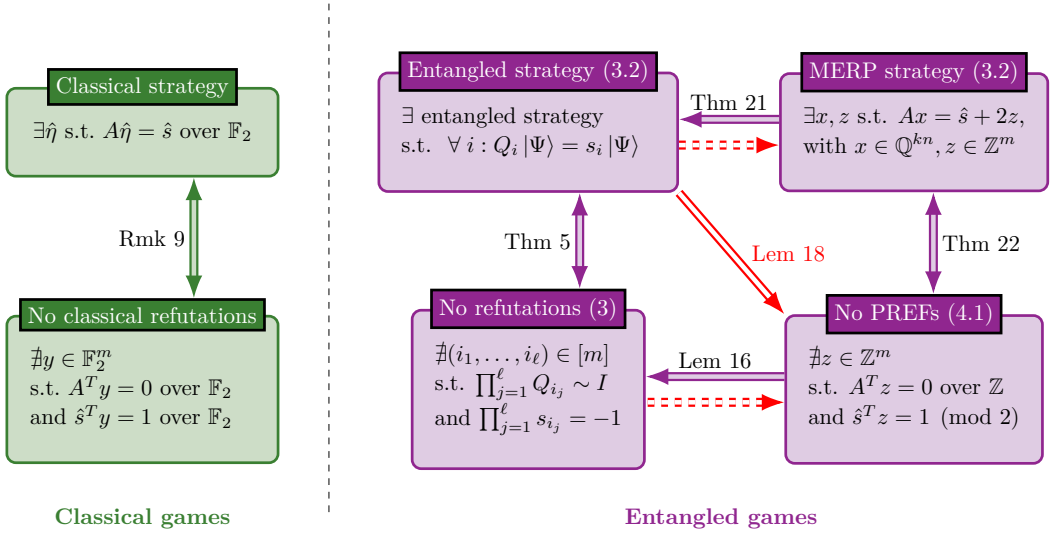


**Figure 1** We extend the well-understood duality relation for classical XOR games (left) to a more complex set of dualities characterizing perfect strategies for entangled XOR games (right). The arrows indicate implications, with the red, unfilled arrows holding for symmetric games only. The dashed red arrows follow from the key lemma for symmetric games. Definitions and notation are developed in the remaining sections.

searching over refutations, is new to this work and makes essential use of the properties of symmetric XOR games. We consider it an interesting open question whether our techniques could be adapted to the Binary Constraint System case.

The symmetry condition on the game is important to our analysis, and it is worth going into some detail as it presents an interesting divergence from the classical case. Classically, any game can be symmetrized as follows: for each clause consisting of questions $(q^{(1)}, \ldots, q^{(k)})$ asked to players $1, \ldots, k$, the referee chooses a random permutation $\pi$ of $\{1, \ldots, k\}$, and sends player $i$ the pair $(\pi(i), q^{(\pi(i))})$. Each player $i$ then follows the same strategy that player $\pi(i)$ would have used in the original, unsymmetrized game. In the quantum setting, this transformation fails to preserve completeness: for instance, if an entangled strategy for a three-player unsymmetrized game requires players 1 and 2 to share entanglement, in the symmetrized version, a player receiving the index 1 does not know which other player received index 2, and thus does not know who to be entangled with. This can be understood as an instance of the phenomenon of *monogamy*, which distinguishes entanglement from classical correlations. It is an interesting question for future work to extend our methods to the nonsymmetric case.

Furthermore, as alluded to earlier, our algorithm yields an understanding of the typical value of a random symmetric XOR game. Classically, research in this direction draws significantly on insights from statistical mechanics and has proven that there exist sharp satisfiable/unsatisfiable thresholds for random $k$-SAT and related games. But these techniques do not carry over to the quantum case. For random classical games, a basic method is to look at the expected number of winning strategies (the "first moment method") or the variance (the "second moment method") as we randomize the referee's payoff function within some family such as random $k$-SAT or random $k$-XOR. This suffices, for example, to show that random 3-XOR games with $n$ variables and $Cn$ clauses are satisfiable with high probability if and only if $C \lesssim 0.92$ [12]. Since quantum strategies do not form a discrete (or even finite-dimensional)

set, these methods are not possible. Nor is it obvious how to use more refined tools such as Shearer's Lemma or the Lovász Local Lemma, which address the question of when sets of overlapping constraints can be simultaneously satisfied. Our duality theory enables us to avoid these obstacles by studying refutations, rather than strategies. Refutations are discrete objects and thus are more amenable to combinatorial and probabilistic techniques. Using our techniques we are able to prove that random quantum XOR games have an unsatisfiable phase above a certain clause density.

▶ **Theorem 2** (Theorem 26 in the body). *For every $k$, there exists a constant $C_k^{unsat}$ depending only on $k$ such that a random $k$-XOR game $G$ with $m \geq C_k^{unsat} n$ clauses has value $\omega^*(G) < 1$ with probability $1 - o(1)$.*

In our overall approach in this paper, we were inspired by the work of Grigoriev [17], who studied the power of SoS refutations for random classical XOR games. We view Theorem 2, together with the ncSoS lower bounds of Theorem 23, as a quantum generalization of Grigoriev's results.

## 2 XOR Games

We begin by defining a $k$-XOR game, along with its classical and quantum values.

▶ **Definition 3.** Define a **clause** $c = (q, s)$ to be any $(k+1)$-tuple consisting of a **query** $q \in [n]^k$ and **parity bit** $s \in \{-1, 1\}$.

In a $k$-**XOR game** $G$ associated with a set of clauses $M$, a verifier selects a clause $c_i = (q_i, s_i)$ uniformly at random from $M$. For all $\alpha \in [k]$, the **question** $q_i^{(\alpha)}$ is then sent to the $\alpha$-th player of the game. Without communicating, the players then each send back a single output $\in \{-1, 1\}$, and win the game if their outputs multiply to $s_i$.

The GHZ game [16] is a canonical example of a 3-XOR game. It is defined by the clauses (here we use the labels $\{x, y\}$ for the questions instead of the typical $\{1, 2\}$):

$$G_{GHZ} := \left\{ \begin{bmatrix} x \\ x \\ x \\ +1 \end{bmatrix}, \begin{bmatrix} y \\ y \\ x \\ -1 \end{bmatrix}, \begin{bmatrix} y \\ x \\ y \\ -1 \end{bmatrix}, \begin{bmatrix} x \\ y \\ y \\ -1 \end{bmatrix} \right\} \begin{matrix} \leftarrow \text{Player A} \\ \leftarrow \text{Player B} \\ \leftarrow \text{Player C} \\ \leftarrow \text{Desired product} \end{matrix} . \tag{1}$$

There is a natural reduction from a $k$-CSP over $\mathbb{F}_2$ to a $k$-XOR game, based on the isomorphism between the groups $(\{0, 1\}, +_{\text{mod } 2})$ and $(\{1, -1\}, \times)$. The XOR game corresponding to a CSP has clauses defined by picking a clause from the CSP at random, sending each player a question corresponding to a random distinct variable from the CSP clause, then choosing a parity bit by demanding that players' answers satisfy the CSP. Games constructed from CSPs in this manner are symmetric over permutation of the players, and are therefore called **symmetric games**.

By excluding communication during the game, the classical game tests whether the players have cooperatively solved the CSP before the questioning began. When the players are given access to quantum resources, the game instead probes "quantum solutions" to the CSP, described by measurements of a shared state in some Hilbert space.

The **value** of that game is defined to be the optimal win probability obtained by the players. We distinguish various possible classes of resources that may be made available to the players in executing a strategy, each of which defines a particular type of value.

▶ **Definition 4.** We define three versions of the value of game $G$.

1. The **classical value** $\omega(G)$ is the value achievable by players using only classical shared information.
2. The **tensor-product value** is the value obtainable by players sharing a quantum state but restricted to making measurements on distinct factors of a tensor-product Hilbert space. Intuitively, this is a no-communication condition described in Hilbert-space language.
3. The **commuting-operator value**[6] $\omega^*(G)$ is the value obtainable by players making commuting measurements on a shared quantum state. Intuitively, this is a weaker form of the no-communication constraint, permitting states living in non-separable Hilbert spaces.

When a CSP is reduced to an XOR game $G$, the classical value roughly corresponds to the fraction of satisfiable constraints, with $\omega(G) = 1$ if and only if the CSP is satisfiable. Sharing a quantum state may allow the players to provide a quantum solution ($\omega^* = 1$) even when $\omega \neq 1$. Famously, the GHZ game is a symmetric game that corresponds to a test of the classically-unsatisfiable, yet quantumly-soluble CSP:

$$x + x + x = 0 \pmod 2 \quad \text{and} \quad x + y + y = 1 \pmod 2. \tag{2}$$

Games (such as the GHZ game) that satisfy $\omega < 1$ and $\omega^* = 1$ are called **pseudo-telepathy games**. Identifying other XOR pseudo-telepathy games is one of the motivating goals of this work.

For a given game, the set of values achievable by tensor-product strategies may not be closed [27]. Whether the closure of this set can differ from the commuting-operator value of the game remains unanswered[7]. In this paper, we focus primarily on a description of the commuting-operator value but in many cases can show that it coincides with the tensor-product value.

## 3 Refutations

A main aim of this paper is to characterize the set of XOR games with commuting-operator value $\omega^* = 1$. In the case of $k = 2$ players, Tsirelson gave an efficient semidefinite program that computes the exact value of $\omega^*$; however, this technique does not generalize easily to $k \geq 3$ [6, 28]. Furthermore, the potentially unbounded size of the players' resource state makes it impossible to upper bound the value of a game via brute force search over strategies.

To avoid these problems, this work introduces a dual characterization that certifies games with value $\omega^* < 1$. This is a natural generalization of a well-understood dual system of equations that certifies games with classical value $\omega < 1$, and employs operator language similar to the quantum satisfying assignments for Binary Constraint System games presented in [9]. The dual pictures in both the classical and commuting-operator cases introduce the notion of a "refutation": intuitively, a sequence of game clauses that together contradict the existence of a value-1 strategy. We show Theorem 5 that refutations are dual to $\omega^* = 1$; our proof can be viewed as a quantum generalization of [17].

[6] $\omega^*(G)$ is often also referred to as the field-theoretic value of $G$.
[7] And hard! For general two-player games this question is known to be equivalent to Connes' embedding conjecture [15].

▶ **Theorem 5** (Strategy-Refutation Duality). *An XOR game $G$ has commuting-operator value $\omega^*(G) = 1$ if and only if it admits no refutations.*

We first take a small detour into the classical duality picture to build intuition and necessary notation (Section 3.1), then describe refutations and outline the proof of duality in the quantum case (Section 3.2).

## 3.1 Classical Strategies and Refutations

Classically, refutations emerge naturally from the linear-algebraic dual to the equations satisfied by a classical value-1 strategy.

For any game, the optimal classical strategy can be specified via a map $[kn] \to \{1, -1\}$ giving a deterministic answer to each possible question given to each player. In order to use linear algebraic tools, we exploit the isomorphism $(\{1,-1\}, \times) \sim (\{0,1\}, +_{\text{mod }2})$ and specify a classical strategy via a vector $\hat{\eta} \in \mathbb{F}_2^{kn}$. Explicitly: $\hat{\eta}_{(\alpha-1)n+j} = 0$ if player $\alpha$ responds to question $j$ with a 1, and equals 1 if the player responds with a $-1$. To clearly specify the player and question, we use the notation

$$\hat{\eta}(\alpha, j) := \hat{\eta}_{(\alpha-1)n+j}.$$

From this point on, we will freely switch back and forth between an additive and a multiplicative representation of strategies, leaving the mapping implicit.

To complete the linear algebraic picture, we also define a vector of desired outputs $\hat{s}$ and a game matrix $A$. The game matrix is defined such that given a strategy vector $\hat{\eta}$, the parities of the outputs for each clause are given by the vector $A\hat{\eta}$.

▶ **Definition 6.** Given a $k$-XOR game with $m$ queries and alphabet size $n$, the **game matrix** $A$ is an $m \times kn$ matrix describing query-player-question incidence, and the length-$m$ **parity bit vector** $\hat{s} \in \mathbb{F}_2^m$ the desired outputs:

$$A_{i,(\alpha-1)n+j} := \begin{cases} 1 \text{ if } q_i^{(\alpha)} = j \\ 0 \text{ otherwise} \end{cases} \qquad \text{and} \qquad \hat{s}_i := \begin{cases} 0 \text{ if } s_i = 1 \\ 1 \text{ if } s_i = -1 \end{cases}. \tag{3}$$

An XOR game $G$ is completely specified by providing the game matrix $A$ and parity bit vector $\hat{s}$, i.e. $G \sim (A, \hat{s})$. For example, we translate the GHZ queries into $A_{GHZ}$ and parity bits into $\hat{s}_{GHZ}$ by:

$$\implies A_{GHZ} := \left(\begin{array}{cc|cc|cc} (A) & & (B) & & (C) & \\ \hline 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 \end{array}\right) \quad \text{and} \quad \hat{s}_{GHZ} := \begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \tag{4}$$

A linear-algebraic constraint for achieving classical value 1 can then be defined by asking that a strategy exists that outputs the desired $\hat{s}$.

▶ **Definition 7.** The **classical constraint equation** for strategy $\hat{\eta}$ on game $G \sim (A, \hat{s})$ is

$$A\hat{\eta} = \hat{s} \qquad (\text{over } \mathbb{F}_2). \tag{5}$$

The solutions to (5) are exactly the classical strategies achieving value 1 on game $G \sim (A, \hat{s})$. In other words, a game $G$ has classical value 1 iff (5) has a solution. Gaussian elimination can be used to check for a solution to (5), so we can decide whether a game has $\omega = 1$ in P. Even so, transforming to the dual picture provides a useful analog to the quantum case.

▶ **Definition 8.** Define a **classical refutation** $y \in \mathbb{F}_2^m$ as any vector satisfying the equation dual to (5),

$$\begin{bmatrix} A^T \\ \hat{s}^T \end{bmatrix} y = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \qquad (\text{over } \mathbb{F}_2). \tag{6}$$

▶ **Theorem 9.** *Either a classical refutation $y$ satisfying* (6) *or a classical strategy $\hat{\eta}$ satisfying* (5) *must exist.*

**Proof.** Immediate from the observation that equations (5) and (6) are dual. ◀

A refutation $y$ has a direct interpretation as a certificate that $\omega < 1$: collecting constraints from clauses $i$ corresponding to non-zero entries $y_i$ produces a contradiction to the value-1 hypothesis. To understand this, note that satisfying clause $i$ requires that the $i$th row of (5) is true under strategy $\hat{\eta}$,

$$\begin{aligned} [A\hat{\eta}]_i &= \hat{s}_i \\ \Leftrightarrow \sum_\alpha \hat{\eta}(\alpha, q_i^{(\alpha)}) &= \hat{s}_i \pmod 2. \end{aligned}$$

Summing the value-1 constraints selected by $y$ (left-multiplication by $y$) produces the desired contradiction when $y$ satisfies (6),

$$\begin{aligned} \sum_{i:y_i=1} \sum_\alpha \hat{\eta}(\alpha, q_i^{(\alpha)}) &= \sum_{i:y_i=1} \hat{s}_i \\ \Leftrightarrow y^T A\hat{\eta} &= y^T \hat{s} \pmod 2 \\ \Longrightarrow 0 &= 1 \pmod 2. \end{aligned}$$

This interpretation is key to generalizing refutations to the commuting-operator value of XOR games.

## 3.2 Commuting-Operator Strategies and Refutations

Whereas classical strategies are specified by assigning deterministic output to every player-question pair, commuting-operator strategies are specified by assigning a $\pm 1$ valued quantum measurement to every player-question pair and fixing some entangled state shared by the players. Each player then executes a commuting-operator strategy by selecting the measurement corresponding to the question they receive, then returning the result of applying it to the shared state.

Using the Naimark dilation theorem, we can restrict the measurements in the players' strategies to be Projection-Valued Measures (PVMs). This is the quantum mechanical analogue of the statement that the optimal classical strategy can be taken to be deterministic. In the case of XOR games, this means the measurements can be chosen to be a pair of projectors $\{O_1, O_{-1}\}$ that partition the space into two subspaces, corresponding to outputs 1 and $-1$. We make this restriction for the remainder of the paper.

All that remains is to enforce the no communication requirement on the quantum players. This is done in one of two possible ways. In tensor product strategies the Hilbert space in which $|\Psi\rangle$ lives is taken to be separable, with different players' measurements acting on disjoint parts of the state. In commuting-operator strategies no restriction is placed on the Hilbert space, but the Hermitian matrices corresponding to different players' measurements are forced to commute. (These two restrictions are distinct only in the case of an infinite dimensional Hilbert space). In this paper we work exclusively with the commuting-operator definition, though all the explicit strategies we construct are also valid tensor product strategies.

Putting this all together, we can define for any strategy the observables corresponding to the players' measurements.

▶ **Definition 10.** Given a commuting-operator strategy consisting of measurements $\{O_1^\alpha(j), O_{-1}^\alpha(j)\}$ for each possible question $j$ and player $\alpha$, define the Hermitian **strategy observable** for player $\alpha$ upon receiving question $j$

$$O^\alpha(j) := O_1^\alpha(j) - O_{-1}^\alpha(j). \tag{7}$$

Operators $O^\alpha(j)$ can equivalently be chosen without reference to particular PVMs by taking any set of Hermitian operators that satisfy the constraints (for all players $\alpha \neq \beta$ and questions $j, j'$)

$$[O^\alpha(j), O^\beta(j')] = 0 \qquad \text{(operators held by distinct players commute)} \tag{8a}$$

$$(O^\alpha(j))^2 = I \qquad \text{(square identity, enforcing } \pm 1 \text{ eigenvalues).} \tag{8b}$$

This abstract definition of strategy observables will be the one most frequently referenced in the remainder of this paper.

Given Hermitian observables, the condition for commuting-operator strategies to achieve value 1 is an eigenvalue condition, generalizing (5).

▶ **Definition 11.** For a $k$-XOR game $G$, define the **commuting-operator constraint equations**:

$$\forall\, i \in [m]: \qquad Q_i \left|\Psi\right\rangle = s_i \left|\Psi\right\rangle \tag{9}$$

where the **query observable** $Q_i := \prod_\alpha O^\alpha(q_i^{(\alpha)})$ is the product of all players' observables for the $i$th query.

A strategy achieving value $\omega^* = 1$ must be played on a state $|\Psi\rangle$ which is an eigenvector of every query observable, with appropriate eigenvalue, to ensure zero probability of outputting an incorrect response to some query. This eigenvalue criterion guarantees that the players win all queries. A game $G$ therefore has commuting operator value $\omega^* = 1$ iff there exists some state and strategy observables that satisfy (8) and (9).

While there is an efficient algorithm to solve the classical constraint equations, no such algorithm is known to exist for the commuting-operator constraint equations. Indeed, there is no known upper bound on the dimension of the Hilbert space required to optimally play an entangled game, meaning the search space of the commuting-operator equations is not finite, and the equations themselves may be undecidable. To work around this we develop refutations to characterize the commuting-operator value of XOR games. This technique gives a search space which is still infinite, but is at least discrete, allowing for some progress to be made via combinatorial analysis.

We would like to construct a dual to the commuting-operator constraint equations, meaning a certificate for the unsatisfiability of (9). As there is no immediate analogue to the linear algebraic methods used in the classical case, we instead generalize the view of a refutation as a collection of clauses producing a contradiction (similar to (6)). At a high level, refutations are obtained by multiplying together constraints of the form (9) and applying the known operator identities of (8a) and (8b) to arrive at an equation of the form $I\left|\Psi\right\rangle = -\left|\Psi\right\rangle$ which cannot be true for a normalized quantum state.[8]

[8] Importantly, the order in which the constraint equations are multiplied matters, as two distinct commuting-operator strategy observables with the same player label may not commute. Further, the same constraint equation may need to be incorporated multiple times before one can arrive at a contradiction.

To formally define a refutation, we use an equivalence relation between possible strings of strategy observables on the LHS of an equation of the form (9).

▶ **Definition 12.** Let $Z_1$ and $Z_2$ be two operators formed from products of strategy observables. We say $Z_1$ is equivalent to $Z_2$, written $Z_1 \sim Z_2$, if $Z_1 = Z_2$ is an identity for all strategy observables satisfying (8).

Definitions 11 and 12 then allow us to precisely define a (quantum) refutation, analogous to Definition 8. From now on, a "refutation" will be a quantum refutation unless otherwise specified.

▶ **Definition 13.** Let $G$ be some $k$-XOR game with $m$ clauses. A **refutation for $G$** is defined to be a sequence of clause indices $(i_1, i_2, \ldots, i_\ell) \in [m]^\ell$ satisfying

$$Q_{i_1} Q_{i_2} \ldots Q_{i_\ell} \sim I \qquad \text{and} \qquad s_{i_1} s_{i_2} \ldots s_{i_\ell} = -1. \tag{10}$$

Assuming the value-1 hypothesis and combining the $\ell$ constraint equations satisfying (10) then gives the desired contradiction, $I|\Psi\rangle = -|\Psi\rangle$. Whether a product of queries $Q_{i_1} Q_{i_2} \ldots Q_{i_\ell}$ is equivalent to $I$ can be efficiently checked by collecting each player's operators using (8a) and repeatedly applying (8b) to greedily cancel operators.

Refutations certify that $\omega^* < 1$ analogously to the way that classical refutations certify that $\omega < 1$. We prove in the full paper that the converse is also true, completing the proof of Theorem 5. The proof of this fact relies on a connection between refutations and the ncSoS hierarchy analogous to a connection made by Grigoriev [17] between classical refutations and the SoS hierarchy. In particular, we show the ncSoS algorithm takes time exponential in the minimum length refutation to prove a game has value $\omega^* < 1$. Theorem 5 then follows from completeness of ncSoS.

It is not obvious that one can find refutations more easily than one can find strategies. The remaining results focus primarily on subclasses of XOR games for which we can apply the refutations picture to exactly characterize the games with $\omega^* = 1$. In particular, we identify an easily-computed stronger refutation condition that is complete for symmetric games, i.e. those naturally corresponding to CSPs. Subsequently we analyze specific families of games that give bounds on the behavior of ncSoS and insight into the structure of the XOR game landscape.

## 4 Symmetric Games

The refutation technology developed above gives surprisingly powerful results when applied to symmetric games. In particular, we show:

▶ **Theorem 14.** *Membership in the set of symmetric games with $\omega^* = 1$ can be efficiently decided via a system of linear Diophantine equations.*

Previously the question of whether symmetric games took value 1 was not known to be decidable. Theorem 14 affirms that it is decidable and in fact in P. We prove the theorem by introducing a simple necessary condition for refutations to exist, then showing the structure of symmetric games ensures this condition is also sufficient for a refutation.

Games that do not satisfy this necessary condition have value $\omega^* = 1$. By returning to duality arguments, we further show that they can be played optimally by a simple family of strategies.

▶ **Theorem 15.** *Any value-1 symmetric game can be played optimally by a single qubit strategy using a GHZ resource state. Furthermore, this strategy can be found in polynomial time.*

## 4.1 A Necessary Condition for Refutation (PREF)

Any valid refutation involves a product of query observables that cancel to $I$ via the square identity (8b). Focusing on the strategy observables corresponding to one player, we see every operator at an even depth in the sequence must cancel with one at an odd depth. Given any sequence of query observables we can count the number of copies of $O^\alpha(j)$ at odd and even depths – if the sequence corresponds to a refutation the counts must be equal. Then for any given game, it is necessary to be able to construct some sequence of queries $Q_{i_1} \dots Q_{i_\ell}$ satisfying this counting equality (with appropriate parity $s_{i_1} \dots s_{i_\ell} = -1$) in order to construct a true refutation. We call such a sequence a **parity-permuted refutation (PREF)**.

To prove properties of such PREFs, we find it useful to introduce a freer equivalence relation $\overset{p}{\sim}$ on strings of queries that allows reordering within the even positions and the odd positions before cancellation (compare to Definition 12). A PREF is then a string of clauses $(i_1, \dots, i_\ell)$ satisfying $Q_{i_1} \dots Q_{i_\ell} \overset{p}{\sim} I$ and the same parity-bit requirement as a regular refutation. In later sections, we carefully define and use a technical version of $\overset{p}{\sim}$, but here simply state that $\overset{p}{\sim}$ is a more inclusive equivalence relation than $\sim$, which immediately gives Lemma 16.

▶ **Lemma 16** (Necessary condition for refutation). *If a game $G$ admits a refutation, it contains a PREF.*

We define **noPREF games** to be those games that do not admit a PREF. These games admit no true refutations by the previous lemma, and so have $\omega^* = 1$. Whether or not a game contains a PREF is efficiently decidable by checking for a solution to a linear system of equations. We sketch the algorithm that decides membership, delegating a rigorous proof to the full paper.

▶ **Lemma 17** (Informal). *Membership in the set of noPREF games can be efficiently decided by a system of linear Diophantine equations.*

**Proof (sketch).** We prove the result by showing that a game $G \sim (A, \hat{s})$ contains a PREF if and only if there is a solution to the set of equations

$$A^T z = 0 \tag{11}$$

$$\hat{s}^T z = 1 \pmod 2 \tag{12}$$

for some $z \in \mathbb{Z}^m$. To prove the forward direction we note that each row of (11) guarantees that a particular player-question pair has equal positive and negative count and (12) ensures the parity bit requirement is met, such that the game $G$ contains a PREF built by interleaving the multisets of clause indices

$$\mathcal{O} = \{i \text{ with multiplicity } |z_i| \; \forall \; i : z_i > 0\} \tag{13a}$$

$$\mathcal{E} = \{i \text{ with multiplicity } |z_i| \; \forall \; i : z_i < 0\} \tag{13b}$$

with their elements placed in odd and even positions, respectively. The reverse direction is proved similarly, but requires a technical lemma relating the even and odd clauses of a PREF. Then standard techniques for solving linear Diophantine equations complete the proof. ◀

Symmetric games have additional structure that allows us to prove a stronger statement.

▶ **Lemma 18** (Informal). *The noPREF characterization is complete for symmetric games. That is, every value 1 symmetric game is in the noPREF set.*

**Proof.** The proof of Lemma 18 is both constructive and purely combinatorial. It involves first showing that symmetric games have enough structure to construct "shuffle gadgets" which let us approximately commute strategy observables past each other. Then, careful application of these shuffle gadgets lets us transform a PREF into a true refutation. The technical proof is presented in the full paper. ◀

Theorem 14 then follows directly from Lemmas 17 and 18 since symmetric games have commuting-operator value 1 if and only if they admit no PREFs and this condition is efficiently checkable.

## 4.2 A Single Qubit Strategy (MERP)

We arrived at the noPREF classification of XOR games by generalizing the classical dual picture to the entangled games case, then finding a simple necessary condition for a dual proof (refutation) to exist. In this section we take the dual of a dual, and give a simple technique to construct a strategy that achieves value 1 for any noPREF game, thus sketching the proof of Theorem 15.

We call the resulting family of strategies **Maximal Entanglement, Relative Phase (MERP)** strategies. These strategies are exactly the class of strategies developed in [31] to optimally solve all games with alphabet size two. They are played on a $k$ qubit $GHZ$ state (one qubit per player), and achieve value one if and only if (Theorem 21) there is a solution to the equation

$$A\hat{\theta} = \hat{s} \pmod 2, \qquad \hat{\theta} \in \mathbb{Q}^{kn}.$$

▶ **Definition 19** (MERP)**.** Given a $k$-XOR game $G$, a **MERP strategy** for $G$ is a tensor-product strategy in which:

1. The $k$ players share the maximally entangled state
$$|\Psi\rangle = \frac{1}{\sqrt{2}}\left[|0\rangle^{\otimes k} + |1\rangle^{\otimes k}\right] \tag{14}$$
with player $\alpha$ having access to the $\alpha$-th qubit of the state.
2. Upon receiving question $j$ from the verifier, player $\alpha$ rotates his qubit by an angle $\theta(\alpha, j)$ about the $Z$ axis, then measures his qubit in the $X$ basis and sends his observed outcome to the verifier.
Equivalently, we define the states
$$|\theta(\alpha,j)_\pm\rangle := \frac{1}{\sqrt{2}}\left[|0\rangle \pm e^{i\theta(\alpha,j)}\,|1\rangle\right] \tag{15}$$
and pick strategy observables
$$O^\alpha(j) := |\theta(\alpha,j)_+\rangle\langle\theta(\alpha,j)_+| - |\theta(\alpha,j)_-\rangle\langle\theta(\alpha,j)_-|\,. \tag{16}$$

Each $Z$ rotation executed by the players introduces a relative phase between the $|0\rangle^{\otimes k}$ and $|1\rangle^{\otimes k}$ components of the GHZ state, and these relative phases add. The measurement is constructed such that a total relative phase that is an even multiple of $\pi$ results in overall output 1 while an odd multiple of $\pi$ results in overall output $-1$. Collecting the rotation angles into a strategy vector

$$\hat{\theta}_{(\alpha-1)n+j} := \frac{1}{\pi}\theta(\alpha,j) \tag{17}$$

results in a useful parallel between MERP strategies and classical strategies:

▶ **Definition 20.** Define the **MERP constraint equations** for game $G$ by

$$A\hat{\theta} = \hat{s} \pmod 2 \tag{18}$$

▶ **Theorem 21.** *The value obtained by a MERP strategy is given by*

$$v^{MERP}(G, \hat{\theta}) := \frac{1}{2} + \frac{1}{2m}\left(\sum_{i=1}^{m} \cos\Big(\pi\left[(A\hat{\theta})_i - \hat{s}_i\right]\Big)\right). \tag{19}$$

*Consequently, there exists a MERP strategy achieving $v^{MERP} = 1$ on a game $G$ iff its MERP constraint equations have a solution over $\mathbb{Q}$.*

The proof of Theorem 21 is computational. For any game, Theorem 21 indicates that a MERP strategy achieves value 1 if and only if the MERP constraint equation (18) is satisfied. Similarly to the classical case, this can be efficiently checked by Gaussian elimination. Here, however, the underlying field is $\mathbb{Q}$ as opposed to $\mathbb{F}_2$.

Both MERP strategies and PREF specifications are defined by linear systems of equations, over $\mathbb{Q}$ and $\mathbb{Z}$ respectively. Remarkably, these systems of equations are dual to each other, in much the same way as classical strategies and refutations. By showing this, we prove Theorem 22.

▶ **Theorem 22.** *Any game $G$ admits either a PREF or a MERP strategy with value 1.*

**Proof.** Technical proof in the style of a Theorem of Alternatives presented in the full paper. ◀

Intuitively, this result (coupled with Lemma 18 and Theorem 21) indicates that the power of quantum solutions to noPREF games is equivalent to promoting the underlying field from $\mathbb{F}_2$ to $\mathbb{Q}$. We expect further advantage to be gained from the non-commuting nature of operator-valued solutions in cases where a game admits a PREF but no true refutation. These classes of games are the main subject of future work.

Figure 1 (given in the Introduction) summarizes the new duality relations presented in this paper. We repeat them here. The general quantum duality provides a complex but complete description of games with $\omega^* = 1$. The PREF conditions are efficient to compute, but are only *necessary* conditions for constructing commuting-operator refutations, and thus the dual, MERP value 1, holds true for only a subset of all $\omega^* = 1$ games. We can make a stronger statement about symmetric games: PREFs are both necessary and sufficient for a symmetric game to have a refutation, so the duality ensures MERP achieves value 1 for all symmetric games with $\omega^* = 1$.

## 5 ncSoS Bounds and the XOR Landscape

The refutations picture also allows us to give worst and average case bounds on the behavior of ncSoS for XOR games, and construct new families of games with interesting properties.

First, we construct a family of games that have $\omega^* < 1$, but are built in such a way that the ncSoS algorithm has a hard time recognizing this. Called Capped GHZ games, games in this family are symmetric and contain PREFs that all are at least exponentially long. The ncSoS algorithm[9] then requires time doubly exponential to prove that these games

[9] Our results imply that in the case of entangled XOR games the ncSoS has runtime within a polynomial factor of a brute force search over refutations up to a certain length.

have commuting-operator value $< 1$. This gives a rare example of a problem whose solution requires a superlinear number of levels of ncSoS, and illustrates the distinction between ncSoS and SoS, the latter always terminating after at most a linear number of levels. On the other hand, Lemma 18 tells us the existence of a PREF for this symmetric game indicates $\omega^* < 1$, which allows us to solve the same problem in polynomial time. This leads us to prove:

▶ **Theorem 23.** *There exists a family of 3-XOR games with $\omega^* < 1$ but for which the minimum refutation length scales exponentially in the number of clauses $m$ and alphabet size $n$. For these games exponentially many levels of ncSoS are needed to witness that $\omega^* < 1$.*

**Proof.** An explicit construction of Capped GHZ games presented in the full paper. These are symmetric games designed such that the PREF condition for these games can only be satisfied by sets containing exponentially many clauses. Since any refutation also gives a PREF, this means all refutations have at least exponential length. Finally, as a symmetric game, the existence of a PREF indicates the game has $\omega* < 1$. By the connection between ncSoS and refutations, this means the ncSoS algorithm requires at least exponentially many levels to detect the game has value less than one. ◀

Returning to pseudo-telepathy, we construct a family of games that generalize the GHZ game, termed the Asymptotically Perfect Difference (APD) family. Members are parameterized by scale $K$, with the $K$-th member having $k = 2^K - 1$ players. The APD family is designed such that any desired parity bits $s_i$ can be produced by some strategy (the game is in the noPREF set regardless of the $s_i$, so $\omega^* = 1$ for all $s_i$). On the other hand, a growing fraction of possible assignments of $s_i$ correspond to low classical value, and the family has perfect difference [4] in the asymptotic limit, $\lim_{K\to\infty} 2(\omega^* - \omega) = 1$. In comparison, there are randomized constructions of families of games whose bias ratio $\frac{\omega^*-1/2}{\omega-1/2}$ diverges for fixed $k \geq 3$ as $n \to \infty$ [25, 4], but these constructions give no guarantee on the difference. Specifically, we prove it is possible to force an upper bound on $\omega$ in terms of the number of players $k$ while preserving $\omega^* = 1$:

▶ **Theorem 24.** *There exists a family of $k$-XOR games, parametrized by $K$, for which $\omega^*(G(K)) = 1$ and the classical value is bounded by*

$$\frac{1}{2} \leq \omega(G(K)) \leq \frac{1}{2} + \sqrt{\frac{K+1}{2^{K+1}}} \leq \frac{1}{2} + \sqrt{\frac{\log k}{k}}. \tag{20}$$

Quicker asymptotic convergence to difference $2(\omega^* - \omega) \to 1$ could be achieved in other ways, for example by the generalized Mermin-GHZ game [2] or by XORing together the answers (aka "XOR repetition") of other pseudo-telepathy XOR games [5]. Although their bias scales in a weaker way, the APD games have the property that perfect entangled strategies exist for *any* choice of target bit strings $s$.

To investigate the incompleteness of the PREF condition, we define an XOR game that contains a PREF, but provably has commuting-operator value 1. This game is solved by a single-qubit strategy employing measurements in the $X$, $Y$, and $Z$ bases. This may be a starting point for stricter necessary criterion, building towards a complete algorithm for deciding the value of entangled XOR games. The existence of this game proves the following theorem.

▶ **Theorem 25.** *The PREF characterization is incomplete. In particular, there exists an XOR game with six players, alphabet size three, for which the entangled value is 1, but the noPREF condition is unable to detect this.*

Finally, we investigate thresholds in $\omega^*$ by considering the behavior of randomly generated XOR games with a large number of clauses. We prove Theorem 26, which shows that (much like the classical case) random XOR games become unsatisfiable with high probability when $m$ is larger than some constant times $n$. Previous techniques could show that $k$-XOR instances were unsatisfiable only in the "dense" regime, i.e. where $m \geq \Omega(n^k)$ [24].

▶ **Theorem 26.** *For every $k$, there exists a constant $C_k^{unsat}$ depending only on $k$ such that a random $k$-XOR game $G$ with $m \geq C_k^{unsat} n$ clauses has value $\omega^*(G) < 1$ with probability $1 - o(1)$.*

**Proof.** Explicit construction of a refutation presented in the full paper. ◀

We also investigate the average case performance of ncSoS. We show that random games with a fixed ratio of $m$ to $n$ have a minimal length refutation that scales like $\Omega(n \log(n)/\log(\log(n)))$, implying that it takes the ncSoS algorithm superexponential time to show that these games have $\omega^* < 1$ (Theorem 27). These results should be thought of as quantum analogues of Grigoriev's [17] integrality gap instances for classical XOR games.

▶ **Theorem 27.** *For any constant $C$, the minimum length refutation of a random 3-XOR game with $m = Cn$ queries on an alphabet of size $n$ has length at least*

$$\frac{en\log(n)}{8C^2\log(\log(n))} - o\left(\frac{n\log(n)}{\log(\log(n))}\right) \tag{21}$$

*with probability $1 - o(1)$ (as $n \to \infty$). Hence, either $\omega^* = 1$ or $\Omega(n\log(n)/\log(\log(n)))$ levels of the ncSoS hierarchy are needed to witness that $\omega^* < 1$ for such games.*

## 6 Future Work

We see four main directions in which our characterization of non-local XOR games could be extended.

First, our linear algebraic characterization of $\omega^* = 1$ games is incomplete: there exist games with $\omega^* = 1$ for which a MERP strategy cannot achieve value 1. We expect a strengthening of the PREF condition may allow us to extend our decidability algorithm to detect these games and develop dual strategies that solve them. Understanding the structure of such games would give further intuition about the behavior of optimal XOR commuting-operator strategies, in particular strategies which may require more entanglement than the simple MERP strategies.

Second, determining whether $\omega^* = 1$ for nonsymmetric XOR games may be outside P or even undecidable. In the realm of Binary Constraint System (BCS) games, [27] shows that determining whether a general BCS game has perfect value is undecidable. The structural similarity between BCS games and XOR games suggests that perhaps some of the group theoretic techniques of that work could be applied to XOR games to arrive at a similar conclusion. An interesting class of games which may serve as an intermediate class between XOR and BCS games are "incomplete" XOR games in which there are $k$ players but queries can involve $< k$ variables, effectively ignoring some players. Even for $k = 2$, Tsirelson's semidefinite programming characterization of $\omega^*$ does not apply to incomplete XOR games, although in this case it is still easy to decide whether $\omega^* = 1$.

Third, while in this work we have focused on computing the entangled game value $\omega^*$, our methods may also be useful from the perspective of Bell inequalities, in which the quantity of interest is the maximal violation achievable by an entangled strategy. While this has

conventionally been measured in terms of the bias ratio $(\omega^* - 1/2)/(\omega - 1/2)$, the difference $2(\omega^* - \omega)$ is an equally natural measure, and we hope that our techniques will render it more amenable to analysis. Indeed, in addition to the construction of Asymptotically Perfect Difference games mentioned above, our results have the following simple consequence: for symmetric games with $\omega^* = 1$, our characterization of the optimal strategies (MERP) together with the Grothendieck-type inequality of [3] imply that the bias ratio and difference are both bounded by constants depending only on $k$, and that for the difference, this constant is strictly less than one.

Finally, our results are almost entirely concerned with the question of determining whether $\omega^* = 1$ or $\omega^* < 1$. However, we note that the MERP family of strategies includes the optimal strategy for the CHSH game [7] and more generally any multiplayer game with question and answer alphabet size two [31], but not for all XOR games [23]. It is an interesting open question to fully characterize when MERP strategies are optimal. In this setting there are still many classical tools which we do not know how to extend to the quantum case. As an example, consider *overconstrained* games in which there are many more constraints than variables and the signs of those constraints are chosen randomly. In the classical case, second moment methods can show that the value is close to 1/2 while in the quantum case we can only conclude that it is $< 1$.

## References


1 Jonathan Barrett, Lucien Hardy, and Adrian Kent. No signaling and quantum key distribution. *Physical review letters*, 95(1):010503, 2005. arXiv:quant-ph/0405101.

2 Gilles Brassard, Anne Broadbent, and Alain Tapp. Recasting mermin's multi-player game into the framework of pseudo-telepathy. *ArXiv e-prints*, 2004. arXiv:quant-ph/0408052.

3 Jop Briët, Harry Buhrman, Troy Lee, and Thomas Vidick. Multipartite Entanglement in XOR Games. *Quantum Info. Comput.*, 13(3-4):334–360, March 2013. arXiv:0911.4007.

4 Jop Briët and Thomas Vidick. Explicit lower and upper bounds on the entangled value of multiplayer XOR games. *Communications in Mathematical Physics*, 321(1):181–207, 2013. arXiv:1108.5647.

5 Anne Lise Broadbent. Quantum pseudo-telepathy games. Master's thesis, Université de Montréal, 2004.

6 Boris S Cirel'son. Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics*, 4(2):93–100, 1980.

7 John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed Experiment to Test Local Hidden-Variable Theories. *Phys. Rev. Lett.*, 23:880–884, October 1969. doi:10.1103/PhysRevLett.23.880.

8 Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and Limits of Nonlocal Strategies. In *CCC '04*, pages 236–249, 2004. arXiv:quant-ph/0404076.

9 Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *International Colloquium on Automata, Languages, and Programming*, pages 320–331. Springer, 2014.

10 Roger Colbeck. *Quantum And Relativistic Protocols For Secure Multi-Party Computation*. PhD thesis, University of Cambridge, 2006. arXiv:0911.3814.

11 Andrew C. Doherty, Yeong-Cherng Liang, Ben Toner, and Stephanie Wehner. The Quantum Moment Problem and Bounds on Entangled Multi-prover Games. In *CCC '08*, pages 199–210, 2008. arXiv:0803.4373.

12 Olivier Dubois and Jacques Mandler. The 3-XORSAT threshold. *Comptes Rendus Mathématique*, 335(11):963–966, 2002.

**13** Artur K Ekert. Quantum cryptography based on Bell's theorem. *Physical review letters*, 67(6):661, 1991.

**14** Joseph Fitzsimons, Zhengfeng Ji, Thomas Vidick, and Henry Yuen. Quantum proof systems for iterated exponential time, and beyond. *arXiv*, 2018. arXiv:1805.12166.

**15** Tobias Fritz, Tim Netzer, and Andreas Thom. Can you compute the operator norm? *Proceedings of the American Mathematical Society*, 142(12):4265–4276, 2014. arXiv:1207.0975.

**16** Daniel M Greenberger, Michael A Horne, Abner Shimony, and Anton Zeilinger. Bell's theorem without inequalities. *American Journal of Physics*, 58(12):1131–1143, 1990.

**17** Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259:613–622, 2001. doi:10.1016/S0304-3975(00)00157-2.

**18** Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

**19** Zhengfeng Ji. Classical verification of quantum proofs. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 885–898. ACM, 2016.

**20** Zhengfeng Ji. Compression of Quantum Multi-Prover Interactive Proofs. *arXiv*, 2016. arXiv:1610.03133.

**21** Anand Natarajan and Thomas Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP for QMA. *arXiv*, 2018. arXiv:1801.03821v2.

**22** Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.*, 10(7):073013, 2008. arXiv:0803.4290.

**23** Dimiter Ostrev and Thomas Vidick. Entanglement of approximate quantum strategies in XOR games, 2016. arXiv:1609.01652.

**24** C. Palazuelos and T. Vidick. Survey on nonlocal games and operator space theory. *Journal of Mathematical Physics*, 57(1):015220, January 2016. doi:10.1063/1.4938052.

**25** David Pérez-García, Michael M Wolf, Carlos Palazuelos, Ignacio Villanueva, and Marius Junge. Unbounded violation of tripartite Bell inequalities. *Communications in Mathematical Physics*, 279(2):455–486, 2008. arXiv:quant-ph/0702189.

**26** Ben W Reichardt, Falk Unger, and Umesh Vazirani. Classical command of quantum systems. *Nature*, 496(7446):456–460, 2013.

**27** William Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games, 2016. arXiv:1606.03140.

**28** Boris S Tsirel'son. Quantum analogues of the Bell inequalities. The case of two spatially separated domains. *Journal of Mathematical Sciences*, 36(4):557–570, 1987.

**29** Umesh Vazirani and Thomas Vidick. Fully device-independent quantum key distribution. *Physical review letters*, 113(14):140501, 2014.

**30** Thomas Vidick. Three-Player Entangled XOR Games Are NP-Hard to Approximate. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, FOCS '13, pages 766–775. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.87.

**31** Reinhard F Werner and Michael M Wolf. All-multipartite Bell-correlation inequalities for two dichotomic observables per site. *Physical Review A*, 64(3):032112, 2001.