Quantum de finetti theorems under local measurements with applications

Quantum de Finetti Theorems under Local Measurements with Applications

Fernando G.S.L. Brandão[*]   Aram W. Harrow [†]

October 24, 2013

Abstract

Quantum de Finetti theorems are a useful tool in the study of correlations in quantum multipartite states. In this paper we prove two new quantum de Finetti theorems, both showing that under tests formed by local measurements in each of the subsystems one can get a much improved error dependence on the dimension of the subsystems. We also obtain similar results for non-signaling probability distributions. We give the following applications of the results to quantum complexity theory, polynomial optimization, and quantum information theory:

- We prove the optimality of the Chen-Drucker protocol for 3-SAT, under the assumption there is no subexponential-time algorithm for SAT. In the protocol a prover sends to a verifier $\sqrt{n}\,\text{polylog}(n)$ unentangled quantum states, each composed of $O(\log(n))$ qubits, as a proof of the satisfiability of a 3-SAT instance with n variables and $O(n)$ clauses. The quantum verifier checks the validity of the proof by performing local measurements on each of the proofs and classically processing the outcomes. We show that any similar protocol with $O(n^{1/2-\varepsilon})$ qubits would imply a $\exp(n^{1-2\varepsilon}\,\text{polylog}(n))$-time algorithm for 3-SAT.

- We show that the maximum winning probability of free games (in which the questions to each prover are chosen independently) can be estimated by linear programming in time $\exp(O(\log|Q| + \log^2|A|/\varepsilon^2))$, with $|Q|$ and $|A|$ the question and answer alphabet sizes, respectively, matching the performance of a previously known algorithm due to Aaronson, Impagliazzo, Moshkovitz, and Shor. This result follows from a new monogamy relation for non-locality, showing that k-extendible non-signaling distributions give at most a $O(k^{-1/2})$ advantage over classical strategies for free games. We also show that 3-SAT with n variables can be reduced to obtaining a constant error approximation of the maximum winning probability under entangled strategies of $O(\sqrt{n})$-player one-round non-local games, in which only two players are selected to send $O(\sqrt{n})$-bit messages.

- We show that the optimization of certain polynomials over the complex hypersphere can be performed in quasipolynomial time in the number of variables n by considering $O(\log(n))$ rounds of the Sum-of-Squares (Parrilo/Lasserre) hierarchy of semidefinite programs. This can be considered an analogue to the hypersphere of a similar known results for the simplex. As an application to entanglement theory, we find a quasipolynomial-time algorithm for deciding multipartite separability.

- We consider a quantum tomography result due to Aaronson – showing that given an unknown n-qubit state one can perform tomography that works well for most observables by

---

[*]Department of Computer Science, University College London, and National Quantum Information Center of Gdansk. email: `fgslbrandao@gmail.com`

[†]Center for Theoretical Physics, Massachusetts Institute of Technology. email: `aram@mit.edu`

1

measuring only $O(n)$ independent and identically distributed (i.i.d.) copies of the state – and relax the assumption of having i.i.d copies of the state to merely the ability to select subsystems at random from a quantum multipartite state.

The proofs of the new quantum de Finetti theorems are based on information theory, in particular on the chain rule of mutual information. The results constitute improvements and generalizations of a recent de Finetti theorem due to Brandão, Christandl and Yard.

1 Background

A central problem in quantum information theory, quantum computation, and physics in general is to understand *entanglement*: quantum correlations with no counterpart in classical probability theory. An important technique in the study of entanglement are quantum versions of the de Finetti theorem. The latter states that the marginal probability distribution $p^{X_1...X_l}$ on l subsystems of a permutation-symmetric probability distribution $p^{X_1...X_k}$ on $k \geq l$ subsystems is close (within $l(l-1)/k$ in variational distance) to a convex combination of independent and identically distributed (i.i.d.) probability distributions [39]. This is a powerful result as it allows us to infer a very particular form for $p^{X_1...X_l}$ merely based on a symmetry assumption on $p^{X_1...X_k}$. Note we can always make sure this assumption holds true by merely forgetting the order of the k subsystems. Quantum versions of the de Finetti theorem state that a l-partite quantum state $\rho^{A_1...A_l}$ that is a reduced state of a permutation-symmetric state on $k \geq l$ subsystems is close (for $k \gg l$) to a convex combination of i.i.d. quantum states, i.e. $\rho^{A_1...A_l} \approx \int \mu(d\sigma)\sigma^{\otimes l}$ for a probability measure μ on quantum states.

The quantum version appears very similar to the original de Finetti theorem, but it is much more remarkable. Not only it says that the correlations are arranged in an organized fashion (as a convex combination of i.i.d. states) but also that the state of l subsystems is close to a *separable*, non-entangled, state. A well-known property of entanglement is that it is monogamous: A quantum system cannot be very much entangled with a large number of other systems. The quantum de Finetti theorems provide a quantitative statement for the monogamy of entanglement; in a symmetric state all the subsystems are equally correlated with all the others and so each of them can only be slightly entangled with a few of the others.

We now know several possible quantum versions of the de Finetti theorem [47, 83, 44, 77, 88, 27, 62, 32, 80, 70, 22]. A natural way to quantify the closeness to convex combinations of i.i.d. states is by the trace norm [1]. In this case Christandl, König, Mitchison, and Renner [32] proved an almost optimal quantum de Finetti theorem: $\rho^{A_1...A_l}$ is $(2d^2l/k)$-close to a convex combination of i.i.d. states in trace norm, with d the dimension of the subsystems, while there are examples where the error is $\Omega(dl/k)$. However in many applications this error is too large to be useful. One possible way forward is therefore to consider other ways of quantifying the approximation rather than the trace norm.

There are two known quantum de Finetti theorems following this idea. The first is the exponential de Finetti theorem of Renner [80], that achieves an exponentially small error in $k - l$, but only shows that $\rho^{A_1...A_l}$ is close to a convex combination of "almost i.i.d." states, a generalization of i.i.d. states having similar properties with respect to certain statistical tests. The second is the de Finetti theorem proved in Ref. [24], which works for $l = 2$ and has an error of $\sqrt{16\ln(d)/k}$,

---

[1]The trace norm gives the maximum probability of distinguishing two quantum states by arbitrary measurements.

an exponential improvement on the dimension dependence. The approximation is quantified by the one-way LOCC[2] norm, a variant of the trace norm for bipartite systems in which only measurements implementable by local operations and one-directional classical communication are allowed. Both results have found interesting applications: The first to quantum key distribution [79], quantum hypothesis testing [23], and quantum state tomography [80]; the second to entanglement testing, where it gives a quasipolynomial-time algorithm for determining if a quantum state is entangled or not [22], and to quantum complexity theory [22]. These two results suggest that more quantum versions of the de Finetti theorem might exist. In this paper we show that this is indeed the case.

It has emerged that some of the properties of entanglement, such as its monogamous character, are shared by more general classes of correlations [67]. A particular interesting example is the class of non-signaling distributions, which are a generalization of the correlations attainable by quantum mechanics. Versions of the de Finetti theorem for non-signaling distributions have also been derived [33, 12], although here again the scaling of the error – linear in the number of possible measurements – has limited the applicability of the results.

Another way to study quantum entanglement is via its role in operational tasks, e.g. in quantum key distribution and quantum computation. One fascinating case is the role of entanglement in *quantum proof systems*. The goal there is to understand how useful are entangled states for convincing a verifier the truth of a mathematical statement. There are many settings, such as interactive or non-interactive protocols, one or multiple provers, and which type of communication is allowed among the provers and the verifier (see e.g. [86]). In this paper we will be concerned with two such settings in particular. The first is MIP*, in which the provers share entanglement (or even general non-signaling correlations) and are only allowed to communicate with the verifier and not with each other [60]. The second is QMA(k), meaning non-interactive multiple proof protocols with the assumption that the proofs are *not* entangled [61]. Here we have the interesting situation where the assumption of not having entanglement among the proofs appears to give extra power to the proof system. Both settings have been extensively studied in the past (see e.g. [35, 87, 71, 40, 57, 56, 53, 52, 34, 58, 54] for work on MIP*/QMIP and [2, 19, 46, 29, 20, 24, 13, 65, 30, 45, 69, 28, 74, 81] for work on QMA(k)), although there are still many interesting open questions concerning them.

2 Results

The main results of this paper are two new quantum versions of the de Finetti theorem, along with extensions to arbitrary non-signaling distributions. Both are based on a coarser notion of approximation to the target state than the trace norm, but as a pay-off their error scales exponentially better with dimension. The notion of approximation used is that two quantum states are close if they have the same statistics under any local measurements on the subsystems. Our results thus extend the de Finetti bound of Ref. [22] to an arbitrary number of subsystems while improving on the error term to depend on the number of measurements instead of the local dimension, generalizing it to general non-signaling distributions, and in some cases providing an explicit rounding scheme. Among the applications of the new quantum de Finetti theorems we address two problems in quantum complexity theory, each concerning one of the proof systems mentioned above.

---

[2]The name LOCC stands for local operations and classical communication. See Eq. (92) for a precise definition of one-way LOCC.

Below we give a brief description of these applications.

Multiple Unentangled Proofs: The first application concerns a protocol due to Chen and Drucker [29] in which a prover sends to a verifier $\sqrt{n}\,\mathrm{polylog}(n)$ unentangled quantum states, each composed of $O(\log(n))$ qubits, as a proof of the satisfiability of a 3-SAT instance with n variables and $O(n)$ clauses. The quantum verifier then checks the validity of the proof by performing local quantum measurements on each of the proofs and post-processing the outcomes. This result (building on [2]), is surprising since one can convince a verifier the satisfiability of a 3-SAT instance by sending only $\sqrt{n}\,\mathrm{polylog}(n)$ qubits! It is a natural question whether the total number of qubits could be decreased even further. As a direct application of one of the new quantum de Finetti theorems we give strong evidence against any further decrease: We show that any similar protocol with $O(n^{1/2-\varepsilon})$ qubits, for any $\varepsilon > 0$, would imply a $\exp(n^{1-2\varepsilon}\,\mathrm{polylog}(n))$-time algorithm for 3-SAT. This proves the optimality of the protocol under the plausible assumption that there are no subexponential-time algorithms for SAT [49].

A related, but harder, problem is whether QMA(2) protocols can give at most a quadratic reduction in proof size with respect to QMA [3,4]. We believe the result we obtain gives evidence that this might be the case and that a suitable quantum version of the de Finetti theorem might be the right tool to show it [5].

Non-local Games: The second application concerns the computational complexity of non-local games. We give two results in this direction. The first is algorithmic and concerns the class of free games, defined as games in which the questions to each prover are chosen independently. We show that the maximum winning probability of such games can be approximated within additive error ε by a linear program in time $\exp(O(\log|Q| + \log^2|A|/\varepsilon^2))$, with $|Q|$ and $|A|$ the question and answer alphabet sizes, respectively. The run-time matches the performance of a different algorithm for the problem due to Aaronson, Impagliazzo, Moshkovitz, and Shor [3][6]. Although this is a purely classical result, we establish it by exploring a connection to non-local games: We show that for any two-player one-round free game, one can find another game on m players such that the maximum winning probability under non-signaling strategies, which can be computed by a linear program [51], gives a $\sqrt{\frac{\ln|A|}{2m}}$-additive approximation to the maximum winning probability of the original game. Note that since non-signaling strategies are at least as powerful as entangled strategies, the same result holds also for games in which the players share entanglement.

Using the observation above for entangled strategies, together with a hardness result for free games from [3], we also show that 3-SAT on n variables can be reduced to obtaining a *constant error* approximation of the maximum winning probability under entangled strategies of $O(\sqrt{n})$-player one-round non-local games, in which the players communicate $O(\sqrt{n})$ bits all together. Finally, we show how one would be able to establish NP-hardness of approximating the maximum winning probability under entangled strategies of a 4-player one-round game if one could strengthen appropriately one of the new quantum de Finetti theorems of this paper. This gives a new approach to this famous problem, which was only recently resolved [85].

---

[3]QMA is the quantum version of NP. QMA(2), in turn, is a version of QMA in which one is given two proofs, with the promise they are not entangled with each other; see section 2.3.

[4]By Ref. [46] we know QMA(2) with constant soundness gives at least a quadratic reduction in proof size relative to QMA, under plausible computational complexity assumptions; see section 2.3.

[5]See [22, 46] for more evidence this might be the case, along with obstacles to prove it.

[6]This algorithm was communicated to us already in 2010, although the result has appeared publicly only in [3].

Polynomial Optimization: We consider the connection [41, 10, 43] between quantum de Finetti theorems and the optimization over separable states, on one hand, and polynomial optimization and the Sum-of-Squares (Parrilo/Lasserre) hierachy, on the other hand, and prove that the optimization of certain degree-d polynomials over the n-dimensional hypersphere can be approximated to error ε in quasipolynomial-time in the number of variables by considering $O(\log(n)d^2\varepsilon^{-2})$ rounds of the Sum-of-Squares hierarchy of semidefinite programs. This result can be considered as an extension to the hypersphere of similar results for the simplex [76]. Moreover employing the result of Chen and Drucker [29], we show that $\Omega(d^2)$ rounds are necessary to obtain even a constant error-approximation, unless there are subexponential-time algorithms for SAT.

Separability Testing: Another application is to give an algorithm for deciding separability of multipartite states which is quasi-polynomial in the local dimensions of the subsystems. Given a multipartite state $\rho_{A_1,\ldots,A_l}$, we prove one can decide whether it is fully separable or ε-away from separable in time $\exp\left(O\left((\sum_k \ln|A_k|)^2 l^2 \varepsilon^{-2}\right)\right)$, with distance measured either by the one-way LOCC norm [21] or by a multipartite version of the Frobenius norm introduced in [63]. This generalizes the findings of [22] from bipartite states to general multipartite states, and vastly improves on the bound of [21].

Efficient State Tomography: A final application of the new de Finetti theorems is to quantum state tomography. The starting point is a result due to Aaronson [1], based on computational learning theory, showing that given an unknown n-qubit state one can perform tomography that allows us to compute to good accuracy the statistics of most observables by measuring only $O(n)$ i.i.d. copies of the state. The new de Finetti theorem we prove allow us to relax the assumption of having i.i.d. copies of the state (which can never be fully certified), showing that essentially the same conclusion holds true for arbitrary quantum states, as long as one can selects a few of its subsystems at random and performs the original scheme on them (weakening however the number of subsystems needed from $O(n)$ to $\text{poly}(n)$, of which only $O(n)$ are measured and the rest discarded).

Notation: Let $\mathcal{D}(\mathcal{H})$ be the set of quantum states on $\mathcal{H}$, i.e. positive semidefinite matrices of unit trace acting on the vector space $\mathcal{H}$. We say $\rho^{AB} \in \mathcal{D}(A \otimes B)$ is a k-extendible state if there is a state $\tilde{\rho}^{AB_1\ldots B_k} \in \mathcal{D}(A \otimes B^{\otimes k})$ such that $\tilde{\rho}^{AB_j} = \rho^{AB}$ for all $j \in [k]$. For a multipartite state such as ρ^{XY}, we use the convention that omitting subscripts corresponds to taking the partial trace over those systems; e.g. $\rho^X = \text{tr}_Y \rho^{XY}$ in the previous example. Let $\text{Sep}(A : B)$ denote the set of separable states in $\mathcal{D}(A \otimes B)$, which is defined to be the convex hull of the states of the form $\rho^A \otimes \rho^B$ (product states). Similarly $\text{Sep}(A^{\otimes l})$ is the convex hull of states of the form $\rho_1 \otimes \cdots \otimes \rho_l$. We say $\rho^{A_1\ldots A_k} \in \mathcal{D}(A^{\otimes k})$ is permutation symmetric if $\rho^{A_{\pi(1)}\ldots A_{\pi(k)}} = \rho^{A_1\ldots A_k}$ for any permutation $\pi \in S_k$ (with S_k the symmetric group of order k).

A quantum measurement (also called a POVM or positive-operator valued measure) is given by a set of matrices $\{M_k\}$ such that $M_k \geq 0$ and $\sum_k M_k = I$. We associate to any measurement a map $\Lambda(X) = \sum_k \text{tr}(M_k X)|k\rangle\langle k|$, with $\{|k\rangle\}$ an orthonormal basis. We denote the set of maps associated to measurements by $\mathcal{M}$. These are also called quantum-classical channels, since they map quantum states to probability distributions.

Let $p(x_1,\ldots,x_k|a_1,\ldots,a_k) \in \mathcal{X}^{\times k} \times \mathcal{A}^{\times k}$ be a conditional probability distribution. We say it is non-signaling if $p(x_j|a_j)$ is independent of a_k for $k \neq j$. We say $p(x,y|a,b)$ is k-extendible if there is a non-signaling distribution $p(x,y_1,\ldots,y_k|a,b_1,\ldots,b_k)$ which is permutation-symmetric in the B systems, i.e. $p(x,\pi^{-1}(y_1),\ldots,\pi^{-1}(y_k)|a,\pi^{-1}(b_1),\ldots,\pi^{-1}(b_k)) = p(x,y_1,\ldots,y_k|a,b_1,\ldots,b_k)$ for

5

all permutations $\pi \in S_k$, and whose marginal is $p(x, y|a, b)$. We call LHV (local hidden variable) the set of conditional probability distributions of the form $p(x, y|a, b) = \sum_l \pi_l q_l(x|a) r_l(y, |b)$ for a probability distribution π and local conditional distributions q_l, r_l.

2.1 Quantum de Finetti Theorems for Local Measurements

By monogamy of entanglement we expect that a k-extendible state ρ^{AB} to be close to a separable state, since the A subsystem is equally correlated to k systems. The next theorem gives a quantitative version of this fact both for entanglement and for non-signaling distributions.

Theorem 1.

1. Let $\rho^{AB} \in \mathcal{D}(A \otimes B)$ be a k-extendible state and $\mu(m)$ a distribution over quantum operations $\{\mathcal{E}_m^{A \to \tilde{A}}\}_m$, with $\mathcal{E}_m^{A \to \tilde{A}} : \mathcal{D}(A) \to \mathcal{D}(\tilde{A})$. Then

$$\min_{\sigma \in \mathrm{Sep}(A:B)} \max_{\Lambda^B \in \mathcal{M}} \mathbb{E}_{m \sim \mu} \left\| \mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB} \right) \right\|_1 \leq \sqrt{\frac{2 \ln |\tilde{A}|}{k}}. \tag{1}$$

2. Let $\rho^{AB} \in \mathcal{D}(A \otimes B)$ be a k-extendible state, $\mu(m)$ a distribution over operators $\{\mathcal{E}_m^{A \to \tilde{A}}\}_m$ from $\mathcal{D}(A) \to \mathcal{D}(\tilde{A})$ and Λ^B a measurement on $\mathcal{D}(B)$. Then in time $\mathrm{poly}(|A|, |B|^k)$ a classical computer can compute $\sigma \in \mathrm{Sep}(A : B)$ such that

$$\mathbb{E}_{m \sim \mu} \left\| \mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB} \right) \right\|_1 \leq \sqrt{\frac{2 \ln |\tilde{A}|}{k}}. \tag{2}$$

3. Let $p(x, y|a, b) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{A} \times \mathcal{B}$ be a k-extendible non-signaling conditional probability distribution and let μ be a distribution over $\mathcal{A}$. Then

$$\min_{q \in LHV} \max_{b \in \mathcal{B}} \mathbb{E}_{a \sim \mu} \|p(x, y|a, b) - q(x, y|a, b)\|_1 \leq \sqrt{\frac{2 \ln |X|}{k}}. \tag{3}$$

The de Finetti bound from Ref. [22] can be recovered (with an improved constant) as a special case of part 1 of Theorem 1 by choosing the singleton distribution composed of the ideal channel on A, since [7]

$$\|\rho^{AB} - \sigma^{AB}\|_{\mathrm{LOCC}^\leftarrow} = \max_{\Lambda \in \mathcal{M}} \|(\mathrm{id} \otimes \Lambda)(\rho^{AB} - \sigma^{AB})\|_1. \tag{4}$$

However, Theorem 1 improves on Ref. [22] in several ways. First and most importantly, the error term is independent of the subsystem dimensions of ρ^{AB}, and only depends on the output dimension of the family of quantum operations $\{\mathcal{E}_m^{A \to \tilde{A}}\}_m$, thus yielding nontrivial bounds even if A is infinite-dimensional. Likewise, for non-signaling distributions the bound in part 3 is independent of the number of measurement settings of $p(x, y|a, b)$. Second, if we think of k-extendable states as a relaxation of Sep, then part 2 provides an "explicit rounding," which did not exist in Ref. [22], although we note the caveat that σ depends on the measurement Λ^B. Third, part 3 generalizes the

---

[7] The one-way LOCC norm is defined as $\|X\|_{\mathrm{LOCC}^\leftarrow} = \max_{0 \leq M \leq I} \mathrm{tr}(XM)$, with the the maximization over all POVMs $\{M, I - M\}$ that can be realized by local operations and one-way classical communication from B to A.

result to non-signaling distributions. Note that in part 1, taking system $\tilde{A}$ to be classical would yield a special case of part 3, but in the more general case where $\tilde{A}$ is a quantum system, parts 1 and 3 are incomparable.

We remark that (4) also follows from the work of Yang [89] using the fact that the entanglement of formation [18] is upper-bounded by the log of either of the local dimensions, together with a variant of the Pinsker inequality adapted to LOCC$^{\leftarrow}$ [75]. It also follows from the recent work of Li and Winter [66]..

The proof of Theorem 1 (found in Section 3) is more direct and general than the proofs in [22, 89, 66], in particular not making use of entanglement measures in any explicit way. This enables us to obtain parts 2 and 3 of the theorem (but see the discussion of Conjecture 5 for an example of how the generality of Theorem 1 limits our abilities to further improve it). We remark that the explicit rounding in part 2 was mostly known only for the variants of the de Finetti theorem requiring $k \geq d$ [80, 32, 70, 43], and the previous de Finetti theorems for non-signaling boxes [33, 12] were similarly inefficient. The main exception to this is [11], which achieves a similar but incomparable bound for measurements with nonnegative matrix elements, together with an efficient rounding scheme. Ref. [11] was also an important source of inspiration for the current work.

The next theorem gives a generalization of the result of [22] to an arbitrary number of subsystems, as well as to non-signaling distributions.

Theorem 2.

1. Let $\rho^{A_1 \ldots A_k} \in \mathcal{D}(A^{\otimes k})$ be a permutation-invariant state. Then for every $0 \leq l \leq k$ there is a measure ν on $\mathcal{D}(A)$ such that

$$\max_{\Lambda_2, \ldots, \Lambda_l \in \mathcal{M}} \left\| (\mathrm{id} \otimes \Lambda_2 \otimes \ldots \otimes \Lambda_l) \left(\rho^{A_1 \ldots A_l} - \int \nu(d\sigma)\sigma^{\otimes l} \right) \right\|_1 \leq \sqrt{\frac{2l^2 \ln |A|}{k - l}}. \tag{5}$$

2. Let $p(X_1 \cdots X_k | A_1 \cdots A_k)$ be a permutation-invariant non-signaling conditional probability distribution (i.e. p is invariant under simultaneous permutation of the X and A systems). Fix a product distribution $\mu = \mu_1 \otimes \cdots \otimes \mu_k$ on $A_1 \times \cdots \times A_k$. Then for every $0 < l < k$ there is a measure ν on single-system conditional probability distributions such that

$$\mathbb{E}_{a_1, \ldots, a_l \sim \mu} \left\| p(X_1 \cdots X_l | a_1, \ldots, a_l) - \mathbb{E}_{q \sim \nu} q(X_1 | a_1) \otimes \cdots \otimes q(X_l | a_l) \right\|_1 \leq \sqrt{\frac{2l^2 \ln |X|}{k - l}} \tag{6}$$

In Ref. [39] Diaconis and Freedman proved that for a permutation-symmetric probability distribution p_k on k subsystems, p_l is $\frac{l(l-1)}{k}$-close (in variational distance) to a convex combination of i.i.d. probability distributions. Theorem 2 can be seen as an analogue of this result to quantum states and non-signaling probability distributions. However instead of having a bound which is independent of the dimension, we only have a bound that depends logarithmic on the dimension (and the notion of approximation is weaker than variational distance). It is an interesting question whether this can be improved. Note however that we give in Section 2.3 a computational complexity argument that the $k \geq \Omega(l^2)$ dependency is optimal.

Just as Theorem 1 yielded a stronger version of the BCY result [22] as a corollary, Theorem 2 leads to a multipartite version of (4). The main difference, apart from considering state on l systems that have symmetric k-partite extensions, is that the bipartite LOCC$^{\leftarrow}$ norm is replaced by

one in which parties $2, \ldots, l$ measure their systems and communicate the outcomes to party 1, who can then choose a measurement adaptively based on these messages. This leads to a norm on states that can be thought of as a multipartite generalization of the LOCC$^{\leftarrow}$ norm. We note that [21] also derived a multipartite generalization of [22] but with an exponentially worse scaling of k with l.

Proof idea: The proofs can be found in Section 3, but here we sketch the intuition behind part 1 of Theorem 1. If ρ^{AB} is k-extendible, then we can treat it instead as part of a state $\rho^{AB_1 \cdots B_k}$ with $\rho^{AB} = \rho^{AB_i}$ for each i. First examine systems A and B_1. If ρ^{AB_1} is approximately product, then ρ^{AB} is approximately separable and we are done. If not, then the correlations between A and B_1 mean that conditioning on B_1 will reduce the entropy of A. Then we can examine the mutual information between A and B_2 conditioned on B_1. Again, if this is small, then we have a nearly separable state and can stop, and if not, then we can condition on B_2 and further reduce the entropy of A. Since the initial entropy of A is at most $\log |A|$, this process is effective as long as $k \gg \log |A|$, which is a benefit of our information-theoretic approach over most previous versions of the de Finetti theorem. The main difficulty is that conditioning on a system does not work (indeed is not defined) if that system is quantum. This introduces the main subtlety, which is that we need to measure all of the B_i systems, but then use the outputs of the measurement to reason about the properties of the state before it was measured. Since the post-measured state is automatically separable, this requires some care.

2.2 Non-Local Games: Algorithms and Hardness Results

One application of Theorem 1 is to the computational complexity of non-local games. A multiprover game is played between a set of cooperative players/provers, who are not allowed to communicate with each other, and a referee/verifier who interrogates the provers to decide if they win the game. In a one-round game, for example, the verifier chooses questions to each prover at random and checks the answers obtained from the provers in order to decide whether to accept or not. Even though the provers cannot communicate with each other, they can agree on a common strategy in order to win the game with the maximum probability possible.

Multiprover games have had a central role in computational complexity theory. In a seminal paper Babai, Fortnow, and Lund proved NEXP = MIP [9], with MIP the class of languages having multi-prover interactive proof with a polynomial number of provers, rounds, and bits exchanged between the provers and the verifier in each round. Building on [9], it was then proven in [6, 7] that it is NP-hard to approximate to constant error the maximum winning probability of a two-player one-round game (with the input size given by the total number of questions to the players and their answers). This hardness result is equivalent to the celebrated PCP theorem [6, 7], which has a pivotal role in hardness of approximation results (see e.g. [4]).

It is natural to allow the players to share correlations that might assist them in winning the game with a higher probability. While it is easy to see that shared randomness is of no help, it has been known since the seminal work of Bell [15] that entanglement might help the players to win with a probability strictly larger than with a purely classical strategy. One can even consider stronger correlations than the ones allowed by quantum mechanics, such as arbitrary non-signaling correlations. Games in which the players can use entanglement (or more general non-signaling correlations) are known as non-local games, since the extra shared resources allow the players to sometimes use strategies that cannot be reproduced by local ones (i.e. strategies only using shared randomness and local actions). Upper bounds on the maximum winning prob-

ability of a one-round game under classical strategies are known as Bell inequalities, and non-local strategies that beat these bounds are known as Bell inequality violations. Such violations of Bell inequalities are central in the foundations of quantum mechanics as they can be implemented experimentally to show that nature cannot be described by a local hidden variable theory [8].

Given the usefulness of multiprover games to computational complexity theory and of non-local games to the foundations of quantum mechanics, it is interesting to study how difficult it is to compute the *entangled* value of the game, defined as the maximum probability of winning the game using entanglement, or the *non-signaling* value of the game, defined as the optimal probability under non-signaling strategies. By contrast, the maximum winning probability under classical strategies is called the *classical* value of the game.

Although *a priori* computing the entangled value of a game requires optimizing over a large set, in some cases this can be easier. Indeed, for unique games, the best known algorithms for the classical value [5] run in time $\exp(n^\varepsilon)$ (with $0 < \varepsilon < 1$ depending on the desired degree of approximation), whereas the entangled value of the game can be estimated in polynomial time using semidefinite programming [57] (or exactly calculated for the special case of XOR games [35]). These two classes of games could be taken as evidence that the estimation of the entangled value is generally easier than of the classical value. However if one is interested in a high-accuracy estimation this turns out not to be true. Kempe, Kobayashi, Matsumoto, Toner, and Vidick proved that it is NP-hard to approximate to an inverse polynomial (in the size of the game) the entangled value of one-round 3-prover games [56] (see also [53, 52, 34]). Recently in a beautiful development Ito and Vidick [54] proved that it is NP-hard, under quasi-polynomial reductions (improved to a polynomial-time reduction in Ref. [85]), to approximate the entangled value of 3-prover games with polynomially many rounds even to constant error. The result [54] has a more elegant formulation in terms of interactive proof systems: It shows that $\text{NEXP} \subseteq \text{MIP}^*$, with MIP^* the analogue of MIP in which the provers share entanglement [60]. The maximum probability of non-signaling strategies, in turn, can always be computed efficiently by linear programming [51].

Probably the biggest open question in this area is to determine the computational complexity of approximating the entangled value of one-round games to constant accuracy (although recent work of Vidick [85] has now resolved this in all but the case of two players). There are two reasons why this is a particular interesting setting. The first is the fact that the PCP theorem can be stated as the NP-hardness of approximating the classical value of one-round games to constant accuracy. Thus an analogous result for the entangled value could be interpreted as a version of the PCP theorem in the presence of entanglement. Second, in Bell inequality violation experiments, which are one-round non-local games, one can only obtain a constant-accuracy approximation to the true violation due to experimental error. Therefore it is important to understand how efficiently one can estimate to constant error the maximum violation of a Bell inequality, since this the most experimentally relevant approximation scale. One of our goals here is to propose a new approach to address this problem.

A particular class of games that we will consider are the so-called *free games*, defined as games in which the questions to each of the players are chosen independently from the questions to the other players [26]. A famous example from physics is the CHSH game. The fact that the verifier cannot coordinate questions suggests that the computation of the maximum winning probability of such games should not be as hard as for general games. And indeed Bellare, Feige and Killian proved that the analogue of MIP for poly-round free games is equal to PSPACE [16], while Aaronson, Impagliazzo, Moshkovitz, and Shor [3] proved that the classical value of one-round

free games with questions to the two provers in $Q \times Q$ and answers in $A_1 \times A_2$ can be simulated to within error ε by AM (Arthur-Merlin) proofs with an $O(\log |Q| + \log(|A_1| \cdot |A_2|)/\varepsilon)$-bit message from Arthur to Merlin and an $O(\log |A_1| \log |A_2|/\varepsilon)$-bit message from Merlin to Arthur. As a result, the value of such games can be estimated in time $\mathrm{poly}(\log |Q|) \exp(\log |A_1| \log |A_2|/\varepsilon)$. They also gave a matching hardness of approximation result for free games, showing that one can reduce 3-SAT on n binary variables to computing $\omega_c(G)$ to within constant additive error for 2-player one-round free games with $\exp(O(\sqrt{n}))$-sized answer alphabet[8].

As a corollary of Theorem 1 we will prove that the classical value of free games can be computed efficiently by linear programming, matching the run-time of the algorithm of [3]. Moreover, we will also derive a non-trivial hardness of approximation result for the *entangled* value of free games by importing to the case of entangled strategies the hardness of approximation result for the classical value of free games from [3] . Finally we will show how a conjectured strengthening of Theorem 1 would yield an alternate proof of the NP-hardness of obtaining a constant error approximation of ω_e for four-player one-round games.

Before we turn to the precise statement of the main result of this section let us give a more formal definition of non-local games.

Definition 3. *We define a m-prover game $G(m, \pi, V)$ by two parameters π and V:*

1. *π is a probability distribution on $Q_1 \times \ldots \times Q_m$ for finite sets $Q_1, \ldots, Q_m$.*

2. *V is a predicate on $Q_1 \times \ldots \times Q_m \times A_1 \times \ldots \times A_m$ for finite sets $A_1, \ldots, A_m$.*

The sets Q_i and A_i consist of the possible questions and answers, respectively, for player i. The predicate $0 \leq V(a_1, \ldots, a_m | q_1, \ldots, q_m) \leq 1$ is the pay-off function of the answer $(a_1, \ldots, a_m)$ given the question $(q_1, \ldots, q_m)$.

The *classical* value of the game G is given by

$$\omega_c(G(m, \pi, V)) := \max_{a_1, \ldots, a_m} \sum_{q_1, \ldots, q_m} \pi(q_1, \ldots, q_m) V(a_1(q_1), \ldots, a_m(q_m) | q_1, \ldots, q_m), \qquad (7)$$

where the maximum is over all functions $a_j : Q_j \to A_j$.

The *entangled* value of the game, in turn, is given by

$$\begin{aligned} &\omega_e(G(m, \pi, V)) \\ &:= \sup \sum_{q_1, \ldots, q_m} \pi(q_1, \ldots, q_m) \sum_{a_1, \ldots, a_m} V(a_1, \ldots, a_m | q_1, \ldots, q_m) \langle \psi | M^1_{a_1 | q_1} \otimes \ldots \otimes M^m_{a_m | q_m} | \psi \rangle, \end{aligned} (8)$$

where the supremum is over states $|\psi\rangle$ of arbitrary dimension and arbitrary POVMs

$$\{M^1_{a_1 | q_1}\}_{a_1 \in A_1}, \ldots, \{M^m_{a_m | q_m}\}_{a_m \in A_m}, \qquad (9)$$

with $\sum_{a_k \in A_k} M^k_{a_k | q_k} = I$ for every $q_k \in Q_k$ and $k \in [m]$.

Finally, the *non-signaling* value of the game G is defined as

$$\begin{aligned} &\omega_{ns}(G(m, \pi, V)) \\ &:= \max \sum_{q_1, \ldots, q_m} \pi(q_1, \ldots, q_m) \sum_{a_1, \ldots, a_m} V(a_1, \ldots, a_m | q_1, \ldots, q_m) p(a_1, \ldots, a_m | q_1, \ldots, q_m), \end{aligned} (10)$$

---

[8]There are suggestive similarities between this result and results about QMA(2) and variants thereof; see Section 2.3 and [46].

10

where the maximum is over all non-signaling probability distributions $p(a_1, \ldots, a_m | q_1, \ldots, q_m)$.

Corollary 4.

1. *Let $G(2, \pi, V)$ be a two-player one-round non-local free game with π a product probability distribution on $R \times Q$ and V a predicate on $R \times Q \times A \times B$. Then there is a $(m+1)$-player one-round non-local game $\overline{G}(m + 1, \overline{\pi}, \overline{V})$ with $\overline{\pi}$ a probability distribution on $R \times Q_1 \times \ldots \times Q_m$, with $|Q_k| = |Q|$ for $k \in [m]$, and $\overline{V}$ a predicate on $R \times Q_1 \times \ldots \times Q_m \times A \times B_1 \times \ldots \times B_m$, with $|B_k| = |B|$ for $k \in [m]$, such that*

$$\omega_c(G) = \omega_c(\overline{G}) \leq \omega_e(\overline{G}) \leq \omega_{ns}(\overline{G}) \leq \omega_c(G) + \sqrt{\frac{\ln|A|}{2m}}. \tag{11}$$

2. *For a free game $G(2, \pi, V)$ there is a linear-programming relaxation of size $|R||A| \, (|Q||B|)^{\frac{\ln|A|}{2\varepsilon^2}}$ for computing $\omega_c(G)$ to within additive error ε.*

3. *One can reduce 3-SAT on n variables to computing $\omega_e(G)$ to within constant additive error for $O(\sqrt{n})$-player one-round non-local games with answer alphabet size of $\exp(O(\sqrt{n}))$ in which only two players are asked questions.*

See Section 5 for the proof.

We note that it is trivial to prove either a version of part 3 of Corollary 4 in which the answer alphabet size is 2^n (in which case even one prover is clearly enough), or one in which the answer alphabet size is constant but one has n provers, or one with $\sqrt{n}$ provers and alphabet size $2^{\sqrt{n}}$ in which *all* provers respond. However, in our result, the total number of bits sent is $O(\sqrt{n})$.

Part 2 of Corollary 4 follows directly from part 1 and the fact that ω_{ns} can be computed by linear-programming. This gives a new algorithm matching the performance of the algorithm due to Aaronson, Impagliazzo, Moshkovitz, and Shor [3]. Part 3 of Corollary 4 follows from part 1 and the hardness of approximation result of Ref. [3] for free games.

Part 1 in turn gives a generic relation between the classical value of a free game, on one hand, and the quantum and non-signaling values of a modified game with more players, one the other hand. The idea of adding more players is to try to immunize the original game from entanglement (or general non-signaling correlations) by adding extra consistency tests that forces the entanglement between the players to have a specific form. Indeed the new game with $m+1$ players consists of playing the original game with player one and one of the remaining m players chosen at random. This essentially allows us to consider a two-player game where the provers can only share an m-extendible state (or m-extendible non-signaling conditional distribution). Then by Theorem 1 we obtain that this m-extendible state cannot be much better than a separable state or a local hidden variable distribution (which themselves are no better than just having shared randomness). The crucial aspect of Theorem 1 used here is that the error term only depends on the number of *outcomes* (which is given by the number of possible answers of the non-local game in question), and not on the dimension of the entangled state or on the number of different POVMs in the family in the quantum case (or the number of measurement settings in the non-signaling case). The idea of immunizing entanglement by introducing more players is not new and was used before by Kempe *et al* [56] to prove the hardness of estimating the entangled value within error inverse polynomial in the size of the game.

11

More generally, it was observed by Terhal, Doherty, and Schwab [84] that m-extendible states cannot violate any Bell inequality with fewer than m measurements for Bob (and an arbitrary number of measurements for Alice). In contrast Theorem 1 shows that a non-signaling m-extendible conditional distribution can violate a Bell inequality associated to a free game (an example of which is the CHSH inequality) with an arbitrary number of measurements, each with M possible outcomes, by at most $\frac{1}{2}\sqrt{\frac{2\ln(M)}{m}}$. This is an instance of the concept of monogamy of entanglement (which is known to hold true for non-signaling distributions as well [33]), in this case to the non-locality of quantum states (i.e. the maximum possible violation of a Bell inequality). Note that to be ε-close to a separable state in trace norm (thus having similar statistics under general quantum measurements) one must consider m-extendible states with $m = \Omega(|B|/\varepsilon)$, with $|B|$ the dimension of the B subsystem [32]. The monogamy of non-locality we find here, in comparison, has a bound that is *independent* of the dimension of the state.

Finally let us mention a conjecture whose validity would imply the NP-hardness of estimating ω_e to within constant error for 4-player one-round games. The conjecture is the following strengthening of Theorem 1.

Conjecture 5. *Let $\rho^{AB} \in \mathcal{D}(A \otimes B)$ be a k-extendible state and $\mu(m)$ a distribution over quantum operations $\{\mathcal{E}_m^{A\to\tilde{A}}\}_m$, with $\mathcal{E}_m^{A\to\tilde{A}} : \mathcal{D}(A) \to \mathcal{D}(\tilde{A})$. Then*

$$\min_{\sigma\in\mathrm{Sep}(A:B)} \mathbb{E}_{m\sim\mu} \max_{\Lambda^B\in\mathcal{M}} \left\|\mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB}\right)\right\|_1 \leq \sqrt{\frac{2\ln|\tilde{A}|}{k}}. \tag{12}$$

The difference with Theorem 1 is that the order of the expectation over μ and the maximization over measurements Λ^B is reversed. It is easy to check that one would be able to carry through the proof of part 1 of Corollary 4 given in Section 5 for general games (of course only for the relation of ω_e and ω_c). The fact that we would be able to prove NP-hardness for 4-player games would then follows from the combination of this stronger version of Eq. (11) with a recent version of the PCP theorem due to Khot and Safra, in the language of two-prover one-round games [59].

We have written (12) in a way that is meant to parallel (1) from Theorem 1, with a consequence that systems A and B are treated very differently. However, the conjecture could equivalently be restated in a more symmetric form. If we explicitly include the maximization over μ, then the LHS (12) becomes $\sup_\mu \min_\sigma [\mathbb{E}_{m\sim\mu} \max_{\Lambda^B} \| \cdots \|_1]$. Observe that the term inside the $[\cdots]$ is linear in μ and convex in σ; indeed, it is a seminorm of σ. Thus, we can use Sion's minimax theorem [82] and reverse the order of the $\sup_\mu$ and $\min_\sigma$. At this point the $\sup_\mu \mathbb{E}_{m\sim\mu}$ become superfluous, and we can replace the pair with simply a maximum over maps $\mathcal{E}^{A\to\tilde{A}}$. Thus, Conjecture 5 could equivalently be stated as

$$\min_{\sigma\in\mathrm{Sep}(A:B)} \max_{\mathcal{E}^{A\to\tilde{A}}} \max_{\Lambda^B\in\mathcal{M}} \left\|\mathcal{E}^{A\to\tilde{A}} \otimes \Lambda^B(\rho^{AB} - \sigma^{AB})\right\|_1 \leq \sqrt{\frac{2\ln|\tilde{A}|}{k}}. \tag{13}$$

Although the conjecture is consistent with all the examples of states we are aware of, we note that a proof would have to follow a very different approach to the one used in Theorem 1, as it cannot apply to non-signaling distributions. The reason is that the quantum version of the conjecture would imply that $\mathsf{NEXP} \subset \mathsf{MIP}^*(4,1)$, and the no-signaling version would imply that

NEXP $\subset$ MIP$^{\text{ns}}(4, 1)$, but this latter class is contained in EXP[9]. A scaled down version of this argument shows that the no-signaling version of Conjecture 5 would imply that P = NP. It is an interesting open question to find a more direct counter-argument, such as an example of a k-extendable no-signaling distribution whose difference from LHV distributions can be detected by correlated measurements.

Thus, despite the superficial similarity of (the quantum version of) Conjecture 5 with our Theorem 1, any proof will need to find features of quantum states that are not shared by no-signaling distributions. In this respect the hypothesis testing approach of Refs. [22, 66] might be a promising route.

2.3 Optimality of Chen and Drucker's Multiple-Proof Protocol for 3-SAT

One first application of Theorem 2 is to unentangled multiple proof systems.

Given a 3-SAT formula with n variables and $O(n)$ clauses, what is the minimum proof that can convince a verifier the formula is satisfiable? Under the exponential time hypothesis [48] – which says 3-SAT cannot be solved in subexponential time – $\Omega(n)$ bits are required, i.e. it is believed one cannot do anything substantially better than just write down the n-bit satisfying assignment. What if we can send a quantum state as a proof to a verifier who has a quantum computer to check its validity? Perhaps we could pack more information into the quantum state so that $o(n)$ qubits would be enough to convince the verifier? It turns out that assuming a quantum version of the exponential time hypothesis – namely that to solve 3-SAT takes exponential time even on a quantum computer (see e.g. [17] for the oracle version of this claim) – $\Omega(n)$ qubits are required [68].

Quantum mechanics allows us to add a new twist to this question. What if we want to convince a quantum verifier by sending a quantum state to her, but with the promise that parts of the quantum state are not entangled with each other? In this case the argument of Ref. [68] does not apply anymore and at least we do not have any implausible consequence for having a sublinear proof. And indeed Aaronson, Beigi, Drucker, Fefferman, and Shor [2] (building on [19]) proved that $\sqrt{n}\,\text{polylog}(n)$ unentangled quantum states, each of $\log(n)$ qubits, are enough to convince a quantum verifier that a 3-SAT instance with n variables and $O(n)$ clauses is satisfiable.

The result of [2] was strengthened in two directions: First Harrow and Montanaro [46] proved that *two* unentangled proofs, each of $\sqrt{n}\,\text{polylog}(n)$ qubits, are sufficient. Second Chen and Drucker [29] showed that $\sqrt{n}\,\text{polylog}(n)$ identical unentangled quantum proofs of $O(\log(n))$ qubits each are sufficient to convince even a verifier who measures each of the proofs separately and postprocesses the outcomes in order to decide whether to accept or not.

To state the main result of this section we define a few quantum complexity classes (see Section 6 for formal definitions). The first is a natural quantum analogue of NP (more precisely of MA). Let QMA$_n(c, s)$ be the class of problems such that: (i) for "yes" instances there is a quantum proof composed of n qubits that makes the verifier, who has access to polynomial quantum computation, to accept with probability at least c; and (ii) for "no" instances every proof is accepted with probability at most c. Let QMA$_n(m, c, s)$ be the analogue of QMA in which instead of one quantum proof the verifier receives m quantum proofs, each of n qubits, with the promise that they are not entangled with each other [61].

---

[9]The proof that MIP$^{\text{ns}}(\text{poly}, \text{poly}) \subseteq$ EXP is an easy application of linear programming (essentially the no-signaling constraints are linear constraints on an exponential-sized prover strategy) which appears not to have been published anywhere. Ref. [56] attribute it to a personal communication from Daniel Preda, and Ref. [51] builds on this approach to show that MIP$^{\text{ns}}(2, 1) \subseteq$ PSPACE by finding a way to parallelize the LP in the 2-prove 1-round case.

Further let $\mathsf{BellQMA}_n(m, c, s)$ be an analogue of $\mathsf{QMA}_n(m, c, s)$ in which the verification procedure is restricted to applying independent measurements to each of the m proofs and then post-processing the outcomes classically [2]. The name of the class comes from the fact that the verifier is basically constrained to apply a Bell test as his verification procedure. Finally let $\mathsf{BellSymQMA}_n(m, c, s)$ be the analogue of $\mathsf{BellQMA}_n(m, c, s)$ in which all the m proofs are promised to be identical.

With this notation the Chen-Drucker result can be stated as showing the containment of 3-SAT with n variables and $O(n)$ clauses in $\mathsf{BellSymQMA}_{\log(n)}(\sqrt{n}\operatorname{polylog}(n), 1 - 2^{-\Omega(\sqrt{n})}, 1/\operatorname{poly}(n))$ [29]. (An analogous, and incomparable, result holds for BellQMA also follows from [29].) A corollary of Theorem 2 is that this is essentially optimal, i.e. the square-root improvement found for the total proof size is all there is if we restrict ourselves to BellSymQMA protocols.

Corollary 6.

1. $\mathsf{BellSymQMA}_n(m, c, s) \subseteq \mathsf{QMA}_{10n^2 m^2/\varepsilon^2}(c, s + \varepsilon)$.

2. *For every $\varepsilon > 0$ and $c - s = \Omega(1)$, there is no* $\mathsf{BellSymQMA}_{O(\log(n))}(n^{\frac{1}{2} - \varepsilon}, c, s)$ *protocol for* 3-SAT *with n variables and $O(n)$ clauses, unless* 3-SAT *can be solved in* $\exp(n^{1 - 2\varepsilon}\operatorname{polylog}(n))$ *time.*

3. $\mathsf{BellQMA}_n(m, c, s) \subseteq \mathsf{QMA}_{10n^2 m^3/\varepsilon^2}(c, s + \varepsilon)$.

4. $\mathsf{QMA}_{\operatorname{poly}(n)}(\frac{2}{3}, \frac{1}{3}) = \mathsf{BellQMA}_{\operatorname{poly}(n)}(\operatorname{poly}(n), \frac{2}{3}, \frac{1}{3})$

See Section 6 for the proof.

In [20, 22] it was shown that $\mathsf{BellQMA}(m)$ is contained in QMA for a constant number of provers m. Corollary 6 strengthens the containment to even to a polynomial number of provers. This gives a new characterization of the class QMA and shows that the only advantage (in the regime where $c - s \geq 1/\operatorname{poly}(n)$) that BellQMA protocols can offer is a polynomial reduction in the proof size, such as in the protocol of [29].

Remark: In fact we can prove something slightly stronger than Corollary 6. Instead of Bell measurements, where k parties individually measure their systems and send the results to a referee, we can handle a slightly larger class of measurements. Our proofs apply equally to the setting where $k - 1$ parties measure their systems and send classical messages to the last party, who can choose a measurement adaptively based on these messages. This will follow from the fact that in part 1 of Theorem 2, we can leave one subsystem unmeasured. To keep the exposition simple, we will not formally state this improved version of Corollary 6.

2.4 Polynomial Optimization and Sum-of-Squares Proofs

Another application of our main theorems is to classical algorithms for maximizing polynomials over $\mathbb{C}^n$. The concepts of k-extendable and separable states turn out to correspond naturally to SDP hierarchies for polynomial optimization, and thus we are able to prove convergence of these hierarchies for polynomials that correspond to LOCC measurements. This connection was first established by Doherty, Parrilo and Spedalieri [41], and was more recently made quantitative for general polynomials over the unit sphere in $\mathbb{R}^n$ by Doherty and Wehner [42, 43].

In this section, we consider the problem of maximizing real-valued polynomial functions over the complex unit sphere $S^{2n-1} \subset \mathbb{C}^n$. More precisely, we consider polynomials of $z_1, \ldots, z_n, \bar{z}_1, \ldots, \bar{z}_n$ that are bihomogenous of degree d, d (i.e. homogenous of degree d in the

$z_1, \ldots, z_n$ and homogenous of degree d in the $\bar{z}_1, \ldots, \bar{z}_n$). This problem is closely related [37] to optimization over the real unit sphere, though not always identical [36]. When $d > 1$, this is generally NP-hard; see [38]. A promising general-purpose approximation scheme is to use an SDP hierarchy invented independently by Parrilo [73] and Lasserre [64]; see also [72] for a recent review of the complexity-theoretic properties of this hierarchy. To define the hierarchy, we introduce some notation. Let $\mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}] := \mathbb{C}[z_1, \ldots, z_n, \bar{z}_1, \ldots, \bar{z}_n]$ denote complex polynomials in n variables, let $\mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{d,d}$ denote the set of bihomogenous polynomials of degree d, d, and let $\mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_d^*$ denote the set of Hermitian linear functionals from $\mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_d$ to $\mathbb{R}$. Here we will consider only *Hermitian* linear functionals L, meaning that $L[\prod_{j=1}^n z_j^{a_j} \bar{z}_j^{b_j}] = L[\prod_{j=1}^n z_j^{b_j} \bar{z}_j^{a_j}]$ for any $a_1, \ldots, a_n, b_1, \ldots, b_n$.

If $p(\mathbf{z}) \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{d,d}$ and $k \geq d$, then we can upper bound $\max_{z \in S^{2n-1}} p(z)$ with the following SDP:

$$\max L(p) \text{ such that} \tag{14a}$$
$$L \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{k,k}^* \tag{14b}$$
$$L(1) = 1 \tag{14c}$$
$$L(q\bar{q}) \geq 0 \qquad \forall q \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{k,0} \tag{14d}$$
$$L((z_1\bar{z}_1 + \ldots + z_n\bar{z}_n)q) = L(q) \qquad \forall q \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{k-1,k-1} \tag{14e}$$

Here (14c) and (14d) are constraints that any collection of moments should satisfy (with (14b) enforcing linearity), while (14e) expresses the $\sum_{i=1}^n |z_i|^2 = 1$ constraint (and can in general be replaced with any polynomial constraint; see [73, 64, 72]). To see that (14) is an SDP, observe that (14e) is a linear constraint and (14d) is equivalent to the constraint that the moment matrix $M(L) \geq 0$, where the entries of $M(L)$ are indexed by monomials in $\mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_k$ and are defined by $M(L)_{\alpha,\beta} := L(\bar{z}^\alpha z^\beta)$. We can interpret this SDP as replacing the maximum over S^{2n-1} by a maximum over probability distributions over S^{2n-1} (which of course changes nothing), and in turn approximating this by considering only the moments of order $\leq k$. The dual of (14) is

$$\min \lambda \text{ such that} \tag{15a}$$
$$\lambda - p = \left(\sum_{i=1}^n z_i \bar{z}_i - 1 \right) q_0 + \sum_{i=1}^m q_i \bar{q}_i \tag{15b}$$
$$q_0 \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{k-1,k-1} \tag{15c}$$
$$q_1, \ldots, q_m \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{k,0} \tag{15d}$$

which can again be seen to be an SDP. This can be thought of as "proving" that $p(x) \leq \lambda$ by using the fact that $p(x) - \lambda$ is a sum of squares of polynomials; hence this SDP is also called the "sum-of-squares" hierarchy.

Under reasonable assumptions, as k grows this SDP converges to $\max_{z \in S^{2n-1}} p(z)$ as k grows [73, 64]. However, since the effort to compute (14) or (15) grows exponentially with k, it is important to determine the rate at which this convergence takes place. This rate is generally well-understood for optimizations over the simplex, but less is known for the sphere [38].

At first glance, the sum-of-squares hierarchy may appear unrelated to the quantum de Finetti theorems studied in this paper. However, the space $\mathbb{C}[\mathbf{z}]_k$ is isomorphic to the symmetric subspace of $(\mathbb{C}^n)^{\otimes k}$. Moreover, the relaxation in (14) is tight in the cases when L approximates the evaluation functional (i.e. $L(p) = p(z)$ for some $z \in \mathbb{C}^n$) on degree-d polynomials which is analogous to the

d-body marginals being approximately product. Indeed, this connection has been explored in [41], where the sum-of-squares hierarchy was used to prove that k-extendable states are approximately separable for sufficiently large k, and in [10], where this connection was used to find cases in which the sum-of-squares hierarchy yielded a good approximation of the $2 \to 4$ norm of a matrix.

To make the connection more explicit, we define, for any convex set K, the support function of K by

$$h_K(x) := \sup_{y \in K} \langle x, y \rangle. \tag{16}$$

For matrices x, y we define $\langle x, y \rangle := \operatorname{tr} x^\dagger y$. Then part 1 of Theorem 2 directly implies the following: Let M be a one-way LOCC operator of the form

$$M = \sum_{i_2,\dots,i_l} P_{i_2,\dots,i_l} \otimes Q_{2,i_2} \otimes \cdots \otimes Q_{l,i_l}, \tag{17}$$

with $0 \le P_{i_2,\dots,i_l} \le I$ for each $i_2,\dots,i_l$ and $0 \le \sum_{i_j} Q_{j,i_j} \le I$ for each $2 \le j \le l$. Defining $k\text{-Ext}(A^{\otimes l})$ to be the set of k-extendable l-partite states, we now have

$$h_{\operatorname{Sep}(A^{\otimes l})}(M) \le h_{k\text{-Ext}(A^{\otimes l})}(M) \le h_{\operatorname{Sep}(A^{\otimes l})}(M) + \sqrt{\frac{2l^2 \ln |A|}{k-l}}. \tag{18}$$

Given such an M and defining $|z\rangle := (z_1,\dots,z_n)$, we observe that $\langle z^{\otimes l}|M|z^{\otimes l}\rangle$ is a degree-l, l polynomial in z. As a result, we immediately obtain a bound on the ability of the sum-of-squares hierarchy to approximate certain polynomials over the complex hypersphere.

Corollary 7. *Let $p \in \mathbb{C}[\mathbf{z}, \bar{\mathbf{z}}]_{l,l}$ be of the form $p(z) = \langle z^{\otimes l}|M|z^{\otimes l}\rangle$ with $|z\rangle := (z_1,\dots,z_n)$ and M described by (17). Then*

$$\max_{\|z\|_2=1} p(z) \tag{19}$$

can be computed to within additive error ε by $O(\log(n)l^2/\varepsilon^2)$ levels of the sum-of-squares hierarchy.

Note that the result of Chen and Drucker [29] implies that $\log(n)l^{2-o(1)}$ levels of the sum-of-squares hierarchy are not sufficient to compute even a constant-error approximation to (19), for general p of the form described in the corollary, unless there is a subexponential time algorithm for 3-SAT.

There is also more direct evidence that Corollary 7 cannot be improved to yield a PTAS for polynomial optimization over the unit sphere. Ref. [25] proved that for any n, there exists a local measurement M (derived from a Bell inequality) on $n \times n$ systems such that

$$\frac{\operatorname{tr}(M\Phi_n)}{h_{\operatorname{Sep}}(M)} \ge \Omega\left(\frac{n}{\log^2(n)}\right), \tag{20}$$

with Φ_n the projector onto the n-dimensional maximally entangled state. Since $\rho := \frac{1}{k}\Phi_n + (1-\frac{1}{k})\frac{I}{n}$ is k-extendable, it follows that the k-extendable approximation can make multiplicative errors as large as $\Omega(\frac{n}{k\log^2(n)})$. Intriguingly, the example of [25] is based on the unique games problem. This suggests that using de Finetti theorems to give algorithms for unique games, as suggested by [10], will need to take advantage of the PPT condition in addition to merely the k-extendability property. The only previous evidence that using the PPT condition gives an asymptotic improvement over mere k-extendability was given by [70].

2.5 Testing Multipartite Separability

Another application of part 1 of Theorem 2, closely related to section 2.4, is to the quantum separability problem, a well-studied problem in quantum information theory of both theoretical and practical interest [50]. Given a multipartite state $\rho^{A_1 \cdots A_l}$ we say it is fully separable if

$$\rho^{A_1 \cdots A_l} = \sum_j p_j \sigma_j^{A_1} \otimes \ldots \otimes \sigma_j^{A_l}, \tag{21}$$

for a probability distribution $\{p_j\}$ and quantum states $\sigma_j^{A_i}$.

The goal in the weak-membership problem for separability is to decide whether a given multipartite state $\rho^{A_1 \cdots A_l}$ is separable or if it is ε-away from any separable state, given the promise that one of the two alternatives holds true. In fact one has a family of problems depending on which norm we choose to quantify the distance of quantum states. We consider two choices of norms. The first is the one-way LOCC norm, defined as

$$\|X\|_{\mathrm{LOCC}^\leftarrow} := \max_{\Lambda_2, \ldots, \Lambda_l} \|\mathrm{id} \otimes \Lambda_2 \otimes \ldots \otimes \Lambda_l(X)\|_1. \tag{22}$$

The name comes from the interpretation of norm as $\max_M \mathrm{tr}(MX)$, with M any POVM element that can be implemented by parties $2, \ldots, l$ measuring their systems locally and communicating the outcome to party 1, who then performs a measurement dependent on the information received. Therefore we have one-directional communication from all the parties to party 1.

The second is a multipartite version of the Forbenius norm recently introduced by Lancien and Winter [63]:

$$\|X\|_{2(l)} := \sqrt{\sum_{I \subseteq [l]} \mathrm{tr} \, |\mathrm{tr}_I X|^2}. \tag{23}$$

Corollary 8. *For some $c > 0$, the Sum-of-Squares hierarchy solves the weak membership problem for separability for the norm $\| * \|_{\mathrm{LOCC}^\leftarrow}$ in time*

$$\exp\left(c \left(\sum_j \log |A_j| \right)^2 l^2 \varepsilon^{-2} \right). \tag{24}$$

In turn, the Sum-of-Squares hierarchy solves the weak membership problem for separability for the norm $\| * \|_{2(l)}$ in time

$$\exp\left(c \left(\sum_j \log |A_j| \right)^2 (18)^{l/2} l^2 \varepsilon^{-2} \right). \tag{25}$$

See Section 7 for the proof.

We note this gives a generalization of the result of [22], which proved the same result for bipartite quantum states. A early generalization of [22] to multipartite states was given in [21]; however there only a bound of

$$\exp\left(c \log |A_1| \cdots \log |A_l| l^{2l-1} \varepsilon^{-2(l-1)} \right) \tag{26}$$

was obtained for the running time of the algorithm.

2.6 Pretty-Good Tomography in Permutation-Symmetric States

A final application of part 1 of Theorem 2 is to quantum state tomography, in which one obtains a description of an unknown quantum system by making measurements on the system. In quantum state tomography one tries to obtain a classical description of an unknown quantum state in the form of a density matrix for the state. By performing sufficiently many measurements of a sufficiently large number of different measurement settings one can obtain an arbitrarily good approximation of the true quantum state. Typically one considers a situation in which one has access to many i.i.d. copies of an unknown quantum state, and one performs measurements on those copies in order to learn the identity of the quantum state. Mathematically we can model this situation as saying that the global quantum state is of the form

$$\omega_n = \int \sigma^{\otimes n} \mu(d\sigma), \tag{27}$$

for an unknown measure μ on quantum states. However the assumption of having many i.i.d. copies of an unknown state cannot always be ensured, and in many situations it simply does not hold true. It is thus an important task to try to relax this requirement. It has long been realized [27] that quantum de Finetti theorems are exactly the right tool here. Instead of having to assume that ω_n has the form given by Eq. (27), one can merely assume that ω_n is the reduced state of a larger permutation-symmetric state ω_{n+k}. Then for k sufficiently large ω_n will be close to a convex combination of i.i.d. states. The point is that one can easily ensure the latter situation by selecting n subsystems at random from the $n + k$ available ones. Our work will allow this i.i.d. assumption to be relaxed. Indeed this was one of the original motivations for quantum de Finetti theorems [27].

The state of affairs is more complicated once complexity is taken into account. Since a quantum state of l qubits has 4^l parameters, reconstructing it generally requires $2^{O(l)}$ different measurement settings. However in many cases most of these parameters do not correspond to relevant questions. For instance, in order to predict expectation values of single-qubit observables, then a linear number of parameters suffices. Is there a way to explore this intuition in order to construct more efficient tomographic schemes?

One beautiful result in this direction was obtained by Aaronson in Ref. [1], using tools from computational learning theory [55], and can be roughly stated as follows: Given an arbitrary distribution $\mathcal{M}$ over measurements and an unknown quantum state on l qubits, $O(l)$ measurements settings are sufficient to get a density matrix which, with high probability over the measurement choice from $\mathcal{M}$, agrees with the expectation of the true quantum state up to small error. Thus a linear – in the number of qubits of the state – number of measurement settings are enough to get a density matrix which gives a good estimate to the statistics of the true state for almost all choices of measurements; one can perform a "pretty-good" tomography just with a linear number of measurement settings. The formal statement of Aaronson's result is as follows, restated slightly in order to facilitate our later extension of the result.

Lemma 9 (Theorem 1.3 of [1]). *Let $\omega_{m+n} \in \mathcal{D}(\mathcal{H}^{\otimes m+n})$ be a state of the form*

$$\omega_{m+n} = \int \nu(d\rho)\rho^{\otimes m+n},$$

for a probability measure ν on $\mathcal{D}(\mathcal{H})$. Let $\mathcal{M}$ be a distribution over two-outcome measurements on $\mathcal{H}$ and $\mathcal{E} = (E_1, \ldots, E_m)$ a training set of independently sampled measurements from $\mathcal{M}$. Suppose we measure

the first m systems of ω according to $\mathcal{E}$ and obtain outcomes $B = (b_1, \ldots, b_m) \in \{0,1\}^m$. For any outcome B, we will choose a hypothesis state

$$\sigma_B := \arg\min_\sigma \sum_{i=1}^m (\operatorname{tr}(E_i\sigma) - b_i)^2. \tag{28}$$

Then there exists a constant $K > 0$ such that if

$$m \geq \frac{K}{\gamma^4\varepsilon^2} \left(\frac{\log|\mathcal{H}|}{\gamma^4\varepsilon^2} \log^2 \frac{1}{\gamma\varepsilon} + \log\frac{1}{\delta} \right), \tag{29}$$

then with probability at least $1 - \delta$ the post-measured state $\tilde{\omega}_n$ satisfies

$$\tilde{\omega}_n = \int \rho^{\otimes n} \mu(d\rho), \tag{30}$$

where the measure μ only has non-zero support on states ρ such that

$$\Pr_{E \in \mathcal{M}} \left[|\operatorname{tr}(E\rho) - \operatorname{tr}(E\sigma_B)| > \gamma \right] \leq \varepsilon. \tag{31}$$

A limitation of Aaronson's result [1], common of other tomographic schemes as well, is the assumption that one is given several i.i.d. copies of the unknown quantum state. Here too one could try to apply the standard quantum de Finetti theorems [62, 32, 80] to find a way around this assumption. However since the error in those depend polynomially on the dimension of the state, one would obtain a non-trivial result only if one would select subsystems at random from a state of $2^{O(l)}$ subsystems, which is not a reasonable assumption. Theorem 2 allows us to circumvent this problem.

Corollary 10. *Let $\omega_{m+n+k} \in \mathcal{D}(\mathcal{H}^{\otimes m+n+k})$ be a permutation-symmetric state, let $\mathcal{M}$ be a distribution over two-outcome measurements on $\mathcal{H}$, and let $\mathcal{E} = (E_1, \ldots, E_m)$ be a training set consisting of m measurements drawn independently from $\mathcal{M}$. Suppose we discard the last k systems, measure the first m systems of ω according to $\mathcal{E}$ and obtain outcomes $B = (b_1, \ldots, b_m) \in \{0,1\}^m$. For any outcome B, we will choose a hypothesis state*

$$\sigma_B := \arg\min_\sigma \sum_{i=1}^m (\operatorname{tr}(E_i\sigma) - b_i)^2. \tag{32}$$

Fix error parameters $\varepsilon, \eta, \gamma, \nu > 0$. Suppose that (for some universal constant $K > 0$) we have

$$m \geq \frac{K}{\gamma^4\varepsilon^2} \left(\frac{\log|\mathcal{H}|}{\gamma^4\varepsilon^2} \log^2 \frac{1}{\gamma\varepsilon} + \log\frac{1}{\delta} \right), \tag{33}$$

$$k \geq \frac{4(m+n)^2 \ln|\mathcal{H}|}{\nu^2}. \tag{34}$$

Then with probability at least $1 - \delta$ the post-measured state $\tilde{\omega}_n$ satisfies

$$\max_{\Lambda_1,\ldots,\Lambda_n} \left\| \Lambda_1 \otimes \ldots \otimes \Lambda_n \left(\tilde{\omega}_n - \int \rho^{\otimes n} \mu(d\rho) \right) \right\|_1 \leq \nu, \tag{35}$$

with the maximum over quantum-classical channels $\Lambda_1, \ldots, \Lambda_n$. Here the measure μ only has non-zero support on states ρ such that

$$\Pr_{E \in \mathcal{M}} \left[|\operatorname{tr}(E\rho) - \operatorname{tr}(E\sigma_B)| > \gamma \right] \leq \varepsilon \tag{36}$$

The proof of Corollary 10 follows immediately from part 1 of Theorem 2 and Lemma 9.

Let us say a few words about the interpretation of the result. Suppose we had Eq. (35) with $\nu = 0$. Then

$$\tilde{\omega}_n = \int \rho^{\otimes n} \mu(d\rho), \tag{37}$$

with μ a measure with non-zero support only on states ρ that, for *most* measurements on $\mathcal{M}$, gives approximately the same statistics as any state σ_B compatible with the observed data (in the sense that it satisfies Eq. (32)). Therefore any state σ_B compatible with the measured data can be used correctly to infer the statistics of future measurements, with high probability over the choice of the observable. For non-zero ν we have a similar situation. While the state $\tilde{\omega}_n$ might be very far away from a convex combination of i.i.d. in trace norm, if we only consider the statistics of local measurements on the n subsystems, then, up to error ν, we have the same conclusions as in the case of $\nu = 0$.

The price we have to pay for being able to relax the assumption of having i.i.d. copies of the state is that instead of starting from $O(\log |\mathcal{H}|) + n$ copies of the state, now we need a global state with $O\left((n + \log |\mathcal{H}|)^2 \log |\mathcal{H}|\right)$ subsystems (of which we only measure $O(\log |\mathcal{H}|)$ of them). The main point is that this is still polynomial in the number of qubits of the unknown state one wants to learn.

We note that while this approach gives an efficient alternative for tomography of states on a large number of qubits in what concerns the number of measurements needed, it says nothing about the computational complexity of finding the hypothesis state σ_B. As noted in [1], it is an interesting problem to determine for which classes of states one can obtain ρ_v efficiently.

3   Proof of Theorem 1

We will prove Theorem 1 by information-theoretic techniques, inspired by [11] and Lemma 4.5 of [78]. Given two quantum states $\rho, \sigma \in \mathcal{D}(\mathcal{H})$, we define the quantum relative entropy (or quantum Kullback-Leibler divergence) as

$$S(\rho||\sigma) := \mathrm{tr}(\rho(\ln(\rho) - \ln(\sigma))). \tag{38}$$

Given a bipartite state $\rho^{AB} \in \mathcal{D}(A \otimes B)$ we define the mutual information as

$$I(A : B)_\rho := S(\rho^{AB}||\rho^A \otimes \rho^B). \tag{39}$$

Given a tripartite qqc state of the form $\rho^{ABK} := \sum_k p_k \rho_k^{AB} \otimes |k\rangle\langle k|^K$ we define the conditional mutual information as

$$I(A : B|X)_\rho := \sum_k p_k I(A : B)_{\rho_k}. \tag{40}$$

The mutual information satisfies the following properties that will be useful in the proof:

Lemma 11.

1. Chain Rule:

$$I(A : BX) = I(A : X) + I(A : B|X) \tag{41}$$

2. *Monotonicity under Local Operations: Let $\pi_{AB} = \mathrm{id} \otimes \Lambda(\rho^{AB})$, then*

$$I(A:B)_\pi \le I(A:B)_\rho \tag{42}$$

3. *Pinsker's Inequality:*

$$I(A:B)_\rho \ge \frac{1}{2}\|\rho^{AB} - \rho^A \otimes \rho^B\|_1^2. \tag{43}$$

(The absence of the usual $\ln(2)$ factor in (43) is because of our convention that entropies are measured in "nats," i.e. with logs taken base e.)

We are now ready to prove Theorem 1:

Theorem 1 (restatement).

1. *Let $\rho^{AB} \in \mathcal{D}(A \otimes B)$ be a k-extendible state and $\mu(m)$ a distribution over quantum operations $\{\mathcal{E}_m^{A\to\tilde{A}}\}_m$, with $\mathcal{E}_m^{A\to\tilde{A}} : \mathcal{D}(A) \to \mathcal{D}(\tilde{A})$. Then*

$$\min_{\sigma\in\mathrm{Sep}(A:B)} \max_{\Lambda^B\in\mathcal{M}} \mathbb{E}_{m\sim\mu} \left\| \mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB} \right) \right\|_1 \le \sqrt{\frac{2\ln|\tilde{A}|}{k}}. \tag{44}$$

2. *Let $\rho^{AB} \in \mathcal{D}(A \otimes B)$ be a k-extendible state, $\mu(m)$ a distribution over operators $\{\mathcal{E}_m^{A\to\tilde{A}}\}_m$ from $\mathcal{D}(A) \to \mathcal{D}(\tilde{A})$ and Λ^B a measurement on $\mathcal{D}(B)$. Then in time $\mathrm{poly}(|A|, |B|^k)$ a classical computer can compute $\sigma \in \mathrm{Sep}(A:B)$ such that*

$$\mathbb{E}_{m\sim\mu} \left\| \mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB} \right) \right\|_1 \le \sqrt{\frac{2\ln|\tilde{A}|}{k}}. \tag{45}$$

3. *Let $p(x,y|a,b) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{A} \times \mathcal{B}$ be a k-extendible non-signaling conditional probability distribution and let μ be a distribution over $\mathcal{A}$. Then*

$$\min_{q\in\mathrm{LHV}} \max_{b\in\mathcal{B}} \mathbb{E}_{a\sim\mu} \|p(x,y|a,b) - q(x,y|a,b)\|_1 \le \sqrt{\frac{2\ln|X|}{k}}. \tag{46}$$

Proof. The three parts of the theorem have similar proofs.
Part 1:

Define the states

$$\pi_{\tilde{A}B_1\ldots B_k M} := \mathbb{E}_{m\sim\mu} \pi_m \otimes |m\rangle\langle m|^M \tag{47}$$

$$\pi_m^{\tilde{A}B_1\ldots B_k} := \left(\mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda_1^{B_1} \otimes \ldots \otimes \Lambda_k^{B_k} \right) (\rho^{AB_1\ldots B_k}),$$

with $\mathcal{E}_m^{A\to\tilde{A}}$ quantum operations from A to $\tilde{A}$, Λ_i quantum-classical channels, and $|m\rangle$ a classical label for which quantum operation $\mathcal{E}_m^{A\to\tilde{A}}$ was applied. Repeatedly applying the chain rule (41), we find

$$I(\tilde{A}:B_1\ldots B_k|M) = I(\tilde{A}:B_1|M) + I(\tilde{A}:B_2|MB_1) + \ldots + I(\tilde{A}:B_k|MB_1\ldots B_{k-1}). \tag{48}$$

21

Now we maximize over measurements and obtain

$$\max_{\Lambda_1,\ldots,\Lambda_k \in \mathcal{M}} I(\tilde{A} : B_1 \ldots B_k | M)_\pi = \tag{49}$$

$$\max_{\Lambda_1,\ldots,\Lambda_{k-1} \in \mathcal{M}} \left(I(\tilde{A} : B_1 | M)_\pi + \ldots + I(\tilde{A} : B_{k-1} | M B_1 \ldots B_{k-2})_\pi + \max_{\Lambda_k \in \mathcal{M}} I(A : B_k | M B_1 \ldots B_{k-1})_\pi \right).$$

Now

$$I(A : B_k | M B_1 \ldots B_{k-1})_\pi = \mathbb{E}_{m \sim \mu} I(A : B_k | B_1 \ldots B_{k-1})_{\pi_m}. \tag{50}$$

Since the $B_1 \ldots B_{k-1}$ systems of π_m are classical, we can write the state of ρ^{AB_k} as an average over them, namely

$$\rho^{AB} = \rho^{AB_k} = \sum_i q_i \rho_i^{AB_k}, \tag{51}$$

where $\{q_i, \rho_i\}$ depend on $\Lambda_1, \ldots, \Lambda_{k-1}$ but not on $\mathcal{E}_m$ and Λ_k. Then define

$$\pi_{i,m}^{AB_k} := \left(\mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^{B_k} \right) (\rho_i^{AB_k}), \tag{52}$$

so that $\pi_m^{AB_k} = \sum_i q_i \pi_{i,m}^{AB_k}$ and

$$I(A : B_k | B_1 \ldots B_{k-1})_{\pi_m} = \sum_i q_i I(A : B_k)_{\pi_{i,m}} \tag{53}$$

By Pinsker's inequality

$$I(A : B_k | B_1 \ldots B_{k-1})_{\pi_m} \geq \frac{1}{2} \sum_i q_i \left\| \mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^{B_k} \left(\rho_i - \rho_i^A \otimes \rho_i^{B_k} \right) \right\|_1^2, \tag{54}$$

where ρ_i^A and $\rho_i^{B_k}$ are the A and B_k reduced states of ρ_i.

By convexity of x^2 and the trace norm

$$I(A : B_k | B_1, \ldots, B_{k-1})_{\pi_m} \geq \frac{1}{2} \left\| \mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^{B_k} \left(\rho^{AB} - \sum_i q_i \rho_i^A \otimes \rho_i^{B_k} \right) \right\|_1^2. \tag{55}$$

Using Eq. (50)

$$\max_{\Lambda_k \in \mathcal{M}} I(A : B_k | M B_1 \ldots B_{k-1})_\pi \geq \frac{1}{2} \max_{\Lambda_k \in \mathcal{M}} \mathbb{E}_{m \sim \mu} \left\| \mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^{B_k} \left(\rho^{AB_k} - \sum_i q_i \rho_i^A \otimes \rho_i^{B_k} \right) \right\|_1^2 \tag{56}$$

$$\geq \frac{1}{2} \min_{\sigma \in \mathrm{SEP}(A:B_k)} \max_{\Lambda_k \in \mathcal{M}} \mathbb{E}_{m \sim \mu} \left\| \mathcal{E}_m^{A \to \tilde{A}} \otimes \Lambda^{B_k} (\rho - \sigma) \right\|_1^2.$$

Note that the second line is independent of $\Lambda_1, \ldots, \Lambda_{k-1}$, since only the ensemble $\{q_i, \rho_i\}$ depended on them.

From (49) and (56),

$$\max_{\Lambda_1,\ldots,\Lambda_k\in\mathcal{M}} I(A:B_1\ldots B_k|M)_\pi \tag{57}$$

$$\geq \max_{\Lambda_1,\ldots,\Lambda_{k-1}\in\mathcal{M}} \sum_{j=1}^{k-1} I(A:B_j|MB_1\ldots B_{j-1})_\pi$$

$$+ \frac{1}{2} \min_{\sigma\in\mathrm{SEP}(A:B_k)} \max_{\Lambda_k\in\mathcal{M}} \mathbb{E}_{m\sim\mu} \left\| \mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda_k (\rho-\sigma) \right\|_1^2.$$

Applying the same argument sequentially to all the remaining conditional mutual informations we find

$$\frac{k}{2} \min_{\sigma\in\mathrm{SEP}(A:B)} \max_{\Lambda\in\mathcal{M}} \mathbb{E}_{m\sim\mu} \left\| \mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB} \right) \right\|_1^2 \leq \max_{\Lambda_1,\ldots,\Lambda_k} I(A:B_1\ldots B_k|M)_\pi \leq \ln|\tilde{A}|, \tag{58}$$

where we used that $\pi^{\tilde{A}} = \mathbb{E}_{m\sim\mu}\left(\mathcal{E}_m^{A\to\tilde{A}}(\rho^A)\right) \in \mathcal{D}(\tilde{A})$. Finally by convexity of x^2,

$$\left(\min_{\sigma\in\mathrm{SEP}(A:B)} \max_{\Lambda^B\in\mathcal{M}} \mathbb{E}_{m\sim\mu} \left\| \mathcal{E}_m^{A\to\tilde{A}} \otimes \Lambda^B \left(\rho^{AB} - \sigma^{AB} \right) \right\|_1 \right)^2 \leq \frac{2\ln|X|}{k}, \tag{59}$$

and we are done with the proof of part 1.

Part 2: The proof of part 2 is mostly the same as that of part 1, and so we only give a brief outline of the changes. The main change is to omit the maximizations over $\Lambda_1,\ldots,\Lambda_k$, instead using only the fixed measurement Λ. We also set $\sigma = \sum_i q_i \rho_i^A \otimes \rho_i^{B_k}$ rather than performing a minimization. As a result, the calculations require only time polynomial in the dimensions of the relevant states.

Part 3: The proof of part 3 is similar to that of part 1, except that we need to make the following replacements.

Part 1	Part 3
quantum states $\rho^{AB_1\ldots B_k}$	non-signaling distributions $p(x,y_1,\ldots,y_k|a,b_1,\ldots,b_k)$
quantum mutual information	classical mutual information maximized over choices of measurements $a,b_1,\ldots b_k$
partial trace	no-signaling condition

For brevity we will use the abbreviations $b^k := (b_1,\ldots,b_k)$, $b^{k-1} = (b_1,\ldots,b_{k-1})$ and so on. In more detail, the analogue of (47) is to define the non-signaling distribution π from $\mathcal{B}^k \to \mathcal{X} \times \mathcal{Y}^k$:

$$\pi(x,y^k,a|b^k) = \mu(a)p(x,y^k|a,b^k) \tag{60}$$

We can also define

$$\pi(x,y^{k-1},a|b^{k-1}) = \mu(a)p(x,y^{k-1}|a,b^k), \tag{61}$$

and, thanks to the no-signaling property of p, this is well-defined, since the RHS does not depend on b_k.

Again the chain rule gives us an analogue of (49).

$$\max_{b^k \in \mathcal{B}^k} I(X : Y^k | A)_{\pi(\cdot|b^k)}$$

$$= \max_{b^{k-1} \in \mathcal{B}^{k-1}} \left(\sum_{j=1}^{k-1} I(X : Y_j | AY^{j-1})_{\pi(\cdot|b^{k-1})} + \max_{b_k \in \mathcal{B}} I(X : Y_k | AY^{k-1})_{\pi(\cdot|b^k)} \right) \qquad (62)$$

Again we focus on the last term of Eq. (62). Define $i := (a, b^k, y^{k-1})$, and compute

$$\max_{b_k \in \mathcal{B}} I(X : Y_k | AY^{k-1})_{\pi(\cdot|b^k)} = \max_{b_k \in \mathcal{B}} \mathop{\mathbb{E}}_{a \sim \mu} I(X : Y_k | Y^{k-1})_{p(\cdot|a, b^k)}$$

$$= \max_{b_k \in \mathcal{B}} \mathop{\mathbb{E}}_{a \sim \mu} \sum_{y^{k-1}} p(y^{k-1}|b^{k-1}) I(X : Y_k)_{p(\cdot|i)}$$

$$= \frac{1}{2} \max_{b_k \in \mathcal{B}} \mathop{\mathbb{E}}_{a \sim \mu} \sum_{y^{k-1}} p(y^{k-1}|b^{k-1}) \left(\sum_{\substack{x \in \mathcal{X} \\ y_k \in \mathcal{Y}}} |p(x, y_k|i) - p(x|i)p(y_k|i)| \right)^2 \qquad \text{Pinsker}$$

$$\geq \frac{1}{2} \max_{b_k \in \mathcal{B}} \mathop{\mathbb{E}}_{a \sim \mu} \left(\sum_{y^{k-1}} p(y^{k-1}|b^{k-1}) \sum_{\substack{x \in \mathcal{X} \\ y_k \in \mathcal{Y}}} |p(x, y_k|i) - p(x|i)p(y_k|i)| \right)^2 \qquad \text{convexity of } x \mapsto x^2$$

$$\geq \frac{1}{2} \max_{b_k \in \mathcal{B}} \mathop{\mathbb{E}}_{a \sim \mu} \left(\sum_{\substack{x \in \mathcal{X} \\ y_k \in \mathcal{Y}}} \left| p(x, y_k|a, b^k) - \sum_{y^{k-1}} p(y^{k-1}|b^{k-1})p(x|i)p(y_k|i) \right| \right)^2 \qquad \text{convexity of } \| \cdot \|_1$$

$$\geq \frac{1}{2} \min_{q \in \text{LHV}} \max_{b \in \mathcal{B}} \mathop{\mathbb{E}}_{a \sim \mu} \|p(X, Y_k|a, b) - q(X, Y_k|a, b)\|_1^2$$

As with part 1, we can repeatedly apply this inequality to (62) in order to prove the theorem.
□

4 Proof of Theorem 2

For a state $\rho^{A_1 \ldots B_k}$ we define the multipartite mutual information

$$I(A_1 : \ldots : A_k) := S(\rho^{A_1 \ldots A_k} \| \rho^{A_1} \otimes \ldots \otimes \rho^{A_k}) = S(A_1) + \ldots + S(A_k) - S(A_1 \ldots A_k). \qquad (63)$$

For a quantum-classical state $\rho^{A_1 \ldots A_k R} = \sum_i p_i \rho_i^{A_1 \ldots A_k} \otimes |i\rangle\langle i|^R$ we define the conditional multipartite mutual information as follows

$$I(A_1 : \ldots : A_k | R)_\rho := \sum_i p_i I(A_1 : \ldots : A_k)_{\rho_i}. \qquad (64)$$

The multipartite mutual information satisfies the following properties:

Lemma 12.

24

1. *Multipartite-to-Bipartite [90]:*

$$I(A_1 : \ldots : A_k|R) = I(A_1 : A_2|R) + I(A_1 A_2 : A_3|R) + \ldots + I(A_1 \ldots A_{k-1} : A_k|R). \quad (65)$$

2. *Monotonicity under Local Operations: Let* $\pi^{A_1 \ldots A_k} = \Lambda^{A_1} \otimes \mathrm{id}^{A_2 \ldots A_k}(\rho^{A_1 \ldots A_k})$, *then*

$$I(A_1 : \ldots A_k)_\pi \leq I(A_1 : \ldots : A_k)_\rho \quad (66)$$

3. *Pinsker's Inequality:*

$$I(A_1 : \ldots : A_k)_\rho \geq \frac{1}{2} \|\rho^{A_1 \ldots A_k} - \rho^{A_1} \otimes \ldots \otimes \rho^{A_k}\|_1^2. \quad (67)$$

Theorem 2 (restatement).

1. *Let* $\rho^{A_1 \ldots A_k} \in \mathcal{D}(A^{\otimes k})$ *be a permutation-invariant state. Then for every* $0 \leq l \leq k$ *there is a measure* ν *on* $\mathcal{D}(A)$ *such that*

$$\max_{\Lambda_2, \ldots, \Lambda_l \in \mathcal{M}} \left\| (\mathrm{id} \otimes \Lambda_2 \otimes \ldots \otimes \Lambda_l) \left(\rho^{A_1 \ldots A_l} - \int \nu(d\sigma)\sigma^{\otimes l} \right) \right\|_1 \leq \sqrt{\frac{2l^2 \ln |A|}{k - l}}. \quad (68)$$

2. *Let* $p(X_1 \cdots X_k|A_1 \cdots A_k)$ *be a permutation-invariant non-signaling conditional probability distribution (i.e.* p *is invariant under simultaneous permutation of the* X *and* A *systems). Fix a product distribution* $\mu = \mu_1 \otimes \cdots \otimes \mu_k$ *on* $A_1 \times \cdots \times A_k$. *Then for every* $0 < l < k$ *there is a measure* ν *on single-system conditional probability distributions such that*

$$\mathbb{E}_{a_1, \ldots, a_l \sim \mu} \left\| p(X_1 \cdots X_l|a_1, \ldots, a_l) - \mathbb{E}_{q \sim \nu} q(X_1|a_1) \otimes \cdots \otimes q(X_l|a_l) \right\|_1 \leq \sqrt{\frac{2l^2 \ln |X|}{k - l}} \quad (69)$$

Proof.
Part 1:
   Let

$$\pi^{A_1 \ldots A_l R} := (\mathrm{id}^{A_1} \otimes \Lambda_2 \otimes \ldots \otimes \Lambda_l \otimes \mathcal{E}^{A_{l+1} \ldots A_k})(\rho^{A_1 \ldots A_k}), \quad (70)$$

with $\Lambda_j : \mathcal{D}(A) \to \mathcal{D}(X)$ and $\mathcal{E} : \mathcal{D}(A^{\otimes k-l}) \to \mathcal{D}(R)$ quantum-classical channels. Then from Eq. (65) of Lemma 12,

$$\min_{\mathcal{E}} \max_{\Lambda_2, \ldots, \Lambda_l} I(A_1 : \ldots : A_l|R)_\pi = \min_{\mathcal{E}} \max_{\Lambda_2, \ldots, \Lambda_l} \sum_{j=2}^{l} I(A_1 \ldots A_{j-1} : A_j|R)_\pi$$

$$\leq \min_{\mathcal{E}} \max_{\Lambda_2, \ldots, \Lambda_l} \sum_{j=2}^{l} I(A_1 \ldots A_{j-1} : A_j|R)_{\pi_j} \quad (71)$$

with

$$\pi_j := (\mathrm{id}^{A_1 \ldots A_{j-1}} \otimes \Lambda_j \otimes \mathrm{id}^{A_{j+1} \ldots A_l} \otimes \mathcal{E}^{A_{l+1} \ldots A_k})(\rho^{A_1 \ldots A_k}). \quad (72)$$

The last inequality in Eq. (71) follows by the monotonicity of the mutual information under local operations (Eq. (66) of Lemma 12). Then

$$
\begin{aligned}
\min_{\mathcal{E}} \max_{\Lambda_2,\ldots,\Lambda_l} I(A_1 : \ldots : A_l | R)_\pi \; &\leq \; \min_{\mathcal{E}} \max_{\Lambda_2,\ldots,\Lambda_l} \sum_{j=2}^{l} I(A_1 \ldots A_{j-1} : A_j | R)_{\pi_j} \\
&= \; \min_{\mathcal{E}} \sum_{j=2}^{l} \max_{\Lambda_j} I(A_1 \ldots A_{j-1} : A_j | R)_{\pi_j} \\
&\leq \; \min_{\mathcal{E}} [(l-1) \max_{\Lambda_l} I(A_1 \ldots A_{l-1} : A_l | R)_{\pi_l}], \qquad (73)
\end{aligned}
$$

where the last inequality follows from the monotonicity of mutual information under tracing out and the permutation invariance of the state $\rho^{A_1 \ldots A_k}$.

We claim that

$$
\min_{\mathcal{E}} \max_{\Lambda_l} I(A_1 \ldots A_{l-1} : A_l | R)_{\pi_l} \leq \frac{(l-1) \ln |A|}{k - l + 1}. \qquad (74)
$$

Indeed, defining $\nu^{A_1 \ldots A_k} := (\mathrm{id}^{A_1 \ldots A_{l-1}} \otimes \Lambda_l \otimes \ldots \otimes \Lambda_k)(\rho^{A_1 \ldots A_k})$, for quantum-classical channels Λ_j, we have

$$
\begin{aligned}
&\max_{\Lambda_l,\ldots,\Lambda_k} I(A_1 \ldots A_{l-1} : A_l \ldots A_k)_\nu \\
&= \max_{\Lambda_l,\ldots,\Lambda_k} \sum_{j=l}^{k} I(A_1 \ldots A_{l-1} : A_j | A_{j+1} \ldots A_k)_\nu \\
&= \max_{\Lambda_{l+1},\ldots,\Lambda_k} \left(\sum_{j=l+1}^{k} I(A_1 \ldots A_{l-1} : A_j | A_{j+1} \ldots A_k)_\nu + \max_{\Lambda_l} I(A_1 \ldots A_{l-1} : A_l | A_{l+1} \ldots A_k)_\nu \right) \\
&\geq \max_{\Lambda_{l+1},\ldots,\Lambda_k} \left(\sum_{j=l+1}^{k} I(A_1 \ldots A_{l-1} : A_j | A_{j+1} \ldots A_k)_\nu + \min_{\mathcal{E}} \max_{\Lambda_l} I(A_1 \ldots A_{l-1} : A_l | R)_{\pi_l} \right), \qquad (75)
\end{aligned}
$$

where the last inequality comes from replacing the specific measurement $\Lambda_{l+1} \otimes \cdots \otimes \Lambda_k$ with the minimum over all measurements $\mathcal{E}$ on systems $A_{l+1} \ldots A_k$. Iterating the argument and exploiting permutation invariance we find

$$
\max_{\Lambda_l,\ldots,\Lambda_k} I(A_1 \ldots A_{l-1} : A_l \ldots A_k)_\nu \geq (k - l + 1) \min_{\mathcal{E}} \max_{\Lambda_l} I(A_1 \ldots A_{l-1} : A_l | R)_{\pi_l}, \qquad (76)
$$

and obtain Eq. (74) from the bound $(l-1) \ln |A| \geq I(A_1 \ldots A_{l-1} : A_l \ldots A_k)_\nu$. Combining it with Eq. (73) we get

$$
\min_{\mathcal{E}} \max_{\Lambda_2,\ldots,\Lambda_l} I(A_1 : \ldots : A_l | R)_\pi \leq \frac{(l-1)^2 \ln |A|}{k - l + 1}. \qquad (77)
$$

We now show how to combine this bound with a few properties of the measure $I(A_1 : \ldots : A_l | R)$ to complete the proof. We have

$$
I(A_1 : \ldots : A_l | R)_\pi = \sum_i p_i I(A_1 : \ldots : A_l)_{\pi_i}, \qquad (78)
$$

with $\pi_i := (\text{id}^{A_1} \otimes \Lambda_2 \otimes \ldots \otimes \Lambda_l)(\rho_i)$, for an ensemble $\{p_i, \rho_i\}$ such that each $\rho_i \in \mathcal{D}(A^{\otimes l})$ is permutation-invariant and $\sum_i p_i \rho_i = \rho^{A_1 \ldots A_l}$. Then, by Pinsker's inequality (Eq. (67)) and the convexity of x^2:

$$\min_{\mathcal{E}} \max_{\Lambda_2, \ldots, \Lambda_l} I(A_1 : \cdots : A_l | R)_\pi \geq \frac{1}{2} \left\| (\text{id}^{A_1} \otimes \Lambda_2 \otimes \cdots \otimes \Lambda_l) \left(\rho^{A_1 \ldots A_l} - \sum_i p_i \rho_i^{A_1} \otimes \ldots \otimes \rho_i^{A_l} \right) \right\|_1^2$$

Part 1 of the theorem follows from Eq. (77).

Part 2: Let $p(x_1, \ldots, x_k | a_1, \ldots, a_k)$ be a permutation-symmetric non-signaling distribution and $\mu = \mu_1 \times \cdots \times \mu_k$ a product distribution on $A_1 \times A_k$. We will use the abbreviations $X_{<l} := X_1 \ldots X_{l-1}$ and $X_{>l} := X_{l+1} \ldots X_k$.

$$\min_{a_{l+1}, \ldots, a_k} \mathbb{E}_{a_1, \ldots, a_l} I(X_1 : \cdots : X_l | X_{>l})_p = \min_{a_{l+1}, \ldots, a_k} \mathbb{E}_{a_1, \ldots, a_l} \sum_{j=2}^{l} I(X_{<j} : X_j | X_{>l})_p \tag{79a}$$

$$= \min_{a_{l+1}, \ldots, a_k} \sum_{j=2}^{l} \mathbb{E}_{a_1, \ldots, a_l} I(X_{<j} : X_j | X_{>l})_p \tag{79b}$$

$$\leq (l-1) \min_{a_{l+1}, \ldots, a_k} \mathbb{E}_{a_1, \ldots, a_l} I(X_{<l} : X_l | X_{>l})_p \tag{79c}$$

To derive the last inequality, observe that $I(X_{<j} : X_j | X_{>l}) = I(X_{<j}; X_l | X_{>l}) \leq I(X_{<l}; X_l | X_{>l})$, where the equality is from the symmetry of p and the inequality is from the monotonicity of mutual information under tracing out systems.

Next,

$$(l-1) \ln |X| \geq \min_{a_{l+1}, \ldots, a_k} \sum_{j=l}^{k} \mathbb{E}_{a_1, \ldots, a_l} I(X_{<l} : X_j | X_{>j})_p \tag{80a}$$

$$= \min_{a_{l+1}, \ldots, a_k} \left(\sum_{j=l+1}^{k} \mathbb{E}_{a_1, \ldots, a_{l-1}} I(X_{<l} : X_j | X_{>j})_p + \mathbb{E}_{a_1, \ldots, a_l} I(X_{<l} : X_l | X_{>l})_p \right) \tag{80b}$$

$$\geq \min_{a_{l+1}, \ldots, a_k} \sum_{j=l+1}^{k} \mathbb{E}_{a_1, \ldots, a_{l-1}} I(X_{<l} : X_j | X_{>j})_p + \min_{a_{l+1}, \ldots, a_k} \mathbb{E}_{a_1, \ldots, a_l} I(X_{<l} : X_l | X_{>l})_p \tag{80c}$$

Iterating, we find that

$$\min_{a_{l+1}, \ldots, a_k} \mathbb{E}_{a_1, \ldots, a_l} I(X_{<l} : X_l | X_{>l})_p \leq \frac{(l-1) \ln |X|}{k-l+1} \tag{81}$$

$$\min_{a_{l+1}, \ldots, a_k} \mathbb{E}_{a_1, \ldots, a_l} I(X_1 : \cdots : X_l | X_{>l})_p \leq \frac{(l-1)^2 \ln |X|}{k-l+1} \quad \text{using (79)} \tag{82}$$

Fix $a_{l+1}, \ldots, a_k$ achieving the minimum in Eq. (82). Using the non-signaling property, we can decompose

$$p(X_{\leq l} | A_{\leq l}) = \sum_{x_{>l}} p(x_{>l} | a_{>l}) p(X_{\leq l} | A_{\leq l}, a_{>l}, x_{>l}). \tag{83}$$

The astute reader will realize that it is now time to deploy Pinsker's inequality (Eq. (67)). Along with Eq. (82) and the convexity of x^2, this concludes the proof of the theorem. $\qquad \square$

5 Proof of Corollary 4

The first lemma is an adaptation of a similar result of Kempe, Kobayashi, Matsumoto, Toner, Vidick [56]. It shows that by symmetrizing the questions and answers of a subset S of the players one can without loss of generality assume that the players follows a symmetric strategy (in the case of classical, entangled, or non-signaling strategies)

Lemma 13. *Let $G(N, \pi, V)$ be a non-signaling-prover game such that $\pi(i_1, \ldots, i_N)$ is symmetric in $i_1, \ldots, i_m$ and V is symmetric under simultaneous permutation of registers $1, \ldots, m$ of the questions $q_{i_1, \ldots, i_N}$ and of the answers $a_{i_1, \ldots, i_N}$ for $m \leq N$. Then given any strategy given by a non-signaling strategy that wins with probability p, there exists a symmetric strategy with respect to provers $1, \ldots, m$.*

The next lemma gives a hardness of approximation result for approximating the classical value of free games.

Lemma 14 (Aaronson-Impagliazzo-Moshkovitz-Shor [3])**.** *3-SAT with n variables can be reduced to the problem of obtaining a constant error approximation to $\omega_c(G)$ for two-player one-round free games with $2^{O(\sqrt{n})}$-sized output alphabet.*

Corollary 4 (restatement).

1. *Let $G(2, \pi, V)$ be a two-player one-round non-local free game with π a product probability distribution on $R \times Q$ and V a predicate on $R \times Q \times A \times B$. Then there is a $(m+1)$-player one-round non-local game $\overline{G}(m + 1, \overline{\pi}, \overline{V})$ with $\overline{\pi}$ a probability distribution on $R \times Q_1 \times \ldots \times Q_m$, with $|Q_k| = |Q|$ for $k \in [m]$, and $\overline{V}$ a predicate on $R \times Q_1 \times \ldots \times Q_m \times A \times B_1 \times \ldots \times B_m$, with $|B_k| = |B|$ for $k \in [m]$, such that*

$$\omega_c(G) = \omega_c(\overline{G}) \leq \omega_e(\overline{G}) \leq \omega_{ns}(\overline{G}) \leq \omega_c(G) + \sqrt{\frac{\ln |A|}{2m}}. \tag{84}$$

2. *For a free game $G(2, \pi, V)$ there is a linear-programming relaxation of size $|R||A| (|Q||B|)^{\frac{\ln |A|}{2\varepsilon^2}}$ for computing $\omega_c(G)$ to within additive error ε.*

3. *One can reduce 3-SAT on n variables to computing $\omega_e(G)$ to within constant additive error for $O(\sqrt{n})$-player one-round non-local games with answer alphabet size of $\exp(O(\sqrt{n}))$ in which only two players are asked questions.*

Proof.
Part 1: Define a game $\overline{G}$ in which the verifier chooses a pair (r, q) from the distribution $\pi(r, q)$ and sends r to the first prover (let us call it Alice) and the q to one of the other m provers chosen at random (let us call them Bob 1 to Bob m). The verifier does not send a question and does not expect an answer from the remaining Bobs. Then the verifier uses the answers obtained from Alice and the chosen Bob to compute $V(a, b|r, q)$. Applying Lemma 13 to the case of non-signaling games we can restrict the parties to use non-signaling distributions which are symmetric on the Bobs. Thus

$$\begin{aligned} \omega_{ns}(\overline{G}) &= \sup_p \sum_{q,r} \pi(r, q) \sum_{a, b_1, \ldots, b_m} \left(\frac{1}{m} \sum_{k=1}^m V(a, b_k|r, q_k) \right) p(a, b_1, \ldots, b_m|r, q_1, \ldots, q_m) \\ &= \sup_{p \in m-\text{Ext}} \sum_{q,r} \pi(r, q) \sum_{a,b} V(a, b|r, q) p(a, b|r, q), \end{aligned} \tag{85}$$

28

where the supremum in the last line is taken over all m-extendible non-signaling distributions p. Then by Theorem 1

$$\sup_{p \in m-\text{Ext}} \sum_{q,r} \pi(r,q) \sum_{a,b} V(a,b|r,q) p(a,b|r,q)$$

$$\leq \sup_{s \in \text{LHV}} \sum_{q,r} \pi(r,q) \sum_{a,b} V(a,b|r,q) s(a,b|r,q) + \frac{1}{2} \sqrt{\frac{2 \ln |A|}{m}}. \tag{86}$$

In more detail, since the game is free we have that $\pi(r,q) = \pi_1(r)\pi_2(q)$. Then

$$\left| \sum_{q,r} \pi_1(r)\pi_2(q) \sum_{a,b} V(a,b|r,q) \left(p(a,b|r,q) - s(a,b|r,q) \right) \right|$$

$$\leq \mathbb{E}_{\pi_1(r)} \mathbb{E}_{\pi_2(q)} \| p(a,b|r,q) - s(a,b|r,q) \|_1$$

$$\leq \mathbb{E}_{\pi_1(r)} \max_{q \in Q} \| p(a,b|r,q) - s(a,b|r,q) \|_1. \tag{87}$$

From theorem 1

$$\min_{s \in \text{LHV}} \mathbb{E}_{\pi_1(r)} \max_{q \in Q} \| p(a,b|r,q) - s(a,b|r,q) \|_1 \leq \frac{1}{2} \sqrt{\frac{2 \ln |A|}{m}}. \tag{88}$$

Part 2: Follows from part 1 and the fact that ω_{ns} can be computed by a linear program [51].

Part 3: Follows from part 1 of this Lemma and part 1 of Lemma 14. □

6 Proof of Corollary 6

We begin with a definition of analogues of QMA with multiple unentangled proofs.

Definition 15. *A language L is in* M-QMA$_n(m,s,c)$ *is there exists a polynomial-time implementable two-outcome measurement $\{M_x, I - M_x\}$ from the class* M *such that*

1. *Completeness: If $x \in L$, there exist m proofs $|\psi_1\rangle, \ldots, |\psi_m\rangle$, each of n qubits, such that*

$$\text{tr} \left(M_x \left(|\psi_1\rangle\langle\psi_1| \otimes \ldots \otimes |\psi_m\rangle\langle\psi_m| \right) \right) \geq c. \tag{89}$$

2. *Soundness: If $x \notin L$, then for any states $|\psi_1\rangle, \ldots, |\psi_k\rangle$,*

$$\text{tr} \left(M_x \left(|\psi_1\rangle\langle\psi_1| \otimes \ldots \otimes |\psi_m\rangle\langle\psi_m| \right) \right) \leq c. \tag{90}$$

If M *is the class of all polynomial-time implementable two-outcome measurements we denote the complexity class simply by* QMA$_n(m,s,c)$.

Some examples of classes of measurements that we consider in this paper are:

29

1. **Bell** is composed of measurements $0 \preceq M \preceq I$ of the form

$$M = \sum_{(i_1,\ldots,i_m)\in S} M_{1,i_1} \otimes \ldots \otimes M_{m,i_m} \tag{91}$$

where $\sum_i M_{j,i} = I$ for all $j \in [m]$, and S is a set of m-tuples of indices. In words the m subsystems are measured locally giving outcomes $(i_1,\ldots,i_m)$ and the verifier accepts if $(i_1,\ldots,i_m) \in S$.

2. **LOCC$_1$** is composed of measurements of the form

$$M = \sum_i M_{1,i} \otimes \ldots \otimes M_{m,i} \tag{92}$$

such that $0 \preceq M_{1,i} \leq I$ for all i, and $0 \leq \sum_i M_{k,i} \leq I$ for every $k \in \{2,\ldots,m\}$.

3. **SEP** is composed of measurements $0 \leq M \leq I$ such that

$$M = \sum_i M_{1,i} \otimes \ldots \otimes M_{1,m}, \tag{93}$$

for positive semi-definite matrices $M_{j,i}$.

See [46] for more examples of classes of measurements as well as relations between them. We will also make use of QMA with multiple identical proofs:

Definition 16. *A language L is in* M-SymQMA$_n(m, s, c)$ *is there exists a polynomial-time implementable two-outcome measurement* $\{M_x, I - M_x\}$ *from the class* M *such that*

1. *Completeness: If $x \in L$, there exist a proof $|\psi\rangle$ of n qubits such that*

$$\mathrm{tr}\left(M_x |\psi\rangle\langle\psi|^{\otimes m}\right) \geq c. \tag{94}$$

2. *Soundness: If $x \notin L$, then for any state $|\psi\rangle$,*

$$\mathrm{tr}\left(M_x |\psi\rangle\langle\psi|^{\otimes m}\right) \leq c. \tag{95}$$

We now turn to the proof of Corollary 6.

Corollary 6 (restatement).

1. BellSymQMA$_n(m, c, s) \subseteq$ QMA$_{10n^2m^2/\varepsilon^2}(c, s + \varepsilon)$.

2. *For every $\varepsilon > 0$ and $c - s = \Omega(1)$, there is no* BellSymQMA$_{O(\log(n))}(n^{\frac{1}{2}-\varepsilon}, c, s)$ *protocol for* 3-SAT *with n variables and $O(n)$ clauses, unless* 3-SAT *can be solved in* $\exp(n^{1-2\varepsilon}\,\mathrm{polylog}(n))$ *time.*

3. BellQMA$_n(m, c, s) \subseteq$ QMA$_{10n^2m^3/\varepsilon^2}(c, s + \varepsilon)$.

4. QMA$_{\mathrm{poly}(n)}(\frac{2}{3}, \frac{1}{3}) =$ BellQMA$_{\mathrm{poly}(n)}(\mathrm{poly}(n), \frac{2}{3}, \frac{1}{3})$

Proof.

Part 1: To simulate a $\mathsf{BellSymQMA}_n(m, c, s)$ protocol in $\mathsf{QMA}_{10n^2m^2/\varepsilon^2}(c, s + \varepsilon)$ the verifier receives the proof of $10n^2m^2/\varepsilon^2$ qubits from the prover and consider it as $10nm^2/\varepsilon^2$ blocks of n qubits. Then he symmetrizes all the blocks, traces out all of them except the first m blocks and runs the original BellQMA protocol on them. It is clear that completeness is not changed. To analyze soundness we use part 1 of Theorem 2.

Part 2: Follows easily from the previous part.

Part 3: To simulate a $\mathsf{BellQMA}_n(m, c, s)$ protocol in $\mathsf{QMA}_{10n^2m^3/\varepsilon^2}(c, s + \varepsilon)$ the verifier receives the proof of $10n^2m^2/\varepsilon^2$ qubits from the prover and consider it as $10nm^2/\varepsilon^2$ blocks of nm qubits. Then he symmetrizes all the blocks, traces out all of them except the first m blocks. Then the divides each of these m blocks into m sub-blocks of n qubits. Let us denote the i-th sub-block of j-th block by $X_{i,j}$. Then the verifier runs the original BellQMA protocol using the state in subsystems $X_{1,1}, X_{2,2}, \ldots, X_{m,m}$ as a proof. It is clear that completeness is not changed. To analyze soundness we use part 1 of Theorem 2.

Part 4: Follows easily from the previous part.

$\square$

7 Proof of Corollary 8

Corollary 8 (restatement). *For some $c > 0$, the Sum-of-Squares hierarchy solves the weak membership problem for separability for the norm $\| * \|_{LOCC^\leftarrow}$ in time*

$$\exp\left(c \left(\sum_j \log |A_j| \right)^2 l^2 \varepsilon^{-2} \right). \tag{96}$$

In turn, the Sum-of-Squares hierarchy solves the weak membership problem for separability for the norm $\| * \|_{2(l)}$ in time

$$\exp\left(c \left(\sum_j \log |A_j| \right)^2 (18)^{l/2} l^2 \varepsilon^{-2} \right). \tag{97}$$

Proof. According to the promise of the weak membership problem, we are given a state $\rho^{A_1,\ldots,A_l} \in \mathcal{D}(A_1 \otimes \cdots \otimes A_l)$ and wish to determine whether it is separable or ε-far from separable in the $\mathrm{LOCC}^\leftarrow$ norm.

The idea of the proof is to approximate the set $\mathrm{Sep}(A_1 : \cdots : A_l)$ with its k-extendible relaxation, for k chosen to give a good approximation guarantee according to Theorem 2. This means introducing systems $X^1, \ldots, X^k$, each of which is composed of l subsystems (i.e. $X^j := X_1^j \cdots X_l^j$ for each j, with $X_i^j \cong A_i$ for each i, j), and searching for a state $\sigma^{X^1 \cdots X^k}$ such that

1. σ is invariant under permutation of the $X^1, \ldots, X^k$ subsystems;

2. and $\sigma^{X_1^1 X_2^2 \ldots X_l^l} = \rho^{A_1 A_2 \ldots A_l}$.

31

Given a separable $\rho^{A_1 \cdots A_l}$, such an extension $\sigma^{X^1 \cdots X^k}$ exists for every $k \geq l$. We can determine whether such a σ exists using semidefinite programming in time polynomial in the overall dimension of σ. If we choose $k = l + 4l^2 \varepsilon^{-2} \sum_j \log |A_j|$, then this will yield the runtime claimed in (96). Moreover, by Theorem 2 we have that there exists a measure ν on $\mathcal{D}(A_1 \otimes \cdots \otimes A_l)$ such that

$$\max_{\Gamma_2, \ldots, \Gamma_l \in \mathcal{M}} \left\| (\mathrm{id} \otimes \Gamma_2 \otimes \cdots \otimes \Gamma_l) \left(\sigma^{X^1 \cdots X^l} - \int \nu(d\omega) \omega^{\otimes l} \right) \right\|_1 \leq \varepsilon, \tag{98}$$

where $\Gamma_2, \ldots, \Gamma_l$ range over all measurements of $X^2, \ldots, X^l$. Restricting to measurements on the $X_2^2, \ldots, X_l^l$ subsystems (which we denote $\Lambda_2, \ldots, \Lambda_l$) and using the monotonicity of the trace norm under partial trace, we obtain

$$\min_{\omega \in \mathrm{Sep}(A_1 : \cdots : A_l)} \max_{\Lambda_2, \ldots, \Lambda_l \in \mathcal{M}} \left\| (\mathrm{id} \otimes \Lambda_2 \otimes \cdots \otimes \Lambda_l) \left(\sigma^{X_1^1 \cdots X_l^l} - \omega \right) \right\|_1 \leq \varepsilon. \tag{99}$$

Of course, $\sigma^{X_1^1 \cdots X_l^l}$ in (99) is equal to $\rho^{A_1 \cdots A_l}$, and so the existence of the symmetric extension σ implies that ρ is no more than ε-far from separable in the one-way LOCC norm. Conversely, if ρ is more than ε-far from separable in the $\mathrm{LOCC}^{\leftarrow}$ norm, then such a σ will not exist, and thus our algorithm will be able to distinguish this case from the one where ρ is separable.

The bound for $\| * \|_{2(l)}$ follows from the reasoning above and the following bound (given by Theorem 5 of [63]):
$$\|X\|_{\mathrm{LOCC}^{\leftarrow}} \geq 18^{-l/2} \|X\|_{2(l)}. \tag{100}$$

$\square$

8 Open Problems

It would be desirable to strengthen several of the results in this work:

1. Conjecture 5 is a proposed improvement of Theorem 1 that would imply that $O(\log(k))$-extendable states cannot be distinguished from separable states by Bell measurements with k outcomes per party. As we discuss in Section 2.2 this would have a very interesting application to the complexity of non-local games.

2. We would also like to improve Theorem 1 to apply to separable measurements[10] instead of merely 1-LOCC measurements. If this were true, it would imply by the results of [46], that $\mathrm{QMA}_n(m, c, s) \subseteq \mathrm{QMA}_{O(mn^2/\varepsilon)}(1, c, s + \varepsilon)$. It would also yield quasipolynomial-time classical algorithms for separability testing and a large number of tensor optimization problems described in [46].

3. One of the few barriers to improving de Finetti theorems is the example of the maximally mixed state on the antisymmetric subspace of $\mathbb{C}^d \otimes \mathbb{C}^d$ [32]. This so-called "universal counterexample" state is d-extendable, and yet is far from separable. However, this distinguishability cannot be achieved by a measurement whose "not separable" outcome is itself a separable measurement operator; aka a "SEP-YES" measurement. As mentioned in the previous open problem, proving a more efficient de Finetti theorem against such measurements

---

[10]Technically, we refer here to approximating $h_{\mathrm{Sep}}(M)$ for "SEP-YES" measurements, meaning that M is of the form $\sum_i A_i \otimes B_i$ for p.s.d. A_i, B_i, but without any such requirement for $I - M$.

would improve the algorithm for approximating $h_{\text{Sep}}(M)$ for general measurements M. Additionally, the antisymmetric state is not PPT, and such examples of highly-extendable far-from-separable states are not known to occur when we add the PPT constraint, as proposed by [41]. Intriguingly, the "worst" known example (i.e. most extendable while being far from separable) of a PPT state is only $O(\log d)$-extendable [24] [11]. It would be of great interest either to prove a better bound on the combination of PPT and k-extendable constraints (see [70] or Section 9.3.2 of [10] for some progress), or to find better counterexample states.

4. It would also be interesting to use our information-theoretic techniques to examine the various extensions of the de Finetti theorem. For example, is there a version of the post-selection technique [31] where the dimension dependence is replaced by a dependence on the number of measurement outcomes? One difficulty here (highlighted by taking the local dimension to be infinite) is in choosing the right test state upon which the channels should act. Another question is whether our techniques can improve the exponential de Finetti theorem [79]. Unfortunately, this theorem is known not to have a classical analogue (due to unpublished work of Christandl and Toner), while our proofs use entropic properties of classical, or classical-quantum, states.

Acknowledgments

References

[1] S. Aaronson. The learnability of quantum states. *Proc. R. Soc. A*, 463:2088, 2007, `arXiv:quant-ph/0608142`.

---

[11]In more detail this follows by considering a variant of the example of [24] (page 6) in which the EPR pair is replaced by a constant-dimensional PPT entangled state. Note that by [14] for every $\varepsilon > 0$ there is a PPT state with trace distance $2-\varepsilon$ from separable states, thus for every $\varepsilon > 0$ one can get a PPT $O(\log(d))$-extendible $d \times d$ state which is $(2-\varepsilon)$-away from any separable state. The same holds true if we use the 1-LOCC version of the trace norm.

[2] S. Aaronson, S. Beigi, A. Drucker, B. Fefferman, and P. Shor. The power of unentanglement. *Theory of Computing*, 5(1):1–42, 2009, `arXiv:0804.0802`.

[3] S. Aaronson, R. Impagliazzo, D. Moshkovitz, and P. Shor. AM with multiple Merlins, 2013. in preparation.

[4] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

[5] S. Arora, B. Barak, and D. Steurer. Subexponential algorithms for unique games and related problems. In *FOCS*, pages 563–572, 2010.

[6] 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, May 1998.

[7] S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of NP. *J. ACM*, 45(1):70–122, Jan. 1998.

[8] A. Aspect, J. Dalibard, and G. Roger. Experimental test of Bell's inequalities using time-varying analyzers. *Phys. Rev. Lett.*, 49:1804–1807, Dec 1982.

[9] L. Babai, L. Fortnow, and C. Lund. Non-deterministic exponential time has two-prover interactive protocols. *computational complexity*, 1:3–40, 1991.

[10] B. Barak, F. G. Brandão, A. W. Harrow, J. Kelner, D. Steurer, and Y. Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *STOC '12*, STOC '12, pages 307–326, 2012, `arXiv:1205.4484`.

[11] B. Barak, J. Kelner, and D. Steurer. Iterative rounding for sum-of-squares relaxations, 2012. unpublished manuscript.

[12] J. Barrett and M. Leifer. The de Finetti theorem for test spaces. *New J. Phys.*, 11:033024, 2009. 0712.2265.

[13] S. Beigi. NP vs QMA_log(2). *Quant. Inf. Comp.*, 10(1&2):0141–0151, 2010, `arXiv:0810.5109`.

[14] S. Beigi and P. Shor. Approximating the set of separable states using the positive partial transpose test. *J. Math. Phys.*, 51:042202, 2010.

[15] J. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195, 1964.

[16] M. Bellare, U. Feige, and J. Kilian. On the role of shared randomness in two prover proof systems. In *ISTCS '95*, pages 199–208, 1995.

[17] C. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. The strengths and weaknesses of quantum computation. *SIAM Journal on Computing*, 26:1510–1523, 1997, `arXiv:quant-ph/9701001`.

[18] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters. Mixed-state entanglement and quantum error correction. *Phys. Rev. A*, 52:3824–3851, 1996, `arXiv:quant-ph/9604024`.

[19] H. Blier and A. Tapp. All languages in NP have very short quantum proofs. In *First International Conference on Quantum, Nano, and Micro Technologies*, pages 34–37, Los Alamitos, CA, USA, 2009. IEEE Computer Society, `arXiv:0709.0738`.

[20] F. Brandão. *Entanglement Theory and the Quantum Simulation of Many-Body Physics*. PhD thesis, Imperial College London, 2008, `arXiv:0810.0026`.

[21] F. Brandão and M. Christandl. Detection of multiparticle entanglement: Quantifying the search for symmetric extensions, 2011, `arXiv:1105.5720`.

[22] F. G. Brandão, M. Christandl, and J. Yard. A quasipolynomial-time algorithm for the quantum separability problem. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pages 343–352, 2011, `arXiv:1011.2751`.

[23] F. G. Brandão and M. B. Plenio. A generalization of quantum Stein's lemma. *Commun. Math. Phys.*, 295:791, 2010, `arXiv:0904.0281`.

[24] F. G. S. L. Brandão, M. Christandl, and J. Yard. Faithful squashed entanglement. *Commun. Math. Phys.*, 306(3):805–830, 2011, `arXiv:1010.1750`.

[25] H. Buhrman, O. Regev, G. Scarpa, and R. de Wolf. Near-optimal and explicit Bell inequality violations. In *CCC '11*, pages 157–166, 2011, `arXiv:1012.5043`.

[26] J.-Y. Cai, A. Condon, and R. Lipton. Playing games of incomplete information. In *STACS '90*, volume 415, pages 58–69, 1990.

[27] C. M. Caves, C. A. Fuchs, and R. Schack. Unknown quantum states: The quantum de Finetti representation. *J. Math. Phys.*, 43(9):4537–4559, 2002, `arXiv:quant-ph/0104088`.

[28] A. Chailloux and O. Sattath. The complexity of the separable Hamiltonian problem, 2011, `arXiv:1111.5247`.

[29] J. Chen and A. Drucker. Short multi-prover quantum proofs for SAT without entangled measurements, 2010, `arXiv:1011.0716`.

[30] A. Chiesa and M. Forbes. Improved soundness for QMA with multiple provers, 2011, `arXiv:1108.2098`.

[31] M. Christandl, R. Koenig, and R. Renner. Post-selection technique for quantum channels with applications to quantum cryptography. *Phys. Rev. Lett.*, 102:020504, 2009, `arXiv:0809.3019`.

[32] M. Christandl, R. König, G. Mitchison, and R. Renner. One-and-a-half quantum de Finetti theorems. *Commun. Math. Phys.*, 273:473–498, 2007, `arXiv:quant-ph/0602130`.

[33] M. Christandl and B. Toner. Finite de Finetti theorem for conditional probability distributions describing physical theories. *J. Math. Phys.*, 50(4):042104, 2009, `arXiv:0712.0916`.

[34] R. Cleve, D. Gavinsky, and R. Jain. Entanglement-resistant two-prover interactive proof systems and non-adaptive PIRs. *Quantum Info. Comput.*, 9(7):648–656, July 2009, `arXiv:0707.1729`.

[35] R. Cleve, P. Hoyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *CCC '04*, pages 236–249, 2004, `arXiv:quant-ph/0404076`.

[36] F. Cobos, T. Kühn, and J. Peetre. Remarks on symmetries of trilinear forms. *Rev. R. Acad. Cienc. Exact. Fis.Nat. (Esp)*, 94(4):441–449, 2000.

[37] J. P. D'Angelo and M. Putinar. Polynomial optimization on odd-dimensional spheres. In M. Putinar and S. Sullivant, editors, *Emerging Applications of Algebraic Geometry*, volume 149 of *The IMA Volumes in Mathematics and its Applications*, pages 1–15. Springer New York, 2009.

[38] E. de Klerk. The complexity of optimizing over a simplex, hypercube or sphere: A short survey. *Central European Journal of Operations Research*, 16(2):111–125, 2008.

[39] P. Diaconis and D. Freedman. Finite exchangeable sequences. *Annals of Probability*, 8:745–764, 1980.

[40] A. C. Doherty, Y.-C. Liang, B. Toner, and S. Wehner. The quantum moment problem and bounds on entangled multi-prover games. In *CCC '08*, pages 199–210, 2008, `arXiv:0803.4373`.

[41] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri. Complete family of separability criteria. *Phys. Rev. A*, 69:022308, Feb 2004, `arXiv:quant-ph/0308032`.

[42] A. C. Doherty and S. Wehner, 2009. personal communication.

[43] A. C. Doherty and S. Wehner. Convergence of SDP hierarchies for polynomial optimization on the hypersphere, 2012, `arXiv:1210.5048`.

[44] M. Fannes, J. T. Lewis, and A. Verbeure. Symmetric states of composite systems. *Lett. Math. Phys.*, 15:255–260, 1988.

[45] S. Gharibian, J. Sikora, and S. Upadhyay. QMA variants with polynomially many provers, 2011, `arXiv:1108.0617`.

[46] A. W. Harrow and A. Montanaro. An efficient test for product states, with applications to quantum Merlin-Arthur games. In *FOCS '10*, pages 633–642, 2010, `arXiv:1001.0017`.

[47] R. L. Hudson and G. R. Moody. Locally normal symmetric states and an analogue of de Finetti's theorem. *Z. Wahrschein. verw. Geb.*, 33:343–351, 1976.

[48] R. Impagliazzo and R. Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.

[49] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? In *FOCS'98*, pages 653–662. IEEE, 1998.

[50] L. Ioannou. Computational complexity of the quantum separability problem. *Quantum Information and Computation*, 7(4):335, 2007.

[51] T. Ito. Polynomial-space approximation of no-signaling provers. In *ICALP'10*, pages 140–151, 2010, `arXiv:0908.2363`.

[52] T. Ito, H. Kobayashi, and K. Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *CCC '09*, pages 217–228, 2009, `arXiv:0810.0693`.

[53] T. Ito, H. Kobayashi, D. Preda, X. Sun, and A. C. C. Yao. Generalized tsirelson inequalities, commuting-operator provers, and multi-prover interactive proof systems. In *CCC '08*, pages 187–198, 2008, `arXiv:0712.2163`.

[54] T. Ito and T. Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. In *FOCS '12*, 2012, `arXiv:1207.0550`.

[55] M. Kearns and U. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, 1994.

[56] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. Entangled games are hard to approximate. *SIAM J. Comput.*, 40(3):848–877, 2011, `arXiv:0704.2903`.

[57] J. Kempe, O. Regev, and B. Toner. Unique games with entangled provers are easy. *SIAM J. Comput.*, 39(7):3207–3229, July 2010, `arXiv:0710.0655`.

[58] J. Kempe and T. Vidick. Parallel repetition of entangled games. In *STOC '11*, pages 353–362, 2011, `arXiv:1012.4728`.

[59] S. Khot and M. Safra. A two prover one round game with strong soundness. In *FOCS '11*, pages 648–657, 2011.

[60] H. Kobayashi and K. Matsumoto. Quantum multi-prover interactive proof systems with limited prior entanglement. *J. Comput. Syst. Sci.*, 66(3):429–450, May 2003, `arXiv:cs/0102013`.

[61] H. Kobayashi, K. Matsumoto, and T. Yamakami. Quantum Merlin-Arthur proof systems: Are multiple Merlins more helpful to Arthur? In *ISAAC*, volume 2906, pages 189–198, 2003, `arXiv:quant-ph/0306051`.

[62] R. Koenig and R. Renner. A de Finetti representation for finite symmetric quantum states. *J. Math. Phys.*, 46(12):122108, 2005, `arXiv:quant-ph/0410229`.

[63] C. Lancien and A. Winter. Distinguishing multi-partite states by local measurements, 2012, `arXiv:1206.2884`.

[64] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Opt.*, 11(3):796–817, 2001.

[65] F. Le Gall, S. Nakagawa, and H. Nishimura. On QMA protocols with two short quantum proofs. *Quant. Inf. Comp.*, 12:0589, 2012, `arXiv:1108.4306`.

[66] K. Li and A. Winter. Relative entropy and squashed entanglement, 2012, `arXiv:1210.3181`.

[67] N. G. Ll. Masanes, A. Acin. General properties of nonsignaling theories. *Phys. Rev. A.*, 73:012112, 2006.

[68] C. Marriott and J. Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005, `arXiv:cs/0506068`.

[69] M. McKague. On the power of quantum computation over real Hilbert spaces, 2011, `arXiv:1109.0795`.

[70] M. Navascues, M. Owari, and M. B. Plenio. The power of symmetric extensions for entanglement detection. *Phys. Rev. A*, 80:052306, 2009, `arXiv:0906.2731`.

[71] M. Navascués, S. Pironio, and A. Acin. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.*, 10(7):073013, 2008, `arXiv:0803.4290`.

[72] R. O'Donnell and Y. Zhou. Approximability and proof complexity, 2013, `arXiv:1211.1958`.

[73] P. A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, MIT, 2000.

[74] A. Pereszlenyi. Multi-prover quantum merlin-arthur proof systems with small gap, 2012, `arXiv:1205.2761`.

[75] M. Piani. Relative entropy of entanglement and restricted measurements. *Phys. Rev. Lett.*, 103:160504, 2009.

[76] V. Powers and B. Reznick. A new bound for Pólya's theorem with applications to polynomials positive on polyhedra. *Journal of Pure and Applied Algebra*, 164(1–2):221–229, 2001.

[77] G. A. Raggio and R. F. Werner. Quantum statistical mechanics of general mean field systems. *Helv. Phys. Acta*, 62:980–1003, 1989.

[78] P. Raghavendra and N. Tan. Approximating CSPs with global cardinality constraints using SDP hierarchies. In *SODA '12*, pages 373–387, 2012, `arXiv:1110.1064`.

[79] R. Renner. *Security of quantum key distribution*. PhD thesis, ETHZ, Zurich, 2005, `arXiv:quant-ph/0512258`.

[80] R. Renner. Symmetry implies independence. *Nature Physics*, 3:645–649, 2007, `arXiv:quant-ph/0703069`.

[81] Y. Shi and X. Wu. Epsilon-net method for optimizations over separable states, 2011, `arXiv:1112.0808`.

[82] M. Sion. On general minimax theorems. *Pacific J. Math*, 8(1):171–176, 1958.

[83] E. Størmer. Symmetric states of infinite tensor products of c-algebras. *J. Funct. Anal.*, 3:48, 1969.

[84] B. M. Terhal, A. C. Doherty, and D. Schwab. Symmetric extensions of quantum states and local hidden variable theories. *Phys. Rev. Lett.*, 90:157903, 2003, `arXiv:quant-ph/0210053`.

[85] T. Vidick. Three-player entangled XOR games are NP-hard to approximate, 2013, `arXiv:1302.1242`.

[86] J. Watrous. Quantum computational complexity, 2008, `arXiv:0804.3401`.

[87] S. Wehner. Entanglement in interactive proof systems with binary answers. In *STACS'06*, pages 162–171, 2006, `arXiv:quant-ph/0508201`.

[88] R. F. Werner. An application of Bell's inequalities to a quantum state extension problem. *Lett. Math. Phys.*, 17:359, 1989.

[89] D. Yang. A simple proof of monogamy of entanglement. *Physics Letters A*, 360(2):249–250, 2006, `arXiv:quant-ph/0604168`.

[90] D. Yang, K. Horodecki, M. Horodecki, P. Horodecki, J. Oppenheim, and W. Song. Squashed entanglement for multipartite states and entanglement measures based on the mixed convex roof. *IEEE Trans. Inf. Th.*, 55(7):3375–3387, July 2009, `arXiv:0704.2236`.