

The Unique Games Conjecture with Entangled Provers is False

Julia Kempe*

Oded Regev[†]

Ben Toner[‡]

October 18, 2007

Abstract

We consider one-round games between a classical verifier and two provers who share entanglement. We show that when the constraints enforced by the verifier are ‘unique’ constraints (i.e., permutations), the value of the game can be well approximated by a semidefinite program. Essentially the only algorithm known previously was for the special case of binary answers, as follows from the work of Tsirelson in 1980. Among other things, our result implies that the variant of the unique games conjecture where we allow the provers to share entanglement is false. Our proof is based on a novel ‘quantum rounding technique’, showing how to take a solution to an SDP and transform it to a strategy for entangled provers.

1 Introduction

Games: For nearly two decades, two-prover one-round games have played a major role in many of the most important developments in theoretical computer science. Such games consist of a verifier and two provers who are unable to communicate with each other. The game starts when the verifier sends two questions, one to each prover, chosen according to some joint distribution. Each prover then replies with an answer chosen from the alphabet $\{1, \dots, k\}$ for some $k \geq 1$. Finally, the verifier decides whether to accept or reject, based on the answers he received. The *value* of such a game is defined as the maximum success probability that the provers can achieve. For example, let us consider the following very simple game known as the CHSH game [CHSH69]: the verifier sends a random bit to each of the provers, who then reply with one bit each (so $k = 2$). The verifier accepts if and only if the XOR of the answers is equal to the AND of his questions. A moment’s reflection shows that the value of this game is $\frac{3}{4}$, and is obtained, say, when the provers always return 0.

*School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848, by an Alon Fellowship of the Israeli Higher Council of Academic Research and by an Individual Research Grant of the Israeli Science Foundation.

[†]School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by the Binational Science Foundation, by the Israel Science Foundation, and by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848.

[‡]Centrum voor Wiskunde en Informatica, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands. Supported by the Dutch BSIK/BRICKS project, by EU project QAP, and by NWO VICI project 639-023-302.

One of the most important breakthroughs in theoretical computer science was the discovery of the PCP theorem in the early 90s [AS98, ALM⁺98]. Combined with Raz’s parallel repetition theorem [Raz98], it implies the following.

Theorem 1.1 ([AS98, ALM⁺98, Raz98]) *For any $\delta > 0$ there exists a $k = k(\delta)$ such that it is NP-hard to determine whether, given a (two-prover one-round) game with answers from a domain of size k , its value is 1 or at most δ .*

This result has led to many important advances in the field, including in particular many tight NP-hardness results. For instance, Håstad [Hås01] showed that it is NP-hard to tell whether a given 3SAT formula is satisfiable, or not more than a $\frac{7}{8} + \epsilon$ fraction of its constraints can be satisfied. This shows that the algorithm that simply assigns random values to the variables is essentially optimal. Other tight NP-hardness results that follow from the PCP theorem include a hardness factor of $\frac{1}{2} + \epsilon$ for E3LIN2 [Hås01], a hardness factor of $n^{1-\epsilon}$ for MAXCLIQUE [Hås99], and a hardness factor of $\ln n$ for SETCOVER [Fei98].

One important special case of games is that of *unique games*. Here, the verifier’s decision is restricted to be of a very specific form. Namely, for any questions s, t sent to the provers, the verifier accepts answers a, b if and only if $b = \sigma_{st}(a)$ where σ_{st} is some permutation on $\{1, \dots, k\}$. In 2002, Khot [Kho02] presented a conjecture known as the unique games conjecture (UGC) that essentially says that it is hard to approximate the value of a unique game, even if we are only interested in distinguishing the almost satisfiable case from the almost completely unsatisfiable case.

Conjecture 1.2 (Unique games conjecture [Kho02]) *For any $\epsilon, \delta > 0$ there exists a $k = k(\epsilon, \delta)$ such that it is NP-hard to determine whether, given a unique game with answers from a domain of size k , its value is at least $1 - \epsilon$ or at most δ .*

It is not hard to see that determining whether the value of a unique game is 1 (i.e., perfectly satisfiable) can be done efficiently using a simple algorithm, and therefore it is crucial that we insist here on $\epsilon > 0$ (cf. Theorem 1.1). Let us also mention that there *exist* $\epsilon, \delta > 0$ for which the problem in the conjecture is known to be NP-hard (even with $k = 2$). This follows from Håstad’s hardness result for MAXCUT [Hås01]. Despite a considerable amount of work in the last few years, the plausibility of the conjecture is still uncertain, and this issue is currently one of the central topics in theoretical computer science.

The tremendous importance of the unique games conjecture stems from the fact that for many fundamental problems, it implies strong, and often tight, inapproximability results that are not known to hold under more conventional assumptions. As an example, let us consider the MAXCUT problem. The best known algorithm for this problem was given in 1994 by Goemans and Williamson, and achieves an approximation factor of ≈ 0.878 [GW95]. It consists of two main steps: first, one writes a semidefinite programming (SDP) relaxation of the given MAXCUT instance, where by ‘relaxation’ we mean that by construction, the value of the SDP is guaranteed to be not smaller than the size of the maximum cut. This SDP can then be solved efficiently using known techniques for convex optimization, such as the ellipsoid algorithm (see, e.g., [BV04]). The second part of their algorithm is a ‘rounding procedure’ in which the solution to the semidefinite program is converted

into a solution to the MAXCUT problem. The name ‘rounding’ comes from the fact that this step can be seen as a way to round the ‘continuous’ SDP solution into a ‘discrete’ solution to MAXCUT.

Despite intensive research, no better algorithm for MAXCUT has been found until this day. The best known NP-hardness result, due to Håstad, shows that obtaining approximation ratio above ≈ 0.941 is NP-hard [Hås01]. The hardness for approximation factors between ≈ 0.878 and ≈ 0.941 was unclear for many years. Recently, it was shown by Khot et al. [KKMO07] that the UGC implies a tight inapproximability result of ≈ 0.878 , thereby giving a partial answer to this long-standing open question.

Another problem for which the UGC implies a tight hardness result is the Vertex Cover problem, where a simple algorithm gives an approximation factor of 2 and the UGC implies a hardness factor of $2 - \varepsilon$ for any $\varepsilon > 0$ [KR03] (whereas the best known NP-hardness result is 1.36 [DS05]). The UGC also implies strong inapproximability results for graph coloring problems [DMR06] and the Sparsest Cut problem [KV05, CKK⁺06].

In another line of work, attempts have been made to disprove the conjecture by means of efficient approximation algorithms for the value of unique games [Tre05, CMM06a, GT06, CMM06b]. So far, however, none of these results was able to disprove this conjecture, and this by itself might be seen by some as evidence in favor of the conjecture. Among the best algorithms is the one by Charikar et al. [CMM06a] that, given a unique game on alphabet size k whose value is $1 - \varepsilon$, outputs a solution of value $1 - O(\sqrt{\varepsilon \log k})$. This does not disprove the conjecture, but instead gives us a lower bound on $k(\varepsilon, \delta)$ for the conjecture to make sense (see also [KKMO07]). Another recent result is by Chlamtac et al. [CMM06b] who show how to compute, given a unique game on alphabet size k whose value is $1 - \varepsilon$, a solution of value $1 - O(\varepsilon \sqrt{\log n \log k})$. This is better than [CMM06a] for small values of ε , but as before, is not enough to disprove the conjecture and only tells us that k should be large enough as a function of ε and δ . We remark that most of these results are based on an SDP relaxation, followed by a (usually quite sophisticated) rounding procedure.

Games with entangled provers: In this paper we consider the model of two-prover one-round games in which the provers are allowed to share entanglement. (The verifier and all communication remain classical, as before.) Such games are sometimes known in the quantum information literature as *nonlocal games* and have their origins in a seminal 1935 paper by Einstein, Podolsky, and Rosen [EPR35] and a 1964 paper by Bell [Bel64]. We define the *entangled value* of a game as the maximum success probability achievable by provers that share entanglement. For instance, it is known that the entangled value of the CHSH game is $\frac{1}{2} + \frac{1}{2\sqrt{2}} \approx 85\%$, which is strictly greater than the 75% achievable without entanglement. This remarkable ability of entanglement to create correlations that are impossible to obtain classically (something Einstein referred to as “spooky”) is one of the most peculiar aspects of quantum mechanics and required many years to be properly understood.

One motivation for this model comes from the fact that although a verifier can guarantee that provers don’t communicate (by separating them in space, say), he has no way to guarantee that they don’t share entanglement. Therefore, in order for a proof system with two provers to be sound in our quantum physical world, we must consider the scenario where the provers share entanglement, even when the verifier is classical. This is especially true for multi-prover crypto-

graphic protocols, where the presence of entanglement could break the security of the protocol. Another purely mathematical motivation for this model comes from the hope that through studying this model, we can reach a better understanding of the non-classical correlations that arise in quantum mechanics, and even obtain some new insights into multi-prover games similar to those obtained from the PCP theorem.

Despite considerable work on this model, our understanding of it is still quite limited. One of the earliest and most important results in this area is due to Tsirelson [Cir80], who showed that, for the special case of unique games with an alphabet of size $k = 2$, the entangled value is given exactly by the optimum of a certain SDP and can therefore be computed efficiently (see also [CHTW04] where this is made explicit and [CSUU07] for a nice application of this SDP). Unique games with $k = 2$ are also known as XOR-games because one can think of the two possible answers as a bit, and then the only possible unique constraints are $a \oplus b = 0$ and $a \oplus b = 1$. This result is in contrast to the (non-entangled) value of an XOR-game, which is NP-hard to compute exactly or even to approximate (as follows from Håstad’s hardness result for MAXCUT). Finally, we note that the CHSH game is an XOR-game, and the way to derive its entangled value of $\frac{1}{2} + \frac{1}{2\sqrt{2}}$ is by computing Tsirelson’s SDP.

Unfortunately, our current understanding of the entangled value of games does not extend much beyond the case of XOR-games. To the best of our knowledge, the only other general result is by Masanes [Mas05], who shows how to compute the entangled value of games with only two possible questions to each prover and $k = 2$.¹ Although restricted to this very special case, his result still allows us to handle a few cases not handled by Tsirelson’s result (namely, non-unique games for $k = 2$ with two questions).

In all other cases, no method is known to compute or even approximate (with provable guarantees) the entangled value of a game. Even for some very small fixed size games, there is still uncertainty regarding their entangled value (see, e.g., [BM04]). One recent attempt to handle more general games was made by Navascues et al. [NPA07], who outlined a hierarchy of SDP relaxations of the entangled value of a game. Unfortunately, there are no known bounds on the quality of their SDP relaxations.

Our results: Our main result is an approximation algorithm for the entangled value of any unique game. More precisely, our main theorem is the following.

Theorem 1.3 *There exists an efficient algorithm that, given a unique game whose entangled value is $1 - \varepsilon$, outputs a value $\varepsilon/6 \leq \varepsilon' \leq \varepsilon$ and a description of an entangled strategy for the provers whose success probability is at least $1 - 6\varepsilon'$.*

This theorem gives, for the first time, a way to approximate the entangled value of games with more than two possible answers. It is also the first provable approximation (as opposed to exact) algorithm for the entangled value of a game.

Our result shows that the analogue of Conjecture 1.2 for entangled provers is *false*. Indeed, as long as, $6\varepsilon + \delta < 1$, our algorithm can efficiently tell whether the entangled value of a game is at least $1 - \varepsilon$ or at most δ . This can be seen as a (modest) contribution to the understanding of the ever-more-mysterious unique games conjecture.

¹His method also handles the case of more than two provers.

It is interesting to compare our algorithm with the approximation algorithms for the (non-entangled) value of unique games. Given a unique game with entangled value $1 - \varepsilon$, our algorithm outputs a strategy whose entangled value is at least $1 - 6\varepsilon$. In contrast, the algorithms in [Tre05, CMM06a, GT06, CMM06b] given a unique game with value $1 - \varepsilon$, output a strategy whose value is at least $1 - f(\varepsilon, k)$ for some function $f(\varepsilon, k)$. The fact that our approximation is independent of k is crucial.

Techniques: The proof of Theorem 1.3 is based on a semidefinite programming (SDP) relaxation of the entangled value. Our SDP turns out to be equivalent to the one used by Khot in [Kho02] as a relaxation of the (non-entangled) value of a game (and, in fact, this SDP originates in the work of Feige and Lovász [FL92]). We note, however, that in the non-entangled case, certain extra constraints are sometimes used that are not known to hold in the entangled case.

The heart of the proof is in the second step, where we show how to take a solution to the SDP and transform it into a strategy for entangled provers. We call this step the ‘quantum rounding’ step in analogy with the rounding procedure used in the non-entangled case. We hope that this novel technique will be useful for other problems as well. The main idea in our rounding step is to use the vectors given by the SDP solution as a quantum measurement performed by the provers on a maximally entangled state shared by them.

Extensions: We present two extensions of our main theorem. The first involves a special case of unique games which we call *uniform* unique games. These are unique games for which there exists an optimal strategy in which each prover’s answer distribution is uniform (for each question). As we show later, any unique game in which the verifier’s decision is based solely on $a - b \pmod k$ is a uniform unique game. This includes XOR-games as well as the unique games constructed in [KKMO07]. For this special case, we show that the factor 6 in our main theorem can be improved to 4. We also extend our main theorem to *d-to-d* games, which are another type of game considered in [Kho02]. Namely, we show that Khot’s conjecture for *d-to-d* games is false in the case that the provers share entanglement.

Discussion: Our work gives for the first time a way to approximate the entangled value of games with more than two possible answers. One open question this raises is whether there exist better algorithms for approximating (or even computing exactly) the entangled value of a unique game. So far we only know of such a result in the case $k = 2$, where Tsirelson’s SDP gives an exact answer [Cir80]. Extending this to $k > 2$ might require improving our quantum rounding procedure, and might also involve the use of a tighter SDP, perhaps taken from the SDP hierarchy outlined in [NPA07]. Another open question is whether our quantum rounding technique can be used for other types of games. One good candidate are games with inequality constraints, such as MAX-K-CUT, as those are relatively well-understood [KKMO07].

One might also hope to extend our results and deal with the case of general games. In fact, for all that we know, it might be the case that there exists an efficient algorithm that can compute *exactly* the entangled value of any given game. This strange state of affairs is (unfortunately) consistent with everything we know so far. The only indication that such an algorithm is unlikely to exist comes from recent work by Kempe et al. [KKM⁺] who show that computing the entangled

value of a game with *three* provers is NP-hard. Strengthening this result to show also hardness of *approximating* the entangled value is one of the most important open questions in this area. Ideally, we would like to prove the analogue of Theorem 1.1 for the case of entangled provers, or determine that it is not true.

Strong violation of Bell inequalities: Games exhibiting a gap between their entangled value and non-entangled value are of great interest to physicists, for possible use in experiments whose goal is to demonstrate the presence of quantum entanglement (see, e.g., [WW01] and references therein). Such games are said in the physics literature to exhibit a ‘violation of Bell inequalities’. By combining our main result with a remarkable construction by Khot and Vishnoi [KV05], we can obtain unique games whose entangled value is very close to 1 even though their value is very close to 0. Previously, such large gaps were known only for non-unique games (such as the parallel repetition of the Magic Square game). The simpler structure of unique games might be an advantage in certain circumstances. A related result was recently established for three-prover games with binary answers [PGWP⁺].

In more detail, Khot and Vishnoi constructed for any $k \geq 1$ and $\eta > 0$, a unique game with $2^k/k$ questions to each prover and answer alphabet of size k for which the value of our SDP relaxation is at least $1 - 9\eta$ and whose (non-entangled) value is at most $2/k^\eta$.² (We note that the existence of unique games whose SDP value is close to 1 and whose value is close to 0 follows from the UGC, but Khot and Vishnoi’s result is unconditional and also gives explicit parameters.) By combining their result with Theorem 1.3, we obtain that for any $k \geq 1$ and $\eta > 0$, there exists a unique game G with $2^k/k$ questions to each prover and answer alphabet of size k for which the entangled value is at least $1 - 54\eta$ and whose (non-entangled) value is at most $2/k^\eta$.

2 Preliminaries

We study *one-round two-prover cooperative games of incomplete information*, also known in the quantum information literature as *nonlocal games*. In such a game, a referee (also called the verifier) asks questions to two provers, Alice and Bob, who cooperate with each other. A game $G = G(\pi, V)$ is specified by a set Q and a number $k \geq 1$, a probability distribution $\pi : Q \times Q \rightarrow [0, 1]$, and a predicate $V : [k] \times [k] \times Q \times Q \rightarrow \{0, 1\}$. The game proceeds as follows: the referee samples $(s, t) \in Q \times Q$ according to π and sends question s to Alice and question t to Bob. Alice replies with an answer $a \in [k]$, and Bob with an answer $b \in [k]$. The provers win if and only if $V(a, b | s, t) = 1$.³ The provers are allowed to agree on a strategy before the game starts, but are not allowed to communicate with each other after receiving their questions. The *value* of a game is the maximum probability with which the provers can win. The provers may share randomness, but it is easy to see that this does not increase the value of the game.

The provers can also share an entangled state, which can sometimes increase their winning probability (for background on quantum information see, e.g., [NC00]). We therefore define the

²Strictly speaking, their construction gives a general constraint graph, and not a two-prover game as needed in our case. In order to derive a two-prover game from their construction, simply choose a random constraint and then randomly send one question to each prover.

³We write $V(a, b | s, t)$ for $V(a, b, s, t)$ to distinguish variables for questions from variables for answers.

entangled value of a game to be the highest winning probability of entangled provers. Let us define this more explicitly. In general, a strategy for entangled provers is described by a shared (possibly mixed) quantum state, as well as a general measurement on Alice’s part of the state for each of her questions, and a general measurement on Bob’s part of the state for each of his questions. On obtaining question s , Alice performs the measurement corresponding to s on her part of the state and returns as answer the result; Bob’s behavior is similar. By standard arguments, we can assume without loss of generality that Alice and Bob share a *pure* quantum state $|\psi\rangle \in \mathbb{C}^{d \times d}$ for some $d \geq 1$, and that, moreover, they use *projective measurements*, i.e., for each s Alice’s measurement is described by $\{A_a^s\}_a$ where the A_a^s are orthogonal projectors and $\sum_a A_a^s = I$, and similarly Bob uses measurements $\{B_b^t\}_b$. By definition, the probability that on questions s, t Alice answers a and Bob answers b is given by $\langle \psi | A_a^s \otimes B_b^t | \psi \rangle$. Therefore, the entangled value of G can be written as

$$\omega^*(G) = \lim_{d \rightarrow \infty} \max_{|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d} \max_{A_a^s, B_b^t} \sum_{abst} \pi(s, t) V(a, b | s, t) \langle \psi | A_a^s \otimes B_b^t | \psi \rangle.$$

We shall be concerned with games of a specific form.

Definition 2.1 *A game is termed unique if we can associate a permutation σ_{st} on $[k]$ with each pair of questions (s, t) such that $V(a, b | s, t) = 1$ if and only if $b = \sigma_{st}(a)$.*

We also want to consider games in which the answers of each prover in an optimal strategy are distributed uniformly in $[k]$.

Definition 2.2 *A game is termed uniform if there exists an optimal strategy for entangled provers in which, for each prover and for each question, the marginal distribution of his answers is uniform over $[k]$.*

As an example, consider a game in which the verifier accepts answers a, b solely as a function of $a - b \pmod k$. In such a game, Alice and Bob, by using their shared randomness (or entanglement), can choose a number c uniformly from $[k]$ and add it modulo k to their responses. This does not change their probability of winning the game, but causes each prover’s output to be uniformly distributed in $[k]$. Therefore, any such game is a uniform game.

Finally, we consider more general games known as *d -to- d' games*.

Definition 2.3 *A game has the d -to- d' property if for each pair of questions (s, t) and each answer a of the first prover, there are at most d answers b of the second prover for which $V(a, b | s, t) = 1$, and similarly, for each answer b of the second prover, there are at most d' answers a for which $V(a, b | s, t) = 1$.*

In [Kho02], Khot conjectured that for any $\delta > 0$ there exists a $k = k(\delta)$ such that it is NP-hard to determine whether, given a 2-to-1 game with answers from a domain of size k , its value is exactly 1 (i.e., perfectly satisfiable) or at most δ . This conjecture was used in proving the hardness of graph coloring problems [DMR06]. In Theorem 4.6 below we show that if the provers are allowed entanglement, then this conjecture, as well as its extension to d -to- d' games, is false.

3 SDP Relaxation

We use the following SDP relaxation for the entangled value of an arbitrary two-prover one-round game. The SDP maximizes over the real vectors $\{u_a^s\}$, $\{v_b^t\}$, and z .

SDP 1

$$\begin{aligned}
\text{Maximize:} & \quad \sum_{abst} \pi(s, t) V(a, b | s, t) \langle u_a^s, v_b^t \rangle \\
\text{Subject to:} & \quad \|z\| = 1 \\
& \quad \forall s, t, \sum_a u_a^s = \sum_b v_b^t = z \\
& \quad \forall s, t, \forall a \neq b, \langle u_a^s, u_b^s \rangle = 0 \text{ and } \langle v_a^t, v_b^t \rangle = 0 \\
& \quad \forall s, t, a, b, \langle u_a^s, v_b^t \rangle \geq 0
\end{aligned}$$

Note that the second constraint is, strictly speaking, not an SDP constraint. However, it is easy to see that there is an equivalent formulation in SDP language. For instance, we can replace the first two constraints by $\sum_{a,b} \langle u_a^s, v_b^t \rangle = 1$, $\sum_a \langle u_a^s, u_a^s \rangle = 1$ and $\sum_b \langle v_b^t, v_b^t \rangle = 1$.

For a game G , let $\omega_{\text{sdp1}}(G)$ be the value of SDP 1. We start by showing that it is indeed a relaxation of the entangled value of the game.

Lemma 3.1 *Let $G = G(\pi, V)$ be a (not necessarily unique) one-round two-prover game. Then $\omega^*(G) \leq \omega_{\text{sdp1}}(G)$.*

Proof: Consider any strategy for the entangled provers, specified by a state $|\psi\rangle \in \mathbb{C}^{d \times d}$ and projectors $\{A_a^s\}$ and $\{B_b^t\}$. Define the vectors $\tilde{u}_a^s = (A_a^s \otimes I)|\psi\rangle$ and $\tilde{v}_b^t = (I \otimes B_b^t)|\psi\rangle$ in $\mathbb{C}^{d \times d}$. Consider now the *real* $2d^2$ -dimensional vectors defined by $u_a^s = \text{Re}(\tilde{u}_a^s) \oplus \text{Im}(\tilde{u}_a^s)$, $v_b^t = \text{Re}(\tilde{v}_b^t) \oplus \text{Im}(\tilde{v}_b^t)$ and $z = \text{Re}(|\psi\rangle) \oplus \text{Im}(|\psi\rangle)$. Note that because $\langle \tilde{u}_a^s, \tilde{v}_b^t \rangle = \langle \psi | A_a^s \otimes B_b^t | \psi \rangle$ is real, we have that

$$\langle u_a^s, v_b^t \rangle = \text{Re}(\tilde{u}_a^s) \text{Re}(\tilde{v}_b^t) + \text{Im}(\tilde{u}_a^s) \text{Im}(\tilde{v}_b^t) = \text{Re}(\langle \tilde{u}_a^s, \tilde{v}_b^t \rangle) = \langle \psi | A_a^s \otimes B_b^t | \psi \rangle \geq 0.$$

The other constraints follow from the observations that $\sum_a \tilde{u}_a^s = |\psi\rangle = \sum_b \tilde{v}_b^t$, that $\langle z, z \rangle = \text{Re}(\langle \psi | \psi \rangle) = 1$, and that for $a \neq b$,

$$\langle u_a^s, u_b^s \rangle = \text{Re}(\langle \tilde{u}_a^s, \tilde{u}_b^s \rangle) = \text{Re}(\langle \psi | A_a^s A_b^s \otimes I | \psi \rangle) = 0,$$

since $A_a^s A_b^s = 0$, and similarly $\langle v_a^t, v_b^t \rangle = 0$. ■

In the case of uniform games there exists an optimal strategy in which the provers' output distribution is uniform on $[k]$. This allows us to add the following constraint to the original SDP:

$$\text{Additional constraint for SDP 2: } \quad \forall s, t, a, b, \|u_a^s\| = \|v_b^t\| = 1/\sqrt{k}.$$

This gives a more constrained SDP relaxation for uniform games, which we call SDP 2 and whose value we denote by $\omega_{\text{sdp2}}(G)$. To see that this is indeed a relaxation of uniform games, note that with the notation of the proof of Lemma 3.1,

$$\langle u_a^s, u_a^s \rangle = \langle \tilde{u}_a^s, \tilde{u}_a^s \rangle = \langle \psi | A_a^s \otimes I | \psi \rangle,$$

which is equal to $\frac{1}{k}$ since the last expression is exactly Alice's marginal distribution, and similarly for v_b^t . This extra constraint will allow us to slightly improve our quantum rounding procedure.

4 Quantum Rounding

In this section we describe how to round the solution of our SDP to a quantum strategy. We start with an informal outline of the rounding algorithm for the special case of *uniform* unique games, and then describe how to modify it for general unique games. The formal description of the rounding algorithm is given as Algorithm 1 below, and it uses a particular measurement, given here as Measurement 1, as a subroutine.

Our goal is the following. The SDP relaxation of a game gives us a solution $\{u_a^s\}, \{v_b^t\}$ where for fixed s, t the inner products $\langle u_a^s, v_b^t \rangle$ can be interpreted as a joint probability distribution on (a, b) (note that they are non-negative and sum to 1). The marginal distribution on a is given by $\|u_a^s\|^2$ and on b by $\|v_b^t\|^2$; in particular, for SDP 2 these marginal distributions are *uniform*. The value of the SDP then represents the winning probability in the corresponding game (given by π and V), when the provers answer according to this probability distribution. Hence we would like to design a quantum strategy that *reproduces* this probability distribution as closely as possible. Then also its winning probability will be close to the value of the SDP solution.

The basic idea is to use the solution to the SDP to define a measurement for Alice and Bob on the *maximally entangled* state $|\psi\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^n |i, i\rangle$. This state has the property that for any orthonormal basis of *real* vectors $\{|u_i\rangle\}_{i=1}^n$ it can be written as $|\psi\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^n |u_i, u_i\rangle$. This implies that if Alice measures in such a basis and obtains outcome i , then Bob's state collapses to $|u_i\rangle$. If Bob then measures in a basis $\{|v_j\rangle\}$ in which one of the vectors, say $|v_j\rangle$, is close to $|u_i\rangle$, then Bob's measurement outcome is likely to be j .

We now describe our rounding algorithm for the simpler case of uniform unique games, using SDP 2. Consider a solution of SDP 2 and assume that it lies in \mathbb{R}^n for some $n \geq 1$. Assume moreover that the value of this solution is $1 - \varepsilon$ for some small $\varepsilon > 0$. This means that for a typical pair s, t the sum $\sum_{a=1}^k \langle u_a^s, v_{\sigma_{st}(a)}^t \rangle$ is close to 1, and hence, since the norms of all these vectors are $1/\sqrt{k}$, u_a^s is typically close to $v_{\sigma_{st}(a)}^t$. We now use this solution to define local projective measurements on the n -dimensional maximally-entangled state $|\psi\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^n |i, i\rangle$. For fixed s , the k vectors u_a^s are orthogonal, and so, after normalization, define part of a basis. We complete this to a basis of \mathbb{R}^n in an arbitrary way. When Alice is asked question s , she measures her half of $|\psi\rangle$ in this basis, outputting a if her measurement result corresponds to the basis element u_a^s , and outputting (for the moment) nothing if she obtains one of the extra basis elements. Similarly, when Bob is asked question t , he measures his half of $|\psi\rangle$ in a basis that contains the vectors $\{v_b^t\}_b$, outputting b if his measurement result corresponds to the vector v_b^t , and nothing otherwise.

This rounding scheme has two properties – one good, the other bad. First the good property: if Alice outputs a , then Bob will output $\sigma_{st}(a)$ with high probability, since the vector u_a^s is close to the vector $v_{\sigma_{st}(a)}^t$. Hence we have that if Alice does give an answer, Bob's answer will be correct with high probability. The problem is that the probability that Alice does give an answer is k/n , which is typically very small.

Luckily, the solution to this problem is easy. If Alice doesn't obtain a suitable outcome, she just starts over, performing her measurement on a fresh maximally entangled state. She keeps performing her measurement on fresh states until she obtains an outcome corresponding to a vector u_a^s . Bob does likewise, performing his measurement on fresh states, making sure to use them in the same order as Alice. This fixes the problem above, since now Alice answers with probability

1. In doing so, however, it seems we have created a new problem: it is entirely possible that Alice and Bob's measurements will succeed on different copies of the maximally-entangled state, in which case their answers won't be correlated at all. But here we are saved by the good property of our measurement: if Alice does obtain an outcome a in a given round, then Bob will obtain a correct outcome b with high probability, assuming he hasn't already answered. Conversely, if Bob obtains an outcome b in a given round, then Alice will obtain a correct outcome a with high probability, assuming she hasn't already answered. All this means that, with high probability, their measurements succeed on the same copy of $|\psi\rangle$.

Let us now consider general (not necessarily uniform) unique games. Our starting point now is a solution to SDP 1 of value $1 - \varepsilon$. Again, for a typical pair s, t we have that $\sum_a \langle u_a^s, v_{\sigma_{st}(a)}^t \rangle$ is close to 1. However, in our rounding algorithm we have to add an extra step to account for the fact that the vectors u_a^s, v_b^t might not all be of the same length. The projective measurement that we have constructed does not take this into account, because all basis vectors after renormalization are equally likely to occur. Recall that our goal is to reproduce the joint probability distribution on a, b given by $\langle u_a^s, v_b^t \rangle$, and, in particular, its marginal distributions on a and b , given by $\|u_a^s\|^2$ and $\|v_b^t\|^2$: Our rounding algorithm has to ensure that outcomes that correspond to short vectors u_a^s or v_b^t should be less likely than those corresponding to longer vectors. To this end we use a rejection sampling technique, as follows. Alice samples λ uniformly from $[0, 1]$. If Alice's measurement outcome is a , she outputs it if and only if $\lambda \leq \|u_a^s\|^2$. This ensures that the probability that Alice answers a is the same as that given by the SDP relaxation. Similarly, if Bob's measurement outcome is b , he outputs it if and only if $\lambda \leq \|v_b^t\|^2$. Again a problem arises that, even if Alice's and Bob's outcomes are otherwise correlated, the rejection sampling could make Alice accept and Bob reject (or vice versa) on the same copy of $|\psi\rangle$. Luckily, we are helped again by the fact that on average, u_a^s and $v_{\sigma_{st}(a)}^t$ are close, and in particular have comparable length. Therefore, to coordinate the rejection sampling procedure, Alice and Bob will use a *shared* random variable λ in each step, which means that with high probability they will either both accept or both reject. (Note that it is easy to obtain a shared random variable from shared entanglement.)

Algorithm 1 Quantum rounding for unique games.

Setup:	Alice and Bob share many copies of an n -dimensional maximally entangled state $ \psi\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^n i, i\rangle$, for some fixed basis $\{ i\rangle\}$ of \mathbb{C}^n , as well as a sequence $\Lambda = (\lambda_1, \lambda_2, \dots)$ of real numbers, where the λ_i are independent and each is sampled uniformly from $[0, 1]$.
Alice:	On input s , performs the measurement $\text{MEASURE}(u_s^1, u_s^2, \dots, u_s^k)$ on her share of the maximally entangled states and the sequence Λ .
Bob:	On input t , performs the measurement $\text{MEASURE}(v_t^1, v_t^2, \dots, v_t^k)$ on his share of the maximally entangled states and the sequence Λ .

4.1 Analysis of the measurement procedure

Lemma 4.1 *Let x_1, \dots, x_k and y_1, \dots, y_k be two sequences of orthogonal vectors in \mathbb{R}^n such that $\sum_{i=1}^k \|x_i\|^2 = \sum_{i=1}^k \|y_i\|^2 = 1$. Assume Alice and Bob apply Measurement 1, Alice using (x_i) and Bob using (y_i) . For*

Measurement 1 The measurement $\text{MEASURE}(x_1, x_2, \dots, x_k)$ used in Algorithm 1.

Input: A state on a Hilbert space $\mathcal{H} = \bigotimes_{r=1}^{\infty} \mathcal{H}_r$, where each $\mathcal{H}_r \cong \mathbb{C}^n$, and a sequence of real numbers $\Lambda = (\lambda_1, \lambda_2, \dots)$, where each $\lambda_r \in [0, 1]$.

Parameters: k orthogonal vectors $x_1, x_2, \dots, x_k \in \mathbb{R}^n$.

Output: An integer $m \in \{1, 2, \dots, k\}$.

Measurement: Define a POVM on \mathbb{C}^n with elements

$$P_i = \left| \frac{x_i}{\|x_i\|} \right\rangle \left\langle \frac{x_i}{\|x_i\|} \right| \text{ for } i = 1, 2, \dots, k \text{ and } P_0 = I - \sum_{i=1}^k P_i,$$

where for a vector $w \in \mathbb{R}^n$ we write $|w\rangle = \sum_i (w)_i |i\rangle$ for its embedding into \mathbb{C}^n .

For $r = 1, 2, \dots$ **do:**

Measure \mathcal{H}_r using POVM (P_0, \dots, P_k) , obtaining outcome m .

If $(m \neq 0 \text{ and } \lambda_r \leq \|x_m\|^2)$ **then** output m and exit.

any $i, j \in \{1, \dots, k\}$ define

$$q_{i,j} := \left\langle \frac{x_i}{\|x_i\|}, \frac{y_j}{\|y_j\|} \right\rangle^2 \min(\|x_i\|^2, \|y_j\|^2)$$

and let $q_{\text{total}} := \sum_{i,j} q_{i,j}$. Then for any $i, j \in \{1, \dots, k\}$, the probability that Alice outputs i and Bob outputs j is at least

$$\frac{q_{i,j}}{2 - q_{\text{total}}}.$$

Proof: We start by analyzing one round of the measurement, i.e., Alice and Bob share a maximally entangled n -dimensional state and a random number $\lambda \in [0, 1]$. Each performs a measurement given by his or her input vectors, and outputs the outcome m if $m \neq 0$ and $\lambda \leq \|x_m\|^2$ (resp. $\lambda \leq \|y_m\|^2$), or nothing otherwise.

A round can end in one of four possible ways, to which we assign probabilities as follows:

- (p_{done}) Both Alice and Bob give an output;
- ($p_{\text{fail},1}$) Alice gives an output while Bob does not;
- ($p_{\text{fail},2}$) Bob gives an output while Alice does not;
- (p_{retry}) Neither Alice nor Bob gives an output.

Hence $p_{\text{done}} + p_{\text{fail},1} + p_{\text{fail},2} + p_{\text{retry}} = 1$. Let us also define $p_{i,j}$ for $i, j \in \{1, \dots, k\}$ as the probability that Alice outputs i and Bob outputs j in one round. Notice that $p_{\text{done}} = \sum_{i,j=1}^k p_{i,j}$.

We now compute each of these probabilities. By construction, the probability that Alice obtains an outcome $i \neq 0$ from her POVM is exactly $1/n$. Conditioned on that happening, Bob's state collapses to the pure state given by the vector $|x_i\rangle / \|x_i\|$. Therefore, the conditional probability that he obtains an outcome $j \neq 0$ in his POVM is given by $\left\langle \frac{x_i}{\|x_i\|}, \frac{y_j}{\|y_j\|} \right\rangle^2$. Finally, conditioned on Alice measuring $i \neq 0$ and Bob measuring $j \neq 0$, the probability that both actually output their

values is $\min(\|x_i\|^2, \|y_j\|^2)$. Hence we see that for any $i, j \in \{1, \dots, k\}$, the probability that in one round of the measurement Alice outputs i and Bob outputs j is

$$p_{i,j} := \frac{1}{n} \left\langle \frac{x_i}{\|x_i\|}, \frac{y_j}{\|y_j\|} \right\rangle^2 \min(\|x_i\|^2, \|y_j\|^2) = \frac{1}{n} q_{i,j}.$$

Moreover, it is easy to see that the probability that Alice gives an output is

$$\sum_{i=1}^m \frac{1}{n} \|x_i\|^2 = \frac{1}{n}$$

and similarly for Bob. This implies that

$$p_{\text{fail},1} = p_{\text{fail},2} = \frac{1}{n} - p_{\text{done}}.$$

To complete the proof, let us consider the probability that in Measurement 1, Alice outputs i and Bob outputs j . This probability is lower bounded by the probability that Alice outputs i and Bob outputs j in the same round. The latter probability is given by

$$\sum_{r=0}^{\infty} (p_{\text{retry}})^r p_{i,j} = \frac{p_{i,j}}{1 - p_{\text{retry}}} = \frac{p_{i,j}}{p_{\text{done}} + p_{\text{fail},1} + p_{\text{fail},2}} = \frac{p_{i,j}}{\frac{2}{n} - p_{\text{done}}} = \frac{q_{i,j}}{2 - q_{\text{total}}}.$$

■

Corollary 4.2 *Let V be a subset of $\{1, \dots, k\}^2$. Then, in the setting of Lemma 4.1, the probability that Alice's output i and Bob's output j are such that $(i, j) \in V$ is at least*

$$\frac{p_V}{2 - p_V} \geq 1 - 2(1 - p_V),$$

where

$$p_V := \sum_{i,j \in V} \left\langle \frac{x_i}{\|x_i\|}, \frac{y_j}{\|y_j\|} \right\rangle^2 \min(\|x_i\|^2, \|y_j\|^2).$$

4.2 Analysis of the quantum rounding

We first analyze the easier case of uniform unique games. We remark that we could have slightly simplified the algorithm for this case by avoiding the rejection sampling step, but for convenience we keep it since it does not affect our results.

Theorem 4.3 (Uniform unique games) *Let G be a uniform unique game. Suppose that $\omega_{\text{sdp}2}(G) = 1 - \varepsilon$. Then $\omega^*(G) \geq 1 - 4\varepsilon$.*

Proof: Fix a solution $\{u_a^s\}, \{v_b^t\}, z$ to SDP 2 with value $1 - \varepsilon$ and consider the strategy of Alice and Bob given by Algorithm 1. Our goal is to show that this strategy has success probability at least $1 - 4\varepsilon$. In order to show this, it suffices to show that for any questions s, t , the success probability of Alice and Bob on these questions is at least $1 - 4(1 - \sum_{ab} V(a, b | s, t) \langle u_a^s, v_b^t \rangle)$.

So from now on fix a pair of questions s, t and let σ be the permutation corresponding to the constraint between s and t , i.e., $V(a, b | s, t) = 1$ if and only if $b = \sigma(a)$. For $i = 1, \dots, k$, define

$u_i = u_i^s$ and $v_i = v_{\sigma(i)}^t$. Suppose that $\sum_i \langle u_i, v_i \rangle \geq 1 - \tilde{\epsilon}$ for some $\tilde{\epsilon} \geq 0$ and recall that our goal is to show that Alice and Bob succeed with probability at least $1 - 4\tilde{\epsilon}$. By Corollary 4.2, their success probability is at least

$$p_{\text{succ}} \geq 1 - 2(1 - p'_{\text{succ}}),$$

where

$$p'_{\text{succ}} = \sum_{i=1}^k \left\langle \frac{u_i}{\|u_i\|}, \frac{v_i}{\|v_i\|} \right\rangle^2 \min(\|u_i\|^2, \|v_i\|^2). \quad (1)$$

It therefore suffices to show that $p'_{\text{succ}} \geq 1 - 2\tilde{\epsilon}$. Using the extra constraints in SDP 2 and the Cauchy-Schwartz inequality,

$$p'_{\text{succ}} = k \sum_{i=1}^k \langle u_i, v_i \rangle^2 \geq \left(\sum_{i=1}^k \langle u_i, v_i \rangle \right)^2 \geq 1 - 2\tilde{\epsilon}.$$

■

Theorem 4.4 (Unique games) *Let G be a unique game. Suppose that $\omega_{\text{sdp1}}(G) = 1 - \epsilon$. Then $\omega^*(G) \geq 1 - 6\epsilon$.*

Proof: As in the proof of Theorem 4.3, we have vectors u_i, v_i , this time coming from SDP 1, satisfying $\sum_i \langle u_i, v_i \rangle \geq 1 - \tilde{\epsilon}$. Our goal now is to show that $p'_{\text{succ}} \geq 1 - 3\tilde{\epsilon}$ where p'_{succ} is defined as in Eq. (1).

Let $F := \sum_i \|u_i\| \|v_i\|$. We first notice that

$$F \leq \left(\sum_i \|u_i\|^2 \right)^{1/2} \left(\sum_i \|v_i\|^2 \right)^{1/2} = 1.$$

Define

$$p''_{\text{succ}} := \sum_i \left\langle \frac{u_i}{\|u_i\|}, \frac{v_i}{\|v_i\|} \right\rangle^2 \|u_i\| \|v_i\|.$$

Then, by convexity,

$$\begin{aligned} p''_{\text{succ}} &= F \sum_i \frac{\|u_i\| \|v_i\|}{F} \left\langle \frac{u_i}{\|u_i\|}, \frac{v_i}{\|v_i\|} \right\rangle^2 \\ &\geq F \left(\sum_i \frac{\|u_i\| \|v_i\|}{F} \left\langle \frac{u_i}{\|u_i\|}, \frac{v_i}{\|v_i\|} \right\rangle \right)^2 \\ &= \frac{1}{F} \left(\sum_i \langle u_i, v_i \rangle \right)^2 \\ &\geq 1 - 2\tilde{\epsilon}. \end{aligned}$$

Moreover, using the fact that for any nonnegative $a, b \in \mathbb{R}$, $\sqrt{ab} - \min(a, b) \leq |a - b|/2$,

$$\begin{aligned}
p''_{\text{succ}} - p'_{\text{succ}} &\leq \frac{1}{2} \sum_i \left\langle \frac{u_i}{\|u_i\|}, \frac{v_i}{\|v_i\|} \right\rangle^2 |\|u_i\|^2 - \|v_i\|^2| \\
&\leq \frac{1}{2} \sum_i |\|u_i\|^2 - \|v_i\|^2| \\
&= \frac{1}{2} \sum_i |\langle u_i, z \rangle - \langle v_i, z \rangle| \\
&= \frac{1}{2} \sum_i \left| \sum_{j \neq i} \langle u_i, v_j \rangle - \sum_{j \neq i} \langle v_i, u_j \rangle \right| \\
&\leq \frac{1}{2} \left(\sum_i \sum_{j \neq i} \langle u_i, v_j \rangle + \sum_i \sum_{j \neq i} \langle v_i, u_j \rangle \right) \leq \tilde{\epsilon}.
\end{aligned}$$

■

Remark: The above analysis is ‘locally tight’ in the following sense. For any small $\tilde{\epsilon} > 0$ and large enough k , there exist two sequences of orthogonal vectors u_1, \dots, u_k and v_1, \dots, v_k satisfying (i) $\sum_i u_i = \sum_i v_i$ has norm 1, (ii) for all i, j , $\langle u_i, v_j \rangle \geq 0$, (iii) $\sum_{i=1}^k \langle u_i, v_i \rangle = 1 - \tilde{\epsilon}$, and (iv) the probability that the quantum rounding procedure produces a pair (i, j) with $i = j$ is roughly $1 - 6\tilde{\epsilon}$. Let $a = \sqrt{(1 - \tilde{\epsilon})/k}$ and $b = \sqrt{2\tilde{\epsilon}/k}$, and let $e_1, \dots, e_k, f_1, \dots, f_{k/2}$ be orthonormal unit vectors. Our vectors are given by $u_i = ae_i + bf_i$, $v_i = ae_i$ for $i = 1, \dots, \frac{k}{2}$, and $u_i = ae_i$, $v_i = ae_i + bf_{i - \frac{k}{2}}$ for $i = \frac{k}{2} + 1, \dots, k$.

Our final theorem deals with d -to- d games, and uses the following combinatorial claim.

Claim 4.5 *Let (V, E) be a directed acyclic graph with non-negative weights associated to its vertices, and let $V' \subseteq V$ denote the set of vertices with outdegree zero. Assume, moreover, that all indegrees are at most D , and that the weight of each vertex in $V \setminus V'$ is smaller by a factor of at least $2D$ than the sum of weights of its out-neighbors. Then the total weight in V' is at least half the total weight in V .*

Proof: Assume without loss of generality that $V = \{1, \dots, n\}$, that all edges are facing forward (i.e., are of the form (i, j) with $j > i$), and that $V' = \{n - m + 1, \dots, n\}$ for $m = |V'| \geq 1$. Consider the following process. Initially, for $i = 1, \dots, n$, set a_i to be the weight of vertex i and b_i to be zero. Then, for $i = 1, \dots, n - m$ do the following: for each edge (i, j) leaving i , add a_j/D to b_j , and then set $a_i = b_i = 0$.

Notice the following two properties. First, since the in-degrees are at most D , we always have $b_i \leq a_i$ for all i . Second, $\sum_{i=1}^n (a_i + b_i)$ never decreases during this process. The reason is that although we decrease the sum by $a_i + b_i$ when we set $a_i = b_i = 0$, we also increase it by the sum of weights of i 's out-neighbors divided by D , which is by assumption at least $2a_i \geq a_i + b_i$. Now consider the situation at the end of the process. One one hand,

$$\sum_{i=1}^n (a_i + b_i) = \sum_{i=n-m+1}^n (a_i + b_i) \leq 2 \sum_{i=n-m+1}^n a_i,$$

which is exactly twice the weight of vertices in V' . On the other hand, this same quantity is at least the weight of vertices in V , and hence the claim follows. ■

Theorem 4.6 (*d*-to-*d* games) Let G be a d -to- d game for some $d \geq 2$, and assume that $\omega_{\text{sdp1}}(G) = 1$. Then $\omega^*(G) \geq \frac{1}{20(d-1)}$.

Proof: Fix a solution $\{u_a^s\}, \{v_b^t\}, z$ to SDP 1 with value 1 and consider the strategy of Alice and Bob given by Algorithm 1. Our goal is to show that this strategy has success probability at least $\frac{1}{20(d-1)}$. Clearly, it suffices to show this for any fixed questions s, t . So from now on fix a pair of questions s, t , and let $a_{ij} = \langle u_i^s, v_j^t \rangle$. Let $V \subseteq \{1, \dots, k\}^2$ be the set of allowed answers from the provers, and notice that $\sum_{i,j \in V} a_{ij} = 1$ and a_{ij} is zero for all $(i, j) \notin V$. By Corollary 4.2, the success probability is at least $p_V/2$ where

$$p_V = \sum_{i,j \in V} a_{ij} \frac{a_{ij}}{\max(\sum_{i'} a_{i'j}, \sum_{j'} a_{ij'})}.$$

Let $V' \subseteq V$ be the set of all pairs $(i, j) \in V$ for which $\max(\sum_{i'} a_{i'j}, \sum_{j'} a_{ij'}) \leq 5(d-1)a_{ij}$. Clearly, $p_V \geq \frac{1}{5(d-1)} \sum_{i,j \in V'} a_{ij}$, and hence it suffices to lower bound $\sum_{i,j \in V'} a_{ij}$ by $\frac{1}{2}$.

Consider the directed graph on vertex set V defined as follows. We assign the weight a_{ij} to each vertex $(i, j) \in V$ so that the total weight of vertices is 1. Let (i, j) be some vertex $V \setminus V'$. If $\sum_{i'} a_{i'j} > 5(d-1)a_{ij}$, then we put an edge from (i, j) to (i', j) for all i' such that $a_{i'j} > a_{ij}$. Otherwise, it must be the case that $\sum_{j'} a_{ij'} > 5(d-1)a_{ij}$, and we proceed similarly, placing an edge from (i, j) to (i, j') for all j' such that $a_{ij'} > a_{ij}$.

The graph obtained is clearly acyclic. Moreover, the sum of weights of the out-neighbors of each vertex $(i, j) \in V \setminus V'$ is at least $5(d-1)a_{ij} - (d-1)a_{ij} = 4(d-1)a_{ij}$, the worst case being when $d-1$ elements in the sum $\sum_{j'} a_{ij'}$ (or $\sum_{i'} a_{i'j}$) are equal to a_{ij} . Also, the vertices in V' are exactly those with outdegree zero, and all indegrees are at most $2(d-1)$. We can therefore apply Claim 4.5 with $D = 2(d-1)$ to obtain that the weight of vertices in V' is at least $\frac{1}{2}$. This implies that $p_V \geq \frac{1}{10(d-1)}$, as required. ■

Acknowledgments

We thank Stephanie Wehner for useful discussions.

References

- [ALM⁺98] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [AS98] S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of np. *J. ACM*, 45(1):70–122, 1998.
- [Bel64] J. S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.
- [BM04] H. Buhrman and S. Massar. Causality and Cirel’son bounds, 2004. Quant-ph/0409066.
- [BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004. ISBN 0-521-83378-7.

- [CHSH69] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23:880–884, 1969.
- [CHTW04] R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Conference on Computational Complexity (CCC 2004)*, pages 236–249. 2004.
- [Cir80] B. S. Cirel’son. Quantum generalizations of Bell’s inequality. *Lett. Math. Phys.*, 4:93–100, 1980.
- [CKK⁺06] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Comput. Complexity*, 15(2):94–114, 2006. ISSN 1016-3328.
- [CMM06a] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games (extended abstract). In *STOC’06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 205–214. ACM, New York, 2006.
- [CMM06b] E. Chlamtac, K. Makarychev, and Y. Makarychev. How to play unique games using embeddings. In *Proc. 47th IEEE Symp. on Foundations of Computer Science*, pages 687 – 696. 2006.
- [CSUU07] R. Cleve, W. Slofstra, F. Unger, and S. Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Proc. of 22nd IEEE Annual Conference on Computational Complexity (CCC 2007)*, pages 109–114, 2007. ISSN 1093-0159. doi: <http://doi.ieeecomputersociety.org/10.1109/CCC.2007.24>.
- [DMR06] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. In *Proc. 38th ACM Symp. on Theory of Computing (STOC)*, pages 344–353. 2006.
- [DS05] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Ann. of Math. (2)*, 162(1):439–485, 2005. ISSN 0003-486X.
- [EPR35] A. Einstein, P. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, 1935.
- [Fei98] U. Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998. ISSN 0004-5411.
- [FL92] U. Feige and L. Lovász. Two-prover one-round proof systems: Their power and their problems. In *Proc. 24th ACM Symp. on Theory of Computing*, pages 733–741. 1992.
- [GT06] A. Gupta and K. Talwar. Approximating unique games. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 99–106. 2006.
- [GW95] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6):1115–1145, 1995. ISSN 0004-5411.

- [Hås99] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.*, 182(1):105–142, 1999. ISSN 0001-5962.
- [Hås01] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. ISSN 0004-5411.
- [Kho02] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 767–775. ACM, New York, 2002.
- [KKM⁺] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. On the power of entangled provers: Immunizing games against entanglement. ArXiv:0704.2903.
- [KKMO07] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for max-cut and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. ISSN 0097-5397.
- [KR03] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. In *Proc. of 18th IEEE Annual Conference on Computational Complexity (CCC)*, pages 379–386. 2003.
- [KV05] S. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into l_1 . In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 53–62. 2005.
- [Mas05] L. Masanes. Extremal quantum correlations for N parties with two dichotomic observables per site, 2005. Quant-ph/0512100.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, New York, 2000.
- [NPA07] M. Navascues, S. Pironio, and A. Acín. Bounding the set of quantum correlations. *Phys. Rev. Lett.*, 98(1):010401, 2007. doi:10.1103/PhysRevLett.98.010401.
- [PGWP⁺] D. Perez-Garcia, M. Wolf, C. Palazuelos, I. Villanueva, and M. Junge. Unbounded violation of tripartite Bell inequalities. Quant-ph/0702189.
- [Raz98] R. Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. ISSN 0097-5397.
- [Tre05] L. Trevisan. Approximation algorithms for unique games. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 197–205. 2005.
- [WW01] R. F. Werner and M. M. Wolf. Bell inequalities and entanglement. *Quantum Information and Computation*, 1(3):1–25, 2001.