

Quantum hedging in two-round prover-verifier interactions

Srinivasan Arunachalam*, Abel Molina†, and Vincent Russo†

March 14, 2017

Abstract

We consider the problem of a particular kind of quantum correlation that arises in some two-party games. In these games, one player is presented with a question they must answer, yielding an outcome of either “win” or “lose”. Molina and Watrous [MW12] studied such a game that exhibited a perfect form of *hedging*, where the risk of losing a first game can completely offset the corresponding risk for a second game. This is a non-classical quantum phenomenon, and establishes the impossibility of performing strong error-reduction for quantum interactive proof systems by parallel repetition, unlike for classical interactive proof systems. We take a step in this article towards a better understanding of the hedging phenomenon by giving a complete characterization of when perfect hedging is possible for a natural generalization of the game in [MW12]. Exploring in a different direction the subject of quantum hedging, and motivated by implementation concerns regarding loss-tolerance, we also consider a variation of the protocol where the player who receives the question can choose to restart the game rather than return an answer. We show that in this setting there is no possible hedging for any game played with state spaces corresponding to finite-dimensional complex Euclidean spaces.

1 Overview and motivation

The interactions we study consist of parallel repetitions of a game played between players Alice and Bob, also referred to as the *verifier* and *prover* respectively. The setting of the game is:

1. Alice prepares a question, and sends this question to Bob.
2. Bob generates an answer, and sends it back to Alice.
3. Alice evaluates this answer and decides if Bob wins or loses.

It is assumed that Bob has complete knowledge of Alice’s specification, including both the method used to determine Alice’s question and the criteria that she uses to determine whether Bob has won or lost the game.

Molina and Watrous [MW12] consider a specific instance of this setting where Alice sends half of a 2-qubit Bell state $\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$ to Bob. Bob replies with a qubit

*QuSoft, CWI

†Institute for Quantum Computing & School of Computer Science, University of Waterloo.

and Alice evaluates Bob’s answer by measuring his qubit and the second half of the Bell state against the state $\cos(\pi/8)|00\rangle + \sin(\pi/8)|11\rangle$. A victory for Bob corresponds to the outcome of Alice measurement corresponding to $\cos(\pi/8)|00\rangle + \sin(\pi/8)|11\rangle$. When Alice and Bob play two repetitions of this game in parallel, the results in [MW12] show that there exists a strategy for Bob that guarantees he wins *at least one* of the two repetitions with probability 1. However, when the game is played once, the probability that Bob wins is at most $\cos(\pi/8)^2 \approx 0.8536$. Playing two repetitions in parallel leads then to a *hedging* phenomenon, where if Bob wants to decrease his chance of losing both repetitions, he can do so by not playing each game independently and optimally. This hedging is also *perfect*, in the sense that Bob can completely offset the risk of losing both games.

This is a completely quantum phenomenon, with no classical counterpart. Indeed, when classical information is considered, and for any game that fits the setting we study, it is immediate to show that when Bob wants to win at least k out of n parallel repetitions, it is optimal for him to play independently (however, this is not the case when considering multiple provers [For89, Fei91, Raz98, Hol07, BG15]). This establishes the non-triviality of the set of outcome distributions that are possible to obtain from parallel repetition of the games that we study, when compared to the classical case. In particular, it immediately illustrates that the technique of parallel repetition cannot be used to trivially achieve strong error reduction for the complexity class QIP(2), a class studied for example in [Raz05, Weh06, JUW09, HMW14]. The quantum hedging phenomenon is also an example where the quantum version of a game produces outcomes unachievable by its classical counterpart. Most famously considered by Bell [Bel64], this type of violation has been observed in a number of game-like frameworks [CHSH, Mer90, Per90, CHTW04, BRSdW12, Reg12, CJPPG15].

It is natural then to ask how general is the hedging phenomenon, both qualitatively and quantitatively. A complete understanding of this question would allow us to characterize the outcome distributions that can arise from Alice and Bob playing n parallel repetitions of a prover-verifier game in our setting. Consequently, it could lead to a protocol for achieving error reduction via parallel repetition for QIP(2) simpler than the one currently known [JUW09]. The techniques used to achieve such an understanding could conceivably also extend to the analysis of prover-verifier games involving further rounds of communication, and more generally to other kinds of multi-party quantum interactions. This would lead to results for the corresponding complexity classes (and likely also for their classical parallels) about error reduction by parallel repetition. Taking a step towards such a complete understanding, we consider in Section 3 a 2-parameter generalization of the game in [MW12], and characterize when Bob can guarantee that he wins at least 1 out of n parallel repetitions, for every n . We also give optimal strategies for Bob to win at least 1 out of n parallel repetitions, both when perfect hedging is possible and not possible. We believe these findings are a valuable stepping stone towards a more complete understanding of hedging behaviors for fully arbitrary initial states, fully arbitrary quantum measurements, and k -out-of- n settings, as well as highly non-trivial from a mathematical point of view. The formulas that we obtain also open the door for connections between the hedging phenomenon and recent work [BJOP14] involving generalizations of the PBR game [PBR], as we will discuss further in Section 5.

Exploring in a different direction the subject of quantum hedging, it also seems natural to consider the possibility of implementing a game that exhibits quantum hedging using existing

quantum information processing devices. One possible choice would be to use optical quantum devices, but the immediate concern arises [Smi11] of how to account for the fact that photon losses will often occur, leading to a communication error between Alice and Bob. Even if one chose another implementation method where communication is more reliable, one would still need to consider the general fact that communication errors can occur. More generally, the consideration of implementation inaccuracies is a standard direction in which to extend results concerning quantum information protocols – see for example recent work regarding loss-tolerant protocols for quantum coin-flipping [AMS10] and QKD, [TCK⁺14] and noise-tolerant protocols for quantum money [PYJ⁺12], quantum coin-flipping [ZZ15] and quantum randomness amplification [BRG⁺16].

Along this direction, we consider a loss-tolerant formalism in Section 4, and prove that under our formalism quantum hedging is not possible. To model communication errors, we assume that Alice cannot distinguish a communication error from Bob choosing not to return an answer. Therefore, our formalism simply allows for the possibility that Bob chooses not to return an answer, in which case the game is repeated. Bob choosing in our formalism a random whether to return an answer or not would correspond to a genuine disruption of communication, while Bob strategizing about when to return an answer would correspond to Bob using communication errors as an excuse to avoid a losing outcome. Our particular choice of framework can also be seen as adding postselection to two-round quantum prover-verifier interactions. This addition of post-selection has been previously considered in the case of single-party quantum computation [Aar05, SDLF10, YS11, MdW15], but not to our knowledge in the context of quantum prover-verifier interactions.

The techniques used to obtain our results in Section 4 are inspired by the techniques in [Fiu04], which studies a particular case of quantum cloning. The connection between quantum cloning and semidefinite programming was observed in [ADM02], and has been used to obtain results regarding quantum cloning (see the review in [CF06]). However, this is the first time to our knowledge that this connection with semidefinite programming acts as a bridge to apply ideas about optimal quantum cloning to the context of fully general two-round quantum prover-verifier interactions.

Both of our results leave room for further progress. In particular, one can consider hedging in a wider context than the setting in Section 3, and consider formalisms that model communication errors in a different way than in Section 4. We give some suggestions in Section 5 concerning corresponding choices for further exploration.

2 Notation

We will denote the set of binary strings with length n as $\{0,1\}^n$. These strings will be indexed from 0 to $n - 1$. Therefore, we will denote the n successive binary symbols or bits in $a \in \{0,1\}^n$ as a_0, \dots, a_{n-1} . $\wedge r$, $\vee r$, and $\oplus r$ refer to the logical AND, OR, and XOR of the bits of $r \in \{0,1\}^n$, respectively, while $|r|$ refers to its Hamming weight.

Vector spaces associated with a quantum system are defined as complex Euclidean spaces. We denote these spaces by the capital script letters \mathcal{X} , \mathcal{Y} , and \mathcal{Z} . The dual of a vector x in a complex Euclidean vector space \mathcal{X} will be the linear functional $A : \mathcal{X} \rightarrow \mathbb{C}$ that maps y to $\langle x, y \rangle$. For a d -dimensional complex Euclidean space, we will often fix a standard *computational* basis and, using bra-ket notation, address its elements and their duals as

$\{|0\rangle, \dots, |d-1\rangle\}$ and $\{\langle 0|, \dots, \langle d-1|\}$, respectively. The encoding of the label inside a bra or a ket will often be done in binary for ease of explanation.

The complex vector space of linear operators of the form $A : \mathcal{X} \rightarrow \mathcal{Y}$ is denoted by $L(\mathcal{X}, \mathcal{Y})$. We write $A \in L(\mathcal{X})$ as a shorthand for $A : \mathcal{X} \rightarrow \mathcal{X}$. The adjoint X^* of an operator $X \in L(\mathcal{X})$ is the operator such that for all $u, v \in \mathcal{X}$, $\langle u, Xv \rangle = \langle X^*u, v \rangle$. An operator $H \in L(\mathcal{X})$ is *Hermitian* if $H = H^*$. We write $\text{Herm}(\mathcal{X})$ to denote the set of all Hermitian operators. The inner product $\langle A, B \rangle = \text{Tr}(AB)$ between two operators $A, B \in \text{Herm}(\mathcal{X})$ is real and satisfies $\langle A, B \rangle = \langle B, A \rangle$. If an operator $P \in \text{Herm}(\mathcal{X})$, and all eigenvalues of P are non-negative, then we call P *positive semidefinite*, and refer to all such operators as $P \in \text{Pos}(\mathcal{X})$. For a Hermitian operator H , $\|H\|$ denotes the operator norm of H , that is, the largest absolute value of an eigenvalue. If for an operator $\rho \in \text{Pos}(\mathcal{X})$ it is the case that $\text{Tr}(\rho) = 1$, then ρ is said to be a *density operator*, and is referred to as $\rho \in D(\mathcal{X})$. We adopt the convention of writing $\mathcal{I}_{\mathcal{X}}$ as opposed to \mathcal{I} to indicate that the identity is acting on the space \mathcal{X} when convenient to do so. We will define the $\text{vec} : L(\mathcal{X}, \mathcal{Y}) \rightarrow \mathcal{X} \otimes \mathcal{Y}$ mapping to be the one that takes yx^* to $x \otimes y$, for x and y elements of the standard/computational basis of \mathcal{X} and \mathcal{Y} . This can be seen as flattening a matrix into a vector. For any two operators $A, B \in L(\mathcal{X}, \mathcal{Y})$, it will hold that $\langle A, B \rangle = \langle \text{vec}(A), \text{vec}(B) \rangle$.

We also consider linear mappings of the form $\Phi : L(\mathcal{X}) \rightarrow L(\mathcal{Y})$. The space of all such mappings is denoted as $T(\mathcal{X}, \mathcal{Y})$. For each $\Phi \in T(\mathcal{X}, \mathcal{Y})$, a unique adjoint mapping $\Phi^* \in T(\mathcal{Y}, \mathcal{X})$ is defined by the property that $\langle Y, \Phi(X) \rangle = \langle \Phi^*(Y), X \rangle$ for all $X \in L(\mathcal{X})$ $Y \in L(\mathcal{Y})$. Throughout this work, we define quantum states by the set of density operators $\rho \in D(\mathcal{X})$, with \mathcal{X} a complex Euclidean space. Associated with the space \mathcal{X} one may consider a *register* denoted X in which the state ρ is contained. We consider measurements of a register X as being described by a set of positive semidefinite operators $\{P_a : a \in \Sigma\}$ indexed by a finite non-empty set Σ of measurement outcomes which satisfies the constraint $\sum_{a \in \Sigma} P_a = \mathcal{I}_{\mathcal{X}}$. By performing a measurement on X in state ρ , the outcome $a \in \Sigma$ results with probability $\langle P_a, \rho \rangle$. These measurements are known as POVMs. We can also consider quantum states stored across n registers $(\mathsf{X}_1, \mathsf{X}_2, \dots, \mathsf{X}_n)$. We can describe the joint state of those registers by a density operator $\sigma \in D(\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_n)$.

A linear mapping $\Phi : L(\mathcal{X}) \rightarrow L(\mathcal{Y})$ is said to be *completely positive* if $\Phi \otimes \mathcal{I}_{\mathcal{Z}}$ is a map that preserves positive semidefiniteness for every complex Euclidean space \mathcal{Z} and Φ is said to be *trace-preserving* if $\text{Tr}(\Phi(X)) = \text{Tr}(X)$ for all $X \in L(\mathcal{X})$. We define a *quantum channel* as a linear mapping $\Phi : L(\mathcal{X}) \rightarrow L(\mathcal{Y})$ that is completely positive and trace preserving. A channel transforms some state ρ stored in register X into the state $\Phi(\rho)$ of another register Y . The set of all channels between such two registers is denoted by $C(\mathcal{X}, \mathcal{Y})$, and is a compact and convex set. Note that the channel corresponding to an unitary operator U is the one that maps a quantum state σ to $U\sigma U^*$.

For spaces \mathcal{X} and \mathcal{Y} , one may define the Choi representation of an operator $\Phi \in T(\mathcal{X}, \mathcal{Y})$ as $J(\Phi) = \sum_{i,j} \Phi(|i\rangle\langle j|) \otimes |i\rangle\langle j|$, where $J : T(\mathcal{X}, \mathcal{Y}) \rightarrow L(\mathcal{Y} \otimes \mathcal{X})$, and i and j iterate over the computational basis for \mathcal{X} . Note that the mapping J is linear, bijective, and multiplicative with respect to the tensor product. The Choi representation has a number of more complex properties, three of which will be useful to us:

Lemma 1.

1. The mapping Φ is completely positive if and only if $J(\Phi) \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X})$.

2. The mapping Φ is trace preserving if and only if $\text{Tr}_{\mathcal{Y}}(J(\Phi)) = \mathcal{I}_{\mathcal{X}}$
3. $\Phi(Z) = \text{Tr}_{\mathcal{X}} [J(\Phi) (\mathcal{I}_{\mathcal{Y}} \otimes Z^T)]$

We refer the reader to [Wat11] for the proof of Lemma 1 and further details on the notation.

3 Hedging to win 1 out of n parallel repetitions of a game

Let G denote the following game:

1. Alice prepares the 2-qubit state $\rho_{\alpha} = u_{\alpha}u_{\alpha}^* \in \text{D}(\mathcal{X} \otimes \mathcal{Z})$ in registers (X, Z) where

$$u_{\alpha} = \alpha |00\rangle + \sqrt{1 - \alpha^2} |11\rangle \in \mathcal{X} \otimes \mathcal{Z}, \quad (1)$$

for $\alpha \in (0, 1]$. Alice sends register X to Bob.

2. Bob applies a channel $\Phi \in \text{C}(\mathcal{X}, \mathcal{Y})$ to the contents of X . This results in a state $\sigma \in \text{D}(\mathcal{Y} \otimes \mathcal{Z})$, contained in registers (Y, Z) . Register Y is sent back to Alice.
3. Alice performs a measurement on the state σ . This measurement is $\{P_{0,\theta}, P_{1,\theta}\}$ for $\theta \in [0, 2\pi)$, with

$$\begin{aligned} P_{1,\theta} &= v_{\theta}v_{\theta}^*, \quad P_{0,\theta} = \mathcal{I} - P_{1,\theta}, \\ v_{\theta} &= \cos(\theta) |00\rangle + \sin(\theta) |11\rangle \in \mathcal{Y} \otimes \mathcal{Z}. \end{aligned} \quad (2)$$

An outcome of “0” or “1” denotes a losing or winning outcome for Bob, respectively.

One can imagine repeating the game G n times in parallel. This is denoted as G^n , and illustrated in Figure 1. In this setting, Alice prepares n states $\rho_{1,\alpha}, \dots, \rho_{n,\alpha}$ in registers $((\mathsf{X}_1, \mathsf{Z}_1), \dots, (\mathsf{X}_n, \mathsf{Z}_n))$ where

$$\rho_{1,\alpha} \in \text{D}(\mathcal{X}_1 \otimes \mathcal{Z}_1), \dots, \rho_{n,\alpha} \in \text{D}(\mathcal{X}_n \otimes \mathcal{Z}_n). \quad (3)$$

Alice sends the registers $(\mathsf{X}_1, \dots, \mathsf{X}_n)$ to Bob and he applies his quantum channel,

$$\Phi_n \in \text{C}(\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_n, \mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n). \quad (4)$$

The resulting states are sent back to Alice and she performs a series of n projective measurements with respect to the operators $P_{0,\theta}, P_{1,\theta}$. These give n outcomes of either 0 or 1, loss or win. Since Bob’s actions are not required to respect the independence of the measurements, they may cause correlations between the n measurement outcomes.

Indeed, in [MW12], Molina and Watrous analyzed G^n for $n = 2$ where $\alpha = 1/\sqrt{2}$ and $\theta = \pi/8$, and found that Bob wins one out of the two games with certainty if he applies a specific correlated strategy. If on the other hand, Bob treated each repetition independently, it would *not* be guaranteed that Bob would win at least one of the games.

We consider G^n for any $n \geq 1$ and ask for what values of α and θ is it true that Bob can make sure to win with certainty *at least* one out of the n games in G^n . Let $p_{n,\alpha,\theta}(\Phi_n) \in [0, 1]$

be the probability that Bob loses all n outcomes of G^n using the strategy defined by Φ_n . This is given by:

$$p_{n,\alpha,\theta}(\Phi_n) = \left\langle P_{0,\theta}^{\otimes n}, (\Phi_n \otimes \mathcal{I}_{\mathcal{Z}_1 \otimes \dots \otimes \mathcal{Z}_n}) \left(\bigotimes_{i=1}^n \rho_{i,\alpha} \right) \right\rangle. \quad (5)$$

Let $m_{n,\alpha,\theta} \in [0, 1]$ be $\min_{\Phi_n} p_{n,\alpha,\theta}(\Phi_n)$. We refer to a quantum channel Φ_n that minimizes $m_{n,\alpha,\theta}$ as an *optimal strategy*. That is, equal to the *minimum probability* with which Bob loses each game over all choices of quantum channels Φ_n of the form in (4). If $m_{n,\alpha,\theta}$ evaluates to 0, then there exists a Φ_n that ensures Bob wins at least one game.

The quantity $m_{n,\alpha,\theta}$ is expressible as the optimal value of a semidefinite program. Let $Q_{0,\alpha,\theta} \in \text{Pos}(\mathcal{Y}_i \otimes \mathcal{X}_i)$ be defined as

$$Q_{0,\alpha,\theta} = (\mathcal{I}_{\mathcal{Y}_i} \otimes \Psi_{\rho_\alpha})(P_{0,\theta}), \quad (6)$$

where the mapping $\Psi_{\rho_\alpha} : L(\mathcal{Z}) \rightarrow L(\mathcal{X})$ is defined by $J(\Psi_{\rho_\alpha}) = \overline{\rho_\alpha}$ (the entry-wise complex conjugate of ρ_α). This makes $Q_{0,\alpha,\theta}$ a function of both $P_{0,\theta}$ and ρ_α .

It follows from Lemma 1 of [MW12] that Q_0 is positive semidefinite, and that for any channel $\Phi : L(\mathcal{X}) \rightarrow L(\mathcal{Y})$, we have $\langle P_{0,\theta}, (\Phi \otimes \mathcal{I})(\rho_{i,\alpha}) \rangle = \langle Q_{0,\theta}, J(\Phi) \rangle$. This can be proved by considering the case where $\rho_{i,\alpha}$ corresponds to a rank-1 operator that transforms a state of the computational basis into another one, and then using the linearity properties of the inner product (see Appendix A.1) for more details of this derivation). Putting this together with facts 1 and 2 about the Choi representation in Lemma 1, and the bijective property of the $J(\cdot)$ map, we obtain that the following primal and dual pair gives a semidefinite program to compute $m_{n,\alpha,\theta}$:




Figure 1: The parallel repetition G^n of n copies of a game G of the type we study.

$$\begin{aligned}
& \underline{m_{n,\alpha,\theta}: \text{ Primal problem}} \\
\text{minimize: } & \left\langle Q_{0,\alpha,\theta}^{\otimes n}, X \right\rangle \\
\text{subject to: } & \text{Tr}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n}(X) = \mathcal{I}_{\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_n}, \\
& X \in \text{Pos}(X_1 \otimes \mathcal{Y}_1 \otimes \dots \otimes \mathcal{X}_n \otimes \mathcal{Y}_n).
\end{aligned} \tag{7}$$

$$\begin{aligned}
& \underline{m_{n,\alpha,\theta}: \text{ Dual problem}} \\
\text{maximize: } & \text{Tr}(Y) \\
\text{subject to: } & \pi(\mathcal{I}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n} \otimes Y) \pi^* \leq Q_{0,\alpha,\theta}^{\otimes n}, \\
& Y \in \text{Herm}(\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_n).
\end{aligned} \tag{8}$$

where π is a unitary permutation operator defined by the action

$$\pi(y_1 \otimes \dots \otimes y_n \otimes x_1 \otimes \dots \otimes x_n) = y_1 \otimes x_1 \otimes \dots \otimes y_n \otimes x_n$$

for all $y_1 \in \mathcal{Y}_1, \dots, y_n \in \mathcal{Y}_n$ and $x_1 \in \mathcal{X}_1, \dots, x_n \in \mathcal{X}_n$. Note that strong duality holds for the above semidefinite program, by choosing the primal and dual feasible solutions (X, Y) for the application of Slater's theorem as a scalar multiple of the identity. The derivation to obtain this semidefinite program is similar to that in [MW12], and previously in [GW07] and [Gut10]. We point the reader to [AMR13] for MATLAB code that solves SDPs (7) and (8), using the CVX convex optimization package [GBY08].

We present now for fixed n and α the range of θ which characterizes the measurements for which Bob can make sure he wins at least 1 parallel repetition in G^n . That is, it characterizes when is Bob able to perform perfect hedging. Furthermore, we present strategies that give Bob an optimal probability to win at least 1 out of n games, both when Bob is able to perform perfect hedging and when he is not.

Theorem 2. *Let*

$$\begin{aligned}
\theta_{n,\alpha} &= \tan^{-1} \left(\sqrt{\frac{1}{\alpha^2} - 1} (2^{1/n} - 1) \right), \\
\gamma_{n,\alpha} &= \tan^{-1} \left(\sqrt{\frac{1}{\alpha^2} - 1} \left(\frac{1}{2^{1/n} - 1} \right) \right),
\end{aligned} \tag{9}$$

where the trigonometric domain is restricted to $[0, \pi/2]$. If and only if Alice's rank-1 projective measurement $\{P_0, P_1\}$ is parametrized by $\theta \in [\theta_{n,\alpha}, \gamma_{n,\alpha}]$, then there exists a strategy for Bob to perform perfect hedging.

We see then that the angle $\pi/8$ used for θ in [MW12] corresponds to the lower bound $\theta_{2,1/\sqrt{2}} = \pi/8$ from Theorem 2, but also that perfect hedging can be attained for this setting up to $\gamma_{2,1/\sqrt{2}} = 3\pi/8$. Note that as the number of games n increases, the size of this range increases. Moreover, for any choice of θ in $(0, \pi/2)$, there is an n large enough for perfect hedging to be possible. As one can see in our plot of $\theta_{n,\alpha}$ and $\gamma_{n,\alpha}$, the cases where perfect hedging are possible are symmetric with respect to the case where the initial state and the desired final state are the same (i.e., $\theta = \tan^{-1}(\sqrt{1 - \alpha^2}/\alpha)$). Note also that the size of the range where perfect hedging is possible is minimized for for $\theta = 0$ and $\theta = \pi/2$, which correspond to a standard basis measurement done by Alice.




Figure 2: $\gamma_{2,\alpha}$ and $\theta_{2,\alpha}$ as a function of $\tan^{-1}\left(\frac{\sqrt{1-\alpha^2}}{\alpha}\right)$

The proof of Theorem 2 follows immediately from Lemma 5 and Lemma 6, stated below. Theorem 2 results in the following corollary:

Corollary 3. *For a fixed n , perfect hedging occurs for the largest range of θ angles when Alice initially prepares a maximally entangled state (that is, when $\alpha = \frac{1}{\sqrt{2}}$).*

The proof for the corollary follows from directly maximizing $\gamma_{n,\alpha} - \theta_{n,\alpha}$ over all α , by taking derivatives with respect to α . This corollary tells us then that an assumption that the initial state is maximally entangled represents an extremal case in our quantum hedging context. One might be able to use when trying to generalize our results, as we will further discuss in Section 5.

In the following lemmas, we define an optimal choice for Bob of the channel Φ that he applies to the part of the state he receives from Alice:

Lemma 4. *Let $n \geq 2$ be a positive integer, let $\alpha \in (0, 1]$, let $\theta_{n,\alpha}$ and $\gamma_{n,\alpha}$ be angles defined as in Theorem 2, and let*

$$\begin{aligned}\Lambda_n &= \sum_{r \in \{0,1\}^n} (-1)^{\wedge r + \oplus r} |r\rangle \langle r|, \\ \Xi_n &= \sum_{r \in \{0,1\}^n} (-1)^{\vee r + \oplus r} |r\rangle \langle r|,\end{aligned}\tag{10}$$

be unitary operators that Bob applies as his strategy in G^n . Then it holds that

$$p_{n,\alpha,\theta_{n,\alpha}}(\Lambda_n) = 0 = p_{n,\alpha,\gamma_{n,\alpha}}(\Xi_n).\tag{11}$$

This shows the existence of strategies $\{\Lambda_n, \Xi_n\}$ for Bob at $\{\theta_{n,\alpha}, \gamma_{n,\alpha}\}$ that achieve a value of 0 for the SDP (7). The next lemma proves that for all points within these two bounds there exists such a strategy as well. Note that Λ_n and Ξ_n do not depend on α . Also, note that when $n = 2$, Bob's unitary Λ_2 on the two qubits that he receives is

$$\Lambda_2 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix},\tag{12}$$

which gives us the same strategy as in [MW12]. The proof of the lemma follows from observing that the final state after Bob applies Λ_n/Ξ_n has zero overlap with the state corresponding to Bob losing all the repetitions. The details of the derivation are included in Appendix A.2.

Lemma 5. *In the scenario where the projective measurements are parametrized by $\theta \in [\theta_{n,\alpha}, \gamma_{n,\alpha}]$ for $\theta_{n,\alpha}$ and $\gamma_{n,\alpha}$ defined as in Theorem 2, Bob can apply the strategy corresponding to the following unitary operator to achieve perfect hedging for 1 out of n games:*

$$(-1)^n |0^n\rangle\langle 0^n| - |1^n\rangle\langle 1^n| + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} (-1)^{n+i} k_r |r\rangle\langle r|, \quad (13)$$

where for a fixed choice of $|r| = i$, the corresponding k_r are $\binom{n}{i}$ complex numbers with the following properties

$$k_r = \begin{cases} s_{\theta,\alpha,n} + i\sqrt{1 - s_{\theta,\alpha,n}^2} & \text{for } \lfloor \binom{n}{i}/2 \rfloor \text{ values of } r, \\ s_{\theta,\alpha,n} - i\sqrt{1 - s_{\theta,\alpha,n}^2} & \text{for } \lfloor \binom{n}{i}/2 \rfloor \text{ values of } r, \\ -1 & \text{for the remaining values of } r \text{ when } \binom{n}{i} \text{ is} \\ & \text{odd and } \tan(\theta) \geq \sqrt{\frac{1}{\alpha^2} - 1}, \\ 1 & \text{for the remaining values of } r \text{ when } \binom{n}{i} \text{ is odd} \\ & \text{and } \tan(\theta) < \sqrt{\frac{1}{\alpha^2} - 1}, \end{cases}$$

where $s_{\theta,\alpha,n}$ is a real number $\in [-1, 1]$ whose existence we guarantee in the proof of this lemma.

Since Bob has complete knowledge of the game, for any $\theta \in [\theta_{n,\alpha}, \gamma_{n,\alpha}]$ Bob can apply the strategy corresponding to the angle θ selected by Alice. It is clear that the optimal strategy for Bob is not unique, since our definition does not uniquely specify which coefficients k_r correspond to which values of r . This lemma is derived by performing a computation (similar to the one for Lemma 4) that computes the overlap between the resulting state after Bob applies the strategy we describe and the state corresponding to Bob losing all n repetitions. Then, we consider the cases $s_{\theta,\alpha,n} = -1$ and $s_{\theta,\alpha,n} = 1$ and obtain through continuity arguments that there must be a value of $s_{\theta,\alpha,n}$ in the $[-1, 1]$ range that results in perfect hedging. The details of the corresponding derivation are included in Appendix A.3.

We have thus far considered the case when perfect hedging is possible. The following result deals with characterizing the scenario when perfect hedging is not possible, and provides a corresponding strategy for Bob to play optimally.

Lemma 6. *For $n \geq 2$ and for $\theta \in [0, \theta_{n,\alpha}) \cup (\gamma_{n,\alpha}, \pi/2]$ perfect hedging cannot occur, and the strategies Λ_n and Ξ_n mentioned in Lemma 4 are respective optimal strategies for Bob.*

The proof of this lemma is obtained by using SDP complementary slackness [Wat] to obtain a candidate solution for the dual SDP (8) with the same objective value as the chance of achieving 1-out-of- n hedging for Λ_n/Ξ_n . Then, one can use a direct sum decomposition of the matrices involved in the SDP constraint to prove the feasibility of this candidate solution.

The details of the corresponding calculations are available in Appendix A.4. Note that the strategy Bob adopts is independent of the parameter θ , implying that when perfect hedging is not possible the strategy is optimal regardless of the projective measurements chosen by Alice.

It can also be observed from Lemma 5 and Lemma 6 that a unitary (and in fact, a diagonal in the computational basis) strategy is always sufficient for Bob to win at least once with optimal probability. Note that it intuitively makes sense that Bob’s strategy is a diagonal unitary, since switching a $|0\rangle$ to a $|1\rangle$ or vice-versa on his side will produce a state with no overlap with the target state $\cos(\theta)|00\rangle + \sin(\theta)|11\rangle$.

4 (Lack of) Hedging in a Loss-Tolerant Prover-Verifier Model

We consider a variation of the prover-verifier setting where Bob has the choice to not respond to Alice, in order to model communication errors, as described in Section 1. If Bob chooses not to respond, and therefore Alice does not receive an answer, the game is repeated again, and this goes on until an answer is returned by Bob. Bob might want to do this whenever using his complete knowledge of the game, he can predict that an answer will result in Alice obtaining a negative outcome in her measurement. Indeed, to see how this variation can change the result of an interaction, consider the following game where Bob is always forced to return an answer:

1. Alice prepares the maximally entangled state $\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$ and sends the second qubit to Bob.
2. Bob responds by sending a qubit to Alice.
3. Alice ignores Bob’s answer, and measures the qubit she kept with respect to the projective measurement $\{P_0, P_1\}$, where $P_0 = |1\rangle\langle 1|$ and $P_1 = |0\rangle\langle 0|$.

It is clear that the maximum probability for Bob to win the game is 50%. This follows from the fact that the actions of Bob cannot alter the reduced state that Alice holds, and the outcome of the interaction depends only on this state. However, the situation changes drastically when Bob is allowed to return no answer in the second step. In that case, Bob can choose to perform a measurement using the computational basis on the qubit he receives. If the measurement results in the outcome $|0\rangle$, corresponding to P_1 , he will return an answer, and otherwise he will not, and force a restart. The entanglement between the qubit that Alice keeps and the one that Bob receives guarantees then that the outcome will always be the successful one.

It seems clear then that giving Bob the choice to abort the protocol can have significant changes on what optimal behaviors for Bob are like. This motivates the consideration of whether any form of quantum hedging (perfect or not) is still possible in the “repetition after communication error” setting for an arbitrary two-message quantum-verifier interaction (described by an arbitrary finite-dimensional initial quantum state ρ prepared by Alice and an arbitrary finite-dimensional POVM $\{P_i\}$ used to determine the interaction’s outcome.) We ask in this context then whether it will be optimal for Bob to play each interaction independently when trying to optimize his chance of winning at least k out of n parallel interactions.

To answer this question, we will assume in our analysis that Bob always has a nonzero chance of winning a single interaction. If this were not the case, the question of whether or not hedging occurs would be uninteresting. This is because in this case, the optimal probability for Bob to win k out of n parallel repetitions would always be zero. To see why, assume to the contrary that Bob can manage to win $k > 0$ out of $n > 1$ repetitions with non-zero probability. Then, whenever Bob plays a single game with Alice, he could simulate the input for $n - 1$ additional interactions, and since the possibility that he wins $k > 0$ of the n games is greater than zero, and the situation is symmetrical, the possibility that he wins the single “real” game is greater than zero as well, which contradicts our starting premise.

Furthermore, we need to specify how does the “repetition after communication error” aspect of the framework interacts with the “repeating n interactions in parallel” aspect of the framework. For simplicity, we will make in our model the assumption that whenever Alice does not receive an answer to one out of n parallel interactions, she will restart all of the n parallel interactions.

To start our analysis, we consider an intermediate setting where we allow Bob to not give an answer, and Alice does not repeat the interaction when she doesn’t obtain an answer, and instead counts that as a loss for Bob. This means that Bob can return a state with trace less than one. Using the properties of the Choi representation, and following the same analysis as in [MW12] and Section 3, the optimal probability for Bob of achieving outcome a is the value of

$$\begin{aligned}
 & \text{Primal problem} \\
 & \text{maximize: } \quad \langle Q_a, X \rangle \\
 & \text{subject to: } \quad \text{Tr}_{\mathcal{Y}}(X) \leq \mathcal{I}_{\mathcal{X}}, \\
 & \quad \quad \quad X \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X}),
 \end{aligned} \tag{14}$$

where Q_a is defined as in (6), starting from an arbitrary POVM $\{P_i\}$ and a state ρ . Without loss of generality, we assume that Bob wants to achieve quantum hedging with respect to outcome a , and group all other outcomes into a single outcome corresponding to Q_{1-a} .

Now we take into account the fact that the interaction is repeated whenever an answer is not received. To do this, it is enough to divide the objective function, which corresponds to the probability of obtaining outcome a , by the probability that an answer is returned. This is because we can ignore previous rounds of the interaction, since the repeated rounds occur in series, and Alice acts independently between them. Indeed, the way in which previous rounds would be taken into account would be with an additional input for Bob, corresponding to his memory after the previous rounds of the protocol. But the fact that there is no computational restriction on Bob and no hidden information means that for any possible value of that input, Bob could just simulate the previous rounds to generate it, so the additional memory input is not needed, and we can ignore previous rounds.

Note that the division by the probability that Bob returns an answer would not be possible if Bob just chose not to return an answer. However, that strategy can just be ignored as a non-optimal one, since we are assuming Bob can win with non-zero probability.

The probability that an answer is returned is the trace of the state after Bob returns an answer, which is a linear function of the variable X in SDP (14). In particular, the probability

is given by $\langle E, X \rangle$, where

$$\begin{aligned} E &= \sum_i Q_i = \sum_i (\mathcal{I}_{\mathcal{L}(\mathcal{Y})} \otimes \Psi_\rho)(P_i) \\ &= (\mathcal{I}_{\mathcal{L}(\mathcal{Y})} \otimes \Psi_\rho) \mathcal{I}_{\mathcal{Y} \otimes \mathcal{Z}} = \mathcal{I}_{\mathcal{Y}} \otimes \text{Tr}_{\mathcal{Z}}(\bar{\rho}), \end{aligned} \tag{15}$$

and the last step uses the third fact in Lemma 1. Note that since $\sum_i Q_i = E$, $Q_a \leq E$.

This tells us then how to modify the SDP (14) that describes Bob's optimal probability of obtaining outcome a in a way that takes into account our loss-tolerant framework. In particular, we have that the equivalent of SDP (14) is now given by

$$\begin{aligned} &\text{Primal problem} \\ \text{maximize:} & \frac{\langle Q_a, X \rangle}{\langle E, X \rangle} \\ \text{subject to:} & \text{Tr}_{\mathcal{Y}}(X) \leq \mathcal{I}_{\mathcal{X}}, \\ & X \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X}), \langle E, X \rangle \neq 0. \end{aligned} \tag{16}$$

We use now an analysis inspired by the one in [CF06] to obtain a more explicit form for the value of this SDP. First, notice that scaling a solution X by a nonzero constant will not change the value of the objective function. Since the partial trace operation preserves positive semidefiniteness, we can then get rid of the $\text{Tr}_{\mathcal{Y}}(X) \leq \mathcal{I}_{\mathcal{X}}$ constraint:

$$\begin{aligned} &\text{Primal problem} \\ \text{maximize:} & \frac{\langle Q_a, X \rangle}{\langle E, X \rangle} \\ \text{subject to:} & X \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X}), \langle E, X \rangle \neq 0. \end{aligned} \tag{17}$$

At this point, we can assume that X corresponds to a rank-one operator. To see why, consider an X that corresponds to a sum of two solutions, X_1 and X_2 . Then, the value of the objective function will be

$$\frac{\langle Q_a, X_1 \rangle + \langle Q_a, X_2 \rangle}{\langle E, X_1 \rangle + \langle E, X_2 \rangle} \leq \max \left(\frac{\langle Q_a, X_1 \rangle}{\langle E, X_1 \rangle}, \frac{\langle Q_a, X_2 \rangle}{\langle E, X_2 \rangle} \right), \tag{18}$$

where the inequality follows from the fact that all values on the left-hand side are positive. We obtain the problem

$$\begin{aligned} &\text{Primal problem} \\ \text{maximize:} & \frac{x^* Q_a x}{x^* E x} \\ \text{subject to:} & x \in \mathcal{Y} \otimes \mathcal{X}, x^* E x \neq 0. \end{aligned} \tag{19}$$

Note now that we can assume without loss of generality that an optimal solution x is contained within the support of E . In this domain the Moore-Penrose pseudo-inverse of E , E^+ , acts as a bijection. Therefore, we replace x by $(E^+)^{1/2}x$ in the objective function, and obtain

$$\begin{aligned}
& \text{Primal problem} \\
\text{maximize: } & \frac{x^*(E^+)^{1/2}Q_a(E^+)^{1/2}x}{x^*x} \\
\text{subject to: } & x \in \mathcal{Y} \otimes \mathcal{X}, x \perp \ker(E),
\end{aligned} \tag{20}$$

which has the value $\|(E^+)^{1/2}Q_a(E^+)^{1/2}\|$. We denote this as $\|\Lambda\|$.

When Bob wants to be successful in at least k out of n parallel interactions, with Alice acting independently, one just needs to replace Q_a by the sum of tensor products of Q_i 's corresponding to at least k outcomes equal to a . Remembering that the sum of all the Q_i is equal to E , the same analysis that we performed for a single repetition gives us an optimal probability of $\|\Lambda_{k,n}\|$, with $\Lambda_{k,n}$ given by :

$$\left\| (\sqrt{E^+})^{\otimes n} \left(E^{\otimes n} - \sum_{t=0}^{k-1} \pi_t (Q_{1-a}^{\otimes n-t} \otimes Q_a^{\otimes t}) \right) (\sqrt{E^+})^{\otimes n} \right\| \tag{21}$$

where $\pi_t(x)$ is the sum of all $\binom{n}{t}$ unique permutations of x .

As an aside, note that one can assume that ρ corresponds to a pure state ψ . This is because given a protocol where Alice initially prepares a mixed state, we can easily modify it so that Alice prepares a purification of that state instead, and just ignores the extra qubits when performing the final measurement. Using this, we observe an interesting fact about this model, which is that at least when one restricts Bob to perform a rank-one measurement, the optimal success probability for Bob does not depend on the Schmidt coefficients of ψ . This is proved by letting the initial state that Alice holds be given by $\sum_i \sqrt{p_i} a_i \otimes b_i$, and the state corresponding to Bob's projection by $\sum_i \sqrt{q_i} c_i \otimes d_i$. Using algebraic manipulations we obtain that the optimal probability of winning for Bob in a single parallel repetition is

$$\left\| \sum_{i,j,k,l} \sqrt{q_j q_l} b_i^* d_l d_j^* b_k a_i a_k^* \otimes c_j c_l^* \right\|, \tag{22}$$

with no dependence on the p_i .

This suggests that the example we gave at the beginning of this section might capture all the additional power Bob has in this model. In particular, it suggests that an optimal strategy for Bob might always consist of performing an orthogonal measurement on the qubits he is given, and then refusing to give an answer except when he obtains the "best" outcome.

As for our main subject of concern (quantum hedging), it turns out that in the model we just described quantum hedging is not possible. One can interpret this as saying that Bob is already so powerful in one single repetition (since he can choose not to return an answer) than the power to entangle several answers does not add anything in comparison. More precisely, we have the following theorem:

Theorem 7. *Consider a two-message prover-verifier interaction characterized by an arbitrary initial state ρ and an arbitrary POVM $\{P_i\}$, both on a finite number of qubits. Then, under the loss-tolerant setting described in this section, it is optimal for Bob to play independently in order to maximize his chance of winning at least k out of n parallel interactions.*

The statement of the theorem results from a straightforward spectral analysis of the $\Lambda_{k,n}$ operator by induction on n and then k . The details of the corresponding computation are included in Appendix A.5.

5 Discussion

We have analyzed generalizations of a specific prover-verifier interaction where the verifier can use a quantum hedging strategy to win at least one of n parallel repetitions with a higher probability than what would have been possible playing each game independently. This interesting phenomenon was originally described in [MW12], where the authors illustrated an explicit example of perfect hedging when two repetitions of the game were carried out. It was previously unknown how the perfect hedging phenomenon generalizes to the case when n repetitions of the game are performed. We resolved this question for a generalization of the game in [MW12], and provided strategies for Bob that allow him to achieve perfect hedging whenever it is possible.

We also analyzed a variant of this setting where Bob is not obligated to return an answer to Alice. In a practical sense, Bob's refusal to respond to Alice can be viewed in terms of an experimental setup where the lack of a response could correspond to a communication error [Smi11]. This consideration led to a different semidefinite program that characterized the interaction between Alice and Bob. We then used this SDP (16) to ask whether or not Bob still had the ability to take advantage of hedging behavior, with a negative answer.

While we have considered this hedging behavior in a number of settings, there are still many questions remaining. As mentioned, we have characterized the conditions that allow Bob to win 1 out of n repetitions in a framework that generalizes the game in [MW12]. However, it still remains open to determine the conditions under which Bob can always win at least k out of n repetitions for some $k > 1$. It would be interesting to determine the threshold of k for which perfect hedging occurs, and to also provide a characterization in regards to the strategy that Bob uses to achieve this result. Running numerical instances for higher values of k and n using a simple formulation in CVX [GBY08] quickly becomes computationally infeasible, as can be observed from the software we have provided in [AMR13]. It is possible that this code could be optimized to consider further cases, leading to conjectures regarding the behavior for arbitrary k and n that could be then proved analytically. Based on our current numerical evidence, it is possible that Bob cannot perfectly hedge more than $k = n/2$ games. Note also that when $k \leq n/2$ one can design a strategy for the goal of winning k out of n repetitions by dividing the n parallel repetitions into several smaller groups, and then using the strategies described in this paper in order to always win at least one repetition in each group. It is left as an open question (whose solution we believe to be a significant task) whether the range of parameters in which the resulting strategy always wins k out of n repetitions is the optimal one. Motivated by our results in Corollary 3, one could also look into the subject of reducibility between different games in our framework, asking for example whether there is a procedure with an intuitive operational description that transforms a game with an arbitrary shared initial state between Alice and Bob to one where the initial shared state is now maximally entangled, while the possibilities of achieving k -out-of- n hedging remains the same.

It is also worth noting that the problem of conclusive state exclusion, which was recently considered in [BJOP14], seems to be connected to the interaction we have analyzed in this work. In this problem, Alice prepares a mixed state from a given distribution and sends it to Bob, and for Bob to win, he has to accurately discard at least one of the possible options. In [BJOP14] the PBR game, originally formulated in [PBR], was analyzed in terms of an semidefinite program using the conclusive state exclusion framework. Some of the

formulas we obtain in Section 3 are similar to the ones [BJOP14] derive in their analysis of the PBR game, specifically equations (9) and (10). Looking at the SDPs involved in their work and in ours, it seems clear that the similarity arises from the fact that diagonal unitaries happen to be optimal for hedging. The fact that they are optimal means that the optimization problem we examine in SDP (7) is equivalent to that of optimizing along complex vectors where each entry of the vector is a unit. Then, to establish the connection with the PBR setting, one would establish an equivalence between these types of vectors and highly symmetrical projective measurements like those obtained as optimal solutions in the corresponding PBR state exclusion setting. However, in a setting with initial states outside the $\alpha |00\rangle + \sqrt{1 - \alpha^2} |11\rangle$ family we consider in Section 3, there is no reason why the optimal channel for winning 1 out of n parallel interactions should correspond to a diagonal unitary. It remains then to see whether any similar connections can be established between such a setting and a state exclusion setting. It seems plausible that further work clarifying these connections could be used to apply existing results concerning the conclusive state exclusion framework to the hedging framework, and vice versa.

One could also further consider the setting in which protocol errors are considered. Here, we have assumed that Bob can delay returning an answer for as many iterations of the protocol as he desires. An obvious follow-up question then is to determine whether an advantage from hedging behavior is possible when this is not the case. One might restrain Bob to behaviors where on average he will return an answer within a fixed number of iterations, or introduce constraints be of the form “After X iterations, Bob’s probability of having return an answer must be at least equal to Y ”. A special case of those constraints that might be particularly interesting is when Bob is required to return an answer within a fixed number of iterations. We could also modify the way in which the “repeating after failure” and “repeating in parallel” frameworks interact. In particular, we could have Alice repeat only a subset of interactions if answers corresponding to the other interactions have been obtained from Bob.

Note that when trying to analyze more general models (in both the ideal and loss-tolerant cases) along the lines described in this section, it might be fruitful to look into whether it is possible to again use ideas from the quantum cloning literature, as we did here in Section 4. It is possible as well that progress can be made using representation theory tools to simplify or avoid the analysis of semidefinite programs, as done for example in [GP04, CLP07, CSW, KGDdS15].

Acknowledgements

We would like to thank Devin Smith for the question that led to Section 4 in this paper. A significant amount of thanks is also due to John Watrous for numerous insightful discussions and suggestions. We thank Ronald de Wolf for helpful comments. We wish to thank Nathaniel Johnston for helpful suggestions, as well as for the use of his QETLAB quantum entanglement MATLAB package [Joh]. Thanks are also due to Alessandro Cosentino, Gus Gutoski, and Christopher Perry for insightful discussions. This work was partially supported by Canada’s NSERC, the US ARO, the ERC Consolidator Grant QPROGRESS, the Mike and Ophelia Lazaridis Fellowship program, and the David R. Chariton Graduate Scholarship program.

References

- [Aar05] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 461, pages 3473–3482. The Royal Society, 2005.
- [ADM02] Koenraad Audenaert and Bart De Moor. Optimizing completely positive maps using semidefinite programming. *Physical Review A*, 65(3):030302, 2002.
- [AMR13] Srinivasan Arunachalam, Abel Molina, and Vincent Russo. Software for implementing some of the semidefinite programs in this paper. Available at <https://bitbucket.org/vprusso/quantum-hedging>, 2013.
- [AMS10] Nati Aharon, Serge Massar, and Jonathan Silman. Family of loss-tolerant quantum coin-flipping protocols. *Physical Review A*, 82(5):052307, 2010.
- [Bel64] John Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1(3):195–200, 1964.
- [BG15] Mark Braverman and Ankit Garg. Small value parallel repetition for general games. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 335–340. ACM, 2015.
- [BJOP14] Somshubhro Bandyopadhyay, Rahul Jain, Jonathan Oppenheim, and Christopher Perry. Conclusive exclusion of quantum states. *Physical Review A*, 89(2):022336, 2014.
- [BRG⁺16] Fernando GSL Brandão, Ravishankar Ramanathan, Andrzej Grudka, Karol Horodecki, Michał Horodecki, Paweł Horodecki, Tomasz Szarek, and Hanna Wojewódka. Realistic noise-tolerant randomness amplification using finite number of devices. *Nature communications*, 7, 2016.
- [BRSdW12] Harry Buhrman, Oded Regev, Giannicola Scarpa, and Ronald de Wolf. Near-optimal and explicit bell inequality violations. *Theory of Computing*, 8(1):623–645, 2012.
- [CF06] Nicolas Cerf and Jaromir Fiurasek. Optical quantum cloning. *Progress in Optics*, 49:455, 2006.
- [CHSH] John Clauser, Michael Horne, Abner Shimony, and Richard Holt.
- [CHTW04] Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Computational Complexity, 2004. Proceedings. 19th IEEE Annual Conference on*, pages 236–249. IEEE, 2004.
- [CJPPG15] Tom Cooney, Marius Junge, Carlos Palazuelos, and David Pérez-García. Rank-one quantum games. *Computational Complexity*, 24(1):133–196, 2015.

- [CLP07] Andrew M Childs, Andrew J Landahl, and Pablo A Parrilo. Quantum algorithms for the ordered search problem via semidefinite programming. *Physical Review A*, 75(3):032335, 2007.
- [CSW] Matthias Christandl, Norbert Schuch, and Andreas Winter.
- [Fei91] Uriel Feige. On the success probability of the two provers in one-round proof systems. In *Structure in Complexity Theory Conference, 1991., Proceedings of the Sixth Annual*, pages 116–123. IEEE, 1991.
- [Fiu04] Jaromír Fiurášek. Optimal probabilistic cloning and purification of quantum states. *Physical Review A*, 70(3):032308, 2004.
- [For89] Lance Jeremy Fortnow. *Complexity-theoretic aspects of interactive proof systems*. PhD thesis, Massachusetts Institute of Technology, 1989.
- [GBY08] Michael Grant, Stephen Boyd, and Yinyu Ye. CVX: Matlab software for disciplined convex programming. <http://cvxr.com/cvx/>, 2008.
- [GP04] Karin Gatermann and Pablo A Parrilo. Symmetry groups, semidefinite programs, and sums of squares. *Journal of Pure and Applied Algebra*, 192(1):95–128, 2004.
- [Gut10] Gus Gutoski. Quantum strategies and local operations. *arXiv preprint arXiv:1003.0038*, 2010.
- [GW07] Gus Gutoski and John Watrous. Toward a general theory of quantum games. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 565–574. ACM, 2007.
- [HMW14] Patrick Hayden, Kevin Milner, and Mark Wilde. Two-message quantum interactive proofs and the quantum separability problem. *Quantum Information & Computation*, 14(5&6):384–416, 2014.
- [Hol07] Thomas Holenstein. Parallel repetition: simplifications and the no-signaling case. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 411–419. ACM, 2007.
- [Joh] Nathaniel Johnston. QETLAB: MATLAB Software for quantum entanglement. Available at <http://www.qetlab.com/>.
- [JUW09] Rahul Jain, Sarvagya Upadhyay, and John Watrous. Two-message quantum interactive proofs are in PSPACE. In *50th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543. IEEE, 2009.
- [KGDdS15] Hari Krovi, Saikat Guha, Zachary Dutton, and Marcus P da Silva. Optimal measurements for symmetric quantum states with applications to optical communication. *Physical Review A*, 92(6):062333, 2015.

- [MdW15] Urmila Mahadev and Ronald de Wolf. Rational approximations and quantum algorithms with postselection. *Quantum Information & Computation*, 15(3&4):295–307, 2015.
- [Mer90] N David Mermin. Simple unified form for the major no-hidden-variables theorems. *Physical Review Letters*, 65(27):3373, 1990.
- [MW12] Abel Molina and John Watrous. Hedging bets with correlated quantum strategies. In *Proc. R. Soc. A*. The Royal Society, 2012.
- [PBR] Matthew Pusey, Jonathan Barrett, and Terry Rudolph.
- [Per90] Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151(3-4):107–108, 1990.
- [PYJ⁺12] Fernando Pastawski, Norman Y Yao, Liang Jiang, Mikhail D Lukin, and J Ignacio Cirac. Unforgeable noise-tolerant quantum tokens. *Proceedings of the National Academy of Sciences*, 109(40):16079–16082, 2012.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998.
- [Raz05] Ran Raz. Quantum information and the PCP theorem. In *46th Annual IEEE Symposium on Foundations of Computer Science.*, pages 459–468. IEEE, 2005.
- [Reg12] Oded Regev. Bell violations through independent bases games. *Quantum Information & Computation*, 12(1-2):9–20, 2012.
- [SDLF10] Oksana Scegulnaja-Dubrovskaja, Lelde Lāce, and Rūsiņš Freivalds. Postselection finite quantum automata. In *International Conference on Unconventional Computation*, pages 115–126. Springer, 2010.
- [Smi11] Devin Smith. Personal communication, 2011.
- [TCK⁺14] Kiyoshi Tamaki, Marcos Curty, Go Kato, Hoi-Kwong Lo, and Koji Azuma. Loss-tolerant quantum cryptography with imperfect sources. *Physical Review A*, 90(5):052314, 2014.
- [Wat] John Watrous. <https://cs.uwaterloo.ca/~watrous/CS766/ProblemSets/solutions.2.pdf>.
- [Wat11] John Watrous. Theory of quantum information lecture notes. <https://cs.uwaterloo.ca/~watrous/LectureNotes.html>, 2011.
- [Weh06] Stephanie Wehner. Entanglement in interactive proof systems with binary answers. In *STACS 2006*, pages 162–171. Springer, 2006.
- [YS11] Abuzer Yakaryilmaz and AC Say. Probabilistic and quantum finite automata with postselection. *arXiv preprint arXiv:1102.0666*, 2011.
- [ZZ15] Sheng Zhang and Yuexin Zhang. Quantum coin flipping secure against channel noises. *Physical Review A*, 92(2):022313, 2015.

A Mathematical derivations

A.1 Verification of procedure to group the starting state and the final measurement into a single variable

Consider first the case where we have a matrix $A \in L(\mathcal{X} \otimes \mathcal{Z})$ that corresponds to a rank-1 operator that transforms a state of the computational basis into another one. Let it be equal to $|a\rangle\langle c| \otimes |b\rangle\langle d|$, with $|a\rangle\langle c| \in L(\mathcal{X})$, $|b\rangle\langle d| \in L(\mathcal{Z})$. The channel $\Psi_A : L(\mathcal{Z}) \rightarrow L(\mathcal{X})$ such that $J(\Psi_A) = \overline{A}$ is then the one that maps $|b\rangle\langle d| \in L(\mathcal{Z})$ to $|a\rangle\langle c| \in L(\mathcal{X})$, and everything else in the computational basis for $L(\mathcal{Z})$ to 0.

Consider now an operator $M \in L(\mathcal{Y} \otimes \mathcal{Z})$, and a channel $\Phi : L(\mathcal{X}) \rightarrow L(\mathcal{Y})$. We want to verify that

$$\langle M, (\Phi \otimes \mathcal{I})(A) \rangle = \langle (\mathcal{I} \otimes \Psi_A)(M), J(\Phi) \rangle. \quad (23)$$

To do so, consider a computational basis decomposition $M = \sum_{i,j,k,l} m_{i,j,k,l} |i\rangle\langle j| \otimes |k\rangle\langle l|$, with $|i\rangle\langle j| \in L(\mathcal{Y})$, $|k\rangle\langle l| \in L(\mathcal{Z})$. Then, the left hand side of (23) is equal to

$$\left\langle \sum_{i,j,k,l} m_{i,j,k,l} |i\rangle\langle j| \otimes |k\rangle\langle l|, \Phi(|a\rangle\langle c|) \otimes |b\rangle\langle d| \right\rangle = \left\langle \sum_{i,j} m_{i,j,b,d} |i\rangle\langle j|, \Phi(|a\rangle\langle c|) \right\rangle,$$

and the right hand side of (23) is equal to

$$\begin{aligned} \left\langle (\mathcal{I} \otimes \Psi_A) \left(\sum_{i,j,k,l} m_{i,j,k,l} |i\rangle\langle j| \otimes |k\rangle\langle l| \right), J(\Phi) \right\rangle &= \left\langle \sum_{i,j} m_{i,j,b,d} |i\rangle\langle j| \otimes |a\rangle\langle c|, J(\Phi) \right\rangle \\ &= \left\langle \sum_{i,j} m_{i,j,b,d} |i\rangle\langle j|, \Phi(|a\rangle\langle c|) \right\rangle, \end{aligned}$$

so (23) holds.

(23) does extend by linearity to any choice of $A \in L(\mathcal{X} \otimes \mathcal{Z})$. Indeed, assume that it holds for $A, B \in L(\mathcal{X} \otimes \mathcal{Z})$, and consider a linear combination $\lambda_A A + \lambda_B B$, with $\lambda_A, \lambda_B \in \mathbb{C}$. Then, the left hand side of (23) will be given by

$$\begin{aligned} \langle M, (\Phi \otimes \mathcal{I})(\lambda_A A + \lambda_B B) \rangle &= \lambda_A \langle M, (\Phi \otimes \mathcal{I})(A) \rangle + \lambda_B \langle M, (\Phi \otimes \mathcal{I})(B) \rangle \\ &= \lambda_A \langle (\mathcal{I} \otimes \Psi_A)(M), J(\Phi) \rangle + \lambda_B \langle (\mathcal{I} \otimes \Psi_B)(M), J(\Phi) \rangle \\ &= \langle \overline{\lambda_A} (\mathcal{I} \otimes \Psi_A)(M) + \overline{\lambda_B} (\mathcal{I} \otimes \Psi_B)(M), J(\Phi) \rangle. \end{aligned}$$

We want to prove then that

$$\overline{\lambda_A} (\mathcal{I} \otimes \Psi_A)(M) + \overline{\lambda_B} (\mathcal{I} \otimes \Psi_B)(M) = (\mathcal{I} \otimes \Psi_{\lambda_A A + \lambda_B B})(M).$$

To do so, we use the third property of the Choi representation introduced in Lemma 1, and express $\overline{\lambda_A} (\mathcal{I} \otimes \Psi_A)(M) + \overline{\lambda_B} (\mathcal{I} \otimes \Psi_B)(M)$ as

$$\begin{aligned}
& \overline{\lambda}_A \operatorname{Tr}_{\mathcal{Y} \otimes \mathcal{Z}} (J(\mathcal{I}_{\mathcal{L}(Y)} \otimes \Psi_A)(\mathcal{I}_{\mathcal{X} \otimes \mathcal{Z}} \otimes M^T)) + \\
& \overline{\lambda}_B \operatorname{Tr}_{\mathcal{Y} \otimes \mathcal{Z}} (J(\mathcal{I}_{\mathcal{L}(Y)} \otimes \Psi_B)(\mathcal{I}_{\mathcal{X} \otimes \mathcal{Z}} \otimes M^T)) \\
&= \operatorname{Tr}_{\mathcal{Y} \otimes \mathcal{Z}} ((\overline{\lambda}_A J(\mathcal{I}_{\mathcal{L}(Y)} \otimes \Psi_A) + \overline{\lambda}_B J(\mathcal{I}_{\mathcal{L}(Y)} \otimes \Psi_B)) (\mathcal{I}_{\mathcal{X} \otimes \mathcal{Z}} \otimes M^T)) \\
&= \operatorname{Tr}_{\mathcal{Y} \otimes \mathcal{Z}} ((J(\mathcal{I}_{\mathcal{L}(Y)}) \otimes \overline{\lambda}_A \overline{A} + J(\mathcal{I}_{\mathcal{L}(Y)}) \otimes \overline{\lambda}_B \overline{B}) (\mathcal{I}_{\mathcal{X} \otimes \mathcal{Z}} \otimes M^T)) \\
&= \operatorname{Tr}_{\mathcal{Y} \otimes \mathcal{Z}} ((J(\mathcal{I}_{\mathcal{L}(Y)}) \otimes (\overline{\lambda}_A \overline{A} + \overline{\lambda}_B \overline{B})) (\mathcal{I}_{\mathcal{X} \otimes \mathcal{Z}} \otimes M^T)) \\
&= (\mathcal{I} \otimes \Psi_{\lambda_A \overline{A} + \lambda_B \overline{B}})(M).
\end{aligned}$$

A.2 Derivation for Lemma 4

Proof. Given that n parallel repetitions of the game are considered, our claim states that Bob will win *at least* one out of the n repetitions if he adopts Λ_n as his strategy when the projective measurement made by Alice corresponds to the parameter $\theta_{n,\alpha}$. A similar argument also holds for Ξ_n at the corresponding angle $\gamma_{n,\alpha}$. We prove this explicitly for the strategy Λ_n , and the other case follows using the same argument. The proof of this lemma uses a technique of conditioning where we consider the resulting state conditioned on Bob obtaining a losing outcome in the first projective measurement of Alice, and the corresponding probability for such an outcome. Then, we generalize this procedure to the rest of the parallel repetitions. To conclude the proof, we set the probability of the “all-losing state” at the end to zero, which allows us to solve for θ in the final equation.

First, let us define the pure states:

$$\begin{aligned}
v_\theta &= \cos(\theta) |00\rangle + \sin(\theta) |11\rangle, & s_\theta &= |01\rangle, \\
w_\theta &= \sin(\theta) |00\rangle - \cos(\theta) |11\rangle, & t_\theta &= |10\rangle,
\end{aligned} \tag{24}$$

where we recall from Section 3 that $v_\theta \in \mathcal{Y} \otimes \mathcal{Z}$ is the state which corresponds to the winning projective measurement outcome, and w_θ, s_θ , and $t_\theta \in \mathcal{Y} \otimes \mathcal{Z}$ are the states that correspond to the losing projective measurement. Essentially, Bob is trying then to transform the state prepared by Alice to something as close as possible to v_θ , while restricted to operating on one half on the state.

Let Λ_n be the operator defined as

$$\Lambda_n = \sum_{r \in \{0,1\}^n} (-1)^{\wedge r + \oplus r} |r\rangle \langle r|, \tag{25}$$

Λ'_n be the similar operator

$$\Lambda'_n = \sum_{r \in \{0,1\}^n} (-1)^{\oplus r} |r\rangle \langle r|. \tag{26}$$

and define the vector κ_n as

$$\kappa_n = \sum_{a \in \{0,1\}^n} \bigotimes_{i=0}^{n-1} \alpha^{(1-a_i)} (1 - \alpha^2)^{a_i/2} |a_i a_i\rangle. \tag{27}$$

We run now through the parallel repetition of n copies of the game. Since the initial shared state is $u_\alpha^{\otimes n} = (\alpha |00\rangle + \sqrt{1-\alpha^2} |11\rangle)^{\otimes n}$, the state after Bob applies his channel (acting on his qubits for all of the n parallel repetitions) is

$$f_\alpha^0 = (\Lambda_n \otimes \mathcal{I}_{\mathcal{Z}_1 \otimes \dots \otimes \mathcal{Z}_n}) \kappa_n \quad (28)$$

We shall condition now on Bob losing the first out of n parallel repetitions. It should be noted that since Alice starts with the entangled state $u_\alpha^{\otimes n}$ and Bob performs a unitary diagonal operation, the states s_θ and t_θ in (24) do not contribute to the losing projective measurement outcome. Once we condition on Bob losing the first game, the resulting state is then a normalization of

$$\begin{aligned} f_{\alpha,\theta}^1 &= (w_\theta w_\theta^* \otimes \mathcal{I}) f_\alpha^0 \\ &= w_\theta \otimes \alpha \sin(\theta) (\Lambda'_{n-1} \otimes \mathcal{I}_{\mathcal{Z}_2 \otimes \dots \otimes \mathcal{Z}_n}) \kappa_{n-1} \\ &\quad + w_\theta \otimes \sqrt{1-\alpha^2} \cos(\theta) (\Lambda_{n-1} \otimes \mathcal{I}_{\mathcal{Z}_2 \otimes \dots \otimes \mathcal{Z}_n}) \kappa_{n-1}, \end{aligned} \quad (29)$$

with the associated probability being $(f_{\alpha,\theta}^1)^* f_{\alpha,\theta}^1$.

Generalizing this to Bob losing all n games, one can observe that the -1 's for the $\cos(\theta)$ term in w_θ cancel the negative terms from the $(-1)^{\oplus r}$ term in Λ_n , and Equation (29) generalizes to:

$$\begin{aligned} f_{\alpha,\theta}^n &= (w_\theta)^{\otimes n} \left(\alpha^n \sin(\theta)^n + n(\alpha^{n-1} \sqrt{1-\alpha^2}) \sin(\theta)^{n-1} \cos(\theta) + \dots \right. \\ &\quad \left. + n(\alpha(1-\alpha^2)^{(n-1)/2}) \cos(\theta)^{n-1} \sin(\theta) - (1-\alpha^2)^{n/2} \cos(\theta)^n \right) \end{aligned} \quad (30)$$

$$= (w_\theta)^{\otimes n} \left((\alpha \sin(\theta) + \sqrt{1-\alpha^2} \cos(\theta))^n - 2(1-\alpha^2)^{n/2} \cos(\theta)^n \right). \quad (31)$$

In order for Bob to ensure he wins *at least* 1 out of the n games with certainty, we require that $\|f_{\alpha,\theta}^n\| = 0$, which implies:

$$(\alpha \sin(\theta) + \sqrt{1-\alpha^2} \cos(\theta))^n - 2(1-\alpha^2)^{n/2} \cos(\theta)^n = 0. \quad (32)$$

This implies that for the angle $\theta_{n,\alpha} = \tan^{-1} \left(\sqrt{\frac{1}{\alpha^2} - 1} (2^{1/n} - 1) \right)$, the strategy corresponding to Λ_n gives us a perfect hedging strategy. Following the same procedure, using the strategy corresponding to Ξ_n yields the similar condition that:

$$(\alpha \sin(\theta) + \sqrt{1-\alpha^2} \cos(\theta))^n - 2\alpha^n \sin(\theta)^n = 0, \quad (33)$$

giving us as a solution $\gamma_{n,\alpha} = \tan^{-1} \left(\sqrt{\frac{1}{\alpha^2} - 1} \left(\frac{1}{2^{1/n} - 1} \right) \right)$. \square

A.3 Derivation for Lemma 5

Proof. As in the previous proof, to win at least 1 out of n games, Bob needs to avoid the outcome corresponding to the state $(\sin(\theta) |00\rangle - \cos(\theta) |11\rangle)^{\otimes n}$ (other states for the losing

outcome can be ignored since Bob's strategy corresponds to a diagonal matrix). Let us now define a matrix

$$D = \sum_{r \in \{0,1\}^n} (-1)^{|r|} \sin(\theta)^{n-|r|} \cos(\theta)^{|r|} |r\rangle \langle r|, \quad (34)$$

such that $(\sin(\theta) |00\rangle - \cos(\theta) |11\rangle)^{\otimes n} = \text{vec}(D)$. For convenience, we denote $\lambda = \tan(\theta)$, and rewrite D as

$$D = \cos(\theta)^n \sum_{r \in \{0,1\}^n} (-1)^{|r|} \lambda^{n-|r|} |r\rangle \langle r|. \quad (35)$$

We also introduce an operator

$$F = \sum_{r \in \{0,1\}^n} (1 - \alpha^2)^{|r|/2} \alpha^{n-|r|} |r\rangle \langle r|, \quad (36)$$

such that $u_\alpha^{\otimes n} = \text{vec}(F)$, where u_α is again the pure state $\alpha |00\rangle + \sqrt{1 - \alpha^2} |11\rangle$ shared by Alice and Bob at the beginning of a single repetition of the protocol.

From our construction the unitary U that Bob applies in Lemma 5 to his portion of the entangled state $u_\alpha^{\otimes n}$ is

$$U = (-1)^n |0\rangle \langle 0| - |1\rangle \langle 1| + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} (-1)^{n+i} k_r |r\rangle \langle r|. \quad (37)$$

The state that Alice holds before measurement is then $(U \otimes \mathcal{I}_{\mathcal{Z}_{1\dots n}}) u_\alpha^{\otimes n}$. We analyze how successful the application of this channel would be to avoid $(\sin(\theta) |00\rangle - \cos(\theta) |11\rangle)^{\otimes n}$. Upon explicit computation of the formula $\langle \text{vec}(D), (U \otimes \mathcal{I}_{\mathcal{Z}_{1\dots n}}) \text{vec}(F) \rangle$, and using repeatedly the fact that $\text{vec}(V) = (V \otimes \mathcal{I}) \text{vec}(\mathcal{I})$, we obtain $\langle \text{vec}(D), \text{vec}(UF) \rangle$, which is equal to $\langle D, UF \rangle$ by the properties of the vec operator, resulting in the following expression:

$$\begin{aligned}
\langle D, UF \rangle &= \text{Tr} \left((-1)^n \alpha^n \lambda^n |0^n\rangle \langle 0^n| + (1 - \alpha^2)^{n/2} (-1)^{n+1} |1^n\rangle \langle 1^n| \right. \\
&\quad \left. + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} (-1)^n k_r (1 - \alpha^2)^{i/2} \alpha^{n-i} \lambda^{n-i} |r\rangle \langle r| \right) \\
&= (-1)^n \alpha^n \text{Tr} \left(\lambda^n |0^n\rangle \langle 0^n| - \left(\sqrt{\frac{1}{\alpha^2} - 1} \right)^n |1^n\rangle \langle 1^n| \right. \\
&\quad \left. + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} k_r \left(\sqrt{\frac{1}{\alpha^2} - 1} \right)^i \lambda^{n-i} |r\rangle \langle r| \right) \\
&= (-1)^n \alpha^n \left(\lambda^n - \left(\sqrt{\frac{1}{\alpha^2} - 1} \right)^n + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} k_r \left(\sqrt{\frac{1}{\alpha^2} - 1} \right)^i \lambda^{n-i} \right) \\
&= (-1)^n \alpha^n \left(\sqrt{\frac{1}{\alpha^2} - 1} \right)^n \left(\lambda_\alpha^n - 1 + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} k_r \lambda_\alpha^{n-i} \right), \tag{38}
\end{aligned}$$

where $\lambda_\alpha = \lambda \cdot \left(\sqrt{\frac{1}{\alpha^2} - 1} \right)^{-1}$.

Note that for the range of θ we are considering, it holds that $2^{1/n} - 1 \leq \lambda_\alpha \leq \frac{1}{2^{1/n} - 1}$.

Note as well that from our choice of k_r , for all i we have that $\text{Im} \left(\sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} k_r \lambda_\alpha^{n-i} \right) = 0$, and therefore the imaginary part of (38) is equal to 0. It then suffices to prove that for any choice of λ_α and n , there exists an $s_{\theta, \alpha, n} \in [-1, 1]$ such that, when plugged into the definition of k_r in the statement of Lemma 5 we have

$$\lambda_\alpha^n - 1 + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \text{Re}(k_r) \lambda_\alpha^{n-i} = 0. \tag{39}$$

Now, as the left hand side of (39) is an affine function of $s_{\theta, \alpha, n}$ with a positive linear coefficient, to prove the existence of such an $s_{\theta, \alpha, n}$, it suffices to prove that the left hand side of (39) ≤ 0 when $s_{\theta, \alpha, n} = -1$, and that the left hand side of (39) ≥ 0 when $s_{\theta, \alpha, n} = 1$.

We look first into the case when $s = -1$. Then, when $1 \leq \lambda_\alpha \leq \frac{1}{2^{1/n} - 1}$ it holds that:

$$\begin{aligned}
\lambda_\alpha^n - 1 + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \operatorname{Re}(k_r) \lambda_\alpha^{n-i} &= \lambda_\alpha^n - 1 - \sum_{i=1}^{n-1} \binom{n}{n-i} \lambda_\alpha^{n-i} \\
&= 2\lambda_\alpha^n - \lambda_\alpha^n - 1 - \sum_{i=1}^{n-1} \binom{n}{n-i} \lambda_\alpha^{n-i} \\
&= 2\lambda_\alpha^n - (1 + \lambda_\alpha)^n, \tag{40}
\end{aligned}$$

which is ≤ 0 whenever $\lambda_\alpha \leq \frac{1}{2^{1/n} - 1}$. When $2^{1/n} - 1 \leq \lambda_\alpha < 1$, that the left hand side of (39) ≤ 0 follows from two simple facts. First, the fact that $\lambda_\alpha^n < 1$, so $\lambda_\alpha^n - 1 < 0$. Second, the fact that for each $\sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \operatorname{Re}(k_r) \lambda_\alpha^{n-i}$ term, $\sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \operatorname{Re}(k_r) \leq -\binom{n}{i} + 1 \leq 0$.

We look now into the case when $s = 1$. Then, when $2^{1/n} - 1 \leq \lambda_\alpha < 1$ it holds that:

$$\begin{aligned}
\lambda_\alpha^n - 1 + \sum_{i=1}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \operatorname{Re}(k_r) \lambda_\alpha^{n-i} &= \lambda_\alpha^n - 1 + \sum_{i=1}^{n-1} \binom{n}{n-i} \lambda_\alpha^{n-i} \\
&= -2 + \lambda_\alpha^n + 1 + \sum_{i=1}^{n-1} \binom{n}{n-i} \lambda_\alpha^{n-i} \\
&= -2 + (1 + \lambda_\alpha)^n, \tag{41}
\end{aligned}$$

which is ≥ 0 whenever $\lambda_\alpha \geq 2^{1/n} - 1$. When $1 \leq \lambda_\alpha \leq \frac{1}{2^{1/n} - 1}$, that the left hand side of (39) ≥ 0 follows from two simple facts. First, the fact that $\lambda_\alpha^n \geq 1$. Second, the fact that for each $\sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \operatorname{Re}(k_r) \lambda_\alpha^{n-i}$ term, it is the case that $\sum_{\substack{r \in \{0,1\}^n \\ |r|=i}} \operatorname{Re}(k_r) \geq \binom{n}{i} - 1$. □

A.4 Derivation for Lemma 6

Proof. We will consider here the case where $\theta < \theta_{n,\alpha}$. The other case proceeds similarly.

Remember first that we characterized the chance of achieve 1-out-of- n hedging by the following SDP program in Section 3:

$$\begin{aligned}
&\underline{m_{n,\alpha,\theta}: \text{ Primal problem}} \\
\text{minimize:} & \quad \left\langle Q_{0,\alpha,\theta}^{\otimes n}, X \right\rangle \\
\text{subject to:} & \quad \operatorname{Tr}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n}(X) = \mathcal{I}_{\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_n}, \\
& \quad X \in \operatorname{Pos}(X_1 \otimes \mathcal{Y}_1 \otimes \dots \otimes \mathcal{X}_n \otimes \mathcal{Y}_n). \tag{42}
\end{aligned}$$

$m_{n,\alpha,\theta}$: Dual problem

$$\begin{aligned}
& \text{maximize: } \text{Tr}(Y) \\
& \text{subject to: } \pi(\mathcal{I}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n} \otimes Y) \pi^* \leq Q_{0,\alpha,\theta}^{\otimes n}, \\
& Y \in \text{Herm}(\mathcal{X}_1 \otimes \dots \otimes \mathcal{X}_n).
\end{aligned} \tag{43}$$

Then, to prove that perfect hedging is not possible when $\theta < \theta_{n,\alpha}$, we prove the feasibility in the dual SDP (??) of an operator Y with positive objective value. This operator is obtained from applying complementary slackness conditions to the primal solution corresponding to Λ_n . Therefore, it has value for the dual equal to the value in the primal SDP (42) for the solution corresponding to Λ_n . By weak duality, its feasibility proves then the optimality of Λ_n when $\theta < \theta_{n,\alpha}$.

To prove the feasibility of Y , we will express $Q_{0,\alpha,\theta}^{\otimes n} - \pi(\mathcal{I}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n} \otimes Y) \pi^*$ as a direct sum of smaller matrices. This reduces the question about feasibility of Y to a question about the positive-semidefiniteness of these smaller matrices. Each of these smaller matrices will have all proper leading principal minors be positive semi-definite, so by Sylvester's criterion it will suffice to check that their determinant is non-negative. We will then obtain a closed formula for these determinants, and prove that they are indeed non-negative.

We will first consider the case with $\alpha = 1/\sqrt{2}$, and then give an overview of the small changes involved in adapting the proof to other values of α . To simplify our argument, we will incur in a bit of notation abuse in this section, and omit the permutation operators in the definition of the dual SDP (43) that remind us that matrices at the sides of a \leq inequality must have their entries reordered to make the spaces on which they are defined be in the same order at both sides of the inequality.

A.4.1 Study of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$

$Q_{0,\alpha,\theta} \in \text{Pos}(\mathcal{X} \otimes \mathcal{Y})$ is given by $|\psi_0^1\rangle\langle\psi_0^1| + |\psi_0^2\rangle\langle\psi_0^2| + |\psi_0^3\rangle\langle\psi_0^3|$, where the $|\psi_0^i\rangle$ are defined as

$$\begin{aligned}
|\psi_0^1\rangle &= \alpha \sin(\theta) |00\rangle - \sqrt{1 - \alpha^2} \cos(\theta) |11\rangle, \\
|\psi_0^2\rangle &= \alpha |01\rangle, \\
|\psi_0^3\rangle &= \sqrt{1 - \alpha^2} |10\rangle.
\end{aligned} \tag{44}$$

This follows from considering the definition of $P_{0,\theta}$ given in Section 3, and observing that the operator Ψ_{ρ_α} satisfying $J(\Psi_{\rho_\alpha}) = \overline{u_\alpha} u_\alpha^*$ (with $u_\alpha = \alpha |00\rangle + \sqrt{1 - \alpha^2} |11\rangle$ the initial state shared between Alice and Bob) maps a state $\sigma \in \text{D}(\mathcal{Z})$ to $(\alpha|0\rangle\langle 0| + \sqrt{1 - \alpha^2}|1\rangle\langle 1|)\sigma(\alpha|0\rangle\langle 0| + \sqrt{1 - \alpha^2}|1\rangle\langle 1|)$. We can then write $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$ as

$$\begin{aligned}
Q_{0,1/\sqrt{2},\theta}^{\otimes n} &= \left(\frac{1}{2}\right)^n \left((\sin(\theta) |00\rangle - \cos(\theta) |11\rangle)(\sin(\theta) \langle 00| - \cos(\theta) \langle 11|) \right. \\
&\quad \left. + |01\rangle\langle 01| + |10\rangle\langle 10| \right)^{\otimes n} \tag{45} \\
&= \left(\frac{1}{2}\right)^n \sum_{a,b,c,d \in \{0,1\}^n} |a\rangle |b\rangle \langle c| \langle d| \prod_{i=0}^{n-1} \left(\delta_{c_i,1-d_i} \delta_{a_i,c_i} \delta_{b_i,d_i} \right. \\
&\quad \left. + \delta_{a_i,b_i} \delta_{c_i,d_i} \left(\delta_{a_i,1-c_i} (-\sin(\theta) \cos(\theta)) + \delta_{a_i,c_i} \delta_{a_i,1} \cos(\theta)^2 \right) \right. \\
&\quad \left. + \delta_{a_i,c_i} \delta_{a_i,0} \sin(\theta)^2 \right) \\
&= \left(\frac{1}{2}\right)^n \sum_{a,c \in \{0,1\}^n} |a\rangle \langle c| \otimes \sum_{b,d \in \{0,1\}^n} |b\rangle \langle d| \prod_{i=0}^{n-1} \left(\delta_{a_i,1-b_i} \delta_{c_i,1-d_i} \delta_{a_i,c_i} + \right. \\
&\quad \left. \delta_{a_i,b_i} \delta_{c_i,d_i} \left(\delta_{a_i,1-c_i} (-\sin(\theta) \cos(\theta)) + \delta_{a_i,c_i} \delta_{a_i,1} \cos(\theta)^2 \right) \right. \\
&\quad \left. + \delta_{a_i,c_i} \delta_{a_i,0} \sin(\theta)^2 \right). \tag{46}
\end{aligned}$$

The key insight to go ahead with the proof is to notice that this matrix can be written as a direct sum of 3^n smaller matrices. Indeed, observe that (45) can be equivalently written as

$$\frac{1}{2^n} \sum_{w \in \{0,1,2\}^n} \bigotimes_{i=0}^{n-1} |\psi_{w_i}\rangle \langle \psi_{w_i}|, \text{ where } |\psi_{w_i}\rangle = \begin{cases} \sin(\theta) |00\rangle - \cos(\theta) |11\rangle, & \text{if } w_i = 0 \\ |01\rangle, & \text{if } w_i = 1 \\ |10\rangle, & \text{if } w_i = 2 \end{cases}. \tag{47}$$

Then, the coefficient for each $|a\rangle \langle c| \otimes |b\rangle \langle d|$ term in the summation in (46) will receive contribution from at most one of the elements in (47). This element will be the one with

$$w_i = \begin{cases} 0 & \text{if } a_i = b_i \\ 1 & \text{if } (a_i, b_i) = (0, 1) \\ 2 & \text{if } (b_i, a_i) = (1, 0) \end{cases}.$$

Since this only depends on $|ab\rangle$, all elements on the same row of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$ come from the same term in (47). As each row of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$ has at least one non-zero term, (47) implies then a decomposition $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$ into a direct sum of smaller matrices, each of them with rank 1.

We can then identify each of these matrices by the corresponding choice of w in (47). We will do so by writing them as $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w)$. We denote the number of 0s, 1s and 2s in w by $n_0(w)$, $n_1(w)$ and $n_2(w)$, respectively. Also, note that there will be 3^n matrices in our decomposition, with the dimension of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w)$ being given by $2^{n_0(w)}$. Also, note that the number of matrices of size 2^k is given by $\binom{n}{k} 2^{n-k}$. This corresponds to choosing on which k positions $w_i = 0$, and what is the value of w_i for the other ones.

It will be convenient later to have a formula for the restriction to the diagonal of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w)$. Using the description in (47), we have that it is given by

$$\left(\frac{1}{2}\right)^n \sum_{w' \in M_w \subseteq \{0,1\}^n} g(w, w') |w'\rangle |f(w, w')\rangle \langle w'| \langle f(w, w')| \quad (48)$$

where M_w is given by the cartesian product $\times_{i=0}^{n-1} M_{w_i}$, with $\begin{cases} M_0 = \{0, 1\} \\ M_1 = \{0\} \\ M_2 = \{1\} \end{cases}$,

$$g(w, w') = \prod_{i=0}^{n-1} g(w_i, w'_i) \text{ with } \begin{cases} g(0, 0) = \sin^2(\theta) \\ g(0, 1) = \cos^2(\theta) \\ g(1, 0) = 1 \\ g(2, 1) = 1 \end{cases}, f(w, w')_i = \begin{cases} w'_i & \text{if } w_i = 0 \\ 1 - w'_i & \text{if } w_i = 1 \end{cases}.$$

Note that by definition of M_w , it is not necessary to define $g(w_i, w'_i)$ for values of (w_i, w'_i) not included in our definition of g .

A.4.2 Study of our candidate for Y in the $\alpha = 1/\sqrt{2}$ case

We define now our candidate solution Y for the dual problem, given by

$$Y = -\epsilon \left(\left(\frac{1}{\sqrt{2}} \sin(\theta) |0\rangle \langle 0| + \frac{1}{\sqrt{2}} \cos(\theta) |1\rangle \langle 1| \right)^{\otimes n} - 2 \left(\frac{1}{\sqrt{2}} \cos(\theta) |1\rangle \langle 1| \right)^{\otimes n} \right), \quad (49)$$

where ϵ is a value > 0 given by $\left(\frac{1}{2}\right)^{n/2} (2 \cos(\theta)^n - (\cos(\theta) + \sin(\theta))^n)$. Note that the definition of $\theta_{n,1/\sqrt{2}}$ implies that this value is positive indeed for $\theta < \theta_{n,1/\sqrt{2}}$. We can then write Y as

$$\sum_{a \in \{0,1\}^n} \lambda_a |a\rangle \langle a|, \text{ where } \lambda_a = \begin{cases} -\epsilon \left(\frac{1}{2}\right)^{n/2} \sin(\theta)^{n-|a|} \cos(\theta)^{|a|} & \text{for } a \neq 1^n \\ \epsilon \left(\frac{1}{2}\right)^{n/2} \cos(\theta)^n & \text{for } a = 1^n \end{cases} \quad (50)$$

Note that its trace (i.e., its value for the dual program) is given by

$$-\left(\frac{1}{2}\right)^{n/2} \epsilon \left((\sin(\theta) + \cos(\theta))^n - 2 \cos(\theta)^n \right), \quad (51)$$

which will again be positive for $\theta < \theta_{n,1/\sqrt{2}}$ by definition of $\theta_{n,1/\sqrt{2}}$.

This Y has been obtained from the strategy Λ_n in Lemma 4, and its feasibility proves the optimality of Λ_n for $\theta < \theta_{n,1/\sqrt{2}}$. This is an example of complementary slackness behavior, and follows from an observation [Wat] that given a feasible solution X to the primal SDP (42), $\text{Tr}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n} (Q_{0,\alpha,\theta}^{\otimes n} X)$ is an operator with the same objective value for the dual SDP (43). Furthermore, $\text{Tr}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n} (Q_{0,\alpha,\theta}^{\otimes n} X)$ satisfies the feasibility constraints of the dual if and only if

X represents an optimal solution to the primal. Therefore, after we experimentally observed that Λ_n seemed to be optimal for $\theta < \theta_{n,\alpha}$ to obtain our proposed Y we computed the corresponding value of $\text{Tr}_{\mathcal{Y}_1 \otimes \dots \otimes \mathcal{Y}_n} (Q_{0,1/\sqrt{2},\theta}^{\otimes n} X)$. X is given in this computation by the primal solution that represents the channel for the unitary in Λ_n ,

$$X = \sum_{i,j \in \{0,1\}^n} |ii\rangle \langle jj| (-1)^{\wedge i + \oplus i + \wedge j + \oplus j}. \quad (52)$$

A.4.3 Feasibility of Y in the $\alpha = 1/\sqrt{2}$ case

We want to prove that Y is feasible - that is to say, $Q_{0,1/\sqrt{2},\theta}^{\otimes n} - Y \otimes \mathcal{I} \geq 0$. Since Y is diagonal, the direct sum decomposition of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$ corresponds to a direct sum decomposition of Y . Since positive semidefiniteness is preserved by the direct sum operator, it is then enough to prove that each of the $S_w = Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w) - (Y \otimes \mathcal{I})(w)$ matrices are positive semidefinite, where $(Y \otimes \mathcal{I})(w)$ denotes $Y \otimes \mathcal{I}$ restricted to the rows/columns of $Q_{0,1/\sqrt{2},\theta}^{\otimes n}$ assigned to $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w)$.

Consider first the largest of these matrices. This will be S_{0^n} , with size 2^n . Using (47), we have that it is given by

$$S_{0^n} = \left(\frac{1}{2}\right)^n \sum_{a,c \in \{0,1\}^n} |aa\rangle \langle cc| \left(\prod_{i=0}^{n-1} \left(\delta_{a_i,1-c_i} \cdot -\sin(\theta) \cos(\theta) + \delta_{a_i,c_i} \delta_{a_i,1} \cos(\theta)^2 + \delta_{a_i,c_i} \delta_{a_i,0} \sin(\theta)^2 \right) - 2^n \lambda_a \right).$$

For example, for $n = 2$, S_{00} is given by

$$\frac{1}{4} \begin{pmatrix} \sin(\theta)^4 - 4\lambda_{00} & -\sin(\theta)^3 \cos(\theta) & -\sin(\theta)^3 \cos(\theta) & \sin(\theta)^2 \cos(\theta)^2 \\ -\sin(\theta)^3 \cos(\theta) & \sin(\theta)^2 \cos(\theta)^2 - 4\lambda_{01} & \sin(\theta)^2 \cos(\theta)^2 & -\sin(\theta) \cos(\theta)^3 \\ -\sin(\theta)^3 \cos(\theta) & \sin(\theta)^2 \cos(\theta)^2 & \sin(\theta)^2 \cos(\theta)^2 - 4\lambda_{10} & -\sin(\theta) \cos(\theta)^3 \\ \sin(\theta)^2 \cos(\theta)^2 & -\sin(\theta) \cos(\theta)^3 & -\sin(\theta) \cos(\theta)^3 & \cos(\theta)^4 - 4\lambda_{11} \end{pmatrix}$$

Consider now that since $Q_{0,1/\sqrt{2},\theta}^{\otimes n} \geq 0$, and for $a \neq 1^n$, $\lambda_a < 0$, the first $2^n - 1$ principal minors of S_{0^n} are ≥ 0 . By Sylvester's criterion, to prove that $S_{0^n} \geq 0$, it suffices then to prove that $\det(S_{0^n}) \geq 0$. Note that $\det(S_{0^n})$ is a polynomial in ϵ . This polynomial has all the coefficients below the one for ϵ^{2^n-1} equal to 0. This is because $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(0^n)$ has rank 1 - therefore, each minor of it with at least two rows will have determinant equal to zero. Using this, and going through the determinant formula, we see that $\det(S_{0^n})$ is given by

$$\begin{aligned} & \left(\epsilon^{2^n-1} (-1)^{2^n-1} \sum_{a \in \{0,1\}^n} \left(\frac{1}{2}\right)^n \cos(\theta)^{2|a|} \sin(\theta)^{2(n-|a|)} \prod_{\substack{b \in \{0,1\}^n \\ b \neq a}} \frac{\lambda_b}{\epsilon} \right) \\ & + \left(\epsilon^{2^n} (-1)^{2^n} \prod_{a \in \{0,1\}^n} \frac{\lambda_a}{\epsilon} \right) \end{aligned} \quad (53)$$

$$= \epsilon^{2^n-1} \left(\epsilon - \sum_{a \in \{0,1\}^n} \frac{\left(\frac{1}{2}\right)^n \cos(\theta)^{2|a|} \sin(\theta)^{2(n-|a|)}}{\lambda_a/\epsilon} \right) \prod_{a \in \{0,1\}^n} \frac{\lambda_a}{\epsilon} \quad (54)$$

$$= \epsilon^{2^n-1} \left(\epsilon + \sum_{a \in \{0,1\}^n} \left(\frac{1}{2}\right)^{n/2} \cos(\theta)^{|a|} \sin(\theta)^{n-|a|} - 2 \left(\frac{1}{2}\right)^{n/2} \cos(\theta)^n \right) \prod_{a \in \{0,1\}^n} \frac{\lambda_a}{\epsilon} \quad (55)$$

Since all of the λ_a/ϵ except the one for 1^n are negative, we have that $\epsilon^{2^n-1} \prod_{a \in \{0,1\}^n} \frac{\lambda_a}{\epsilon}$ is negative whenever $\epsilon > 0$. Therefore,

$$\det(S_{0^n, 0^n}) \geq 0 \iff \quad (56)$$

$$\epsilon + \sum_{a \in \{0,1\}^n} \left(\frac{1}{2}\right)^{n/2} \cos(\theta)^{|a|} \sin(\theta)^{n-|a|} - 2 \left(\frac{1}{2}\right)^{n/2} \cos(\theta)^n \leq 0 \iff \quad (57)$$

$$\epsilon \leq \left(\frac{1}{2}\right)^{n/2} (2(\cos(\theta))^n - (\cos(\theta) + \sin(\theta))^n), \quad (58)$$

which is true by definition of ϵ . We have then that our proposed feasible solution Y produces a positive-semidefinite S_{0^n} . To verify the feasibility of Y , it remains to prove the positive-semidefiniteness of the rest of the S_w .

To do so, consider an arbitrary S_w , $w \in \{0, 1, 2\}^n - \{0^n\}$, with a corresponding M_w , as defined in (48). Note that M_w is the set of indices such that λ_i appears in the diagonal of S_w , and that each λ_i appears in the diagonal of S_w at most once, as we can see from the expression in (48). If $1^n \notin M_w$, then S_w is trivially positive-semidefinite, since it is obtained by adding a positive-semidefinite diagonal matrix $Y(w)$ to a positive-semidefinite matrix $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w)$. Otherwise, our appeal to Sylvester's criterion from the 0^n case applies again, and it is enough to prove that $\det(S_w) \geq 0$. Also, since $Q_{0,1/\sqrt{2},\theta}^{\otimes n}(w)$ has rank 1, our argument that $\det(S_w)$ is a polynomial of minimum degree $|M_w| - 1$ applies again.

Then, using (48), we have that $\det(S_w)$ is given by

$$\epsilon^{|M_w|-1} \left(\prod_{c \in M_w} \frac{\lambda_c}{\epsilon} \right) \left(\epsilon - \left(\frac{1}{2}\right)^n \sum_{d \in M_w} \frac{g(w, d)}{\lambda_d/\epsilon} \right) \quad (59)$$

Using the recursive definition of M_w in (48), and realizing that $1^n \in M_w$ implies that $n_1(w) = 0$, we have that

$$\sum_{d \in M_w} \frac{g(w, d)}{|\lambda_d/\epsilon|} = \left(\frac{1}{2}\right)^{n/2} (\sin(\theta) + \cos(\theta))^{n_0(w)} \left(\frac{1}{\cos(\theta)}\right)^{n_2(w)}. \quad (60)$$

Now, we have that

$$\frac{1}{\cos(\theta)} \leq \sin(\theta) + \cos(\theta) \iff \frac{1}{\cos(\theta)^2} \leq \tan(\theta) + 1 \quad (61)$$

$$\iff \tan(\theta)^2 \leq \tan(\theta) \iff \theta \leq \pi/4 \quad (62)$$

Since we are looking at the range $\theta < \theta_{n,1/\sqrt{2}} \leq \pi/4$, and $n_0(w) + n_2(w) = n$, we have that

$$(\sin(\theta) + \cos(\theta))^{n_0(w)} \left(\frac{1}{\cos(\theta)}\right)^{n_2(w)} \leq (\sin(\theta) + \cos(\theta))^n. \quad (63)$$

Therefore, since $n_2(w) \leq n$,

$$\left(\frac{1}{2}\right)^n \sum_{d \in M_w} \frac{g(w, d)}{\lambda_d/\epsilon} \geq \left(\frac{1}{2}\right)^{n/2} (2(\cos(\theta))^n - (\cos(\theta) + \sin(\theta))^n). \quad (64)$$

We see then that any ϵ that makes $\det(S_{0^n})$ non-negative will make the determinant of the other S_w non-negative as well.

A.4.4 Generalization to $\alpha \neq 1/\sqrt{2}$

For $\alpha \neq 1/\sqrt{2}$, the changes necessary to make the proof work are limited to arithmetic adjustments. $Q_{0,\alpha,\theta}^{\otimes n}$ will now be given by

$$\begin{aligned} & \sum_{a,c \in \{0,1\}^n} |a\rangle\langle c| \otimes \sum_{b,d \in \{0,1\}^n} |b\rangle\langle d| \prod_{i=0}^{n-1} \left(\delta_{a_i,1-b_i} \delta_{c_i,1-d_i} \delta_{a_i,c_i} \left(\delta_{a_i,1} (1-\alpha^2) + \delta_{a_i,0} \alpha^2 \right) \right. \\ & + \delta_{a_i,b_i} \delta_{c_i,d_i} \left(\delta_{a_i,1-c_i} \cdot -\alpha \sin(\theta) \sqrt{1-\alpha^2} \cos(\theta) + \delta_{a_i,c_i} \delta_{a_i,1} (1-\alpha^2) \cos(\theta)^2 \right. \\ & \left. \left. + \delta_{a_i,c_i} \delta_{a_i,0} \alpha^2 \sin(\theta)^2 \right) \right). \end{aligned} \quad (65)$$

Note that its direct sum decomposition is not affected, since the choice of which terms of $Q_{0,\alpha,\theta}^{\otimes n}$ appear on each term does not depend on α .

Similarly, Y is given now by

$$\begin{aligned} & \sum_{a \in \{0,1\}^n} \lambda_a |a\rangle\langle a|, \text{ where } \lambda_a = \begin{cases} -\epsilon (\alpha \sin(\theta))^{n-|a|} \left(\sqrt{1-\alpha^2} \cos(\theta) \right)^{|a|} & \text{for } a \neq 1^n \\ \epsilon \left(\sqrt{1-\alpha^2} \right)^n \cos(\theta)^n & \text{for } a = 1^n \end{cases} \\ & \text{and } \epsilon = 2 \left(\sqrt{1-\alpha^2} \cos(\theta) \right)^n - \left(\sqrt{1-\alpha^2} \cos(\theta) + \alpha \sin(\theta) \right)^n. \end{aligned} \quad (66)$$

As for the feasibility of Y , we have then that $\det(S_w)$ is given by

$$\epsilon^{|M_w|-1} \left(\prod_{c \in M_w} \frac{\lambda_c}{\epsilon} \right) \left(\epsilon - \sum_{d \in M_w} \frac{g(w, d) \alpha^{2(n-|d|)} (1 - \alpha^2)^{|d|}}{\lambda_d / \epsilon} \right), \quad (67)$$

again non-negative whenever

$$\begin{aligned} \epsilon &\leq \sum_{d \in M_w} \frac{g(w, d) \alpha^{2(n-|d|)} (1 - \alpha^2)^{|d|}}{\lambda_d / \epsilon} \\ &= 2 \left(\sqrt{1 - \alpha^2} \right)^n \cos(\theta)^{2n_0(w) - n} - \sum_{d \in M_w} \frac{g(w, d) \alpha^{2(n-|d|)} (1 - \alpha^2)^{|d|}}{|\lambda_d| / \epsilon}. \end{aligned} \quad (68)$$

Note that we have now that using the recursive definition of M_w in (48),

$$\begin{aligned} &\sum_{d \in M_w} \frac{g(w, d) \alpha^{2(n-|d|)} (1 - \alpha^2)^{|d|}}{|\lambda_d| / \epsilon} \\ &= (\alpha \sin(\theta) + \sqrt{1 - \alpha^2} \cos(\theta))^{n_0(w)} \left(\frac{\sqrt{1 - \alpha^2}}{\cos(\theta)} \right)^{n_2(w)}. \end{aligned}$$

To prove that (68) holds we will need an argument slightly more involved than the corresponding one for the $\alpha = \frac{1}{\sqrt{2}}$ case. First, we consider that for $n_0(w) = n$, the right hand side of (68) is equal to ϵ , by definition of ϵ . Then, we prove that the right hand side of (68) increases as we decrement $n_0(w)$, and increase $n_2(w) = n - n_0(w)$ in parallel. This is because the positive term in the right hand side increases with each decrease of $n_0(w)$, and it does so by a larger factor than the one by which the negative term decreases. More rigorously, consider the expression

$$k = \frac{1}{\cos(\theta)^2} - \frac{\sqrt{1 - \alpha^2}}{(\alpha \sin(\theta) + \sqrt{1 - \alpha^2} \cos(\theta)) \cos(\theta)}. \quad (69)$$

First, note that

$$k \geq 0 \iff \sqrt{1 - \alpha^2} \cos(\theta)^2 \leq (\alpha \sin(\theta) + \sqrt{1 - \alpha^2} \cos(\theta)) \cos(\theta) \quad (70)$$

$$\iff \cos(\theta) \leq \frac{\alpha}{\sqrt{1 - \alpha^2}} \sin(\theta) + \cos(\theta) \quad (71)$$

$$\iff 0 \leq \frac{\alpha}{\sqrt{1 - \alpha^2}} \sin(\theta), \quad (72)$$

which is always true when $0 \leq \theta \leq \pi/2$, which is always the case within the trigonometric domain that we consider. Then, if we denote the right hand side of (68) by $r_{n_0(w)}$, we have the recursive relation

$$r_{n_0(w)} = r_{n_0(w)+1} \frac{1}{\cos(\theta)^2} + k(\alpha \sin(\theta) + \sqrt{1 - \alpha^2} \cos(\theta))^{n_0(w)} \left(\frac{\sqrt{1 - \alpha^2}}{\cos(\theta)} \right)^{n - n_0(w)}$$

We can see indeed that this defines an increasing sequence as we decrease $n_0(w)$, since the second summand is positive, and the first summand multiplies the previous value of r by an amount greater than one. We have then successfully proved that (68) holds in the $\alpha \neq \frac{1}{\sqrt{2}}$ case. \square

A.5 Derivation for Theorem 7

Proof. For didactic purposes, we show our derivation along the line of thought used by us when obtaining it. Therefore, we first consider simple proofs for two particular cases, and then finish with a general proof.

A.5.1 Absence of hedging for the protocol in [MW12]

It is easy to establish that in a generalization of the example in [MW12], the hedging behavior *disappears* if Bob can avoid returning an answer. This generalization considers the set of protocols where the initial quantum state shared between Alice and Bob is a pure state ψ such that $\text{Tr}_{\mathcal{X}}(\psi\psi^*) = \mathcal{I}_{\mathcal{Z}}/\text{dim}(\mathcal{Z})$. It suffices to prove it for one of such states, as the other ones can be obtained from it by Bob applying a unitary. We prove it then for

$$\psi = \frac{1}{\sqrt{\text{dim}(\mathcal{X})}} \sum_i e_i \otimes e_i, \quad (73)$$

with e_i being the computational basis for \mathcal{X} , and corresponding to the case $\text{dim}(\mathcal{X}) = \text{dim}(\mathcal{Z})$.

The reason no hedging behavior is possible is because in this situation, it is always possible for Bob to make sure he obtains the desired outcome. To see this, notice that the operator that we apply to get Q_a from P_a is the identity divided by $\text{dim}(\mathcal{X})$. Similarly, $E = \mathcal{I}_{\mathcal{X} \otimes \mathcal{Y}}/\text{dim}(\mathcal{X})$. Therefore, $(E^+)^{1/2} Q_a (E^+)^{1/2} = P_a$. As this is a projector into a non-empty space (from the assumption that Bob has a nonzero probability of obtaining the desired outcome), the norm of this operator is 1.

A.5.2 Absence of hedging in the classical case

We look now at the behavior when a game is repeated twice in parallel, and the information exchanged between Alice and Bob is classical. This is reflected in the operators ρ and P_a we consider in our model being diagonal matrices. As ρ is a diagonal matrix, then Ψ_ρ maps diagonal matrices to diagonal matrices, so E and the Q_a are diagonal too. Then, if we denote by $\Omega(E)$ the matrix that has a one in a position whenever the corresponding entry of E is nonzero, and a zero otherwise, we have that

$$\|\Lambda_{1,2}\| = \left\| \Omega(E) \otimes \Omega(E) - \left((E^+)^{1/2} \otimes (E^+)^{1/2} \right) (Q_{1-a} \otimes Q_{1-a}) \left((E^+)^{1/2} \otimes (E^+)^{1/2} \right) \right\|.$$

Now, whenever $\Omega(E)$ has a zero entry, $(E^+)^{1/2}Q_{1-a}(E^+)^{1/2}$ has a zero entry as well in that position, as $Q_{1-a} \leq E$. We define now $\lambda_E(X)$ as the minimum entry of a diagonal matrix X , restricted to the positions where E has a nonzero entry. We have then that the value of the game when Bob is trying to win one out of two parallel repetitions is given by:

$$\begin{aligned} & 1 - \lambda_E \left((Q_{1-a} \otimes Q_{1-a}) \left((E^+)^{1/2} \otimes (E^+)^{1/2} \right) \left(Q_{1-a} \otimes Q_{1-a} \right) \right) \\ &= 1 - \lambda_E \left((E^+)^{1/2} Q_{1-a} (E^+)^{1/2} \right)^2. \end{aligned} \quad (74)$$

Since we have that

$$\begin{aligned} \Omega(E) &= (E^+)^{1/2} E (E^+)^{1/2} \\ &= (E^+)^{1/2} (Q_a + Q_{1-a}) (E^+)^{1/2} \\ &= (E^+)^{1/2} Q_{1-a} (E^+)^{1/2} + (E^+)^{1/2} Q_a (E^+)^{1/2} \end{aligned} \quad (75)$$

we have then that

$$\lambda_E \left((E^+)^{1/2} Q_{1-a} (E^+)^{1/2} \right)^2 = 1 - \|(E^+)^{1/2} Q_a (E^+)^{1/2}\| \quad (76)$$

so

$$1 - \lambda_E \left((E^+)^{1/2} Q_{1-a} (E^+)^{1/2} \right)^2 = \|(E^+)^{1/2} Q_a (E^+)^{1/2}\|. \quad (77)$$

Therefore, there is no hedging in this case. Our argument applies similarly to the case where Bob is trying to win k out of n repetitions.

A.5.3 Absence of hedging in the general case

We begin by defining the following operators:

$$A = \Lambda = (E^+)^{1/2} Q_a (E^+)^{1/2}, B = (E^+)^{1/2} E (E^+)^{1/2}. \quad (78)$$

Note that $[Q_a, (E^+)E] = 0$, as $(E^+)E$ is equal to the identity on the support of E and zero outside it, and $Q_a \leq E$, so $E^+ E Q_a = Q_a E^+ E = Q_a$. We have then that $[A, B] = 0$, so A and B are simultaneously diagonalizable. This means that any tensor products of A , B , and \mathcal{I} of the same dimension are simultaneously diagonalizable as well.

We consider first the case where $k = 1$ and $n = 2$, and then use a proof by induction to take care of larger n and k . Using the operators A and B , we can use the fact that $Q_{1-a} = E - Q_a$ to write $\|\Lambda_{1,2}\|$ in terms of A and B as

$$\left\| A \otimes B + B \otimes A - A \otimes A \right\| \leq \left\| A \otimes \mathcal{I} + \mathcal{I} \otimes A - A \otimes A \right\| = 2\|A\| - \|A\|^2, \quad (79)$$

where the inequality follows from the fact that $0 \leq B \leq \mathcal{I}$. The equality follows from considering a basis where A is diagonal, and using the fact that since $Q_a \leq E$, $0 \leq A \leq \mathcal{I}$, so all the eigenvalues of A are at most 1.

We have then that $\|\Lambda_{1,2}\| = 1 - (1 - \|A\|)^2$, since the fact that Bob can just choose to play independently implies $\|\Lambda_{1,2}\| \geq 1 - (1 - \|A\|)^2$. Therefore, we obtain that playing each game independently is an optimal behavior.

In the general case where Bob is trying to win k out of n games, we can again express Q_{1-a} as $E - Q_a$, and thus reduce $\Lambda_{k,n}$ to a sum of tensor products of A and B .

Consider first the case where $k = 1$. Then observe that we can write

$$\Lambda_{1,n} = \Lambda_{1,n-1} \otimes (B - A) + B^{\otimes n-1} \otimes A \leq \Lambda_{1,n-1} \otimes (\mathcal{I} - A) + \mathcal{I}^{\otimes n-1} \otimes A \quad (80)$$

Using as basis the n^{th} tensor product of a basis where A is diagonal, we obtain by induction on n that $\|\Lambda_{1,n}\| = 1 - (1 - \|A\|)^n$. This is because for diagonal positive semidefinite matrices $J \leq \mathcal{I}$ and K , we have $\|J(\mathcal{I} - K) + \mathcal{I} \cdot K\| = \|J\|(1 - \|K\|) + \|K\|$.

Note as well that if x is a largest eigenvalue eigenvector of Λ , a maximum-eigenvalue eigenvector of $\Lambda_{1,n}$ is given by $x^{\otimes n}$. Using this fact, we obtain a proof for the case with $k > 1$. To do this, observe that

$$\Lambda_{k,n} = \Lambda_{k,n-1} \otimes (B - A) + \Lambda_{k-1,n-1} \otimes A \leq \Lambda_{k,n-1} \otimes (\mathcal{I} - A) + \Lambda_{k-1,n-1} \otimes A \quad (81)$$

Then, using again as basis the n^{th} tensor product of a basis where A is diagonal, we obtain by induction that $\|\Lambda_{k,n}\| = 1 - \sum_{t=0}^{k-1} \binom{n}{t} \|A\|^t (1 - \|A\|)^{n-t}$, and that for all choices of k and n , a maximum-eigenvalue eigenvector of $\Lambda_{k,n}$ is given by $x^{\otimes n}$, for x a largest eigenvector of Λ . This is because for diagonal positive semidefinite matrices J, K, H , where J and H share a largest eigenvector, and $\|J\| \leq \|H\|$, we have $\|J(\mathcal{I} - K) + H \cdot K\| = \|J\|(1 - \|K\|) + \|H\|\|K\|$.

We obtain then that in this setting, no quantum advantage can be obtained by correlating Bob's strategy between parallel repetitions. □