Non-Signaling Parallel Repetition Using de Finetti Reductions

Rotem Arnon-Friedman, Renato Renner, and Thomas Vidick

***Abstract*— In the context of multiplayer games, the parallel repetition problem can be phrased as follows: given a game G with optimal winning probability $1-\alpha$ and its repeated version G^n (in which n games are played together, in parallel), can the players use strategies that are substantially better than ones in which each game is played independently? This question is relevant in physics for the study of correlations and plays an important role in computer science in the context of complexity and cryptography. In this paper, the case of multiplayer non-signaling games is considered, i.e., the only restriction on the players is that they are not allowed to communicate during the game. For complete-support games (games where all possible combinations of questions have non-zero probability to be asked) with any number of players, we prove a threshold theorem stating that the probability that non-signaling players win more than a fraction $1-\alpha+\beta$ of the n games is exponentially small in $n\beta^2$ for every $0 \leq \beta \leq \alpha$. For games with incomplete support, we derive a similar statement for a slightly modified form of repetition. The result is proved using a new technique based on a recent de Finetti theorem, which allows us to avoid central technical difficulties that arise in standard proofs of parallel repetition theorems.**

***Index Terms*— Parallel repetition, threshold theorem, multiplayer games, non-signalling players, de Finetti theorems, non-locality, correlations, probability theory, quantum entanglement.**

I. Introduction

A. Multiplayer Games and Parallel Repetition

Multiplayer games are relevant in many areas of both theoretical physics and theoretical computer science. In physics, multiplayer games give an intuitive way to study the role and implications of entanglement and correlations, e.g., in the setting of Bell inequalities [1], [2]. In computer science such games arise in the context of complexity theory [3]–[5] and cryptography [6], [7].

In a game G, a referee asks each of the cooperating players a question chosen according to a given probability distribution. The players then need to supply answers which fulfil a predefined requirement according to which the referee accepts or rejects the answers. In order to do so, they can agree on a strategy beforehand, but once the game begins communication between the players is not allowed. If the referee accepts their answers the players win. The goal of the players is, of course, to maximise their winning probability in the game.

According to the field of interest, one can analyse any game under different restrictions on the players (in addition to not being allowed to communicate). In classical computer science the players are usually assumed to have only classical resources, or strategies. That is, they can use only local operations and shared randomness. In contrast, one can also consider quantum strategies: before the game starts the players create a multipartite quantum state that can be shared among them. When the game begins each player locally measures their own part of the state and bases the answer on their measurement result. It is well-known that sharing quantum entanglement can significantly increase the winning probability in some games [2], [8].

Another, more general, type of strategies are those where the players can use any type of correlations that do not allow them to communicate, also called non-signalling correlations. That is, the *only* restriction on the players is that they are not allowed to communicate (as will be defined formally later).

Considering the non-signalling case is interesting for several reasons. A first reason is to minimise the set of assumptions to the mere necessary one. Indeed, if the players are allowed to communicate by sending signals they can win any game. Minimising the set of assumptions can be useful in cryptography when one wishes to get the strongest result possible, i.e., one where the attack strategies of malicious parties are only restricted minimally (as in [9]–[11] for example). In theoretical physics, non-signalling correlations enable the study of generalised theories possibly beyond quantum theory. It is also important to mention that, due to their linearity, the non-signalling constraints are often easier to analyse than the quantum or the classical constraints. Therefore, even if additional constraints hold, focusing on the non-signalling ones serves as a way to get first insights into a given problem.

One of the most interesting questions regarding multiplayer games is the question of parallel repetition. Given a game G with optimal winning probability $1-\alpha$ (using either classical, quantum, or non-signalling strategies), we are interested in analysing the winning probability in the repeated game G^n.

Manuscript received May 5, 2015; revised September 18, 2015; accepted December 22, 2015. Date of publication January 8, 2016; date of current version February 12, 2016. R. Arnon-Friedman and R. Renner were supported in part by the European Research Council under Grant 258932, in part by the European Commission Strategic Targeted Research Project RAQUEL, in part by the Swiss National Science Foundation via the National Centre of Competence in Research through the QSIT Project, and in part by the CHIST-ERA Project DIQIP. T. Vidick was supported in part by the Simons Institute, Berkeley, CA, USA, in part by the Ministry of Education, Singapore, under the Tier 3 under Grant MOE2012-T3-1-009, and in part by the Perimeter Institute, Waterloo, ON, Canada.

R. Arnon-Friedman and R. Renner are with the Institute for Theoretical Physics, ETH Zürich, Zürich 8092, Switzerland (e-mail: rotema@itp.phys.ethz.ch; renner@phys.ethz.ch).

T. Vidick is with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: vidick@cms.caltech.edu).

Communicated by M. M. Wilde, Associate Editor for Quantum Information Theory.

Digital Object Identifier 10.1109/TIT.2016.2516022

In G^n the referee gives the players n independent tuples of questions at once, to which the players should reply. The most natural winning criterion is that the players answer a certain fraction $1-\alpha+\beta$ of the n game instances correctly, and one can then ask what is the probability that the players succeed as a function of β, as the number of repetitions n increases.

The players can always use the trivial independent and identically distributed (i.i.d.) strategy: they just answer each of the n questions independently according to the optimal one-game strategy. In this case the fraction of successful answers is highly concentrated around $1-\alpha$ (alternatively, the probability to win all games simultaneously is $(1-\alpha)^n$). But can they do significantly better?

If correlated strategies for G^n are not substantially better than independent ones, even in an asymptotic manner, we learn that "one cannot fight independence with correlations". As long as the questions are asked, and the answers are verified, in an independent way, creating correlations between the different answers using a correlated strategy cannot help much. The resulting threshold theorem can then be used, for example, when considering a series of Bell violation experiments performed in parallel, or for hardness amplification in complexity theory and security amplification in classical, quantum and non-signalling cryptography.

B. Related Work

Raz was the first to show in [12] an exponential parallel repetition theorem for classical two-player games. That is, he showed that if the classical optimal winning probability in a game G is smaller than 1, then the probability to win all the games in the repeated game G^n, using a classical strategy, decreases exponentially with the number of repetitions n. Raz's result was then improved and adapted to the non-signalling case by Holenstein [13]. Another improvement was made by Rao in [14], where a threshold theorem for the classical two-player case was proven: Rao showed that the probability to win more than a fraction $1-\alpha+\beta$ of the games for any $\beta > 0$ is exponentially small in the number of repetitions.

Following the same proof technique as [12]–[14], Buhrman, Fehr and Schaffner recently proved in [15] a threshold theorem for multiplayer non-signalling complete-support games (as formally defined in Definition 7). Their threshold theorem was the first result where more than two players were considered.

The question of parallel repetition in the quantum case is less well understood than its classical and non-signalling versions. All currently known results deal with limited classes of two-player games and no general proof is known. The latest results are given in [16]–[18], where different assumptions on the probability distribution over the questions of the game are considered.

C. de Finetti Theorems in the Context of Parallel Repetition

The main difficulty in proving a parallel repetition result comes from the, almost arbitrary, correlations between the different questions-answers pairs in the players' strategy for G^n: as the players get all the n tuples of questions together they can answer them in a correlated way. In most of the known parallel repetition results (e.g., [12]–[15]) the main idea of the proof is to bound the winning probability for some of the questions, *conditioned* on winning the previous questions. However, as the strategy itself introduces correlations between the different tuples of questions, a large amount of technical work is devoted to dealing with the effect of conditioning on the event of winning the previous questions.

When considering the correlations in a strategy for the repeated game there is one type of symmetry which one can take advantage of, but which is usually virtually ignored – permutation invariance. As the game G^n itself is invariant under joint permutation of the tuples of questions and answers, we can restrict our attention to permutation-invariant strategies without loss of generality. Permutation-invariant strategies are strategies which are indifferent to the ordering of the questions given by the referee. That is, the probability of answering a specific set of tuples of questions correctly does not depend on the ordering of the tuples.

Once we restrict our attention to permutation-invariant strategies, de Finetti theorems seem like a natural tool to leverage for the analysis. A de Finetti theorem is any type of theorem which relates permutation-invariant states[1] to a more structured state, having the form of a convex combination of i.i.d. states, called a de Finetti state. The specific relation between the states depends on the type of theorem. The first de Finetti theorem [19] established that the collection of infinitely exchangeable sequences, in other words those distributions on infinite strings that are invariant under all permutations, exactly coincides with the collection of all convex combinations of i.i.d distributions. Subsequent results establish quantitative bounds on the distance of any permutation-invariant state, or subsystems thereof, from some de Finetti state or an approximation of a de Finetti state [20]–[25]. A different form of statement, also called a de Finetti reduction, relates any permutation-invariant state to an explicit de Finetti state by an inequality relation [26], [27]. The common feature of all de Finetti theorems is that they enable a substantially simplified analysis of information-processing tasks by exploiting permutation invariance symmetry. Indeed, quantum de Finetti theorems play a significant role in many quantum information problems such as quantum cryptography [26], [28], tomography [29], channel capacities [30] and complexity [25].

In the context of games and strategies, de Finetti theorems suggest one may be able to reduce the analysis of general permutation-invariant strategies to the analysis of a de Finetti strategy, i.e., a convex combination of i.i.d. strategies. As the behaviour of i.i.d. strategies is trivial under parallel repetition, a reduction of this type could simplify the analysis of parallel repetition theorems and threshold theorems.

Yet, de Finetti theorems were not used in the past in this context, and for a good reason. The many versions of quantum de Finetti theorems (e.g., [23], [26]) could not have been used as they depend on the dimension of the underlying quantum strategies, while in the quantum multiplayer game setting one

[1] Depending on the context, a state can be a probability distribution, a quantum density operator or a conditional probability distribution.

does not wish to restrict the dimension. Non-signalling de Finetti theorems, as in [31] and [32], were also not applicable for non-signalling parallel repetition theorems, as they restrict almost completely the type of allowed correlations in the strategies for the repeated game by assuming very strict non-signalling constraints between the different repetitions, i.e., between the different questions-answers pairs.

In this work we use the recent de Finetti theorem of [27], which imposes no assumptions at all regarding the structure of the strategies (apart from permutation invariance), and is therefore applicable in the context of parallel repetition. This allows us to devise a proof technique which is completely different from the known proofs of parallel repetition results. In particular, at least in the non-signalling case presented here, the conditioning problem described at the beginning of this section disappears completely and the number of players does not play a role in the proof structure.

D. Results and Contributions

The main result presented in this work is a threshold theorem (also called a concentration bound) for the n-fold repetition of any m-player complete-support game (see Definition 7), in which the players are allowed to share any non-signalling strategy. Denote by w_{ns} the optimal non-signalling winning probability in a game G. We prove the following theorem.

Theorem 1: For any complete-support game G with $w_{\text{ns}} = 1 - \alpha$ *there exist* $\mathcal{C}_1(G,n)$ *and* $\mathcal{C}_2(G)$*, where* $\mathcal{C}_1(G,n)$ *is polynomial in the number of repetitions n, such that for every* $0 < \beta \leq \alpha$ *and large enough n, the probability that non-signalling players win more than a fraction* $1 - \alpha + \beta$ *of the n questions in the repeated game* G^n *is at most* $\mathcal{C}_1(G,n)\exp\left[-\mathcal{C}_2(G)n\beta^2\right]$.

That is, for sufficiently many repetitions the probability to win more than a fraction $1 - \alpha + \beta$ of the n games is exponentially small. The constant $\mathcal{C}_1(G,n)$ is such that $\mathcal{C}_1(G,n) < 6m|\mathcal{Q}||\mathcal{A}|\,(n+1)^{|\mathcal{Q}||\mathcal{A}|-1}$ where m is the number of players, and $|\mathcal{Q}|$ and $|\mathcal{A}|$ are the number of possible questions and answers, respectively, in G. $\mathcal{C}_2(G)$ is a finite constant that can be computed by solving the polynomial-size linear program given in Equation (5). A sufficient condition on the number of repetitions for the bound in the theorem to hold is $n = \Omega\big(|\mathcal{Q}||\mathcal{A}|\frac{\mathcal{C}_2}{\beta^2}\ln^2(|\mathcal{Q}||\mathcal{A}|\frac{\mathcal{C}_2}{\beta})\big)$. We refer to Equation (12), and the choice of constants made around Equation (24), for a more precise bound.

There are two main differences between the exponential bound given in the threshold theorem of [15] (Theorem 15 therein) and the bound we give here. First, while our bound suffers from the polynomial dependency on the number of repetitions in $\mathcal{C}_1(G,n)$ (which is inherent to the use of a de Finetti reduction), there is no such dependency in [15]. As the number of repetitions goes to infinity, however, the exponential factor quickly dominates. Both our constant $\mathcal{C}_2(G)$ and the constant μ in [15, Th. 15] depend on the size of the game through a certain linear program (see the proof of Lemma 27 and the discussion that follows it in this paper, and [15, Proof of Proposition 18]), making a direct comparison difficult. Another point of comparison between the bounds is the dependency on β: we obtain the optimal (as follows from optimal formulations of the Chernoff bound) dependency β^2, as compared to β^4 in [15]. As far as we are aware, this is the first threshold theorem where optimal dependency on β is achieved (see also [14]).

Theorem 1 applies to complete-support games. The result is extended in two different directions. First, based on ideas from [33], we show in Appendix A-A that when considering two-player games *without* complete-support Theorem 1 still holds. Second, for general multiplayer games we consider in Appendix A-B a small modification of the repetition procedure. Instead of the usual parallel repetition procedure, in which n tuples of questions are chosen according to the game distribution Q, we change the distribution of questions in the repeated game by sometimes (with small positive probability η) asking the players a tuple of questions q which does not appear in the original game G. We call such questions "dummy questions"; for these questions any answer from the players is accepted. The remaining questions, for which $Q(q) > 0$, are called the "real questions" and the modified game is denoted by $\tilde{G}^n$. We prove the following threshold theorem:

Theorem 2: For any game G with $w_{\text{ns}} = 1 - \alpha$ *there exist* $\mathcal{C}_1(G,n)$ *and* $\mathcal{C}_2(G)$*, where* $\mathcal{C}_1(G,n)$ *is polynomial in the number of repetitions n, such that for every* $0 < \beta \leq \alpha$ *and large enough n, the probability that non-signalling players win more than a fraction* $1-\alpha+\beta$ *of the* real questions *in the modified repeated game* $\tilde{G}^n$ *is at most* $\mathcal{C}_1(G,n)\exp\left[-\mathcal{C}_2(G)n\beta^2\right]$.

The constants $\mathcal{C}_1(G,n)$ and $\mathcal{C}_2(G)$ have the same form as in Theorem 1, but they now depend also on the perturbation η of the original questions distribution. For more details on the definition of $\tilde{G}^n$ and the proof of Theorem 2 see Appendix A-B.

A similar modification was previously considered in both classical [34] and quantum [35] parallel repetition theorems, where the repetitions in which dummy questions are selected were called "confusion rounds". For many applications this modification is harmless, especially as the success probability of "honest" players is not affected by it. However, it is important to note that Theorem 2 only holds for the modified form of repetition of the original game.

In addition to the bounds themselves our, perhaps most important, contribution in this work is the, arguably simpler, proof technique. While most of the known parallel repetition results build on the proof technique of [12] we give a completely different proof, with ideas based on de Finetti theorems and tomography (as explained in the next section). Our proof technique allows us to avoid the usual difficulties which arise in proofs of parallel repetition theorems, such as conditioning on some of the questions and answers or considering an arbitrary number of players. In this sense our proof can be seen as more natural than previous proofs, and therefore more likely to be extendable to the classical and quantum multiplayer cases as well.

E. Proof Idea and Techniques

The goal of this section is to give the reader an intuitive understanding of the proof idea and techniques. The formal and more technical implementations of these ideas are given in the following sections. Nevertheless, the following two definitions are needed.

Definition 3 (Multiplayer Game): An m-player game $G = (\mathcal{Q}, \mathcal{A}, Q, R)$ is defined by a set of possible tuples of questions $\mathcal{Q}$ together with a probability distribution $Q : \mathcal{Q} \to [0, 1]$ (according to which the referee choses the questions) over it, a set of possible tuples of answers $\mathcal{A}$ and a winning condition $R : \mathcal{Q} \times \mathcal{A} \to \{0, 1\}$. An m-tuple of questions $q = (q^1, q^2, \ldots, q^m) \in \mathcal{Q}$ describes the questions given to the different players. Similarly an m-tuple of answers $a = (a^1, a^2, \ldots, a^m) \in \mathcal{A}$ describes the answers given by the different players.

Definition 4 (Strategy): A strategy for an m-player game $G = (\mathcal{Q}, \mathcal{A}, Q, R)$ is a conditional probability distribution $\mathrm{O}_{A|Q} : \mathcal{A} \times \mathcal{Q} \to [0, 1]$, i.e., $\sum_a \mathrm{O}_{A|Q}(a|q) = 1$ for all $q \in \mathcal{Q}$. Similarly, a strategy for a repeated game G^n is a conditional probability distribution denoted by $\mathrm{P}_{\vec{A}|\vec{Q}} : \mathcal{A}^n \times \mathcal{Q}^n \to [0, 1]$.

Throughout the proof strategies for the game G are denoted by $\mathrm{O}_{A|Q}$ and strategies for the repeated game G^n are denoted by $\mathrm{P}_{\vec{A}|\vec{Q}}$.

1) Permutation Invariance and de Finetti Theorems: The first trivial, but crucial, observation made is that when considering strategies for the repeated game, one can concentrate without loss of generality on permutation-invariant strategies. Permutation-invariant strategies are indifferent to the ordering of the tuples of questions given by the referee. That is, the referee can ask the players to answer q_1, q_2, q_3 or q_2, q_3, q_1 (each q_i is an m-tuple); in both cases the winning probability will be the same if the players are using a permutation-invariant strategy. Note that the permutation changes only the order of the tuples of questions. In particular, the players themselves are not being permuted and the questions of all players are permuted in exactly the same way (see Definition 21 and Lemma 22 for the formal argument).

Considering only permutation-invariant strategies allows us to use the de Finetti theorem of [27] which relates any permutation-invariant strategy to a de Finetti strategy. The exact statement of the de Finetti theorem will only be relevant later. For now, using just the intuition of de Finetti theorems, one can think of any permutation-invariant strategy as being a convex combination of i.i.d. strategies. That is,[2]

$$\mathrm{P}_{\vec{A}|\vec{Q}} \approx \int \mathrm{O}_{A|Q}^{\otimes n} \mathrm{dO}_{A|Q} \tag{1}$$

where $\mathrm{dO}_{A|Q}$ is some measure on the space of one-game strategies and $\mathrm{O}_{A|Q}^{\otimes n}$ is a product of n identical strategies $\mathrm{O}_{A|Q}$.

Unfortunately, the convex combination itself (meaning, the measure $\mathrm{dO}_{A|Q}$) is unknown. Moreover, even though we assume that the strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ does not allow the m players to communicate, i.e., it is non-signalling, the convex combination might still include signalling parts, i.e., signalling $\mathrm{O}_{A|Q}$. Indeed, in general, a convex combination of signalling strategies can still be non-signalling.

For the non-signalling parts of the convex combination one can easily prove a strong parallel repetition or threshold theorem. These parts are just i.i.d. non-signalling strategies. The only thing which is left to prove is therefore that the *signalling* part of the convex combination of Equation (1) has an exponentially small weight.[3] We find this question interesting by itself, and of course, the same question can be asked in the classical and quantum case – given a classical or quantum strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$, what is the weight of the non-classical or non-quantum i.i.d. parts in the convex combination?

2) Bounding the Signalling Part: As the convex combination itself in Equation (1) is unknown, one cannot just calculate the weight of the signalling part. We therefore take a more operational approach, following ideas from quantum tomography [29].

Consider a particular (unknown) part $\mathrm{O}_{A|Q}^{\otimes n}$ of the convex combination and divide the n copies of the strategy $\mathrm{O}_{A|Q}$ into two groups – a test group consisting of $n/2$ out of the n copies, and a game group of $n/2$ copies. The general idea is to use the test copies to get an estimation $\mathrm{O}_{A|Q}^{\mathrm{EST}}$ of the strategy $\mathrm{O}_{A|Q}$, which will then help us in proving our claims.

More specifically, we are interested in knowing whether $\mathrm{O}_{A|Q}$ is signalling or not (if it is non-signalling then its winning probability in G is obviously bounded by the optimal non-signalling winning probability $1-\alpha$). For this we define a signalling measure and an operational (and hypothetical) signalling test. Given questions and answers which are distributed according to the $n/2$ copies of $\mathrm{O}_{A|Q}$ and Q, the signalling test will create an estimation $\mathrm{O}_{A|Q}^{\mathrm{EST}}$ and calculate its signalling value. If the signalling value is above a certain threshold the test will accept, and otherwise it will reject.

In order to bound the weight of the signalling part in Equation (1) one can bound the probability that the signalling test accepts. To prove that the acceptance probability is exponentially small we use a combination of two lemmas, which we call the weak and the strong lemma. These lemmas are based on a special guessing game that we construct and on applications of the de Finetti theorem. Both lemmas together show that if the probability of the test accepting is too high, then the original strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ must have been signalling – a contradiction.

3) From Intuition to Practice: In practice, the de Finetti theorem [27] is an inequality relation between any permutation-invariant strategy and a given de Finetti strategy (see Lemma 23) which does not imply Equation (1). As a consequence, considering the test copies and game copies as above does not directly make sense. Nevertheless, we can follow similar ideas by considering the questions-answers pairs in a specific instance of the repeated game. That is, every time the game is played using a strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$, we divide the data, the questions and answers, of this specific run into two groups – test data and game data, consisting of $n/2$ tuples of questions-answers pairs each. Our goal is then to bound

[2] We emphasise once again that this is not a quantitive statement that we claim to be correct. This is just useful as an intuitive way of understanding the proof idea.

[3] As mentioned above, this statement does not hold for an arbitrary decomposition of a non-signalling strategy as a convex combination of other strategies. We will crucially use the fact that each term in the convex combination has an i.i.d. structure.

TABLE I

Constants, Parameters and Their Relations. (g) Next to the Symbol Denotes That This Is a Constant Which Depends on the Considered Game and (t) Denotes a Parameter of the Threshold Theorem. All Other Constants Should Be Chosen Such That All the Requirements in the Last Column of the Table Are Fulfilled

Symbol	Meaning	First appears	Fulfils
m (g)	# of players		
$1-\alpha$ (g)	optimal NS winning probability in G		
κ (g)	bound on an optimal dual solution $y^\star$		$\sum_{j=1}^{d} \lvert y_j^\star \rvert$
W_{ns} (g)	optimal NS winning probability in guessing game		$\max_{q,i} Q(q^i \lvert q^{\bar{i}})$
n (t)	# of repetitions		
β (t)	deviation in the threshold theorem		
ϵ	confidence interval of the test	Equation (8)	$\epsilon \leq \min_q Q(q)$
ζ	signalling threshold of the test	Equation (8)	$7\epsilon \leq \zeta \leq 1$; $\zeta + 2\epsilon \leq \frac{\beta}{\kappa}$
δ	confidence level of the test	Lemma 16	$\delta = (n/2+1)^{\lvert\mathcal{A}\rvert\cdot\lvert\mathcal{Q}\rvert-1} e^{-n\epsilon^2/4}$
ν	signalling threshold	Lemma 18	$\frac{\sqrt{c\delta}}{1-\sqrt{c\delta}} W_{\text{ns}} < \nu < \zeta - 6\epsilon$
c	de Finetti constant	Lemma 23	$c = (n+1)^{\lvert\mathcal{Q}\rvert(\lvert\mathcal{A}\rvert-1)}$
d	# of different signalling tests	Lemma 27	$d < m\lvert\mathcal{Q}\rvert\lvert\mathcal{A}\rvert$

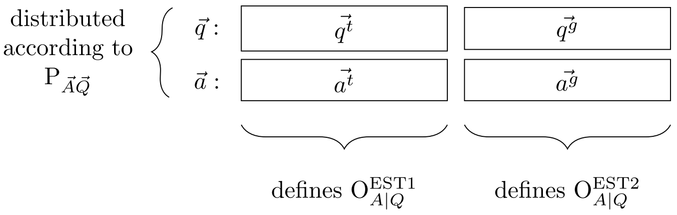
distributed according to
$\mathrm{P}_{\vec{A}\vec{Q}}$
$\vec{q}$:
q^t
q^g
$\vec{a}$:
a^t
a^g
defines $\mathrm{O}^{\mathrm{EST1}}_{A|Q}$
defines $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$

Fig. 1. Division to test and game data.

TABLE II

Symbols of Events Used Throughout the Paper

Symbol	Event
$\mathbf{T}$	passing the signalling test: $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\vec{q^t},\vec{a^t}) = 1$
$\mathbf{in}^{\mathbf{\Sigma}}$	$\mathrm{O}^{\mathrm{EST2}}_{A\lvert Q} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$
$\mathbf{in}^{\sigma}$	$\mathrm{O}^{\mathrm{EST2}}_{A\lvert Q} \in \sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$
$\mathbf{agq}$	all tuples of questions appear in the game questions

the winning frequency in the game data, while the test data is relevant for the hypothetical signalling tests (see also Figure 1 in the following section).

The rest of the paper is organised as follows. We start with some preliminaries in Section II. In Section III we first consider and explain the concept of non-signalling strategies, then define our signalling measures and signalling tests in a formal way and present their important properties. Section IV is devoted to de Finetti and permutation-invariant strategies. Finally, in Section V we connect all the relevant tools together using the weak and the strong lemmas, and prove our main theorem, Theorem 1 (the extension of the theorem to games with incomplete support is relatively straightforward and is given in Appendix A). We conclude in Section VI with open questions and a discussion of possible extensions to the classical and quantum case.

II. Preliminaries

Throughout the proof many constants and parameters are used. For convenience, apart from introducing them when necessary, we list all of them together with their role in Table I. Similarly, a list of relevant events and their symbols appears in Table II.

We use the letters q, r and s to denote tuples of questions and a and b to denote tuples of answers. In the following we define the notation using only q and a. Furthermore, $\wedge$ denotes 'logical-and', $\vee$ 'logical-or', and $\neg$ 'logical-not'.

A. Games and Strategies

In this work we consider a general m-player game $G = (\mathcal{Q}, \mathcal{A}, Q, R)$ as defined in Definition 3 in the previous section. A strategy for a game G is described by a conditional probability distribution $\mathrm{O}_{A|Q} : \mathcal{A} \times \mathcal{Q} \to [0,1]$ as defined in Definition 4. For the joint questions-answers distribution we use $\mathrm{O}_{AQ} = Q \times \mathrm{O}_{A|Q}$.

Definition 5 (Winning Probability): The winning probability of a strategy $\mathrm{O}_{A|Q}$ in game $G = (\mathcal{Q}, \mathcal{A}, Q, R)$ is given by $w\left(\mathrm{O}_{A|Q}\right) = \sum_{q,a} Q(q)R(q,a)\mathrm{O}_{A|Q}(a|q)$.

We use the following definition to measure the distance between two one-game strategies.

Definition 6 (Distance Measure): The distance between $\mathrm{K}_{A|Q}$ and $\mathrm{R}_{A|Q}$ is defined as

$$\left|\mathrm{K}_{A|Q} - \mathrm{R}_{A|Q}\right|_1 = \mathbb{E}_{q\in\mathcal{Q}} \sum_{a\in A} \left|\mathrm{K}_{A|Q}(a|q) - \mathrm{R}_{A|Q}(a|q)\right|$$

where the m-tuples of questions $q \in \mathcal{Q}$ are distributed according to Q defined by the game G.

Note that this is not the standard definition – instead of a maximisation over the tuples of questions as in the usual definition of the trace distance we consider the average over the tuples according to the game distribution. Therefore, the distance between the strategies depends on the specific game G considered.

In the repeated game G^n the referee asks each of the players n questions, all at once. The questions are chosen according to the distribution $Q^{\otimes n}$, i.e., independently using Q. The answers are then checked independently according to the winning condition R. A strategy for the repeated game G^n is denoted by $\mathrm{P}_{\vec{A}|\vec{Q}} : \mathcal{A}^n \times \mathcal{Q}^n \to [0,1]$ and the joint questions-answers distribution is then $P_{\vec{A}\vec{Q}} = Q^{\otimes n} \times \mathrm{P}_{\vec{A}|\vec{Q}}$. When the distributions are clear from the context we sometimes omit the subscripts and write just O and P.

When considering many questions-answers pairs in the repeated game we denote all the questions and answers as vectors $\vec{q}, \vec{a}$. We use a subscript index as in $\vec{q}_j$ to denote the

j'th tuple of questions given to the players. We denote by $\mathrm{O}_{A|Q}^{\otimes n}$ a product of n identical strategies $\mathrm{O}_{A|Q}$. That is, $\mathrm{O}_{A|Q}^{\otimes n}$ is defined according to $\mathrm{O}_{A|Q}^{\otimes n}(\vec{a}|\vec{q}) = \prod_{j=1}^{n} \mathrm{O}_{A|Q}(a_j|q_j)$ for all $\vec{a}, \vec{q}$.

For any m-tuple of questions $q = (q^1, \dots, q^m) \in \mathcal{Q}$ and any $i \in [m] = \{1, \cdots, m\}$ we denote by q^i, using a superscript index, the question given to the i'th player by the referee, and by $q^{\bar{i}} = (q^1, \dots q^{i-1}, q^{i+1}, \cdots, q^m)$ the $(m-1)$-tuple of questions given to all the players but i. Similarly, for a subset $I \subset [m]$, q^I denotes the questions given to all the players $i \in I$ and $q^{\bar{I}}$ denotes the complementary set of questions, i.e., the questions given to all the players $i \notin I$. An analogous notation is used for the answers. Similarly, when considering many questions-answers paris, $\vec{q}^i$ denotes *all* the questions given to the i'th player, and so on.

A tuple of questions $q = (q^1, \dots q^{i-1}, q^i, q^{i+1}, \cdots, q^m)$ can then be also written as $(q^i, q^{\bar{i}})$ where it is understood which player gets which question. Therefore in this notation $Q(q^i, q^{\bar{i}}) = Q(q)$ and similarly $\mathrm{O}(a^i, a^{\bar{i}}|q^i, q^{\bar{i}}) = \mathrm{O}(a|q)$. Moreover, $Q(q^i|q^{\bar{i}}) = \frac{Q(q^i,q^{\bar{i}})}{\sum_{r^i} Q(r^i,q^{\bar{i}})}$ denotes the probability that the i'th player receives question q^i *given* that the other players receive $q^{\bar{i}}$.

In the following we prove Theorem 1, which applies to games with complete-support. A game has complete-support if every possible combination of questions to the players has some non-zero probability according to the question distribution Q. Formally,

Definition 7 (Complete-Support Game): An m-player game has complete-support if for every possible combination of questions to the players $q^1, \dots, q^m$ (i.e., $q^1, \dots, q^m$ such that for all $i \in [m]$ there exist s^i for which $Q\left((s^1, \dots, s^{i-1}, q^i, s^{i+1}, \dots, s^m)\right) > 0$), $Q(q) > 0$.

B. Estimated Strategies

The specific questions and answers in one run of the repeated game $\vec{q}, \vec{a}$ are also called the data of the game. As mentioned in the previous section, the data $\vec{q}, \vec{a}$ is divided into two disjoint sets which we call the test data and the game data. We denote the $n/2$ tuples of test questions and answers by $\vec{q^t}, \vec{a^t}$ respectively and the $n/2$ tuples of game questions and answers by $\vec{q^g}, \vec{a^g}$ respectively. Using this notation $\vec{q}$ is the concatenation of $\vec{q^t}$ and $\vec{q^g}$ and $\vec{a}$ is the concatenation of $\vec{a^t}$ and $\vec{a^g}$. Note that although we denote here the test questions as appearing before the game questions, they are indistinguishable from one another, as they are chosen according to the exact same distribution Q. Had this not been the case, the permutation invariance symmetry would have been broken.

Given the test data $\vec{q^t}, \vec{a^t}$ we create an estimation $\mathrm{O}_{A|Q}^{\mathrm{EST1}}$ of a one-game strategy in the following way. For every tuple of questions $q \in \mathcal{Q}$ and answers $a \in \mathcal{A}$ let f_a^q be the frequency of the tuple of answers a in $\vec{a^t}$ restricted to the indices $j \in [n/2]$ for which $\vec{q^t}_j = q$ (if q did not appear at all set $f_a^q = 0$). Then define $\mathrm{O}_{A|Q}^{\mathrm{EST1}}$ such that $\mathrm{O}_{A|Q}^{\mathrm{EST1}}(a|q) = f_a^q$.

Similarly $\mathrm{O}_{A|Q}^{\mathrm{EST2}}$ is created in the same way, using the game data $\vec{q^g}, \vec{a^g}$ (see Figure 1). Note that since $\mathrm{P}_{\vec{A}|\vec{Q}}$ might be signalling between the different n tuples of questions, both $\mathrm{O}_{A|Q}^{\mathrm{EST1}}$ and $\mathrm{O}_{A|Q}^{\mathrm{EST2}}$ can depend also on the other questions which are not considered in the estimation process.

To evaluate the accuracy of the estimation process described above we will use the following lemma – an application of Sanov's theorem (see, e.g., [36] Section 11.4) to our scenario.

Lemma 8: Let $\delta(l, \epsilon) = (l+1)^{|\mathcal{A}|\cdot|\mathcal{Q}|-1} e^{-l\epsilon^2/2}$. *Then for every i.i.d. strategy* $\mathrm{O}_{A|Q}^{\otimes l}$,

$$\Pr_{\vec{a},\vec{q}\sim \mathrm{O}_{AQ}^{\otimes l}} \left[|\mathrm{O}_{A|Q}^{\mathrm{EST}} - \mathrm{O}_{A|Q}|_1 > \epsilon \right] \leq \delta(l, \epsilon)$$

where $\mathrm{O}_{A|Q}^{\mathrm{EST}}$ *is estimated as above from the data* $\vec{a}, \vec{q}$.

C. Linear Programs

Linear programs (see, e.g., [37]) are a useful tool when considering non-signalling games, as the non-signalling constraints are linear. The following results regarding the sensitivity of linear programs will be of use for us.

Lemma 9 (Sensitivity Analysis of Linear Programs, [37] Section 10.4): Let $\max\{c^T x | Ax \leq b\}$ *be a primal linear program and* $min\{b^T y | A^T y = c, y \geq 0\}$ *its dual. Denote the optimal value of the programs by* ω *and the optimal dual solution by* $y^\star$. *Then the optimal value of the perturbed program* $\omega_e = max\{c^T x | Ax \leq b + e\}$ *for some perturbation* e *is bounded by* $\omega_e \leq \omega + e^T y^\star$.

Lemma 10 (Dual Optimal Solution Bound, [37] Section 10.4): Let A *be an* $r_1 \times r_2$*-matrix and let* Δ *be such that for each non-singular sub-matrix* B *of* A *all entries of* B^{-1} *are at most* Δ *in absolute value. Let* c *be a row vector of dimension* r_2 *and let* $y^\star$ *be the optimal dual solution of* $min\{b^T y | A^T y = c, y \geq 0\}$. *Then*

$$\kappa = \sum_{j=1}^{r_1} |y_j^\star| \leq r_2 \Delta \sum_{j=1}^{r_2} |c_j|.$$

III. Detecting Signalling

A. The Non-Signalling Constraints

We start by defining a non-signalling strategy. To simplify notation we define it using one-game strategies $\mathrm{O}_{A|Q}$. The definition is identical for the strategies $\mathrm{P}_{\vec{A}|\vec{Q}}$.

Definition 11 (Non-Signalling Strategy): An m-player strategy $\mathrm{O}_{A|Q}$ is called non-signalling if for any set of players $I \subset [m]$, for all $a^{\bar{I}}, q^{\bar{I}}, q^I$, and r^I,

$$\mathrm{O}_{A|Q}(\circ, a^{\bar{I}}|q^I, q^{\bar{I}}) = \mathrm{O}_{A|Q}(\circ, a^{\bar{I}}|r^I, q^{\bar{I}})$$

where $\circ$ denotes a marginal, e.g., $\mathrm{O}_{A|Q}(\circ, a^{\bar{I}}|q^I, q^{\bar{I}}) = \sum_{a_i|i\in I} \mathrm{O}_{A|Q}(a|q^I, q^{\bar{I}})$.

Alternatively, one can define a non-signalling strategy using a set of linearly independent non-signalling constraints from which all the constraints in the above definition follow.

Lemma 12 ([38, Lemma 2.7]): An m*-player strategy* $\mathrm{O}_{A|Q}$ *is non-signalling if and only if for any player* $i \in [m]$, *for all* $a^{\bar{i}}, q^{\bar{i}}, q^i$, *and* r^i,

$$\mathrm{O}_{A|Q}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) = \mathrm{O}_{A|Q}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}). \tag{2}$$

From Equation (2) it is clear that for every i and $q^{\bar{i}}$ the marginal states $\mathrm{O}_{A|Q}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}})$ are all equivalent and independent of q^i. Therefore another equivalent formulation of the non-signalling constraints is given by

$$\mathrm{O}_{A|Q}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) = \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}_{A|Q}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})$$

For all $a^{\bar{i}}, q^{\bar{i}}$ and q^i. Here we defined the marginal, which is independent of r^i, as an average over the different $\mathrm{O}_{A|Q}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})$, where the average is taken according to the distribution of the game question Q. It is easy to verify that this condition is equivalent to Equation (2).

We can now write the optimisation problem of finding the optimal winning probability in a complete-support game G using a non-signalling strategy as the following linear program over the variables $\mathrm{O}(a|q)$:

$$\begin{aligned}
\max\ & \sum_{q,a} Q(q)R(q,a)\mathrm{O}(a|q) && \text{(3a)}\\
\text{s.t. } & Q(q^i, q^{\bar{i}})\Big[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) \\
& \qquad - \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})\Big] = 0 \\
& \qquad\qquad \forall i, q^i, q^{\bar{i}}, a^{\bar{i}} && \text{(3b)}\\
& \sum_a \mathrm{O}(a|q) = 1 \quad \forall q && \text{(3c)}\\
& \mathrm{O}(a|q) \geq 0 \quad \forall a, q && \text{(3d)}
\end{aligned}$$

The objective function, Equation (3a), is exactly the winning probability of the game using strategy $\mathrm{O}(a|q)$ as defined in Definition 5. Equations (3c) and (3d) are the normalisation and positivity constraints on the strategy $\mathrm{O}(a|q)$.

In Equation (3b) all the non-signalling constraints are listed, up to a factor of $Q(q)$ which does not change the constraints when considering complete-support games, but will be important later in the following section. We note that the only place in the proof where the complete-support property of the game is used is for writing down the linear program above. In Appendix A we explain the implications of the linear program (3) to games with incomplete support. In particular, in Appendix A-A we show how to modify program (3) for the case of two-player games with incomplete support such that our result still holds. In Appendix A-B we show how one can slightly modify the parallel repetition procedure to derive a general (although modified) threshold theorem for any game.

Next, one can relax the linear program (3) to the following:

$$\begin{aligned}
\max\ & \sum_{q,a} Q(q)R(q,a)\mathrm{O}(a|q) \\
\text{s.t. } & Q(q^i, q^{\bar{i}})\Big[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) \\
& \qquad - \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})\Big] \leq 0 \\
& \qquad\qquad \forall i, q^i, q^{\bar{i}}, a^{\bar{i}} && \text{(4a)}\\
& \sum_a \mathrm{O}(a|q) = 1 \quad \forall q \\
& \mathrm{O}(a|q) \geq 0 \quad \forall a, q
\end{aligned}$$

To see that the relaxation of the non-signalling constraints (3b) to the constraints (4a) does not change the program, i.e., does not change the value of the optimal solution, assume there exists $i, q^i, q^{\bar{i}}, a^{\bar{i}}$ for which

$$Q(q^i, q^{\bar{i}})\Big[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})\Big] < 0.$$

That is, $\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}})$ is smaller than the average $\sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})$, and therefore there must be some s^i for which $\mathrm{O}(\circ, a^{\bar{i}}|s^i, q^{\bar{i}})$ is larger than the average, meaning,

$$Q(s^i, q^{\bar{i}})\Big[\mathrm{O}(\circ, a^{\bar{i}}|s^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}})\Big] > 0,$$

but this contradicts the constraints in (4a).

The dual program of the primal (4) is given below.

$$\begin{aligned}
\min\ & \sum_q z(q) \\
\text{s.t. } & z(q) + \sum_i y_i(q, a^{\bar{i}})Q(q) \\
& - \sum_i \sum_{\substack{r| \\ r^{\bar{i}}=q^{\bar{i}}}} y_i(r, a^{\bar{i}})Q(r)Q(q^i|q^{\bar{i}}) \\
& \geq Q(q)R(q,a) \quad \forall a, q && \text{(5a)}\\
& y_i(q, a^{\bar{i}}) \geq 0 \quad \forall i, q, a^{\bar{i}}
\end{aligned}$$

B. Signalling Measure

Given a general strategy $\mathrm{O}_{A|Q}$ we would like to measure the amount of signalling from every player $i \in [m]$ to all the other players together. Following the linear program (4), we quantify signalling using Definition 13 below.

In the definition we derive all the relevant conditional and marginal distributions from O_{AQ}. Concretely we use the following notation: $\mathrm{O}(\circ, b^{\bar{i}}|s^i, s^{\bar{i}}) = \sum_{b^i} \mathrm{O}(b^i, b^{\bar{i}}|s^i, s^{\bar{i}})$ as before, $\mathrm{O}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}}) = \sum_{b^i, s^i} \mathrm{O}(b^i, b^{\bar{i}}, s^i, s^{\bar{i}})$, and

$$\mathrm{O}(\circ, s^i|b^{\bar{i}}, s^{\bar{i}}) = \sum_{b^i} \mathrm{O}(b^i, s^i|b^{\bar{i}}, s^{\bar{i}}) = \sum_{b^i} \frac{\mathrm{O}(b^i, b^{\bar{i}}, s^i, s^{\bar{i}})}{\mathrm{O}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}})}.$$

Definition 13 (Signalling Measure): The signalling of strategy $\mathrm{O}_{A|Q}$ in direction $i \to \bar{i}$ using outputs $b^{\bar{i}}$ and inputs $s^i, s^{\bar{i}}$ is defined as

$$\begin{aligned}
&\mathrm{Sig}_{(i, b^{\bar{i}}, s^i, s^{\bar{i}})}(\mathrm{O}) \\
&\quad = Q(s^i, s^{\bar{i}})\Big[\mathrm{O}(\circ, b^{\bar{i}}|s^i, s^{\bar{i}}) - \sum_{r^i} Q(r^i|s^{\bar{i}})\mathrm{O}(\circ, b^{\bar{i}}|r^i, s^{\bar{i}})\Big] && \text{(6)}\\
&\quad = \mathrm{O}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}})\Big[\mathrm{O}(\circ, s^i|b^{\bar{i}}, s^{\bar{i}}) - Q(s^i|s^{\bar{i}})\Big]. && \text{(7)}
\end{aligned}$$

That is, we have a signalling measure for every $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$. If $\mathrm{Sig}_{(i, b^{\bar{i}}, s^i, s^{\bar{i}})}(\mathrm{O}) > 0$ we say that the strategy is signalling in direction $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$. A negative signalling value, $\mathrm{Sig}_{(i, b^{\bar{i}}, s^i, s^{\bar{i}})}(\mathrm{O}) < 0$, is not relevant due to the inequality in Equation (4a).

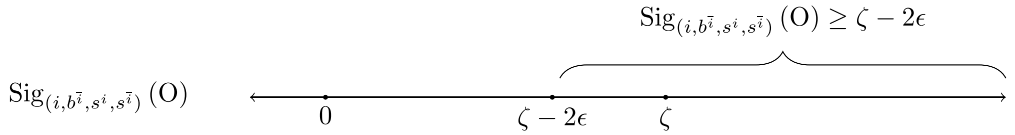
$\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) \geq \zeta - 2\epsilon$
$\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O})$
0
$\zeta - 2\epsilon$
ζ

Fig. 2. The different forms of signalling: every i and every $b^{\bar{i}}, s^i, s^{\bar{i}}$ define a line as in the figure. The value of $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O})$ tells us exactly where we are on the line.

The two forms of $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O})$ given in equations (6) and (7) are equivalent according to Bayes' rule and they will be useful in different places in the proof. Equation (7) for example allows us to quantify the amount by which input s^i is more or less probable given $b^{\bar{i}}$, compared to the prior $Q(s^i|s^{\bar{i}})$ (this will be useful in the proof of Lemma 25).

The following lemma shows that our measure of signalling is continuous. That is, if two strategies are close to one another according to Definition 6 then their signalling values are also close. The proof is given in Appendix B.

Lemma 14 (Continuity of Sig): Let O_1 *and* O_2 *be two one-game strategies such that* $\left|\mathrm{O}_1 - \mathrm{O}_2\right|_1 \leq \epsilon$*. Then for all* i, $b^{\bar{i}}$, s^i *and* $s^{\bar{i}}$,

$$\left|\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}_1) - \mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}_2)\right| \leq 2\epsilon.$$

C. Signalling Tests

In the following we will need an operational way of testing whether a one-game strategy $\mathrm{O}_{A|Q}$ is signalling. This can be done by using many copies of $\mathrm{O}_{A|Q}$ – given data $\vec{q}, \vec{a}$ which is distributed according to many independent copies of O_{AQ} it is possible to create an estimation of $\mathrm{O}_{A|Q}$, $\mathrm{O}^{\mathrm{EST1}}_{A|Q}$, and then evaluate $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\mathrm{O}^{\mathrm{EST1}}\right)$.

To formulate this process we first define an indicator function which will be used in the test. More precisely, for every tuple $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ we define a function $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})} : \mathcal{Q}^t \times \mathcal{A}^t \to \{0, 1\}$:

$$\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\vec{q^t}, \vec{a^t}) = \begin{cases} 1 & \text{if } \mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\mathrm{O}^{\mathrm{EST1}}\right) \geq \zeta - 2\epsilon \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

where $\mathrm{O}^{\mathrm{EST1}}$ is estimated from $\vec{q^t}, \vec{a^t}$ and $\zeta, \epsilon > 0$ are parameters satisfying $\zeta \geq 7\epsilon$ and $\epsilon \leq \min_q Q(q)$. See Figure 2 for a visualisation of the different signalling forms $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ and the signalling values considered in the test.

The following observation will be crucial later on:

Remark 15: According to Definition 13, in order to evaluate $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\mathrm{O}^{\mathrm{EST1}}\right)$ there is no need to know $\mathrm{O}^{\mathrm{EST1}}$ completely; only the marginals $\mathrm{O}^{\mathrm{EST1}}(\circ, b^{\bar{i}}|r^i, s^{\bar{i}})$ for all r^i are needed.

For every $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ we can now consider a signalling test. Given a strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ for the repeated game G^n we sample n tuples of questions $\vec{q}$ using the game distribution $Q^{\otimes n}$ and use them to get n tuples of answers $\vec{a}$ which are distributed according to $\mathrm{P}_{\vec{A}|\vec{Q}}$. Finally, if $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\vec{q^t}, \vec{a^t}) = 1$ the test accepts, and otherwise rejects.[4] Throughout the paper we denote by $\mathbf{T}$ the event of the test accepting (where the index $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ is clear from the context). Note that if a question s does not appear in the test data $\vec{q^t}$ the test $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ rejects by definition.

The following lemma shows that the test is reliable when applied to an i.i.d. strategy $\mathrm{O}^{\otimes n}_{A|Q}$. That is, if $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) \geq \zeta$ the test will detect it with high probability, i.e. the test will accept with high probability, and if $\mathrm{O}_{A|Q}$ is non-signalling then the test will reject with high probability. The proof can be found in Appendix B.

Lemma 16 (Reliable Signalling Test): Assume the players share an i.i.d. strategy $\mathrm{O}^{\otimes n}_{A|Q}$*. For every* $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$,

1) *If* $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) \geq \zeta$ *then* $\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{O}^{\otimes n}_{AQ}}[\mathbf{T}] > 1 - \delta$.
2) *If* $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) = 0$ *then* $\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{O}^{\otimes n}_{AQ}}[\neg\mathbf{T}] > 1 - \delta$.

where $\delta = \delta\left(\frac{n}{2}, \epsilon\right) = \left(\frac{n}{2} + 1\right)^{|\mathcal{A}|\cdot|\mathcal{Q}|-1} e^{-n\epsilon^2/4}$.

Given a specific signalling test $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ we define two relevant sets of one-game strategies:

$$\begin{aligned}\sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})} = \Big\{\mathrm{O}\Big|\forall\bar{\mathrm{O}} \text{ s.t. } |\mathrm{O} - \bar{\mathrm{O}}|_1 \leq \epsilon, \\ \bar{\mathrm{O}} \text{ is } \zeta \text{ signalling or more in } (i, b^{\bar{i}}, s^i, s^{\bar{i}})\Big\}\end{aligned} \tag{9}$$

$$\begin{aligned}\Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})} = \Big\{\mathrm{O}\Big|\exists\bar{\mathrm{O}} \text{ s.t. } |\mathrm{O} - \bar{\mathrm{O}}|_1 \leq \epsilon \\ \wedge \mathrm{Pr}_{\vec{a},\vec{q}\sim\bar{\mathrm{O}}^{\otimes n}_{AQ}}[\mathbf{T}] > \delta\Big\}\end{aligned} \tag{10}$$

The following two lemmas allow us to address these sets also according to the signalling values of the relevant strategies (see also Figure 3).

Lemma 17: For all $\mathrm{O} \notin \sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$, $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) < \zeta + 2\epsilon$.

Lemma 18: Let $\nu > 0$ be any parameter such that $\nu < \zeta - 6\epsilon$. Then

$$\forall\mathrm{O} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}, \quad \mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) > \nu.$$

Lemma 17 follows right away from Lemma 14 and the definition of $\sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$. Lemma 18 is proven in Appendix B.

IV. Using de Finetti Strategies

In this section we start analysing the relation between the test questions-answers $\vec{q^t}, \vec{a^t}$ and the game questions-answers $\vec{q^g}, \vec{a^g}$ in one instance of the repeated game G^n

[4] Note that as $\mathrm{P}_{\vec{A}|\vec{Q}}$ can be signalling between the different n tuples of questions-answers one has to input all the questions before getting the test answers. That is, even though the test considers only the test data, the test questions themselves are not sufficient to get all the necessary information.

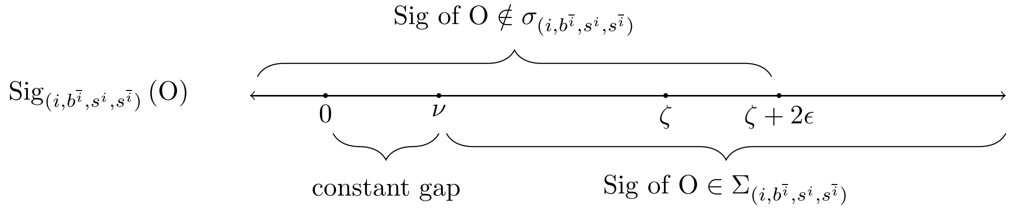
Sig of O $\notin \sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$
$\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O})$
0
ν
ζ
$\zeta + 2\epsilon$
constant gap
Sig of O $\in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$

Fig. 3. Visualization of the signalling values which are relevant for Lemma 18 and the sets $\sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$, $\Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$.

using a strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$. More precisely, we denote the one-game strategy which is estimated from $\vec{q^g}, \vec{a^g}$ by $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$, and we are interested in knowing what is the probability that $\mathrm{O}^{\mathrm{EST2}}_{A|Q} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ or $\mathrm{O}^{\mathrm{EST2}}_{A|Q} \in \sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ given the result of $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\vec{q^t}, \vec{a^t})$. We denote the event $\mathrm{O}^{\mathrm{EST2}}_{A|Q} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ by $\mathbf{in}^{\Sigma}$ and similarly the event $\mathrm{O}^{\mathrm{EST2}}_{A|Q} \in \sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ by $\mathbf{in}^{\sigma}$, where the indices of the sets are clear from the context.

We first do this for any i.i.d. strategy and then extend the results to any permutation-invariant strategy using a de Finetti reduction [27].

A. de Finetti Strategies

As mentioned in Section I, de Finetti strategies are strategies that can be written as a convex combination of i.i.d. strategies. Formally,

Definition 19 (de Finetti Strategy): A de Finetti strategy $\tau_{\vec{A}|\vec{Q}}$ is a strategy of the form

$$\tau_{\vec{A}|\vec{Q}} = \int \mathrm{O}^{\otimes n}_{A|Q} \mathrm{dO}_{A|Q},$$

where $\mathrm{dO}_{A|Q}$ is some measure on the space of one-game strategies.

In the following lemma we are interested in the relation between the test questions-answers $\vec{q^t}, \vec{a^t}$ and the game questions-answers $\vec{q^g}, \vec{a^g}$ in one instance of the repeated game G^n. For i.i.d. strategies (and therefore also for de Finetti strategies) this is simple: $\vec{q^t}, \vec{a^t}$ and $\vec{q^g}, \vec{a^g}$ are independent of each other and conditioning on a property of one of them does not affect the other.

Lemma 20: For a de Finetti strategy $\tau_{\vec{A}|\vec{Q}}$ and every $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$

1) $\Pr_{\vec{a},\vec{q} \sim \tau_{\vec{A}\vec{Q}}} \left[\mathbf{T} \wedge \neg \mathbf{in}^{\Sigma}\right] \leq \delta$
2) $\Pr_{\vec{a},\vec{q} \sim \tau_{\vec{A}\vec{Q}}} \left[\neg \mathbf{T} \wedge \mathbf{in}^{\sigma}\right] \leq \delta$

where $\tau_{\vec{A}\vec{Q}} = Q^{\otimes n} \times \tau_{\vec{A}|\vec{Q}}$.

The proof of this lemma (given in Appendix C) follows from Sanov's theorem stated in Lemma 8. Intuitively, if the event $\mathbf{T}$ holds then $\mathrm{O}^{\mathrm{EST1}}$ is signalling and therefore so should $\mathrm{O}^{\mathrm{EST2}}$ be, and vice versa.

B. de Finetti Reductions

Of course, considering just de Finetti strategies is not interesting by itself. Luckily, we can now use a de Finetti reduction to extend the results of the previous section to any permutation-invariant strategy, where the permutation is performed on the questions-answers pairs (we do not permute the players). As the repeated game G^n is by itself permutation invariant we can restrict the strategies of the players to be permutation invariant without loss of generality.

Definition 21 (Permutation Invariance): Given a strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ and a permutation π of the questions and answers we denote by $\mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi$ the strategy which is defined by

$$\forall \vec{a}, \vec{q} \quad \left(\mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi\right)(\vec{a}|\vec{q}) = \mathrm{P}_{\vec{A}|\vec{Q}}(\pi(\vec{a})|\pi(\vec{q})).$$

$\mathrm{P}_{\vec{A}|\vec{Q}}$ is permutation invariant if for any permutation π, $\mathrm{P}_{\vec{A}|\vec{Q}} = \mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi$.

The following lemma shows that we can restrict our analysis to permutation-invariant strategies without loss of generality.

Lemma 22: For every strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ for the repeated game G^n there exists a permutation-invariant strategy $\tilde{\mathrm{P}}_{\vec{A}|\vec{Q}}$ such that $w\left(\mathrm{P}_{\vec{A}|\vec{Q}}\right) = w\left(\tilde{\mathrm{P}}_{\vec{A}|\vec{Q}}\right)$.

Proof: Given $\mathrm{P}_{\vec{A}|\vec{Q}}$ define its permutation-invariant version to be

$$\tilde{\mathrm{P}}_{\vec{A}|\vec{Q}} = \frac{1}{n!} \sum_{\pi} \mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi.$$

The winning probability of the game is linear in the strategy, therefore we have

$$w\left(\tilde{\mathrm{P}}_{\vec{A}|\vec{Q}}\right) = w\left(\frac{1}{n!} \sum_{\pi} \mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi\right) = \frac{1}{n!} \sum_{\pi} w\left(\mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi\right). \tag{11}$$

Since the tuples of questions in the repeated game are chosen in an i.i.d. manner and the winning condition is checked for each tuple separately, the winning probability is indifferent to the ordering of the questions-answers pairs. As π permutes the tuples of questions and answers together we have $w\left(\mathrm{P}_{\vec{A}|\vec{Q}} \circ \pi\right) = w\left(\mathrm{P}_{\vec{A}|\vec{Q}}\right)$.

Combining this with Equation (11) we get $w\left(\tilde{\mathrm{P}}_{\vec{A}|\vec{Q}}\right) = w\left(\mathrm{P}_{\vec{A}|\vec{Q}}\right)$. □

Lemma 23 (de Finetti Reduction for Conditional Probability Distributions [27]): Let $c = (n+1)^{|\mathcal{Q}|(|\mathcal{A}|-1)}$. There exists a de Finetti strategy $\tau_{\vec{A}|\vec{Q}}$ such that for every permutation-invariant strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$

$$\forall \vec{a}, \vec{q} \quad \mathrm{P}_{\vec{A}|\vec{Q}}(\vec{a}|\vec{q}) \leq c \cdot \tau_{\vec{A}|\vec{Q}}(\vec{a}|\vec{q}).$$

The de Finetti strategy $\tau_{\vec{A}|\vec{Q}}$ is constructed explicitly in [27] but the specific construction is not relevant for our purposes. In some special cases the constant c in Lemma 23 can also be made smaller by taking into account symmetries of the game G itself. For further details see [27].

We now use the de Finetti reduction to show that the properties proven in Lemma 20 for the de Finetti strategy also hold true for permutation-invariant strategies, although with slightly weaker parameters. Concretely, the bound of 2δ in Lemma 20 is replaced by $2c\delta$ in the following lemma. Nevertheless, the bound still decreases exponentially fast with the number of repetitions.[5]

Lemma 24 (Reduction): For every permutation-invariant strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ *and every* $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$

1) $\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[\mathbf{T} \wedge \neg\mathbf{in}^{\Sigma}\right] \leq c\delta$.
2) $\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[\neg\mathbf{T} \wedge \mathbf{in}^{\sigma}\right] \leq c\delta$.

Proof: We prove both of the claims together. Denote the relevant event by $E(\vec{a}, \vec{q})$ and note that for both events we can write

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[E(\vec{a}, \vec{q}) = 1\right] = \sum_{\substack{\vec{a},\vec{q}|\\ E(\vec{a},\vec{q})=1}} \mathrm{P}_{\vec{A}\vec{Q}}(\vec{a}, \vec{q}).$$

From Lemma 23 we get $\mathrm{P}_{\vec{A}\vec{Q}}(\vec{a}, \vec{q}) \leq c \cdot \tau_{\vec{A}\vec{Q}}(\vec{a}, \vec{q})$ and therefore

$$\begin{aligned}
&\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[E(\vec{a}, \vec{q}) = 1\right]\\
&\quad= \sum_{\substack{\vec{a},\vec{q}|\\ E(\vec{a},\vec{q})=1}} \mathrm{P}_{\vec{A}\vec{Q}}(\vec{a}, \vec{q})\\
&\quad\leq c \cdot \sum_{\substack{\vec{a},\vec{q}|\\ E(\vec{a},\vec{q})=1}} \tau_{\vec{A}\vec{Q}}(\vec{a}, \vec{q})\\
&\quad= c \cdot \mathrm{Pr}_{\vec{a},\vec{q}\sim\tau_{\vec{A}\vec{Q}}}\left[E(\vec{a}, \vec{q}) = 1\right].
\end{aligned}$$

Combining this with Lemma 20 proves the lemma. □

V. Threshold Theorem

In this section we prove our threshold theorem, Theorem 1. Before going into the details of the proof, let us explain the high-level idea.

First, to see the connection between what was done so far and a threshold theorem note that the winning probability of $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ in the game G, $w(\mathrm{O}^{\mathrm{EST2}}_{A|Q})$, is exactly the fraction of coordinates in which the game data $\vec{q^g}, \vec{a^g}$ satisfies the winning condition R. Therefore, in order to prove a threshold theorem it is sufficient to prove an upper bound on $w(\mathrm{O}^{\mathrm{EST2}}_{A|Q})$ which holds with high probability.

To do so we use the following sequence of lemmas. The first two lemmas bound the probability that the estimate $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ is significantly signalling[6] in any direction $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ for which $\mathrm{Pr}_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}[\mathbf{T}] \neq 0$. Lemma 25, which we also call the weak lemma, establishes that even conditioned on the test $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\vec{q^t}, \vec{a^t})$ detecting signalling the distribution $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ itself cannot be signalling with very high probability. The proof of the lemma is based on a reduction to a certain guessing game which is used to derive a contradiction between the conclusion that $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ would be signalling and the assumption that the overall distribution $\mathrm{P}_{\vec{A}|\vec{Q}}$ is not. Lemma 26, called the strong lemma, amplifies the conclusion of the weak lemma to show that $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ cannot display too much signalling, even only with small probability. The amplification is obtained by using the properties of permutation-invariant strategies which were proven in Lemma 24 in the previous section.

Having shown that with high probability $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ cannot be too signalling, Lemma 27 derives an upper bound on the winning probability $w(\mathrm{O}^{\mathrm{EST2}}_{A|Q})$. Intuitively, if the strategy $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ does not display strong signalling in any direction it should not lead to a large advantage over strictly non-signalling strategies in the game G. The quantitative argument is based on performing a sensitivity analysis of the appropriate linear program. The three lemmas are brought together in Lemma 28, from which Theorem 1 follows.

In the following lemmas we consider for simplicity the scenarios conditioned on the event in which all tuples of questions appear in the game data, denoted by **agq**, and hence $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ is a strategy. The probability that this event does not hold is exponentially small and will be taken into account in the final bound in Lemma 28. (See Table II to recall all other events which are used in the lemmas of this section).

We are now ready to prove the following lemmas and the threshold theorem.

Lemma 25 (Weak Lemma): Let n *be such that*

$$\frac{n}{\ln(n)} > 20|Q||A|\frac{\ln(2/\epsilon)}{\epsilon^2}, \tag{12}$$

and $\mathrm{P}_{\vec{A}|\vec{Q}}$ *a non-signalling strategy for* G^n*. For any* $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ *denote by* $\mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}$ *the probability distribution* $\mathrm{P}_{\vec{A}\vec{Q}}$ *conditioned on the event* $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\vec{q^t}, \vec{a^t}\right) = 1$*, whenever such a conditional probability distribution is defined. Then,*

$$\mathrm{Pr}_{\vec{a^g},\vec{q^g}\sim\mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}}\left[\mathbf{in}^{\Sigma}|\mathbf{agq}\right] < 1 - \sqrt{c\delta}. \tag{13}$$

Proof: In the proof all the probabilities are conditioned on the event **agq**, i.e., all tuples of questions appear in the game data. To ease notation we do not explicitly write it.

For every signalling test $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ and game questions for players $\bar{i}$ $\vec{q^g}^{\bar{i}}$ such that $\mathrm{Pr}_{\vec{a^g},\vec{q^g}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[\mathbf{T}|\vec{q^g}^{\bar{i}}\right] \neq 0$ we construct a guessing game. Our goal is to derive a contradiction by showing that if Equation (13) is not true, then the guessing game can be won with probability higher than the optimal non-signalling winning probability.

The guessing game is defined as follows. A referee gives the players $n/2$ independent m-tuples of game questions: players $\bar{i}$ get the questions $\vec{q^g}^{\bar{i}}$ and player i gets $\vec{q^g}^{i}$ distributed according to the prior $Q(q^i|q^{\bar{i}})$ (i.e., in $\vec{q^g}$ each tuple is distributed according to the questions distribution Q of the original game G). Players $\bar{i}$ are then allowed to communicate and their goal is to guess and output an index $j \in [n/2]$ such

[5] One would have liked to apply a similar argument to the winning probability of the repeated game right away. That is, $w(\mathrm{P}_{\vec{A}|\vec{Q}}) \leq cw(\tau_{\vec{A}|\vec{Q}})$. This claim is indeed correct, but not useful. A look at the explicit construction of $\tau_{\vec{A}|\vec{Q}}$ itself in [27] will reveal that it is a signalling strategy, hence no non-trivial bound on $w(\tau_{\vec{A}|\vec{Q}})$ holds a priori. For a further discussion see Section VI-B.

[6] In the words of the explanation given in Section I-E, this is where we prove that the signalling weight is exponentially small.

that $\vec{q^g}_j = (s^i, s^{\bar{i}})$ (there is always such an index since we condition on **agq**).

If the players share a non-signalling strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$ then the marginals of players $\bar{i}$ are the same for all $\vec{q^g}^i$. Therefore, their outputs $\vec{a}^{\bar{i}}$ do not give them any information about the question that the i'th player got from the referee (even when players $\bar{i}$ are allowed to communicate among themselves, but not with player i). The best non-signalling strategy is therefore to choose, uniformly at random, an index j for which $\vec{q^g}^{\bar{i}}_j = s^{\bar{i}}$. The winning probability is then given by $W_{\mathrm{ns}} = Q(s^i|s^{\bar{i}}) < 1$ (note that while the questions of players $\bar{i}$ are fixed in a specific instance of the guessing game, the questions of player i are still distributed according to the prior $Q(q^i|q^{\bar{i}})$).

We now show that if the players share $\mathrm{P}_{\vec{A}|\vec{Q}}$ for which

$$\Pr_{\vec{a^g},\vec{q^g}\sim\mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}}\left[\mathbf{in}^{\Sigma}|\vec{q^g}^{\bar{i}}\right] \geq 1-\sqrt{c\delta} \tag{14}$$

then they can win the above guessing game with probability higher than the optimal non-signalling winning probability W_{ns}.

The idea is as follows. The players share many identical copies of $\mathrm{P}_{\vec{A}|\vec{Q}}$. They use the questions given by the referee as the game questions $\vec{q^g}$ in all of the copies and choose, using shared randomness, the test questions $\vec{q^t}$ in each copy (i.e., there are different test questions for each copy). They input the questions into the copies of $\mathrm{P}_{\vec{A}|\vec{Q}}$. Players $\bar{i}$ then look for the first copy of $\mathrm{P}_{\vec{A}|\vec{Q}}$ in which the event $\mathbf{T}$ holds – such a copy exists as long as[7] $\Pr_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[\mathbf{T}|\vec{q^g}^{\bar{i}}\right] \neq 0$; they can find it since they are allowed to communicate among themselves and they know all the inputs for the test questions of player i (as they were chosen using shared randomness which is available to all the players). Recalling Remark 15, they have all the information they need. Player i does not need to know in which copy the test holds.

Using the chosen copy, the players choose a random index $j \in [n/2]$ such that $\vec{q^g}^{\bar{i}}_j = s^{\bar{i}}$ and $\vec{a^g}^{\bar{i}}_j = b^{\bar{i}}$.

Let us show that, as long as $\Pr_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[\mathbf{T}|\vec{q^g}^{\bar{i}}\right] \neq 0$, this strategy achieves a winning probability which is higher than W_{ns}. For the chosen copy the event $\mathbf{T}$ holds and hence $\vec{q^g}, \vec{a^g}$ can be seen as data which is distributed according to $n/2$ identical copies of $\mathrm{O}^{\mathrm{EST2}}$, which is with high probability in $\Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ according to Equation (14). From Lemma 18 this implies

$$\Pr_{\vec{a^g},\vec{q^g}\sim\mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}}\left[\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}^{\mathrm{EST2}}) > \nu|\vec{q^g}^{\bar{i}}\right] \geq 1-\sqrt{c\delta}, \tag{15}$$

where $\nu > 0$ is any parameter satisfying $\nu < \zeta - 6\epsilon$ (recall Lemma 18).

Using the definition of Sig in Equation (7) we know that if indeed $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}^{\mathrm{EST2}}) > \nu$ then $\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}}) > 0$ and

$$\begin{aligned}\mathrm{O}^{\mathrm{EST2}}(\circ, s^i|b^{\bar{i}}, s^{\bar{i}}) &> \frac{\nu}{\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}})} + Q(s^i|s^{\bar{i}})\\ &= \frac{\nu}{\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}})} + W_{\mathrm{ns}}.\end{aligned} \tag{16}$$

That is, by choosing an index for which $a^{\bar{i}} = b^{\bar{i}}$ players $\bar{i}$ increase the winning probability.

On the other hand, if $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}^{\mathrm{EST2}}) \leq \nu$, which can happen with probability $\sqrt{c\delta}$, then the players might decrease their winning probability. In the worst case the winning probability is 0.

Therefore, for the chosen copy (for which $\mathcal{T}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\vec{q^t}, \vec{a^t}) = 1$) we get the following winning probability

$$W \geq (1-\sqrt{c\delta})\left(\frac{\nu}{\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}})} + W_{\mathrm{ns}}\right) + \sqrt{c\delta}\cdot 0. \tag{17}$$

Thus, $W > W_{\mathrm{ns}}$ for

$$\nu > \frac{\sqrt{c\delta}}{1-\sqrt{c\delta}}W_{\mathrm{ns}} \geq \frac{\sqrt{c\delta}}{1-\sqrt{c\delta}}W_{\mathrm{ns}}\cdot\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}}). \tag{18}$$

Using $W_{\mathrm{ns}}\cdot\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}}) \leq 1$ and $\sqrt{c\delta} \leq (n+1)^{|Q||A|}e^{-n\epsilon^2/8}$ (see Table I), we see that as long as $n/\ln(n) > 20|Q||A|\epsilon^{-2}\ln(2/\epsilon)$ the quantity $\sqrt{c\delta}W_{\mathrm{ns}}\mathrm{O}^{\mathrm{EST2}}(\circ, b^{\bar{i}}, \circ, s^{\bar{i}})/(1-\sqrt{c\delta})$ is strictly less than ϵ. Assuming $\zeta \geq 7\epsilon$, there is a choice of ν that satisfies both (18) and the earlier condition that $\nu < \zeta - 6\epsilon$.

We get that Equation (14) must not hold for any $\vec{q^g}^{\bar{i}}$ and hence cannot hold also when we omit the conditioning on $\vec{q^g}^{\bar{i}}$. The lemma therefore follows. □

The bound given in Equation (13) is weak for two reasons. First, the game data $\vec{q^g}, \vec{a^g}$ is distributed according to the conditional distribution $\mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}$ and not according to $\mathrm{P}_{\vec{A}\vec{Q}}$ itself. Second, it only tells us that $\Pr_{\vec{a^g},\vec{q^g}\sim\mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}}\left[\mathrm{O}^{\mathrm{EST2}}_{A|Q} \notin \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\right] \geq \sqrt{c\delta}$, i.e., the probability that $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ has a small value of signalling is higher than $\sqrt{c\delta}$. We show how the statement in the weak lemma can be amplified using the de Finetti reduction from Lemma 24.

Lemma 26 (Strong Lemma): Let $\mathrm{P}_{\vec{A}|\vec{Q}}$ *be a permutation-invariant non-signalling strategy for* G^n *and* n *such that Equation* (18) *is satisfied. Then for any* $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$ *such that* $Q(s^i, s^{\bar{i}}) \neq 0$ *and* $Q(s^i|s^{\bar{i}}) \neq 1$,

$$\Pr_{\vec{a},\vec{q}\sim\mathrm{P}_{\vec{A}\vec{Q}}}\left[\mathbf{in}^{\sigma}|\mathbf{agq}\right] \leq 2\sqrt{c\delta}.$$

Proof: In the proof all the probabilities are conditioned on the event **agq**, i.e., all tuples of questions appear in the game data. To ease notation we do not explicitly write it. Note that while Lemma 24 is stated without conditioning on **agq**, the proof is easily adapted to show that the same statement holds while conditioning on it.

[7] To see this note that since the strategy is non-signalling between player i and players $\bar{i}$, players $\bar{i}$ can check in which copy the test passes even before player i inputs his question. Therefore, the probability to pass the test is independent of the game questions of player i and hence must be non-zero for any of them.

From Lemma 24 part 1 we get

$$\begin{aligned}\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}[\mathbf{T}] \;>\; \sqrt{c\delta} \\ \Rightarrow \mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}}\left[\neg\mathbf{in}^{\Sigma}\right] \leq \sqrt{c\delta}\end{aligned}$$

and this can be rewritten as

$$\begin{aligned}\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}[\mathbf{T}] \;>\; \sqrt{c\delta} \\ \Rightarrow \mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}|\mathcal{T}=1}}\left[\mathbf{in}^{\Sigma}\right] \geq 1-\sqrt{c\delta}.\end{aligned}$$

According to Lemma 25, this implies

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}[\mathbf{T}] > \sqrt{c\delta} \Rightarrow \mathrm{P}_{\vec{A}|\vec{Q}} \text{ is signalling.}$$

Therefore it must be that

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}[\mathbf{T}] \leq \sqrt{c\delta} \tag{19}$$

or alternatively,

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}[\neg\mathbf{T}] \geq 1-\sqrt{c\delta} \tag{20}$$

Next, combining Lemma 24 part 2 with Equation (20) we get

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{AQ|\mathcal{T}=0}}\left[\mathbf{in}^{\sigma}\right] \leq \sqrt{c\delta}.$$

Using Equation (19) we get

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[\mathbf{in}^{\sigma}\right] \leq 2\sqrt{c\delta}. \qquad \square$$

Lemma 26 tells us that if $\mathrm{P}_{\vec{A}|\vec{Q}}$ is a permutation-invariant non-signalling strategy then the probability that $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ is in a given set $\sigma_{(i,b^{\bar{i}},s^{i},s^{\bar{i}})}$ is exponentially small in the number of games. In the next lemma we use this property to get a bound on the winning probability of $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ in the game G.

Lemma 27: Let $\kappa = \sum_{j=1}^{d}|y^{\star}_{j}|$ *where* d *is the number of signalling tests and* $y^{\star}$ *is an optimal solution of the dual program* (5). *Let* $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ *be a strategy (i.e., we assume that the event* **agq** *holds) such that for all* $(i,b^{\bar{i}},s^{i},s^{\bar{i}})$, $\mathrm{O}^{\mathrm{EST2}}_{A|Q} \notin \sigma_{(i,b^{\bar{i}},s^{i},s^{\bar{i}})}$. *Then* $w(\mathrm{O}^{\mathrm{EST2}}_{A|Q}) \leq 1-\alpha+(\zeta+2\epsilon)\kappa$.

Proof: According to Lemma 17, if $\mathrm{O}^{\mathrm{EST2}}_{A|Q} \notin \sigma_{(i,b^{\bar{i}},s^{i},s^{\bar{i}})}$ for every $\sigma_{(i,b^{\bar{i}},s^{i},s^{\bar{i}})}$ then

$$\mathrm{Sig}_{(i,a^{\bar{i}},q^{i},q^{\bar{i}})}(\mathrm{O}^{\mathrm{EST2}}) < \zeta+2\epsilon \tag{21}$$

for every i and every $b^{\bar{i}}, s^{i}, s^{\bar{i}}$. That is, $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ is not "too signalling" in any direction. This can be used to bound the winning probability of $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ in the game G.

The following linear program describes the optimal winning probability of a strategy $\mathrm{O}_{A|Q}$ which fulfils Equation (21):

$$\begin{aligned}\max \;\; & \sum_{q,a} Q(q)R(q,a)\mathrm{O}(a|q) \\ \text{s.t. } & Q(q^{i},q^{\bar{i}})\Big[\mathrm{O}(\circ,a^{\bar{i}}|q^{i},q^{\bar{i}}) \\ & \qquad -\sum_{r^{i}} Q(r^{i}|q^{\bar{i}})\mathrm{O}(\circ,a^{\bar{i}}|r^{i},q^{\bar{i}})\Big] \leq \zeta+2\epsilon \\ & \qquad\qquad\qquad \forall i, q^{i}, q^{\bar{i}}, a^{\bar{i}} \\ & \sum_{a}\mathrm{O}(a|q) = 1 \quad \forall q \\ & \mathrm{O}(a|q) \geq 0 \quad \forall a,q\end{aligned} \tag{22}$$

As $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ is a strategy we have $\sum_{a}\mathrm{O}^{\mathrm{EST2}}(a|q) = 1$ for all q. Hence $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ satisfies all the constraints of the above program and therefore its winning probability in G is bounded by the optimal value of the program. Program (22) can be seen as a perturbation of the linear program (4), we can therefore bound its optimal value by using known tools for sensitivity analysis of linear programs, stated in Lemmas 9 and 10.

Denote by $y^{\star}$ an optimal solution of the dual program[8] (5) and let $\kappa = \sum_{j=1}^{d}|y^{\star}_{j}|$ where d is the number of signalling tests. That is, κ is the sum of all the dual variables which are associated to the non-signalling constraints.

According to Lemma 9 the perturbed winning probability is then bounded by

$$w_{e} \leq 1-\alpha+(\zeta+2\epsilon)\kappa. \qquad \square$$

To get κ in the above lemma, one can use any of the following:

1) Given a description of a game one can easily get κ by solving the dual program (5).[9]
2) If the game involves only 2 players, then following [33] one can get $\kappa \leq d$ where d is the number of different signalling tests ($d < m|\mathcal{Q}||\mathcal{A}|$).
3) Otherwise, the general bound of Lemma 10 can be used. In our case the bound reads $\kappa \leq |\mathcal{A}|^{2}|\mathcal{Q}|\Delta$, where Δ depends only on the game.[10]

Finally we are ready to prove the last lemma:

Lemma 28 (Main Lemma): Let $w(G) = 1-\alpha$ *be the optimal winning probability of a non-signalling strategy in G. Let* $0 < \beta \leq \alpha$ *be some constant and* n *the number of repetitions such that Equation* (18) *is satisfied. Then for any non-signalling strategy* $\mathrm{P}_{\vec{A}|\vec{Q}}$ *of the repeated game,*

$$\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[w(\mathrm{O}^{\mathrm{EST2}}_{A|Q}) > 1-\alpha+\beta\right] \leq 3d\sqrt{c\delta}.$$

Proof: Let $\zeta,\epsilon > 0$ be such that $\zeta+2\epsilon \leq \frac{\beta}{\kappa}$, $\epsilon \leq \min_{q} Q(q)$ and $7\epsilon \leq \zeta \leq 1$.

If all tuples of questions s appear at least once in the game data, i.e., the event **agq** holds, then we can apply Lemma 27 in combination with Lemma 26 and get

$$\begin{aligned}&\mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[w(\mathrm{O}^{\mathrm{EST2}}_{A|Q}) > 1-\alpha+\beta\middle|\mathbf{agq}\right] \\ &\quad\leq \mathrm{Pr}_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[\exists\sigma_{(i,b^{\bar{i}},s^{i},s^{\bar{i}})} \text{ s.t. } \mathbf{in}^{\sigma}\middle|\mathbf{agq}\right] \\ &\quad\leq d\cdot 2\sqrt{c\delta}.\end{aligned}$$

The probability that some tuple of questions does not appear in the game data (i.e., **agq** does not hold) is upper bounded by

$$\begin{aligned}|\mathcal{Q}|\left(1-\min_{s} Q(s)\right)^{n/2} &\leq |\mathcal{Q}|e^{-\min_{s} Q(s)n/2} \\ &\leq |\mathcal{Q}|e^{-\epsilon n/2} \leq d\delta\end{aligned}$$

[8] We are only interested in the value of $y^{\star}$ as $z^{\star}$ will not affect the bound.

[9] Solving the linear program is anyhow usually necessary for knowing the optimal non-signalling value $1-\alpha$.

[10] A similar bound was also used in [15].

and therefore all together we have

$$\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[w(\mathrm{O}^{\mathrm{EST2}}_{A|Q}) > 1-\alpha+\beta\right] \leq 3d\sqrt{c\delta}. \qquad \square$$

Our threshold theorem, Theorem 1, follows from Lemma 28:

Proof of Theorem 1: Let f^g, f^t and f denote the winning frequency in the game data, test data and the entire data respectively (i.e., the fraction of coordinates in which the players win the game). From Lemma 28 we know that $\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[f^g > 1-\alpha+\beta\right] \leq 3d\sqrt{c\delta}$, as the winning frequency in the game questions is exactly $w(\mathrm{O}^{\mathrm{EST2}}_{A|Q})$. As the game data and test data are symmetric (i.e., there is no difference between them except for the name we gave them), the same result also holds for f^t, $\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[f^t > 1-\alpha+\beta\right] \leq 3d\sqrt{c\delta}$.

Finally, as the winning frequency in the entire data is given by $f = \frac{1}{2}\left(f^t + f^g\right)$ we have

$$\begin{aligned}&\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[f > 1-\alpha+\beta\right]\\ &\quad\leq \Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[\left(f^t > 1-\alpha+\beta\right)\vee\left(f^g > 1-\alpha+\beta\right)\right]\\ &\quad\leq 6d\sqrt{c\delta}. \qquad (23)\end{aligned}$$

$\square$

The relations between all the constants and parameters of the theorem and the proofs are listed in Table I. Note that for any game and choice of parameters the bound $6d\sqrt{c\delta}$ is exponentially decreasing with the number of repetitions n.

To get a better feeling of the result, without trying to optimise it, one can make the following choices. Let $\epsilon = \frac{\beta}{10\kappa}$, $\zeta = 8\epsilon$ and $\nu = \epsilon$ (assuming $\min_q Q(q) > \frac{\beta}{10\kappa}$). Using these choices, our proof holds for n and β such that

$$\frac{n}{\ln(n)} > 20|\mathcal{Q}||\mathcal{A}|\frac{\ln(20\kappa/\beta)}{(\beta/10\kappa)^2} \qquad (24)$$

with the following constants in Theorem 1:

$$\begin{aligned}\mathcal{C}_1(G,n) &= 6m|\mathcal{Q}||\mathcal{A}|\,(n+1)^{|\mathcal{Q}||\mathcal{A}|-1}\,,\\ \mathcal{C}_2(G) &= (30\kappa)^{-2}. \qquad (25)\end{aligned}$$

The theorem then reads

$$\begin{aligned}&\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{Q}}}\left[f > 1-\alpha+\beta\right]\\ &\quad\leq 6m|\mathcal{Q}||\mathcal{A}|\,(n+1)^{|\mathcal{Q}||\mathcal{A}|-1}\,e^{-n\left(\frac{\beta}{30\kappa}\right)^2}. \qquad (26)\end{aligned}$$

A different choice of parameters can improve the dependency of the constants on the game G.

VI. Conclusions and Open Questions

A. Current Work and Possible Extensions

In this work a threshold theorem for multiplayer non-signalling games was proven. The threshold theorem given in Theorem 1 is applicable to any multiplayer complete-support game and for every two-player game (not necessarily with complete-support, as proven in Appendix A-A). Hence, all cases for which parallel repetition was already known prior to our work [13], [15] are covered by our proof. For multiplayer-games with incomplete support we considered a small modification of the parallel repetition procedure which results in Theorem 2. We believe a similar modification can be considered to extend the result of [15].

In both theorems it might be possible to improve the dependency of the result on the parameters of the considered game, i.e., improve the constants $\mathcal{C}_1(G,n)$ and $\mathcal{C}_2(G)$. The polynomial dependency of $\mathcal{C}_1(G,n)$ on the number of repetitions, on the other hand, is inherent to the use of the de Finetti theorem. Moreover, further investigation of the dual program (5) could lead to an explicit bound on $\mathcal{C}_2(G)$. This could then be used to extend Theorem 1 to games with incomplete support, as done for two-player games.

The most important contribution of this work is a new proof technique for parallel repetition theorems, based on ideas of de Finetti theorems and tomography. de Finetti theorems seem like a natural tool for parallel repetition theorems, yet, this is the first time that such a result is proven using a de Finetti theorem.

Apart from allowing a different point of view on parallel repetition questions, and the study of correlations in general, the new proof technique has several advantages over the previous proofs.

For instance, note that in the standard proofs of parallel repetition theorems, i.e., proofs following the approach of [12] such as [13]–[15], most of the difficulties arise due to the effect of conditioning on the event of winning some of the game repetitions. As this event is one that depends on the structure of the game and we have no control over it, it can introduce arbitrary correlations between the questions used in different repetitions of the game, a major source of difficulty for the remainder of the argument. In our proof we also need to analyse the effect of conditioning on a certain event, the event of the non-signalling test accepting, and this is done in Lemma 25, the weak lemma. However, the key advantage of our approach is that the test has a very specific structure, and in particular conditioning on the test passing can be done locally by the players in a way that respects the non-signalling constraints. As a result it is almost trivial to deal with the conditioning in the remainder of the proof. This shift from conditioning on an uncontrolled event, success in the game, to a highly controlled one, a non-signalling test that we design ourselves, is a key simplification that we expect to play an important role in any extension of our method to other scenario such as classical or quantum strategies. More specifically, by finding appropriate “non-classicality” and “non-quantumness” measures which can replace our signalling measure in Definition 13 one may be able to adapt the proof to the multiplayer classical and quantum cases as well. The results of Sections III and IV should follow easily for most “non-classicality” and “non-quantumness” measures of one-game strategies. The main difficulty, however, is finding a measure for which Lemma 25 can be proven.

B. What Parallel Repetition Tells us About de Finetti Theorems

In the light of the de Finetti reduction stated in Lemma 23, it is tempting to try and prove a parallel repetition theorem by claiming that for every permutation-invariant strategy $\mathrm{P}_{\vec{A}|\vec{Q}}$,

$$w\left(\mathrm{P}_{\vec{A}|\vec{Q}}\right) \leq c \cdot w\left(\tau_{\vec{A}|\vec{Q}}\right). \tag{27}$$

This claim is correct but, unfortunately, not very useful as $\tau_{\vec{A}|\vec{Q}}$ itself is signalling according to the explicit construction given in [27], hence, no non-trivial bound holds on $w\big(\tau_{\vec{A}|\vec{Q}}\big)$.

One might hope that this is just a technical problem; perhaps a different de Finetti reduction can be proven, where both $\mathrm{P}_{\vec{A}|\vec{Q}}$ and $\tau_{\vec{A}|\vec{Q}}$ can be taken to be non-signalling (or analogously, quantum or classical). Such a de Finetti reduction, if it existed, would have implied a strong parallel repetition theorem (up to the polynomial factor c) for any game right away using Equation (27). This however will stand in contradiction to known impossibility results, such as the result of [39].

We therefore learn an interesting fact about de Finetti reductions by considering parallel repetition theorems: in order to prove a general de Finetti reduction as in Lemma 23, the de Finetti strategy must have some signalling parts. Fortunately, as shown by our result, this does not render a proof for the non-signalling case impossible.

Appendix A
Extending the Result to General Games

Before we show how to extend the threshold theorem to games with incomplete support, let us explain why the proof given for Theorem 1 holds only for complete-support games.

As mentioned in the main text, the reason lies in the linear program (3), and more specifically, in the non-signalling constraints given in Equation (3b). Indeed, if for some q we have $Q(q) = 0$ then the relevant constraint in Equation (3b) is vacuous. It is therefore clear that in this case the constraints given in Equation (3b) are in fact relaxations of the standard non-signalling constraints given in Equation (2).

For some games, this relaxation of the non-signalling constraints is strict. For example,[11] consider a game of 3 players where the questions are uniformly distributed over $\mathcal{Q} = \{(0,0,1),(0,1,0),(1,0,0)\}$ and the winning condition is given by the following predicate:

$$R(q,a) = \begin{cases} 1 & \text{if } q = (0,0,1) \text{ and } a^1 = a^2 \\ 1 & \text{if } q = (0,1,0) \text{ and } a^1 = a^3 \\ 1 & \text{if } q = (1,0,0) \text{ and } a^2 \neq a^3 \\ 0 & \text{otherwise} \end{cases}$$

The optimal non-signalling winning probability in this game is $\frac{2}{3}$ (as can be shown by solving a linear program). However, in the linear program (3) there are no non-trivial constraints (i.e., all the constraints in Equation (3b) are of the form $0 = 0$). Hence, the optimal solution of the program (3) is 1, which is strictly larger than $\frac{2}{3}$. Thus even though the non-signalling conditions are enforced over all "relevant" questions, this does not suffice to guarantee that there exists a strategy achieving the resulting optimum success probability 1 and that can be extended to a non-signalling strategy defined on all questions.

For games with incomplete support in which the optimal value of program (3) is not trivial (i.e., it is smaller than 1), it follows that our proof can be applied as is to derive a non-trivial threshold theorem. Irrespectively of whether this is the case or not one might also elect to work with the weaker definition of non-signalling strategies that is implied by the constraints in (3b), where the behaviour of the strategy is not required to be well-defined for questions which do not appear in the game. In this case the linear program (3) exactly describes the optimal winning probability of such strategies and Theorem 1 holds without any modification.

In other cases, on the other hand, we have to slightly modify the linear program in order to derive a correct threshold theorem. In the following sections we show how to do this.

A. Two-Player Games

For two-player games we consider the modification of the linear program (3) described in (28), as shown at the top of next page, where $\eta > 0$ is some small constant that will be chosen later.

Following the analysis proposed in [33] (Section 4 therein), one can show that the program (28) can be relaxed to the program described in (29).

Moreover, following [33] it can also be shown that the dual variables $y^\star$ which are associated with the primal constraints of Equations (29a) and (29b), as shown at the top of next page, are all upper bounded by 1, independently of the value of η. This implies that $\kappa = \sum_{j=1}^{d} |y_j^\star| \leq d$ is also independent of η (where d is now the total number of constraints in Equations (29a) and (29b) together).

When applying our proof using the linear program (29) we get the perturbed linear program (30) in Lemma 27 (instead of the one given in Equation (22)).

The estimated strategy $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ fulfils the constraints of Equation (30a), as shown at the top of next page, as in the proof in the main text. Moreover, it fulfils Equation (30c), as shown at the top of next page, by definition (see Section II-B). Therefore, in order to ensure that the winning probability of $\mathrm{O}^{\mathrm{EST2}}_{A|Q}$ is bounded by the optimal value of the program (30) we only need to choose $\eta \leq \zeta + 2\epsilon$ such that the constraints of Equation (30b), as shown at the top of next page, will hold as well.

To see that this is possible, recall that the values of ζ and ϵ are chosen such that $\zeta + 2\epsilon \leq \frac{\beta}{\kappa}$. As both β and κ are independent of η we can just choose $\eta \leq \zeta + 2\epsilon$. The rest of the proof then follows in the same way as in the main text and Theorem 1 is derived (without any dependence on η).

B. General Games

As the technique of the previous section is relevant only for two-player games,[12] the aim of this section is to explain how our proof can be adapted to derive a useful result for multiplayer games which do not have complete-support, as stated in Theorem 2. To do so we slightly modify the parallel repetition procedure.

Instead of considering the usual parallel repetition, in which n tuples of questions are chosen according to the game

[11] This example was communicated to us by Christian Schaffner.

[12] To be more precise, it holds for any game where κ can be bounded by a constant independent of the questions distribution Q.

$$
\begin{aligned}
\max \ & \sum_{q,a} Q(q) R(q,a) \mathrm{O}(a|q) \\
\text{s.t. } & Q(q^i, q^{\bar{i}}) \left[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}}) \mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}) \right] = 0 \quad \forall i, a^{\bar{i}}, \forall q^i, q^{\bar{i}} \text{ s.t. } Q(q) \neq 0 \\
& \eta \left[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}}) \mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}) \right] = 0 \quad \forall i, a^{\bar{i}}, \forall q^i, q^{\bar{i}} \text{ s.t. } Q(q) = 0 \\
& \sum_a \mathrm{O}(a|q) = 1 \quad \forall q \\
& \mathrm{O}(a|q) \geq 0 \quad \forall a, q
\end{aligned} \tag{28}
$$

$$
\begin{aligned}
\max \ & \sum_{q,a} Q(q) R(q,a) \mathrm{O}(a|q) \\
\text{s.t. } & Q(q^i, q^{\bar{i}}) \left[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}}) \mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}) \right] \leq 0 \quad \forall i, a^{\bar{i}}, \forall q^i, q^{\bar{i}} \text{ s.t. } Q(q) \neq 0 & \text{(29a)} \\
& \eta \left[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}}) \mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}) \right] \leq 0 \quad \forall i, a^{\bar{i}}, \forall q^i, q^{\bar{i}} \text{ s.t. } Q(q) = 0 & \text{(29b)} \\
& \sum_a \mathrm{O}(a|q) \leq 1 \quad \forall q \\
& \mathrm{O}(a|q) \geq 0 \quad \forall a, q
\end{aligned}
$$

$$
\begin{aligned}
\max \ & \sum_{q,a} Q(q) R(q,a) \mathrm{O}(a|q) \\
\text{s.t. } & Q(q^i, q^{\bar{i}}) \left[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}}) \mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}) \right] \leq \zeta + 2\epsilon \quad \forall i, a^{\bar{i}}, \forall q^i, q^{\bar{i}} \text{ s.t. } Q(q) \neq 0 & \text{(30a)} \\
& \eta \left[\mathrm{O}(\circ, a^{\bar{i}}|q^i, q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}}) \mathrm{O}(\circ, a^{\bar{i}}|r^i, q^{\bar{i}}) \right] \leq \zeta + 2\epsilon \quad \forall i, a^{\bar{i}}, \forall q^i, q^{\bar{i}} \text{ s.t. } Q(q) = 0 & \text{(30b)} \\
& \sum_a \mathrm{O}(a|q) \leq 1 \quad \forall q & \text{(30c)} \\
& \mathrm{O}(a|q) \geq 0 \quad \forall a, q
\end{aligned}
$$

distribution Q, we change the distribution of questions in the repeated game by sometimes (with small positive probability) asking the players a tuple of questions q for which $Q(q) = 0$. We call such questions "dummy questions"; for these questions any answer from the players is accepted. The remaining questions, for which $Q(q) > 0$, are called the "real questions". We denote the modified repeated game by $\tilde{G}^n$.

It is important to note that the standard definition of the non-signalling constraints implies that a non-signalling strategy should have a well-defined behaviour for all possible inputs. As the referee ignores the players' answers to the additional questions, the specific behaviour of the strategy on dummy questions is irrelevant. Therefore, if the optimal non-signalling winning probability in G is 1, then the winning probability in both G^n and $\tilde{G}^n$ is also 1: our modification does not harm the success probability of "honest" players.

To prove Theorem 2 we proceed in two steps: we make a small change in the linear program (4) and then apply our proof using the modified program.

1) Changing the Linear Program: As a first step we define $\tilde{Q}$ to be a complete-support version of Q in the following way.[13]

Let $\mathrm{I}(q)$ be the indicator function such that $\mathrm{I}(q) = 1$ if q is a dummy question, i.e., if $Q(q) = 0$, and 1 otherwise. Denote by $\mathcal{D}$ the number of dummy questions $\mathcal{D} = |\{q|\mathrm{I}(q) = 1\}|$.

Let $\eta > 0$ be some small constant (which can be later chosen to optimise the bound obtained in the final result). We define the following joint probability distribution of q and $d \in \{0, 1\}$:

$$
\mathrm{P}_{\tilde{Q}D}(q,d) = \begin{cases} \frac{\eta}{\mathcal{D}} & \text{if } \mathrm{I}(q) = 1 \text{ and } d = 1 \\ Q(q)(1-\eta) & \text{if } \mathrm{I}(q) = 0 \text{ and } d = 0 \\ 0 & \text{otherwise} \end{cases}
$$

[13] In [34] and [35] a subset of indices in which dummy, or "confusion", questions are asked is chosen. We choose to make a small modification in the questions distribution instead, such that permutation invariance is maintained.

Then $\tilde{Q}(q) = \sum_{d\in\{0,1\}} \mathrm{P}_{\tilde{Q}D}(q,d)$ and we have

$$\mathrm{P}_{\tilde{Q}|D=0}(q) = \frac{\mathrm{P}_{\tilde{Q}D}(q,0)}{\sum_q \mathrm{P}_{\tilde{Q}D}(q,0)} = \frac{\mathrm{P}_{\tilde{Q}D}(q,0)}{1-\eta} = Q(q).$$

That is, when conditioning on the event of a question not being a dummy question we retrieve Q from $\tilde{Q}$.

Next, we use $\tilde{Q}$ to write the non-signalling constraints (but keep Q in the objective function):

$$\begin{aligned}
\max\ & \sum_{q,a} Q(q)R(q,a)\mathrm{O}(a|q) \\
\text{s.t. } & \tilde{Q}(q^i,q^{\bar{i}})\Big[\mathrm{O}(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) \\
& \qquad - \sum_{r^i} \tilde{Q}(r^i|q^{\bar{i}})\mathrm{O}(\circ,a^{\bar{i}}|r^i,q^{\bar{i}})\Big] \le 0 \\
& \forall i,q^i,q^{\bar{i}},a^{\bar{i}} \\
& \sum_a \mathrm{O}(a|q) = 1 \quad \forall q \\
& \mathrm{O}(a|q) \ge 0 \quad \forall a,q
\end{aligned} \tag{31}$$

This linear program replaces the program (4). The distance measure in Definition 6 and the signalling measure in Definition 13 should now be defined with respect to $\tilde{Q}$ as well.

2) Deriving Theorem 2: Following the proof of Theorem 1 with the above changes we get the following statement in the main Lemma, Lemma 28:

$$\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{\tilde{Q}}}}\left[w(\mathrm{O}^{\mathrm{EST2}}_{A|Q}) > 1-\alpha+\beta\right] \le 3d\sqrt{c\delta}. \tag{32}$$

where the data $\vec{a},\vec{q}$ is now distributed according to $\mathrm{P}_{\vec{A}\vec{\tilde{Q}}} = \tilde{Q}^{\otimes n} \times \mathrm{P}_{\vec{A}|\vec{\tilde{Q}}}$ and the parameter δ now depends on the change we did in the question distribution $\tilde{Q}$ (through κ which depends on the solution of the dual program of program (31), and thus has an implicit dependence on η).

As the objective function of program (31) is given using Q and not $\tilde{Q}$, $w(\mathrm{O}^{\mathrm{EST2}}_{A|Q})$ in Equation (32) is the winning probability with respect to the original question distribution Q. It is therefore equal to the winning frequency in the real questions (i.e. it does not take the indices where dummy questions were asked into account). Hence, it leads to the desired statement:

$$\Pr_{\vec{a},\vec{q}\sim \mathrm{P}_{\vec{A}\vec{\tilde{Q}}}}\left[f > 1-\alpha+\beta\right] \le 6d\sqrt{c\delta},$$

where f is the winning frequency in the real questions. This proves Theorem 2.

The parameter η can be optimised in different ways, depending on the application. If one is interested in the bound itself and is not concerned by the modification of the repeated game the precise value of η should be chosen in order to optimise the constants $\mathcal{C}_1(G,n)$ and $\mathcal{C}_2(G)$ appearing in the bound. Alternatively, if one does not wish to change the game by too much, small values for η will ensure that $\tilde{G}^n$ is relatively close to G^n (due to the definition of $\tilde{Q}$ above). A smaller η will lead to a smaller fraction of dummy questions, but could result in worse constants $\mathcal{C}_2(G)$.

Appendix B
Proofs of Section III

In this section we present all the proofs which are relevant to the signalling measures and signalling tests.

The first proof is a proof of Lemma 14 which shows that the signalling measure given in Definition 13 is continuous. We repeat Lemma 14 here:

Lemma 14: Let O_1 *and* O_2 *be two one-game strategies such that* $\left|\mathrm{O}_1 - \mathrm{O}_2\right|_1 \le \epsilon$. *Then for all* $i, a^{\bar{i}}, q^i$ *and* $q^{\bar{i}}$,

$$\left|\mathrm{Sig}_{(i,a^{\bar{i}},q^i,q^{\bar{i}})}(\mathrm{O}_1) - \mathrm{Sig}_{(i,a^{\bar{i}},q^i,q^{\bar{i}})}(\mathrm{O}_2)\right| \le 2\epsilon.$$

Proof: We prove a stronger result from which the lemma follows. We prove that for all i,

$$\sum_{a^{\bar{i}},q}\left|\mathrm{Sig}_{(i,a^{\bar{i}},q^i,q^{\bar{i}})}(\mathrm{O}_1) - \mathrm{Sig}_{(i,a^{\bar{i}},q^i,q^{\bar{i}})}(\mathrm{O}_2)\right| \le 2\epsilon.$$

To do so first note the following,

$$\begin{aligned}
\left|\mathrm{O}_1 - \mathrm{O}_2\right|_1 &= \mathbb{E}_q \sum_a \left|\mathrm{O}_1(a|q) - \mathrm{O}_2(a|q)\right| \\
&\ge \mathbb{E}_q \sum_{a^{\bar{i}}}\left|\sum_{a^i}\Big(\mathrm{O}_1(a^i,a^{\bar{i}}|q) - \mathrm{O}_2(a^i,a^{\bar{i}}|q)\Big)\right| \\
&= \mathbb{E}_q \sum_{a^{\bar{i}}}\left|\mathrm{O}_1(\circ,a^{\bar{i}}|q) - \mathrm{O}_2(\circ,a^{\bar{i}}|q)\right| \\
&= \sum_{a^{\bar{i}},q} Q(q)\left|\mathrm{O}_1(\circ,a^{\bar{i}}|q) - \mathrm{O}_2(\circ,a^{\bar{i}}|q)\right|,
\end{aligned}$$

therefore if $\left|\mathrm{O}_1 - \mathrm{O}_2\right|_1 \le \epsilon$ then

$$\sum_{a^{\bar{i}},q} Q(q)\left|\mathrm{O}_1(\circ,a^{\bar{i}}|q) - \mathrm{O}_2(\circ,a^{\bar{i}}|q)\right| \le \epsilon. \tag{33}$$

Next, using Equation (6) we get the derivation (34), as shown at the top of next page.

where the last inequality follows from Equation (33). □

Next we give the proof of Lemma 16:

Lemma 16: Assume the players share an i.i.d. strategy $\mathrm{O}^{\otimes n}_{A|Q}$ *and let* $\zeta,\epsilon > 0$ *be the the parameters defined as in Equation* (8). *For every* $(i,b^{\bar{i}},s^i,s^{\bar{i}})$,

1) *If* $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) \ge \zeta$ *then*

$$\Pr_{\vec{a},\vec{q}\sim \mathrm{O}^{\otimes n}_{AQ}}[\mathbf{T}] > 1-\delta \tag{35}$$

2) *If* $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) = 0$ *then*

$$\Pr_{\vec{a},\vec{q}\sim \mathrm{O}^{\otimes n}_{AQ}}[\neg\mathbf{T}] > 1-\delta \tag{36}$$

where $\delta = \delta\left(\frac{n}{2},\epsilon\right) = \left(\frac{n}{2}+1\right)^{|\mathcal{A}|\cdot|\mathcal{Q}|-1} e^{-n\epsilon^2/4}$.

Proof: For the first part of the lemma assume that $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) \ge \zeta$. Then

$$\begin{aligned}
&\Pr_{\vec{a},\vec{q}\sim \mathrm{O}^{\otimes n}_{AQ}}[\neg\mathbf{T}] \\
&\quad = \Pr_{\vec{a},\vec{q}\sim \mathrm{O}^{\otimes n}_{AQ}}\left[\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\mathrm{O}^{\mathrm{EST1}}\right) < \zeta - 2\epsilon\right] \\
&\quad \le \Pr_{\vec{a},\vec{q}\sim \mathrm{O}^{\otimes n}_{AQ}}\left[|\mathrm{O}^{\mathrm{EST1}} - \mathrm{O}|_1 > \epsilon\right] \\
&\quad \le \delta
\end{aligned}$$

$$
\begin{aligned}
&\sum_{a^{\bar{i}},q} \left|\mathrm{Sig}_{(i,a^{\bar{i}},q^i,q^{\bar{i}})}\left(\mathrm{O}_1\right) - \mathrm{Sig}_{(i,a^{\bar{i}},q^i,q^{\bar{i}})}\left(\mathrm{O}_2\right)\right| \\
&\quad = \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}_1(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) - \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) + \sum_{r^i} Q(r^i|q^{\bar{i}})\mathrm{O}_2(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) \Big| \\
&\quad = \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) + \sum_{r^i} Q(r^i|q^{\bar{i}}) \Big(\mathrm{O}_2(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) - \mathrm{O}_1(\circ,a^{\bar{i}}|r^i,q^{\bar{i}})\Big) \Big| \\
&\quad \leq \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) \Big| + \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \sum_{r^i} Q(r^i|q^{\bar{i}}) \Big(\mathrm{O}_2(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) - \mathrm{O}_1(\circ,a^{\bar{i}}|r^i,q^{\bar{i}})\Big) \Big| \\
&\quad \leq \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) \Big| + \sum_{a^{\bar{i}},q} \sum_{r^i} Q(r^i|q^{\bar{i}}) Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_2(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) - \mathrm{O}_1(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) \Big| \\
&\quad = \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) \Big| + \sum_{a^{\bar{i}},q^{\bar{i}}} \sum_{r^i} Q(r^i|q^{\bar{i}}) Q(q^{\bar{i}}) \Big| \mathrm{O}_2(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) - \mathrm{O}_1(\circ,a^{\bar{i}}|r^i,q^{\bar{i}}) \Big| \\
&\quad = \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) \Big| + \sum_{a^{\bar{i}},q} Q(q^i,q^{\bar{i}}) \Big| \mathrm{O}_2(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) - \mathrm{O}_1(\circ,a^{\bar{i}}|q^i,q^{\bar{i}}) \Big| \\
&\quad \leq 2\epsilon
\end{aligned}
\tag{34}
$$

where the first inequality is due to Lemma 14 and the second due to Lemma 8. This implies Equation (35). Equation (36) can be proven in an analogous way. □

Next, we give the proof of Lemma 18:

Lemma 18: Let $\nu > 0$ *be any parameter such that* $\nu < \zeta - 6\epsilon$. *Then for every* $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$,

$$\forall \mathrm{O} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}, \quad \mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) > \nu.$$

Proof: Assume by contradiction that there exists $\mathrm{O} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ such that $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}(\mathrm{O}) \leq \nu$. Since $\mathrm{O} \in \Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ there exists $\bar{\mathrm{O}}$ such that $|\mathrm{O} - \bar{\mathrm{O}}|_1 \leq \epsilon$ and

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \bar{\mathrm{O}}_{AQ}^{\otimes n}}[\mathbf{T}] > \delta. \tag{37}$$

Using Lemma 14 we get $\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\bar{\mathrm{O}}\right) \leq \nu + 2\epsilon$.

From Lemma 8 we know that $\mathrm{Pr}_{\vec{a},\vec{q} \sim \bar{\mathrm{O}}_{AQ}^{\otimes n}}\left[|\bar{\mathrm{O}}^{\mathrm{EST1}} - \bar{\mathrm{O}}|_1 > \epsilon\right] \leq \delta$ and therefore, using Lemma 14 again,

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \bar{\mathrm{O}}_{AQ}^{\otimes n}}\left[\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\bar{\mathrm{O}}^{\mathrm{EST1}}\right) > \nu + 4\epsilon\right] \leq \delta.$$

Since $\nu < \zeta - 6\epsilon$ this implies

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \bar{\mathrm{O}}_{AQ}^{\otimes n}}\left[\mathrm{Sig}_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}\left(\bar{\mathrm{O}}^{\mathrm{EST1}}\right) > \zeta - 2\epsilon\right] \leq \delta$$

and therefore, according to the definition of the test,

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \bar{\mathrm{O}}_{AQ}^{\otimes n}}[\mathbf{T}] \leq \delta,$$

which contradicts Equation (37). □

Appendix C
Proofs of Section IV

In this section we prove the relevant properties of the de Finetti strategy. We prove Lemma 20:

Lemma 20: For a de Finetti strategy $\tau_{\vec{A}|\vec{Q}}$ *and every* $(i, b^{\bar{i}}, s^i, s^{\bar{i}})$

1) $\mathrm{Pr}_{\vec{a},\vec{q} \sim \tau_{\vec{A}\vec{Q}}}\left[\mathbf{T} \wedge \neg \mathbf{in}^{\Sigma}\right] \leq \delta$
2) $\mathrm{Pr}_{\vec{a},\vec{q} \sim \tau_{\vec{A}\vec{Q}}}\left[\neg \mathbf{T} \wedge \mathbf{in}^{\sigma}\right] \leq \delta$

Proof: Since a de Finetti strategy is a convex combination of i.i.d. strategies, it is sufficient to prove this for i.i.d. strategies $\mathrm{O}_{A|Q}^{\otimes n}$ and the lemma will follow. We start by proving the first part of the lemma.

If $\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}[\mathbf{T}] \leq \delta$ then we are done. Consider therefore states $\mathrm{O}_{A|Q}$ such that

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}[\mathbf{T}] > \delta.$$

For such states

$$
\begin{aligned}
&\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}\left[\neg \mathbf{in}^{\Sigma}\right] \\
&\quad \leq \mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}\left[|\mathrm{O}_{A|Q}^{\mathrm{EST2}} - \mathrm{O}_{A|Q}|_1 > \epsilon\right] \\
&\quad \leq \delta
\end{aligned}
$$

where the first inequality follows from the definition of $\Sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ and the second from Lemma 8.

All together we get $\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}\left[\mathbf{T} \wedge \neg \mathbf{in}^{\Sigma}\right] \leq \delta$ as required for the first part of the lemma.

We now proceed to the second part of the lemma.

If $\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}[\mathbf{in}^{\sigma}] \leq \delta$ then we are done. Consider therefore states $\mathrm{O}_{A|Q}$ such that

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}\left[\mathbf{in}^{\sigma}\right] > \delta.$$

Using Lemma 8 we know that there exists a state $\mathrm{O}_{A|Q}^{\mathrm{EST2}} \in \sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ such that $|\mathrm{O}_{A|Q}^{\mathrm{EST2}} - \mathrm{O}_{A|Q}|_1 \leq \epsilon$ and according to the definition of $\sigma_{(i,b^{\bar{i}},s^i,s^{\bar{i}})}$ this implies that $\mathrm{O}_{A|Q}$ is ζ signalling or more. Therefore, according to Lemma 16, $\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}[\neg \mathbf{T}] \leq \delta$. All together we get

$$\mathrm{Pr}_{\vec{a},\vec{q} \sim \mathrm{O}_{AQ}^{\otimes n}}\left[\neg \mathbf{T} \wedge \mathbf{in}^{\sigma}\right] \leq \delta.$$

□

Acknowledgment

The authors would like to thank Anthony Leverrier, Yeong-Cherng Liang, and David Sutter for helpful discussions, and Christian Schaffner for pointing out a mistake in a preliminary version of this work, as well as to the anonymous referees for their helpful comments.

References

[1] J. S. Bell *et al.*, "On the Einstein–Podolsky–Rosen paradox," *Physics*, vol. 1, no. 3, pp. 195–200, 1964.

[2] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed experiment to test local hidden-variable theories," *Phys. Rev. Lett.*, vol. 23, no. 15, pp. 880–884, 1969.

[3] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy, "Interactive proofs and the hardness of approximating cliques," *J. ACM*, vol. 43, no. 2, pp. 268–292, 1996.

[4] U. Feige and L. Lovász, "Two-prover one-round proof systems: Their power and their problems (extended abstract)," in *Proc. 24th Annu. ACM Symp. Theory Comput. (STOC)*, 1992, pp. 733–744.

[5] I. Dinur, "The PCP theorem by gap amplification," *J. ACM*, vol. 54, no. 3, Jun. 2007, Art. ID 12. [Online]. Available: http://doi.acm.org/10.1145/1236457.1236459

[6] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson, "Multi-prover interactive proofs: How to remove intractability assumptions," in *Proc. 20th Annu. ACM Symp. Theory Comput. (STOC)*, 1988, pp. 113–131.

[7] Y. T. Kalai, R. Raz, and R. D. Rothblum, "Delegation for bounded space," in *Proc. 45th STOC*, 2013, pp. 565–574. [Online]. Available: http://doi.acm.org/10.1145/2488608.2488679

[8] P. K. Aravind. (2002). "The magic squares and Bell's theorem." [Online]. Available: http://arxiv.org/abs/quant-ph/0206070

[9] E. Hänggi, R. Renner, and S. Wolf. (2009). "Quantum cryptography based solely on Bell's theorem." [Online]. Available: http://arxiv.org/abs/0911.4171

[10] L. Masanes, "Universally composable privacy amplification from causality constraints," *Phys. Rev. Lett.*, vol. 102, no. 14, p. 140501, 2009.

[11] L. Masanes, R. Renner, M. Christandl, A. Winter, and J. Barrett, "Full security of quantum key distribution from no-signaling constraints," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 4973–4986, Aug. 2014.

[12] R. Raz, "A parallel repetition theorem," *SIAM J. Comput.*, vol. 27, no. 3, pp. 763–803, 1998.

[13] T. Holenstein, "Parallel repetition: Simplifications and the no-signaling case," in *Proc. 39th Annu. ACM Symp. Theory Comput.*, 2007, pp. 411–419.

[14] A. Rao, "Parallel repetition in projection games and a concentration bound," *SIAM J. Comput.*, vol. 40, no. 6, pp. 1871–1891, 2011.

[15] H. Buhrman, S. Fehr, and C. Schaffner. (2013). "On the parallel repetition of multi-player games: The no-signaling case." [Online]. Available: http://arxiv.org/abs/1312.7455

[16] R. Jain, A. Pereszlényi, and P. Yao. (2013). "A parallel repetition theorem for entangled two-player one-round games under product distributions." [Online]. Available: http://arxiv.org/abs/1311.6309

[17] I. Dinur, D. Steurer, and T. Vidick. (2013). "A parallel repetition theorem for entangled projection games." [Online]. Available: http://arxiv.org/abs/1310.4113

[18] A. Chailloux and G. Scarpa. (2014). "Parallel repetition of free entangled games: Simplification and improvements." [Online]. Available: http://arxiv.org/abs/1410.4397

[19] B. de Finetti, "Sulla proseguibilità di processi aleatori scambiabili," *Int. J. Math.*, vol. 1, pp. 53–67, 1969.

[20] P. Diaconis and D. Freedman, "Finite exchangeable sequences," *Ann. Probab.*, vol. 8, no. 4, pp. 745–764, 1980.

[21] G. A. Raggio and R. F. Werner, "Quantum statistical mechanics of general mean field systems," *Helvetica Phys. Acta*, vol. 62, no. 8, pp. 980–1003, 1989.

[22] C. M. Caves, C. A. Fuchs, and R. Schack, "Unknown quantum states: The quantum de-Finetti representation," *J. Math. Phys.*, vol. 43, p. 4537, Aug. 2002.

[23] R. Renner, "Symmetry of large physical systems implies independence of subsystems," *Nature Phys.*, vol. 3, no. 9, pp. 645–649, 2007.

[24] M. Christandl, R. König, G. Mitchison, and R. Renner, "One-and-a-half quantum de Finetti theorems," *Commun. Math. Phys.*, vol. 273, no. 2, pp. 473–498, 2007.

[25] F. G. S. L. Brandao and A. W. Harrow, "Quantum de Finetti theorems under local measurements with applications," in *Proc. 45th Annu. ACM Symp. Theory Comput.*, 2013, pp. 861–870.

[26] M. Christandl, R. König, and R. Renner, "Postselection technique for quantum channels with applications to quantum cryptography," *Phys. Rev. Lett.*, vol. 102, no. 2, p. 020504, 2009.

[27] R. Arnon-Friedman and R. Renner, "de Finetti reductions for correlations," *J. Math. Phys.*, vol. 56, no. 5, p. 052203, 2015.

[28] A. Leverrier. (2014). "Composable security proof for continuous-variable quantum key distribution with coherent states." [Online]. Available: http://arxiv.org/abs/1408.5689

[29] M. Christandl and R. Renner, "Reliable quantum state tomography," *Phys. Rev. Lett.*, vol. 109, no. 12, p. 120403, 2012.

[30] M. Berta, M. Christandl, and R. Renner, "The quantum reverse Shannon theorem based on one-shot information theory," *Commun. Math. Phys.*, vol. 306, no. 3, pp. 579–615, 2011.

[31] J. Barrett and M. Leifer, "The de Finetti theorem for test spaces," *New J. Phys.*, vol. 11, no. 3, p. 033024, 2009.

[32] M. Christandl and B. Toner, "Finite de Finetti theorem for conditional probability distributions describing physical theories," *J. Math. Phys.*, vol. 50, no. 4, p. 042104, 2009.

[33] T. Ito, "Polynomial-space approximation of no-signaling provers," in *Automata, Languages and Programming*. Berlin, Germany: Springer, 2010, pp. 140–151.

[34] U. Feige and J. Kilian, "Two prover protocols: Low error at affordable rates," in *Proc. 26th Annu. ACM Symp. Theory Comput.*, 1994, pp. 172–183.

[35] J. Kempe and T. Vidick, "Parallel repetition of entangled games," in *Proc. 43rd Annu. ACM Symp. Theory Comput.*, 2011, pp. 353–362.

[36] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley, 2012.

[37] A. Schrijver, *Theory of Linear and Integer Programming*. New York, NY, USA: Wiley, 1998.

[38] E. Hänggi, "Device-independent quantum key distribution," M.S. thesis, ETH Zurich, Zürich, Switzerland, 2010.

[39] J. Kempe and O. Regev, "No strong parallel repetition with entangled and non-signaling provers," in *Proc. IEEE 25th Annu. Conf. Comput. Complex. (CCC)*, 2010, pp. 7–15.

Rotem Arnon-Friedman received the B.Sc. degree in Physics and Computer Science and M.Sc in Computer Science from Tel-Aviv University in 2010 and 2012. She is currently pursuing a Ph.D. degree at the Institute for Theoretical Physics at ETH Zurich. Her research interests include quantum information theory and quantum cryptogrphy.

Renato Renner was born on December 11, 1974, in Lucerne (Switzerland). He studied physics, first at EPF Lausanne and later at ETH Zurich (Switzerland), where he graduated in theoretical physics in 2000. He then moved to the Computer Science Department of ETH to work on a thesis in the area of quantum cryptography. He received his Ph.D. degree in 2005. Between 2005 and 2007, he held a HP research fellowship in the Department for Applied Mathematics and Theoretical Physics at the University of Cambridge (UK). Since 2007, he is with the ETH Zurich Physics Department, where he is a Professor for Theoretical Physics. His research interests are ranging from quantum information science to foundations of quantum mechanics and thermodynamics.

Thomas Vidick received a Ph.D. in Computer Science from UC Berkeley in 2011, where his advisor was Umesh Vazirani. His thesis focused on the study of quantum entanglement in multi-prover interactive proof systems and in quantum cryptography. After a postdoctoral scholarship at MIT under the supervision of Scott Aaronson, he moved back to California where he is currently an assistant professor in the department of Computing and Mathematical Sciences at the California Institute of Technology. His research interests are in quantum complexity theory, cryptography, and algorithms.