Parallel repetition and concentration for (sub-)no-signalling games via a flexible constrained de Finetti reduction

Cécilia Lancien[1,2,*] and Andreas Winter[1,3,†]

[1]*Física Teòrica: Informació i Fenòmens Quàntics, Universitat Autònoma de Barcelona, 08193 Bellaterra, Barcelona, Spain*

[2]*Institut Camille Jordan, Université Claude Bernard Lyon 1, 43 boulevard du 11 novembre 1918, 69622 Villeurbanne Cedex, France*

[3]*ICREA – Institució Catalana de Recerca i Estudis Avançats, Pg. Lluís Companys, 23, 08010 Barcelona, Spain*

(Dated: 23 June 2015)

We use a recently discovered constrained de Finetti reduction (aka "Post-Selection Lemma") to study the parallel repetition of multi-player non-local games under no-signalling strategies. Since the technique allows us to reduce general strategies to independent plays, we obtain parallel repetition (corresponding to winning all rounds) in the same way as exponential concentration of the probability to win a fraction larger than the value of the game.

Our proof technique leads us naturally to a relaxation of no-signalling (NS) strategies, which we dub *sub-no-signalling (SNOS)*. While for two players the two concepts coincide, they differ for three or more players. Our results are most complete and satisfying for arbitrary number of sub-no-signalling players, where we get universal parallel repetition and concentration for any game, while the no-signalling case is obtained as a corollary, but only for games with "full support".

I. NON-LOCAL GAMES AND NO-SIGNALLING STRATEGIES

A multi-player non-local game is played between cooperating but non-communicating players. Each player receives an input from some input alphabet and has to produce an output in some output alphabet. The common goal of the players is to satisfy some pre-defined predicate on their inputs and outputs. For that, they may agree on a strategy before the game starts, but are then not allowed to communicate anymore. Such games are especially relevant in theoretical physics in the context of the foundations of quantum mechanics and quantum information, and in computer science where they arise in multi-prover interactive proof systems. Indeed, they may provide an intuitive and quantitative understanding of the role played by various degrees of correlations in global systems which are composed of several local subsystems. These games also arise in complexity theory, under the formulation of multi-provers with some shared resources producing a protocol that should convince a referee, or in cryptography as attacks from malicious parties having a more or less restricted physical power.

The *value* of a game is the maximum winning probability of the players, over all allowed joint strategies, using possibly some prescribed correlation resource such as shared randomness, quantum entanglement or no-signalling correlations. It has been a subject of considerable study how the availability of different resources affects the values of certain games [3, 4, 11, 27, 31].

In this context, a natural question is how the value of a game behaves when n independent instances of the game are played simultaneously, i.e. each player gets n independent inputs and has to provide n outputs such that each game instance is won (or a large fraction of them); this is the parallel repetition problem. Playing independently the optimal single-game strategy on all

* lancien@math.univ-lyon1.fr
† andreas.winter@uab.cat

n game instances will result in an exponentially decreasing winning probability. But although that was found paradoxical at first, this is in general not optimal [16, 17]. For classical two-player games, Raz [29], later simplified and improved by Holenstein [20], established the first general parallel repetition theorem, showing that the value of n repetitions decreases exponentially for every game. Holenstein [20] also proved an analogous parallel repetition theorem for the no-signalling value of general two-player games. Only recently, parallel repetition theorems were proved for the entangled value of two-player games: for general games, nothing better than a polynomial decay result is known up to now [24], while exponential decay results have been established in several special cases (perfect parallel repetition for XOR games [12], exponential decrease under parallel repetition for unique games [23], projection games [13], free games [8, 22]).

Multi-player games have received little attention until recently. And apart from the result in [10] (containing both classical and quantum statements), only in the no-signalling setting [2, 7]. The present work has the same focus on multiple no-signalling players, albeit we will find that the theory becomes much more satisfying for *sub-no-signalling* players.

Specifically, we will consider here ℓ-player games G with input alphabets $\mathcal{X}_1, \dots, \mathcal{X}_\ell$ and output alphabets $\mathcal{A}_1, \dots, \mathcal{A}_\ell$. By way of notation, $\underline{\mathcal{X}} := \bigtimes_{i=1}^{\ell} \mathcal{X}_i$ and $\underline{\mathcal{A}} := \bigtimes_{i=1}^{\ell} \mathcal{A}_i$. Furthermore, for a subset $I \subset [\ell]$ of indices, $\mathcal{A}_I := \bigtimes_{i\in I} \mathcal{A}_i$ and $\mathcal{X}_I := \bigtimes_{i\in I} \mathcal{X}_i$. We will be interested in making minimal a priori assumptions on how powerful the ℓ players may be. This will naturally lead us to considering that their common strategy to win the game G could be any no-signalling (or even sub-no-signalling) strategy, which we define now.

Definition 1 *The sets of* no-signalling *and* sub-no-signalling *correlations, denoted $NS(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ and $SNOS(\underline{\mathcal{A}}|\underline{\mathcal{X}})$, respectively, consist of non-negative densities $P(\underline{a}|\underline{x}) \geq 0$ defined as follows:*

$$P \in NS(\underline{\mathcal{A}}|\underline{\mathcal{X}}) :\Leftrightarrow \forall I \subsetneq [\ell]\ \exists Q(\cdot|x_I) \text{ p.d.'s on } \mathcal{A}_I \text{ s.t. } \forall \underline{x}, a_I\ P(a_I|\underline{x}) = Q(a_I|x_I), \tag{1}$$

$$P \in SNOS(\underline{\mathcal{A}}|\underline{\mathcal{X}}) :\Leftrightarrow \forall I \subsetneq [\ell]\ \exists Q(\cdot|x_I) \text{ p.d.'s on } \mathcal{A}_I \text{ s.t. } \forall \underline{x}, a_I\ P(a_I|\underline{x}) \leq Q(a_I|x_I). \tag{2}$$

Here, $P(a_I|\underline{x})$ denotes the marginal density,

$$P(a_I|\underline{x}) = \sum_{a_{I^c} \in \mathcal{A}_{I^c}} P(\underline{a} = a_I a_{I^c}|\underline{x}).$$

Remark Note that under this definition, $\mathrm{NS}(\underline{\mathcal{A}}|\underline{\mathcal{X}}) \subset \mathrm{SNOS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$, but the latter is a strictly larger set (e.g. it always contains the all-zero density). Furthermore, $P \in \mathrm{NS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ iff $P \in \mathrm{SNOS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ and P is *normalized* in the sense that for all $\underline{x} \in \underline{\mathcal{X}}$, $\sum_{\underline{a}} P(\underline{a}|\underline{x}) = 1$. Indeed, NS consists of conditional probability distributions, while SNOS allows, given each input, a total "probability" of less than or equal to 1.

Also, it can be shown that in Eq. (1), only sets of the form $I = [\ell] \setminus i$ need to be considered. This is because the no-signalling conditions take the form of equations and this subset spans the set of all equations required (cf. [19], Lemma 2.7). The analogous statement for sub-no-signalling is not known and likely false. Nevertheless, one might in other contexts consider to relax the conditions of Eq. (2) to hold only for a selected family of subsets $I \subset [\ell]$. □

An ℓ-player game G is given by a game distribution $T(\underline{x})$ on the queries $\underline{\mathcal{X}}$, and a binary predicate $V(\underline{a}, \underline{x}) \in \{0, 1\}$ on $\underline{\mathcal{A}} \times \underline{\mathcal{X}}$. The no-signalling [sub-no-signalling] value of the game, denoted $\omega_{\mathrm{NS}}(G)$ [$\omega_{\mathrm{SNOS}}(G)$] is the maximum of the winning probability

$$\Pr\{\text{win}\} = \mathbb{E}V(\underline{A}, \underline{X}) = \sum_{\underline{a},\underline{x}} T(\underline{x}) V(\underline{a}, \underline{x}) P(\underline{a}|\underline{x})$$

over all $P \in \text{NS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ $[P \in \text{SNOS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})]$, where the distribution of $\underline{X} = X_1 \dots X_\ell$ and $\underline{A} = A_1 \dots A_\ell$ is as expected,

$$\Pr\{\underline{X} = \underline{x},\, \underline{A} = \underline{a}\} = T(\underline{x})P(\underline{a}|\underline{x}).$$

In words, the (sub-)no-signalling value of a game is the maximal probability of winning it when no limitation is assumed on the power of the players, apart from the fact that they cannot signal information instantaneously from one another. In the no-signalling case, players are forced to always produce an output, while in the sub-no-signalling case they are even allowed to abstain from doing so. In Section V we briefly discuss other kinds of restrictions that one may put on the players' physical power, such as shared randomness or shared quantum entanglement only.

A. Two-player SNOS $\equiv$ NS. Not surprisingly, the no-signalling and sub-no-signalling values of games are related. We start by showing that for any two-player game G, they are identical, $\omega_{\text{NS}}(G) = \omega_{\text{SNOS}}(G)$. As NS $\subset$ SNOS, the inequality "$\leq$" is evident, and we only need to prove the opposite inequality "$\geq$". This follows from the following structural lemma.

Lemma 2 (cf. Ito [21]) *Let $P \in SNOS(\mathcal{A}\times\mathcal{B}|\mathcal{X}\times\mathcal{Y})$ be a two-player sub-no-signalling correlation. Then there exists a no-signalling correlation $P' \in NS(\mathcal{A}\times\mathcal{B}|\mathcal{X}\times\mathcal{Y})$ with $P \leq P'$ pointwise, i.e. $P(ab|xy) \leq P'(ab|xy)$ for all a, b, x, y.*

Proof If P is normalized, i.e. if for all x and y, $\sum_{ab} P(ab|xy) = 1$, there is nothing to prove because P is already no-signalling.

Otherwise, there is a pair $xy \in \mathcal{X}\times\mathcal{Y}$ with weight $\sum_{ab} P(ab|xy) =: w < 1$. By sub-no-signalling, we have distributions $Q(a|x)$ and $Q(b|y)$ dominating the marginals:

$$P(a|xy) \leq Q(a|x), \quad P(b|xy) \leq Q(b|y) \quad \forall a, b.$$

As the total weight of both marginals of $P(\cdot|xy)$ is $w < 1$, we can find a and b such that

$$P(a|xy) < Q(a|x), \quad P(b|xy) < Q(b|y),$$

so we can increase $P(ab|xy)$ by some $\epsilon > 0$ to $P'(ab|xy) = P(ab|xy) + \epsilon$ and still satisfy the sub-no-signalling conditions. By choosing ϵ maximally so, we can reduce the total number of strict inequality signs in the SNOS conditions. Iterating this procedure we arrive at a sub-no-signalling correlation P' with all inequalities met with equality, i.e. a no-signalling correlation.

A more "mathematical" presentation of this argument appeals to compactness: Consider all sub-no-signalling P' dominating P and with marginals dominated by $Q(a|x)$ and $Q(b|y)$. Because of compactness, the maximum over such P' of

$$\sum_{xy}\sum_{ab} P'(ab|xy)$$

is attained. If it were less than $|\mathcal{X}\times\mathcal{Y}|$, we could use the procedure above to increase the objective function, contradicting that it is a maximum. $\square$

B. Multi-player SNOS vs. NS. Clearly, $\omega_{\text{NS}}(G) \leq \omega_{\text{SNOS}}(G)$ for every game, and there are examples of games (with game distribution T having strictly smaller than full support) where $\omega_{\text{NS}}(G) < 1$ but $\omega_{\text{SNOS}}(G) = 1$, for instance the *anticorrelation game* (and likewise a number of games where frustration prohibits extension of a winning sub-no-signalling strategy to a normalized, no-signalling one).

Example (Buhrman *et al.* [7]) Consider the three-player *anti-correlation game* A_3, which has binary input and output for all players and game distribution T supported on $\{0,1\}^3 \setminus \{111\}$, i.e. 111 does not occur as a triple of questions. The winning predicate is that if any two inputs are 1, say $x_i = x_j = 1$, then the corresponding outputs must be different, $a_i \neq a_j$. If there are no or only a single 1 in $\underline{x} = x_1x_2x_3$, outputs may be arbitrary.

It is straightforward to verify that the following correlation is in SNOS and wins the game with certainty, i.e. $\omega_{\text{SNOS}}(A_3) = 1$:

$$P(a_1a_2a_3|x_1x_2x_3) = \begin{cases} 0 & \text{if } x_1x_2x_3 = 111, \\ \frac{1}{8} & \text{if } x_1 + x_2 + x_3 \leq 1, \\ \frac{1}{4}\delta_{a_i,1-a_j} & \text{if } x_i = x_j = 1,\ (i \neq j). \end{cases}$$

On the other hand, for, say, T uniform on $\{011, 101, 110\}$, one can check by elementary means that $\omega_{\text{NS}}(A_3) = 2/3$. □

However, for a game distribution T having full support, a simple reasoning shows that $\omega_{\text{NS}}(G) < 1$ implies $\omega_{\text{SNOS}}(G) < 1$. Indeed, we show the contrapositive, assuming that $\omega_{\text{SNOS}}(G) = 1$. Because of the full support of T, this implies that for the optimal sub-no-signalling strategy P and every $\underline{x}$,

$$1 = \sum_{\underline{a}} V(\underline{a}, \underline{x}) P(\underline{a}|\underline{x}) \leq \sum_{\underline{a}} P(\underline{a}|\underline{x}) \leq 1,$$

hence equality (i.e. normalization) holds for all $\underline{x}$. Thus, P is really a no-signalling correlation and so $\omega_{\text{NS}}(G) = 1$. In fact, we can show something stronger, namely the following quantitative relationship.

Lemma 3 *Consider a game distribution T with full support on $\underline{\mathcal{X}}$. Then there exists a constant Γ, which only depends on T, such that for every game G with query distribution T,*

$$\omega_{SNOS}(G) \geq 1 - \epsilon \quad \Longrightarrow \quad \omega_{NS}(G) \geq 1 - (\Gamma + 1)\epsilon.$$

The constant Γ can be taken from [7] or [2].

Proof Take an optimal strategy $P \in \text{SNOS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$, so that $P(a_I|\underline{x}) \leq Q(a_I|x_I)$ for all I, a_I, $\underline{x}$. Then,

$$\sum_{\underline{a},\underline{x}} P(\underline{a}|\underline{x}) T(\underline{x}) \geq \omega_{\text{SNOS}}(G) \geq 1 - \epsilon,$$

and so we get, for all I,

$$\big\| P_{\mathcal{A}_I|\underline{\mathcal{X}}} T_{\underline{\mathcal{X}}} - Q_{\mathcal{A}_I|\mathcal{X}_I} T_{\underline{\mathcal{X}}} \big\|_1 = \sum_{a_I,\underline{x}} \big(Q(a_I|x_I) - P(a_I|\underline{x}) \big) T(\underline{x}) \leq \epsilon,$$

because the difference term in the sum is non-negative. Now simply "bump up" the sub-normalized distributions $P_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$ to a properly normalized conditional probability distribution $P'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$, adding at most a total probability of ϵ, and hence, for all I,

$$\frac{1}{2} \big\| P'_{\mathcal{A}_I|\underline{\mathcal{X}}} T_{\underline{\mathcal{X}}} - Q_{\mathcal{A}_I|\mathcal{X}_I} T_{\underline{\mathcal{X}}} \big\|_1 \leq \epsilon.$$

At this point we can invoke the stability of linear programmes, used in [7] and [2] to conclude that there is $\Gamma = \Gamma(T)$ such that there is a no-signalling correlation $P''_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} \in \text{NS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ with

$$\frac{1}{2}\left\| P''_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} T_{\underline{\mathcal{X}}} - P'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} T_{\underline{\mathcal{X}}} \right\|_1 \leq \Gamma\epsilon.$$

This gives

$$\begin{aligned}
\omega_{\text{NS}}(G) &\geq \sum_{\underline{a},\underline{x}} T(\underline{x})V(\underline{a},\underline{x})P''(\underline{a}|\underline{x}) \\
&\geq \sum_{\underline{a},\underline{x}} T(\underline{x})V(\underline{a},\underline{x})P'(\underline{a}|\underline{x}) - \Gamma\epsilon \\
&\geq \sum_{\underline{a},\underline{x}} T(\underline{x})V(\underline{a},\underline{x})P(\underline{a}|\underline{x}) - \Gamma\epsilon \\
&\geq 1 - (\Gamma + 1)\epsilon,
\end{aligned}$$

where we have used the total variational bound on $P'' - P'$, the fact that P' dominates P and the assumption on the value of P. □

The rest of the paper is structured as follows: In Section II we introduce parallel repetition of games, and state our main results, which improve upon, and partly clarify, earlier findings by Holenstein [20], Buhrman *et al.* [7] and Arnon-Friedman *et al.* [2]. In Section III, we present the main technical tool, a constrained de Finetti reduction, adapted to our present needs, followed by the proofs of the main theorems and corollaries in Section IV. We conclude in Section V.

II. PARALLEL REPETITION AND MAIN RESULTS

Given an ℓ-player game G, with game distribution $T(\underline{x})$ on $\underline{\mathcal{X}}$ and binary predicate $V(\underline{a},\underline{x}) \in \{0,1\}$ on $\underline{\mathcal{A}} \times \underline{\mathcal{X}}$, we are interested in playing the same game n times independently in parallel, and in looking at the probability of winning all n or a subset of t of them.

Formally, the *n-fold parallel repetition of G* is the ℓ-player game G^n having the product game distribution $T^{\otimes n}(\underline{x}^n) = T(\underline{x}^{(1)}) \cdots T(\underline{x}^{(n)})$ on $\underline{\mathcal{X}}^n$ and the product binary predicate $V^{\otimes n}(\underline{a}^n, \underline{x}^n) = V(\underline{a}^{(1)}, \underline{x}^{(1)}) \cdots V(\underline{a}^{(n)}, \underline{x}^{(n)}) \in \{0,1\}$ on $\underline{\mathcal{A}}^n \times \underline{\mathcal{X}}^n$. The no-signalling [sub-no-signalling] value of this n-fold parallel repetition game, denoted $\omega_{\text{NS}}(G^n)$ [$\omega_{\text{SNOS}}(G^n)$] is thus the maximum of the winning probability

$$\Pr\{\text{win}\} = \sum_{\underline{a}^n,\underline{x}^n} T^{\otimes n}(\underline{x}^n)V^{\otimes n}(\underline{a}^n,\underline{x}^n)P(\underline{a}^n|\underline{x}^n)$$

over all $P \in \text{NS}(\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n)$ [$P \in \text{SNOS}(\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n)$].

In words, the players win G^n if they win all n instances of G played in parallel. So we obviously always have (for the allowed set of strategies being $X \in \{\text{NS}, \text{SNOS}\}$)

$$\omega_X(G)^n \leq \omega_X(G^n) \leq \omega_X(G). \tag{3}$$

However, in the case where $\omega_X(G) < 1$, the gap between the lower and upper bounds in Eq. (3) grows exponentially with n, making it very little informative. The parallel repetition problem is thus the following: If none of the players' allowed strategies can make them win 1 instance of G with probability 1, does it necessarily imply that they have an exponentially decaying probability of winning n of them at the same time? And if so at which rate?

More generally, we can study the game $G^{t/n}$, whose winning predicate is defined as winning any t (or more) out of n repetitions [7]:

$$V^{t/n}(\underline{a}^n, \underline{x}^n) := \left\{ \sum_{i=1}^{n} V\big(\underline{a}^{(i)}, \underline{x}^{(i)}\big) \geq t \right\} = \begin{cases} 1 & \text{if } \sum_{i=1}^{n} V\big(\underline{a}^{(i)}, \underline{x}^{(i)}\big) \geq t, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $G^n = G^{n/n}$.

The main results of the present paper are the following ones.

Theorem 4 (Parallel repetition for the sub-no-signalling value of ℓ-player games) *Let G be an ℓ-player game such that $\omega_{SNOS}(G) \leq 1 - \delta$ for some $0 < \delta < 1$. Then, for any $n \in \mathbb{N}$,*

$$\omega_{SNOS}(G^n) \leq \left(1 - \frac{\delta^2}{5C_\ell^2}\right)^n,$$

where $C_\ell = 2^{\ell+1} - 3$, and for $t \geq (1 - \delta + \alpha)n$,

$$\omega_{SNOS}(G^{t/n}) \leq \exp\left(-n\frac{\alpha^2}{5C_\ell^2}\right).$$

As immediate consequences or refinements of Theorem 4, we can get parallel repetition results for the no-signalling value of multiplayer games in some particular instances.

Corollary 5 (Parallel repetition for the no-signalling value of full support ℓ-player games) *Let G be an ℓ-player game whose distribution T has full support, and such that $\omega_{NS}(G) \leq 1 - \delta$ for some $0 < \delta < 1$. Then, for any $n \in \mathbb{N}$, and for $t \geq (1 - \delta + \alpha)n$,*

$$\omega_{NS}(G^n) \leq \left(1 - \frac{\delta^2}{5C_\ell^2(\Gamma + 1)^2}\right)^n,$$
$$\omega_{NS}(G^{t/n}) \leq \exp\left(-n\frac{\alpha^2}{5C_\ell^2(\Gamma + 1)^2}\right).$$

where Γ is the constant from Lemma 3, which only depends on T.

Note that the constant Γ in this corollary depends on the game, and in the worst case carries a heavy dependence on the players' alphabet sizes. This is in contrast to Holenstein's two-player result for no-signalling games, which has no alphabet dependence at all [20]. This is generalized in our Theorem 4, since for two players we know by Lemma 2 that NS $\equiv$ SNOS, and we could directly read off bounds with constants already improving on Holenstein's. Looking a little into the proof allows us to optimize the constants even more, which we record as follows.

Theorem 6 (Parallel repetition for the no-signalling value of 2-player games, cf. Holenstein [20]) *Let G be a 2-player game with $\omega_{NS}(G) \leq 1 - \delta$ for some $0 < \delta < 1$. Then, for any $n \in \mathbb{N}$, and for $t \geq (1 - \delta + \alpha)n$,*

$$\omega_{NS}(G^n) \leq \left(1 - \frac{\delta^2}{27}\right)^n,$$
$$\omega_{NS}(G^{t/n}) \leq \exp\left(-n\frac{\alpha^2}{33}\right).$$

III. CONSTRAINED DE FINETTI REDUCTION

De Finetti reductions are a useful tool when trying to understand any permutation-invariant information processing task. Indeed, these enable to restrict the analysis to that of i.i.d. scenarios, which are usually trivially understood. In the context of multi-player games played n times in parallel, one would like to use the fact that the numbering of the n instances of the repeated game is irrelevant to reduce the study of strategies for the latter to the study of so-called *de Finetti strategies* (i.e. convex combinations of n i.i.d. strategies).

The seminal de Finetti reduction (aka post-selection) lemma was stated in [9], later finding applications in many areas of quantum information theory, from quantum cryptography [26] to quantum Shannon theory [5]. Our proofs though, will rely on two more recently established de Finetti reduction results, which are stated below. Just to fix some definitions: we will say that a (sub-)probability distribution $P_{\mathcal{Z}^n}$ [a conditional probability distribution $P_{\mathcal{B}^n|\mathcal{Y}^n}$] is n-symmetric if for any permutation π of n elements, $\forall\ z^n,\ P(\pi(z^n)) = P(z^n)$ [$\forall\ b^n, y^n,\ P(\pi(b^n)|\pi(y^n)) = P(b^n|y^n)$].

Lemma 7 (de Finetti reduction for conditional p.d.'s [1]) *Let $\mathcal{B}, \mathcal{Y}$ be finite alphabets. There exists a probability measure $dR_{\mathcal{B}|\mathcal{Y}}$ on the set of conditional probability distributions $R_{\mathcal{B}|\mathcal{Y}}$ such that, for any n-symmetric conditional probability distribution $P_{\mathcal{B}^n|\mathcal{Y}^n}$,*

$$P_{\mathcal{B}^n|\mathcal{Y}^n} \leq poly(n) \int_{R_{\mathcal{B}|\mathcal{Y}}} R_{\mathcal{B}|\mathcal{Y}}^{\otimes n}\, dR_{\mathcal{B}|\mathcal{Y}},$$

where the polynomial pre-factor may be upper-bounded as $poly(n) \leq (n+1)^{|\mathcal{B}||\mathcal{Y}|}$.

Lemma 8 (Constrained de Finetti reduction for (sub-)p.d.'s [14, 25]) *Let $\mathcal{Z}$ be a finite alphabet. There exists a probability measure $dQ_{\mathcal{Z}}$ on the set of probability distributions $Q_{\mathcal{Z}}$ on $\mathcal{Z}$ such that, for any n-symmetric (sub-)probability distribution $P_{\mathcal{Z}^n}$ on $\mathcal{Z}^n$,*

$$P_{\mathcal{Z}^n} \leq poly(n) \int_{Q_{\mathcal{Z}}} F\left(P_{\mathcal{Z}^n}, Q_{\mathcal{Z}}^{\otimes n}\right)^2 Q_{\mathcal{Z}}^{\otimes n}\, dQ_{\mathcal{Z}},$$

where the polynomial pre-factor may be upper-bounded as $poly(n) \leq (n+1)^{3|\mathcal{Z}|^2}$.

We are now ready to present the technical lemma that will allow us in Section IV to reduce the study of strategies for repeated games to the study of so-called *de Finetti strategies*, and hence prove our main results.

Lemma 9 (de Finetti reduction for sub-no-signalling correlations) *There exists a probability measure dQ on the set of probability distributions Q on $\underline{\mathcal{A}} \times \underline{\mathcal{X}}$ such that for any probability distribution T on $\underline{\mathcal{X}}$ and any $P \in SNOS(\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n)$ an n-symmetric sub-no-signalling correlation, it holds*

$$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n} \leq poly(n) \int_{Q_{\underline{\mathcal{A}\mathcal{X}}}} \widetilde{F}\left(Q_{\underline{\mathcal{A}\mathcal{X}}}\right)^{2n} Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n}\, dQ_{\underline{\mathcal{A}\mathcal{X}}}, \tag{4}$$

where we defined

$$\widetilde{F}\left(Q_{\underline{\mathcal{A}\mathcal{X}}}\right) := \min_{\emptyset \neq I \subsetneq [\ell]} \max_{R_{\mathcal{A}_I|\mathcal{X}_I}} F\left(T_{\underline{\mathcal{X}}} R_{\mathcal{A}_I|\mathcal{X}_I}, Q_{\mathcal{A}_I \underline{\mathcal{X}}}\right).$$

We mention for the sake of completeness that the $poly(n)$ pre-factor in Eq. (4) may be upper bounded by $(n+1)^{3|\underline{\mathcal{A}}|^2|\underline{\mathcal{X}}|^2 + 2|\underline{\mathcal{A}}||\underline{\mathcal{X}}|}$.

Proof Since $T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ is an n-symmetric sub-probability distribution on $(\underline{\mathcal{A}\mathcal{X}})^n$, we first of all have by Lemma 8 that

$$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n} \leq \text{poly}(n) \int_{Q_{\underline{\mathcal{A}\mathcal{X}}}} F\left(T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}, Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n}\right)^2 Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n}\, dQ_{\underline{\mathcal{A}\mathcal{X}}}.$$

Notice next that, for any $\emptyset \neq I \subsetneq [\ell]$,

$$F\left(T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}, Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n}\right) \leq F\left(T_{\underline{\mathcal{X}}}^{\otimes n} P_{\mathcal{A}_I^n|\underline{\mathcal{X}}^n}, Q_{\mathcal{A}_I\underline{\mathcal{X}}}^{\otimes n}\right) \leq F\left(T_{\underline{\mathcal{X}}}^{\otimes n} P'_{\mathcal{A}_I^n|\mathcal{X}_I^n}, Q_{\mathcal{A}_I\underline{\mathcal{X}}}^{\otimes n}\right).$$

The first inequality is by monotonicity of the fidelity under stochastic maps (in particular taking marginals). While the second inequality is because $P \in \text{SNOS}(\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n)$, so that $P_{\mathcal{A}_I^n|\underline{\mathcal{X}}^n} \leq P'_{\mathcal{A}_I^n|\mathcal{X}_I^n}$ for some conditional p.d. $P'_{\mathcal{A}_I^n|\mathcal{X}_I^n}$, and because the fidelity is order-preserving.

What is more, since for any $\emptyset \neq I \subsetneq [\ell]$, $P'_{\mathcal{A}_I^n|\mathcal{X}_I^n}$ is an n-symmetric conditional probability distribution, we have by Lemma 7 that

$$P'_{\mathcal{A}_I^n|\mathcal{X}_I^n} \leq \text{poly}(n) \int_{R_{\mathcal{A}_I|\mathcal{X}_I}} R_{\mathcal{A}_I|\mathcal{X}_I}^{\otimes n}\, dR_{\mathcal{A}_I|\mathcal{X}_I},$$

and subsequently, using first, once more, that the fidelity is order-preserving, and second that it is multiplicative on tensor products,

$$\begin{aligned} F\left(T_{\underline{\mathcal{X}}}^{\otimes n} P'_{\mathcal{A}_I^n|\mathcal{X}_I^n}, Q_{\mathcal{A}_I\underline{\mathcal{X}}}^{\otimes n}\right) &\leq \text{poly}(n) \max_{R_{\mathcal{A}_I|\mathcal{X}_I}} F\left(T_{\underline{\mathcal{X}}}^{\otimes n} R_{\mathcal{A}_I|\mathcal{X}_I}^{\otimes n}, Q_{\mathcal{A}_I\underline{\mathcal{X}}}^{\otimes n}\right) \\ &= \text{poly}(n) \max_{R_{\mathcal{A}_I|\mathcal{X}_I}} F\left(T_{\underline{\mathcal{X}}} R_{\mathcal{A}_I|\mathcal{X}_I}, Q_{\mathcal{A}_I\underline{\mathcal{X}}}\right)^n. \end{aligned}$$

Recapitulating, we get

$$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n} \leq \text{poly}(n) \int_{Q_{\underline{\mathcal{A}\mathcal{X}}}} \left(\min_{\emptyset\neq I\subsetneq[\ell]} \max_{R_{\mathcal{A}_I|\mathcal{X}_I}} F\left(T_{\underline{\mathcal{X}}} R_{\mathcal{A}_I|\mathcal{X}_I}, Q_{\mathcal{A}_I\underline{\mathcal{X}}}\right)\right)^{2n} Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n}\, dQ_{\underline{\mathcal{A}\mathcal{X}}},$$

as announced. □

IV. PROOFS OF THEOREMS 4 AND 6, AND COROLLARY 5

We need first of all the following extension of Holenstein's [20, Lemma 9.5]:

Lemma 10 *For $\underline{\mathcal{Z}} = \bigtimes_{j=1}^m \mathcal{Z}_j$ and $\underline{\mathcal{B}} = \bigtimes_{j=1}^m \mathcal{B}_j$, consider probability distributions T on $\underline{\mathcal{Z}}$ and P on $\underline{\mathcal{B}} \times \underline{\mathcal{Z}}$ with*

$$\frac{1}{2}\left\|P_{\underline{\mathcal{Z}}} - T_{\underline{\mathcal{Z}}}\right\|_1 \leq \epsilon_0.$$

If for each $j \in [m]$ there exists a conditional probability distribution $Q(b_j|z_j)$ such that

$$\left\|P_{\mathcal{B}_j\underline{\mathcal{Z}}} - Q_{\mathcal{B}_j|\mathcal{Z}_j} T_{\underline{\mathcal{Z}}}\right\|_1 \leq \epsilon_j,$$

then there exists a conditional probability distribution $P'(\underline{b}|\underline{z})$ such that for each $j \in [m]$,

$$P'(b_j|\underline{z}) = P'(b_j|z_j),$$

and

$$\frac{1}{2}\left\|P'_{\underline{\mathcal{B}}|\underline{\mathcal{Z}}} T_{\underline{\mathcal{Z}}} - P_{\underline{\mathcal{B}\mathcal{Z}}}\right\|_1 \leq \epsilon_0 + \sum_{j=1}^m 2\epsilon_j.$$

Proof This works exactly as the proof of Lemma 9.5 in [20], which is a successive application (m times) of Lemma 9.4 in the same paper. □

Note that the conditions enforced in Lemma 10 are not enough to ensure no-signalling of P' for three or more players. They would be sufficient though to guarantee that P' satisfies the relaxed no-signalling constraints considered in [30], namely that any group of $\ell - 1$ players together cannot signal to the remaining player. Nevertheless, we can leverage this result to approximate the given no-signalling correlation by a sub-no-signalling correlation.

Lemma 11 *Let P be a probability distribution on $\underline{\mathcal{A}} \times \underline{\mathcal{X}}$ and T be a probability distribution on $\underline{\mathcal{X}}$. If the no-signalling conditions (1) hold approximately, namely*

$$\forall I \subsetneq [\ell]\ \exists Q(\cdot|x_I) \text{ p.d.'s on } \mathcal{A}_I \text{ s.t. } \frac{1}{2}\left\| P_{\mathcal{A}_I \underline{\mathcal{X}}} - Q_{\mathcal{A}_I|\mathcal{X}_I} T_{\underline{\mathcal{X}}} \right\|_1 \leq \epsilon_I,$$

then there exists a sub-no-signalling correlation $P' \in SNOS(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ that approximates P, in the sense that

$$\frac{1}{2}\left\| P'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} T_{\underline{\mathcal{X}}} - P_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \right\|_1 \leq \epsilon_\emptyset + \sum_{\emptyset \neq I \subsetneq [\ell]} 2\epsilon_I.$$

In the two-player case $\ell = 2$, P' can be chosen to be no-signalling itself, $P' \in NS(\underline{\mathcal{A}}|\underline{\mathcal{X}})$.

Proof We will apply Lemma 10, with $n = 2^\ell - 2$, the index j identifying a non-empty and non-full set $\emptyset \neq I \subsetneq [\ell]$ (for instance via the expansion of j into ℓ binary digits). The local input and output alphabets are

$$\mathcal{Z}_j = \prod_{i \in I} \mathcal{X}_i, \quad \mathcal{B}_j = \prod_{i \in I} \mathcal{A}_i,$$

and the distribution we apply it to is

$$\widehat{P}(\underline{b}\underline{z}) = \begin{cases} P(\underline{a}\underline{x}) & \text{if } \forall j\ b_j = (a_i : i \in I),\ z_j = (x_i : i \in I), \\ 0 & \text{otherwise.} \end{cases}$$

Likewise, the prior distribution on $\underline{\mathcal{Z}}$ is given by

$$\widehat{T}(\underline{z}) = \begin{cases} T(\underline{x}) & \text{if } \forall j\ z_j = (x_i : i \in I), \\ 0 & \text{otherwise,} \end{cases}$$

and we use the conditional distributions $Q(b_j|z_j) = Q(a_I|x_I)$.

Now, the prerequisites of Lemma 10 are given, with $\epsilon_j = \epsilon_I$, and thus we get a conditional probability distribution $\widehat{P}'$ with $\widehat{P}'(b_j|\underline{z}) = \widehat{P}'(b_j|z_j)$ for all j, and

$$\frac{1}{2}\left\| \widehat{P}'_{\underline{\mathcal{B}}|\underline{\mathcal{Z}}} \widehat{T}_{\underline{\mathcal{Z}}} - \widehat{P}_{\underline{\mathcal{B}}\underline{\mathcal{Z}}} \right\|_1 \leq \epsilon_0 + \sum_{j=1}^{n} 2\epsilon_j =: \epsilon.$$

We would like to conclude here by "pulling back" this conditional distribution to a (it would seem: no-signalling) correlation on $\underline{\mathcal{A}} \times \underline{\mathcal{X}}$, except that P' has support outside the image of the diagonal embedding

$$\begin{aligned} \Delta : \underline{\mathcal{A}} &\longrightarrow \underline{\mathcal{B}} \\ \underline{a} &\longmapsto \underline{b} \text{ s.t. } \forall j\ b_j = (a_i : i \in I), \end{aligned}$$

and likewise for $\Delta : \underline{\mathcal{X}} \longrightarrow \underline{\mathcal{Z}}$.

To resolve this issue, we simply remove this part of the distribution, and define the desired sub-normalized conditional densities by letting

$$P'(\underline{a}|\underline{x}) := \widehat{P}'\big(\Delta(\underline{a})|\Delta(\underline{x})\big).$$

From this we see directly that

$$\frac{1}{2}\big\|P'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}T_{\underline{\mathcal{X}}} - P_{\underline{\mathcal{A}}\underline{\mathcal{X}}}\big\|_1 \leq \epsilon,$$

because $\widehat{P}(\underline{b}, \underline{z}) = P(\underline{a}, \underline{x})$ for $\underline{b} = \Delta(\underline{a})$ and $\underline{z} = \Delta(\underline{x})$, and it is 0 outside the image of Δ.

It remains to check that P' is sub-no-signalling. Let $\emptyset \neq I \subsetneq [\ell]$ be a subset with corresponding index $1 \leq j \leq 2^\ell - 2$. Let also $\underline{x} \in \underline{\mathcal{X}}$, $a_I \in \mathcal{A}_I$ be tuples, and set $\underline{z} = \Delta(\underline{x})$, $\underline{b} = \Delta(\underline{x})$ (so that $z_j = x_I \in \mathcal{X}_I = \mathcal{Z}_j$, $b_j = a_I \in \mathcal{A}_I = \mathcal{B}_j$). Then,

$$\begin{aligned}
P'(a_I|\underline{x}) &= \sum_{a_{I^c} \in \mathcal{A}_{I^c}} P'(\underline{a}|\underline{x}) \\
&= \sum_{a_{I^c} \in \mathcal{A}_{I^c}} \widehat{P}'\big(\Delta(\underline{a})|\Delta(\underline{x})\big) \\
&\leq \sum_{b_k \in \mathcal{B}_k : k \neq j} \widehat{P}'\big(\underline{b}|\underline{z}\big) \\
&= \widehat{P}'\big(b_j|\underline{z}\big) \\
&= \widehat{P}'(b_j|z_j) =: Q'(a_I|x_I).
\end{aligned}$$

Here, we have used the definition of the marginal and of P'. The inequality in the third line is because we enlarge the domain of the summation, and the equality in the last line is by the marginal property of $\widehat{P}'$.

The last claim, regarding $\ell = 2$ players, is the original Lemma 9.5 in [20]. □

We are now ready to prove our main theorem, namely the parallel repetition and concentration results for the sub-no-signalling value of multi-player games.

Proof of Theorem 4 Let $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ be a sub-no-signalling correlation which is optimal to win the game G^n. The distribution $T_{\underline{\mathcal{X}}}^{\otimes n}$ and the predicate $V_{\underline{\mathcal{A}}\underline{\mathcal{X}}}^{\otimes n}$ of G^n being n-symmetric, we can assume without loss of generality that $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ is also n-symmetric. Hence, by Lemma 9,

$$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n} \leq \operatorname{poly}(n) \int_{Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}} \widetilde{F}\left(Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}\right)^{2n} Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}^{\otimes n}\, dQ_{\underline{\mathcal{A}}\underline{\mathcal{X}}}.$$

Now, fix $0 < \epsilon < 1$ and define

$$\mathcal{P}_\epsilon := \left\{ Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \ : \ \max_{\emptyset \neq I \subsetneq [\ell]} \min_{R_{\mathcal{A}_I|\mathcal{X}_I}} \frac{1}{2}\|T_{\underline{\mathcal{X}}} R_{\mathcal{A}_I|\mathcal{X}_I} - Q_{\mathcal{A}_I \underline{\mathcal{X}}}\|_1 \leq \epsilon \right\}.$$

Observe that, by well-known relations between fidelity and trace-distance [18], if $Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \notin \mathcal{P}_\epsilon$, then $\widetilde{F}\left(Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}\right)^2 \leq 1 - \epsilon^2$. Hence,

$$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n} \leq \operatorname{poly}(n) \left(\int_{Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \in \mathcal{P}_\epsilon} Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}^{\otimes n}\, dQ_{\underline{\mathcal{A}}\underline{\mathcal{X}}} + (1-\epsilon^2)^n \int_{Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \notin \mathcal{P}_\epsilon} Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}^{\otimes n}\, dQ_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \right).$$

On the other hand, if $Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \in \mathcal{P}_\epsilon$, then by definition

$$\forall\, \emptyset \neq I \subsetneq [\ell] \;\exists\, R_{\mathcal{A}_I|\mathcal{X}_I} : \; \frac{1}{2}\| T_{\underline{\mathcal{X}}} R_{\mathcal{A}_I|\mathcal{X}_I} - Q_{\mathcal{A}_I \underline{\mathcal{X}}} \|_1 \leq \epsilon.$$

By Lemma 11, the latter condition implies that there exists a sub-no-signalling correlation $R'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$ such that

$$\frac{1}{2}\| T_{\underline{\mathcal{X}}} R'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} - Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \|_1 \leq C_\ell \epsilon, \text{ where } C_\ell = 1 + 2(2^\ell - 2) = 2^{\ell+1} - 3.$$

Yet, the winning probability when playing G with a strategy $R'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} \in \text{SNOS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ is, by assumption on G, at most $1-\delta$. So the average of the predicate of G over $Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}} \in \mathcal{P}_\epsilon$ is at most $1-\delta+2C_\ell\epsilon$. Putting everything together, we eventually get that the winning probability when playing G^n with strategy $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ is upper-bounded as

$$\Pr\{\text{win}\} \leq \text{poly}(n) \left((1-\delta+2C_\ell\epsilon)^n + (1-\epsilon^2)^n \right). \tag{5}$$

Choosing in Eq. (5)

$$\epsilon = C_\ell \left(\left(1 + \frac{\delta}{C_\ell^2} \right)^{1/2} - 1 \right) \geq \frac{99\delta}{200 C_\ell}, \text{ so that } \epsilon^2 \geq \frac{\delta^2}{5C_\ell^2},$$

and recalling that $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ is, by hypothesis, an optimal sub-no-signalling strategy, we obtain

$$\omega_{\text{SNOS}}(G^n) \leq \text{poly}(n) \left(1 - \frac{\delta^2}{5C_\ell^2} \right)^n. \tag{6}$$

In order to conclude, we have to remove the polynomial pre-factor. So assume that there exists a constant $C > 0$ such that for some $N \in \mathbb{N}$, $\omega_{\text{SNOS}}(G^N) \geq C\left(1 - \delta^2/5C_\ell^2\right)^N$. Then, for any $n \in \mathbb{N}$, we would have

$$\omega_{\text{SNOS}}(G^{Nn}) \geq \left(\omega_{\text{SNOS}}(G^N)\right)^n \geq C^n \left(1 - \frac{\delta^2}{5C_\ell^2} \right)^{Nn}.$$

On the other hand, however, we still have by Eq. (6)

$$\omega_{\text{SNOS}}(G^{Nn}) \leq \text{poly}(Nn) \left(1 - \frac{\delta^2}{5C_\ell^2} \right)^{Nn}.$$

Letting n grow, we see that the only option to make these two conditions compatible is to have $C \leq 1$, which is precisely what we wanted to show.

Following the exact same lines as above, we also get the concentration bound. Indeed, we now have in place of Eq. (5) that, for any $0 < \epsilon < 1$,

$$\omega_{\text{SNOS}}(G^{t/n}) \leq \text{poly}(n) \left(\exp\left[-2n(\alpha - 2C_\ell\epsilon)^2\right] + \exp\left[-n\epsilon^2\right] \right), \tag{7}$$

by Hoeffding's inequality (and because $e^{-x} \geq 1 - x$ for any $x > 0$).

The announced upper-bound follows from choosing in Eq. (7)

$$\epsilon = \frac{(4C_\ell - \sqrt{2})\alpha}{8C_\ell^2 - 1} \geq \frac{5(20-\sqrt{2})\alpha}{199 C_\ell}, \text{ so that } \epsilon^2 \geq \frac{\alpha^2}{5C_\ell^2},$$

and removing the polynomial pre-factor by the same trick as before. □

Proof of Corollary 5 By Lemma 3, we know that if G is an ℓ-player game with full support satisfying $\omega_{\mathrm{NS}}(G) \leq 1-\delta$, then $\omega_{\mathrm{SNOS}}(G) \leq 1-\delta/(\Gamma+1)$. And thus by Theorem 4,

$$\omega_{\mathrm{NS}}(G^n) \leq \omega_{\mathrm{SNOS}}(G^n) \leq \left(1 - \frac{\delta^2}{5C_\ell^2(\Gamma+1)^2}\right)^n .$$

The concentration bound for $\omega_{\mathrm{NS}}(G^{t/n})$ follows analogously. □

Proof of Theorem 6 We follow the exact same reasoning as in the proof of Theorem 4, and keep the same notations. In the case $\ell = 2$, we have by Lemma 11 that, for any $0 < \epsilon < 1$,

$$Q_{\underline{\mathcal{A}\mathcal{X}}} \in \mathcal{P}_\epsilon \;\Rightarrow\; \exists\, R'_{\underline{\mathcal{A}\mathcal{X}}} \in \mathrm{NS}(\underline{\mathcal{A}}|\underline{\mathcal{X}}) :\; \frac{1}{2}\|T_{\underline{\mathcal{X}}} R'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} - Q_{\underline{\mathcal{A}\mathcal{X}}}\|_1 \leq 5\epsilon.$$

Yet, if the winning probability when playing G with a strategy $R'_{\underline{\mathcal{A}}|\underline{\mathcal{X}}} \in \mathrm{NS}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ is, by assumption on G, at most $1-\delta$, then the average of the predicate of G over $Q_{\underline{\mathcal{A}\mathcal{X}}} \in \mathcal{P}_\epsilon$ is at most $1-\delta+5\epsilon$. This is because we are here dealing with normalised probability distributions. Hence, for any $0 < \epsilon < 1$,

$$\begin{aligned} \omega_{\mathrm{NS}}(G^n) &\leq \mathrm{poly}(n)\left((1-\delta+5\epsilon)^n + (1-\epsilon^2)^n\right), \\ \omega_{\mathrm{SNOS}}(G^{t/n}) &\leq \mathrm{poly}(n)\left(\exp\left[-2n(\alpha-5\epsilon)^2\right] + \exp\left[-n\epsilon^2\right]\right). \end{aligned}$$

We can now choose $\epsilon = (\sqrt{29}-5)\delta/2$ in the parallel repetition estimate and $\epsilon = (10-\sqrt{2})\alpha/49$ in the concentration bound one, and argue as in the proof of Theorem 4 to remove the polynomial pre-factor, which yields the two advertised results. □

V. DISCUSSION

Our main contribution in the present paper is a concentration result for the sub-no-signalling value of multi-player games under parallel repetition. In fact, we believe that our work is the first to recognize the intrinsic interest of the class of sub-no-signalling correlations, which appears naturally as a relaxation of the no-signalling ones.

Specifically, if an ℓ-player game G has SNOS value $1-\delta$, then the probability for SNOS players to win a fraction at least $1-\delta+\alpha$ of n instances of G played in parallel is at most $\exp(-nC_\ell\alpha^2)$, where $C_\ell > 0$ is a constant which only depends on the number ℓ of players. As mentioned in [7], such result, valid for games involving strictly more than 2 players and where not all queries are asked, could potentially find applications in position-based cryptography [6, 15]. In the case $\ell = 2$, this is actually equivalent to the analogous concentration result for the no-signalling value of G, thus with a universal constant $c = C_2$ in the exponential bound. And we know we cannot hope for a better dependence in α than the obtained one, even in the special case $\alpha = \delta$ (see e.g. [20] for an example of 2-player game G which is such that $\omega_{\mathrm{NS}}(G) < 1$ but $\omega_{\mathrm{NS}}(G^2) = \omega_{\mathrm{NS}}(G)$, hence proving that strong parallel repetition in general does not hold for no-signalling players). In the case $\ell > 2$, our result implies a concentration bound for the no-signalling value of G, but only if its input distribution has full support. Besides, the constant in the exponential bound is this time highly game dependent (dependence on the sizes of the input and output alphabets, and on the smallest weight occurring in the input distribution). This is fully comparable to previous work in this direction due to Buhrman, Fehr and Schaffner [7], and Arnon-Friedman, Renner and Vidick [2].

Hence, the most immediate open problem at that point is regarding games with non-full support in the case of three or more players (e.g. the anti-correlation game): does a parallel repetition result hold for the no-signalling value of such multi-player games? Answering this question probably requires to understand first whether in Corollary 5, the presence of the game parameter Γ is really necessary or is just an artifact of the proof technique. In other words, does the rate at which the no-signalling value of a game decays under parallel repetition truly depends on the game distribution?

Another issue that would be worth investigating is whether constrained de Finetti reductions could also be used to establish parallel repetition results for the classical or quantum value of multi-player games. Formally, the sets of classical correlations $\mathrm{C}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ and quantum correlations $\mathrm{Q}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ are defined as follows:

$$P \in \mathrm{C}(\underline{\mathcal{A}}|\underline{\mathcal{X}}) :\Leftrightarrow \forall \underline{x}, \underline{a} \; P(\underline{a}|\underline{x}) = \sum_{m \in \mathcal{M}} Q(m) P_1(a_1|x_1\, m) \cdots P_\ell(a_\ell|x_\ell\, m),$$

for some p.d. Q on $\mathcal{M}$ and some p.d.'s $P_i(\cdot|x_i\, m)$ on $\mathcal{A}_i$.

$$P \in \mathrm{Q}(\underline{\mathcal{A}}|\underline{\mathcal{X}}) :\Leftrightarrow \forall \underline{x}, \underline{a} \; P(\underline{a}|\underline{x}) = \langle \psi | M(x_1)_{a_1} \otimes \cdots \otimes M(x_\ell)_{a_\ell} | \psi \rangle,$$

for some state $|\psi\rangle$ on $\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_\ell$ and some POVMs $M(x_i)$ on $\mathcal{H}_i$.
And the classical [quantum] value of an ℓ-player game G with distribution T and predicate V, denoted $\omega_{\mathrm{C}}(G)$ [$\omega_{\mathrm{Q}}(G)$], is then naturally defined as the maximum [supremum] of the winning probability

$$\Pr\{\text{win}\} = \sum_{\underline{a}, \underline{x}} T(\underline{x}) V(\underline{a}, \underline{x}) P(\underline{a}|\underline{x})$$

over all $P \in \mathrm{C}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$ [$P \in \mathrm{Q}(\underline{\mathcal{A}}|\underline{\mathcal{X}})$].

In the classical case, the first parallel repetition result for two-player games was established by Raz [29], and later improved by Holenstein [20], while Rao [28] gave a concentration bound. However, the proof techniques are arguably not as straightforward as via de Finetti reductions, and do not generalise directly to any number ℓ of players. In the quantum case, even less is known. The best parallel repetition result up to now is the one established by Chailloux and Scarpa [8] (subsequently improved by Chung, Wu and Yuen [10]), which applies to two-player (ℓ-player) free games, and from there to games with full support. That is why being able to export ideas from the de Finetti approach to these two cases would be of great interest. Roughly speaking, the problem we are facing is the following: Given an n-symmetric correlation $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$, we can always write the first step in the proof of Lemma 9, i.e.

$$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n} \leq \mathrm{poly}(n) \int_{Q_{\underline{\mathcal{A}\mathcal{X}}}} F\left(T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}, Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n} \right)^2 Q_{\underline{\mathcal{A}\mathcal{X}}}^{\otimes n} \, dQ_{\underline{\mathcal{A}\mathcal{X}}}. \tag{8}$$

Now, we would like to argue that if $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ is a classical [quantum] correlation, then the p.d.'s $Q_{\underline{\mathcal{A}\mathcal{X}}}$ for which the fidelity weight in the r.h.s. of Eq. (8) is not exponentially small are necessarily close to being of the form $T_{\underline{\mathcal{X}}} R_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$ for some classical [quantum] correlation $R_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$. This was precisely our proof philosophy in the no-signalling case. However, the fact that the classical and quantum conditions are not properties that one can read off on the marginals, contrary to the no-signalling one, seems to be a first obstacle to surmount.

One related legitimate question would be the following: is it possible to make an even stronger statement than the one that, as explained above, we either are looking for (in the classical and quantum cases) or already have (in the no-signalling case)? Namely, could we upper-bound

$T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ by a de Finetti distribution analogous to that in the r.h.s. of Eq. (8), but with weight strictly 0 on p.d.'s $Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}$ which are not of the form $T_{\underline{\mathcal{X}}} R_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$, for $R_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$ belonging to the same class as $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$? The answer to this question is no. Indeed, such improved de Finetti reduction would imply a strong parallel repetition result, which we know does not hold (see [2] for a similar discussion). So the best we can hope for is really to show that the fidelity weight in our upper-bounding de Finetti distribution is exponentially small on the p.d.'s which are too far from being of the desired form.

Finally, let us briefly comment on the main spirit difference between the present work and the one by Arnon-Friedman *et al.* [2]. Our approach consists in using a more "flexible" de Finetti reduction, in which the information on the correlation $P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}$ and the p.d. $T_{\underline{\mathcal{X}}}^{\otimes n}$ of interest are kept in the upper-bounding de Finetti distribution, through the fidelity weight $F(T_{\underline{\mathcal{X}}}^{\otimes n} P_{\underline{\mathcal{A}}^n|\underline{\mathcal{X}}^n}, Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}^{\otimes n})^2$. Whereas in [2], any initial correlation is first upper-bounded by the same universal de Finetti correlation, on which a test (specifically tailored to the considered game distribution) is performed in a second step, that has the property of letting pass, resp. rejecting, with high probability the strategies which are no-signalling, resp. too signalling. So it seems in the end that both approaches are quite closely related: in our case, the "signalling test" which is applied to a given p.d. $Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}$ is nothing else than the maximal fidelity of $Q_{\underline{\mathcal{A}}\underline{\mathcal{X}}}$ to the set of p.d.'s of the form $T_{\underline{\mathcal{X}}} R_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$, with $R_{\underline{\mathcal{A}}|\underline{\mathcal{X}}}$ no-signalling, being above or below a certain threshold value. Also, it would be interesting (and potentially fruitful) to investigate whether one could combine in some way the techniques yielding Lemmas 7 and 8, to get a de Finetti reduction result that would have the advantages of both: namely, that is designed for conditional p.d.'s while at the same carrying the relevant information on the conditional p.d. it is applied to.

ACKNOWLEDGEMENTS

We thank Rotem Arnon-Friedman, Harry Buhrman and Renato Renner for insightful discussions on parallel repetition and de Finetti theorems.

This research was supported by the European Research Council (Advanced Grant IRQUAT, ERC-2010-AdG-267386), the European Commission (STREP RAQUEL, FP7-ICT-2013-C-323970), the Spanish MINECO (projects FIS2008-01236 and FIS2013-40627-P), with the support of FEDER funds, the Generalitat de Catalunya (CIRIT project 2014-SGR-966), and the French CNRS (ANR projects OSQPI 11-BS01-0008 and Stoq 14-CE25-0033).

---

[1] R. Arnon-Friedman, R. Renner, "de Finetti reductions for correlations", *J. Math. Phys.* **56**:052203 (2015); arXiv[quant-ph]:1308.0312.

[2] R. Arnon-Friedman, R. Renner, T. Vidick, "Non-signalling parallel repetition using de Finetti reductions"; arXiv[quant-ph]:1411.1582.

[3] J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, D. Roberts, "Non-local correlations as an information theoretic resource", *Physical Review A* **71**(2):022101 (2005); arXiv:quant-ph/0404097.

[4] J. Bell, "On the Einstein Podolsky Rosen paradox", *Physics* **1**:195 (1964).

[5] M. Berta, M. Christandl, R. Renner, "The quantum reverse Shannon theorem based on one-shot information theory", *Commun. Math. Phys.* **306**, 579 (2011); arXiv[quant-ph]:0912.3805.

[6] H. Buhrman, N. Chandran, S. Fehr, R. Gelles, V. Goyal, R. Ostrovsky, C. Schaffner, "Position-Based Quantum Cryptography: Impossibility and Constructions"; arXiv[quant-ph]:1009.2490.

[7] H. Buhrman, S. Fehr, C. Schaffner, "On the Parallel Repetition of Multi-Player Games: The No-Signaling Case", *Proc. 9th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC14)*, vol. 27, pp. 24-35 (2014); arXiv[quant-ph]:1312.7455.

[8] A. Chailloux, G. Scarpa, "Parallel repetition of free entangled games: simplification and improvements"; arXiv[quant-ph]:1410.4397.

[9] M. Christandl, R. König, R. Renner, "Post-selection technique for quantum channels with applications to quantum cryptography", *Phys. Rev. Lett.* **102**:020504 (2009); arXiv[quant-ph]:0809.3019.

[10] M.K. Chung, X. Wu, H. Yuen, "Parallel repetition for entangled k-player games via fast quantum search"; arXiv[quant-ph]:1501.00033 (2015).

[11] J.F. Clauser, M.A. Horne, A. Shimony, R.A. Holt, "Proposed experiment to test local hidden-variable theories", *Phys. Rev. Lett.* **23**(15), pp. 880-884 (1969).

[12] R. Cleve, W. Slofstra, F. Unger, S. Upadhyay, "Strong parallel repetition theorem for quantum XOR proof systems", *Comput. Complex.* **17**(2), pp. 282-299 (2008); arXiv:quant-ph/0608146.

[13] I. Dinur, D. Steurer, T. Vidick, "A parallel repetition theorem for entangled projection games"; arXiv[quant-ph]:1310.4113.

[14] R. Duan, S. Severini, A. Winter, "On zero-error communication via quantum channels in the presence of noiseless feedback"; arXiv[quant-ph]:1502.02987.

[15] S. Fehr, J. Kaniewski, M. Tomamichel, S. Wehner, "A Monogamy-of-Entanglement Game with Applications to Device-Independent Quantum Cryptography", *New J. Phys.* **15**:103002; arXiv[quant-ph]:1210.4359.

[16] U. Feige, "On the success probability of the two provers in one- round proof systems", *Proc. 6th IEEE Symposium on Structure in Complexity Theory*, pp. 116-123 (1991).

[17] U. Feige, O. Verbitsky, "Error Reduction by Parallel Repetition – A Negative Result", *Combinatorica* **2**(4), pp. 461-478 (2002).

[18] C. A. Fuchs, J. van de Graaf, "Cryptographic Distinguishability Measures for Quantum-Mechanical States", *IEEE Trans. Inf. Theory* **45**(4), pp. 1216-1227 (1999).

[19] E. Hänggi, *Device-independent quantum key distribution*, PhD thesis, ETH Zürich (2010); arXiv[quant-ph]:1012.3878.

[20] T. Holenstein, "Parallel repetition: simplifications and the no-signaling case", *Theory of Computing* **5**(1), pp. 141-172 (2009); arXiv:cs/0607139.

[21] T. Ito, "Polynomial-space approximation of no-signaling provers", *Automata, Languages and Programming* **6198**, pp. 140-151 (2010); arXiv[cs.CC]:09082363.

[22] R. Jain, A. Pereszlenyi, P. Yao, "A parallel repetition theorem for entangled two-player one-round games under product distributions"; arXiv[quant-ph]:1311.6309.

[23] J. Kempe, O. Regev, B. Toner, "Unique games with entangled provers are easy", *Proc. 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS08)*, pp. 457-466 (2008); arXiv[quant-ph]:0710.0655.

[24] J. Kempe, T. Vidick, "Parallel repetition of entangled games", *Proc. 43rd Annual ACM Symposium on Theory of Computing (STOC11)*, pp. 353-362 (2011); arXiv[quant-ph]:1012.4728.

[25] C. Lancien, A. Winter, "Flexible constrained de Finetti reductions and applications", in preparation.

[26] A. Leverrier, R. García-Patrón, R. Renner, N.J. Cerf, "Security of continuous-variable quantum key distribution against general attacks", *Phys. Rev. Lett.* **110**:030502 (2013); arXiv[quant-ph]:1208.4920.

[27] S. Popescu, D. Rohrlich, "Nonlocality as an axiom", *Foundations of Physics* **24**(3), pp. 379385 (1994).

[28] A. Rao, "Parallel repetition in projection games and a concentration bound", *SIAM J. Comput.* **40**(6), pp. 1871-1891 (2011).

[29] R. Raz, "A parallel repetition theorem", *SIAM J. Comput.* **27**(3), pp. 763-803 (1998).

[30] R. Rosen, "A k-provers parallel repetition theorem for a version of no-signaling model", *Discrete Math., Alg. and Appl.* **2**(4), pp. 457-468 (2010).

[31] B.S. Tsirelson, "Quantum generalizations of Bell's inequality", *Lett. Math. Phys.* **4**(2), pp. 93-100 (1980).