Natural Computing manuscript No.
(will be inserted by the editor)

Classical, quantum and non-signalling resources in bipartite games

Gilles Brassard [1], Anne Broadbent [1], Esther Hänggi [2], André Allan Méthot [3], Stefan Wolf [2]

[1] Département d'informatique et de recherche opérationnelle
Université de Montréal
CP 6128, succursale centre-ville
Montréal (Québec), H3C 3J7 Canada
{brassard, broadbea}@iro.umontreal.ca

[2] ETH Zürich
8092 Zürich, Switzerland
{esther.haenggi, wolfst}@inf.ethz.ch

[3] Group of Applied Physics
Université de Genève
Rue de l'École-de-Médecine 20
1211 Genève, Switzerland
andre.methot@physics.unige.ch

4 December 2007

Abstract We study bipartite games that arise in the context of nonlocality with the help of graph theory. Our main results are alternate proofs that deciding whether a no-communication classical winning strategy exists for certain games (called forbidden-edge and covering games) is NP-complete, while the problem of deciding if these games admit a non-signalling winning strategy is in P. We discuss relations between quantum winning strategies and orthogonality graphs. We also show that every pseudo-telepathy game yields both a proof of the Bell-Kochen-Specker theorem and an instance of a two-prover interactive proof system that is classically sound, but that becomes unsound when provers use shared entanglement.

Keywords: Game Theory, Graph Theory, Nonlocality, Bell Theorems, Interactive Proof Systems

1 Introduction

There exist particular measurement scenarios on *entangled* particles that cannot be simulated within a gedanken world in which the particles have predefined outcomes to any measurement [5]. This phenomenon is nowadays called nonlocality. Theoretical proofs of this fact are usually set up in the following paradigm: the entangled particles (we shall restrict ourselves to the case of two particles since it is the most studied case and the subject of this paper) are measured according to a measurement chosen from a given set of measurements. The measurements are timed in such a way that it is impossible for either particle to send a signal that would influence the measurement outcome on the other. The probability distribution of the joint outcomes is then studied. The purpose is to show that no local realistic theory can reproduce this distribution.

Bipartite games are of particular interest in the study of quantum nonlocality. We view the particles as *players*, the measurements as *questions* and the outcomes as *answers*. A proof of nonlocality is then nothing more than showing that quantum players—players that have access to quantum information—can fare better than classical players, who do not have access to this resource.

Recently, the community has studied the amount of resources one needs to give to classical players in order to have them on par with quantum players [10, 15, 21, 34, 39, 45]. The purpose is to help characterize the power of entanglement. This line of thinking has led to many interesting results, such as "if quantum mechanics were too nonlocal, communication complexity would collapse to a single bit for any distributed Boolean function" [19, 9] and "entanglement and nonlocality are incomparable resources" [12]. However, the general question of whether a given quantum probability distribution can be simulated by classical means in different scenarios remains open.

In this paper, we introduce a novel approach to the study of nonlocality: graph theory. Thanks to our new general framework for bipartite games, we give new solutions to known results and also provide some new contributions. Our work paves the way for future research in this general direction. Previous connections between graph theory and nonlocality were established in [14, 24].

We investigate bipartite games and study the cases in which the participants are allowed 1) two-way communication, 2) one-way communication, 3) just local resources, 4) non-signalling resources, or 5) quantum resources. We also establish links between these games and the Bell-Kochen-Specker theorem [6, 30], as well as with interactive proof systems [3, 25].

In particular, we give alternate proofs that deciding whether a particular game is winnable by classical players is NP-complete and deciding whether a particular game is winnable by players sharing non-signalling resources is in P. The first result was originally established by Uriel Feige and László Lovász [20], while the second by Daniel Preda [43] and Ben Toner [44].

In Section 2, we formalize what we mean by bipartite games and introduce the graph theory paradigm. We then study the different types of resources one can give the players in Section 3. The links with the Bell-Kochen-Specker theorem and with interactive proof systems are covered in Sections 4 and 5, respectively.

2 Bipartite Games

A *bipartite game* $\mathbb{G} = (X, Y, W)$ is a set of *inputs* $X = X_A \times X_B$, a set of *outputs* $Y = Y_A \times Y_B$ and a relation $W \subseteq X \times Y$. The relation W is called the *winning condition*. We explain how such a game is played in Section 2.1. Given a bipartite game $\mathbb{G}$, we represent it as a bipartite graph $G_{\mathbb{G}} = (V, E)$ with $A = X_A \times Y_A$ and $B = X_B \times Y_B$ being the classes of the bipartition, meaning that V is the *disjoint* union of A and B. There is an edge between $x_A y_A \in A$ and $x_B y_B \in B$ if and only if $((x_A, x_B), (y_A, y_B)) \in W$. To distinguish vertices coming from class A from those coming from class B, such an edge will be denoted $(x_A y_A, x_B y_B)$ even

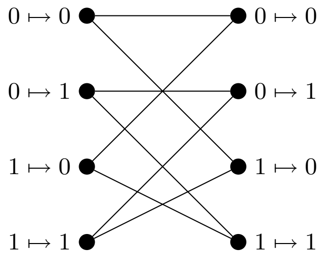

Fig. 1 $G_{\mathbb{G}_1}$

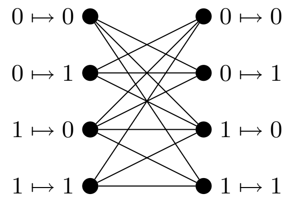

Fig. 2 $G_{\mathbb{G}_2}$

though, formally, edges in the graph are not oriented. This is our graph-theoretical representation for bipartite games. We now give two examples.

Example 1 Game $\mathbb{G}_1$ is given by (X, Y, W) where $X_A = X_B = Y_A = Y_B = \{0, 1\}$ with

$$((x_A, x_B), (y_A, y_B)) \in W \Leftrightarrow y_A \oplus y_B = x_A \wedge x_B \,.$$

The corresponding graph, $G_{\mathbb{G}_1}$, is given in Figure 1. Note that in order to emphasize the structure of the graph in relation to the game, we have labelled the vertices according to the form $x_A \mapsto y_A$ and $x_B \mapsto y_B$. Similarly, edges such as $(x_A y_A, x_B y_B)$ are sometimes denoted $(x_A \mapsto y_A, x_B \mapsto y_B)$. We shall refer to game $\mathbb{G}_1$ in Section 3.5, as it is closely related to the *nonlocal* box [37], inspired by the CHSH game [13, 17].

Example 2 Game $\mathbb{G}_2$ is given by (X, Y, W) where $X_A = X_B = Y_A = Y_B = \{0, 1\}$ with the following edges:

$$(0 \mapsto 0, 0 \mapsto 1), (0 \mapsto 0, 1 \mapsto 0), (0 \mapsto 0, 1 \mapsto 1), (0 \mapsto 1, 0 \mapsto 0), (0 \mapsto 1, 0 \mapsto 1),$$
$$(0 \mapsto 1, 1 \mapsto 0), (1 \mapsto 0, 0 \mapsto 0), (1 \mapsto 0, 0 \mapsto 1), (1 \mapsto 0, 1 \mapsto 0), (1 \mapsto 0, 1 \mapsto 1),$$
$$(1 \mapsto 1, 0 \mapsto 0), (1 \mapsto 1, 1 \mapsto 0), (1 \mapsto 1, 1 \mapsto 1)\,.$$

The corresponding graph, $G_{\mathbb{G}_2}$, is given in Figure 2. The game $\mathbb{G}_2$ is closely linked to the Hardy game [27], which we shall discuss in Section 3.5.

Given a bipartite game $\mathbb{G}$ and its corresponding graph $G_{\mathbb{G}}$, there is a natural partition of each class A and B of the bipartition, which is induced by fixing an element of either X_A or X_B. We refer to these as *Alice's natural partition* (or the *natural partition* of A),

$$P_A = \{\{x_A y_A \mid y_A \in Y_A\} \mid x_A \in X_A\}$$

and *Bob's natural partition* (or the *natural partition* of B),

$$P_B = \{\{x_B y_B \mid y_B \in Y_B\} \mid x_B \in X_B\}\,.$$

For instance, the game of Example 1 (Figure 1) has P_A containing vertices in class A: $P_A = \{\{0 \mapsto 0, 0 \mapsto 1\}, \{1 \mapsto 0, 1 \mapsto 1\}\}$ and P_B containing vertices in class B: $P_B = \{\{0 \mapsto 0, 0 \mapsto 1\}, \{1 \mapsto 0, 1 \mapsto 1\}\}$.
(Note: Despite appearances, P_A and P_B are distinct in this example because recall that the set V of vertices is the *disjoint* union of classes A and B.)

2.1 Bipartite games as cooperative games

We study bipartite games as cooperative games in two scenarios: the *forbidden-edge* games and the *covering* games. In each *round* of game $\mathbb{G}$, Alice and Bob are individually presented with a *question*, $x_A \in X_A$ for Alice and $x_B \in X_B$ for Bob. They must produce an *answer* chosen in $y_A \in Y_A$ for Alice and $y_B \in Y_B$ for Bob. Alice and Bob *win* this round of $\mathbb{G}$ if and only if $(x_A y_A, x_B y_B) \in E$ (in which case, we say that the edge E is *covered*). Whether or not Alice and Bob have a winning strategy for a game depends on the type of game they are playing:

Definition 1 *In a* forbidden-edge *game, a winning strategy for Alice and Bob is such that they win each round.*

The case of a covering game is more complicated: to each covering game $\mathbb{G}$, we associate a probability $p(\mathbb{G})$, which, intuitively, is used to formalize the fact that each possible answer must be given, in turn, with probability at least $p(\mathbb{G})$:

Definition 2 *In a* covering *game, a winning strategy for Alice and Bob is such that they win each round and for a fixed round, each edge is covered with probability at least* $p(\mathbb{G})$.

Players are allowed resources: in all cases, at the onset of the game, they can discuss a common strategy and flip an unlimited number of coins. If these are the only allowed resources, we say that the strategy is *classical*. In some cases, we also allow the players to establish shared quantum entanglement. During the execution of the game, we may also allow communication or the use of non-signalling probability distributions. A forbidden-edge game is called a *pseudo-telepathy* game if

Alice and Bob have a winning strategy using shared entanglement, yet no such classical strategy exists, while a covering game with the same features is called a *Bell theorem without inequalities* (BTWI), a term coined in [26]. For a discussion on the differences between these types of games, see [35].

Bell's theorem [5], which Henry Stapp designated "the most profound discoveries of science" [42] states that quantum mechanics is not a local realistic theory. There is a direct connection between Bell's theorem, pseudo-telepathy and BTWI due to the fact that any such game is a proof of Bell's theorem. This is easily seen by the fact that Alice and Bob (who are unable to communicate, thus are restricted to act in a local realistic world to anyone who doesn't believe in quantum mechanics) have a quantum winning strategy, yet they do not have a classical winning strategy.

In the next section, we study winning strategies according to the available shared resources. We then make connections between pseudo-telepathy games, the Bell-Kochen-Specker theorem (Section 4) and multi-prover interactive proofs (Section 5).

3 Bipartite games and resources

In this section, we give necessary and sufficient conditions for forbidden-edge as well as covering games to exhibit a winning strategy depending on the available resources. We also give a necessary and sufficient condition for a game to exhibit a winning strategy, regardless of the resources available to Alice and Bob. We start with the latter.

3.1 Winnable games

A certain class of bipartite games is rather uninteresting for our purposes; these are the games that *do not* have a winning strategy, no matter the resources that Alice and Bob share (even with unlimited communication!). Intuitively, a forbidden-edge game or a covering-game has a winning strategy (with unlimited resources) if and only if there is a way to win each round (and also, for a covering game, each possible answer must be given with a minimum probability). A game that has a winning strategy with unlimited resources is called *winnable*.

Theorem 1 *Let $\mathbb{G}$ be a bipartite game (either a forbidden-edge game or a covering game) with bipartite graph $G_{\mathbb{G}} = (V, E)$, whose classes are A and B, and let P_A be the natural partition of A and P_B be the natural partition of B. Then $\mathbb{G}$ is winnable if and only if each subgraph induced by an element of P_A and an element of P_B has at least one edge. In addition, in the case of a covering game, the number of edges in the induced subgraph must be at most $1/p(\mathbb{G})$.*

Proof If $\mathbb{G}$ is a forbidden-edge game, then it is winnable if and only if Alice and Bob (who have access to unlimited resources) can win every round. But this is possible if and only if there is at least one answer for each possible question that causes Alice and Bob to win. This is what is formally stated in the lemma. If $\mathbb{G}$ is a covering game, then in each round, each possible edge must be covered with probability at least $p(\mathbb{G})$, which is possible if and only if there are at most $1/p(\mathbb{G})$ edges to cover. □

Theorem 2 *The problem of deciding if a game (forbidden-edge or covering) is winnable is in* P.

Proof As stated in Theorem 1, at most, we need to count the number of edges in each bipartite graph induced by a pair of elements, one in P_A and one in P_B. This can be done in $O(n^3)$, where n is the number of vertices. □

3.2 Two-way communication

The first resource that one probably thinks of is communication. What type of game can players win if they are allowed to communicate? If two-way communication is allowed, the results are simple.

Theorem 3 *Let* $\mathbb{G}$ *be a bipartite game (either a forbidden-edge game or a covering game) with bipartite graph* $G_{\mathbb{G}} = (V, E)$*, whose classes are* A *and* B*, and let* P_A *be the natural partition of* A *and* P_B *be the natural partition of* B*. Then* $\mathbb{G}$ *is winnable with two-way communication if and only if it is a winnable game.*

Proof The strategy for a winnable game is easy. Alice and Bob discuss which questions they receive and jointly decide which edge they want to cover. The other direction of the proof is even simpler. If a game is not winnable, then it is of course not winnable with two-way communication, since winnable has been defined independently of resources. □

Corollary 1 *The problem of deciding if a game (forbidden-edge or covering) is winnable with two-way communication is in* P.

3.3 One-way communication

A more interesting scenario is to allow communication, but to restrict it to being one-way only.

Theorem 4 *Let $\mathbb{G}$ be a forbidden-edge game with bipartite graph $G_{\mathbb{G}} = (V, E)$, whose classes are A and B, and let P_A be the natural partition of A and P_B be the natural partition of B. Then $\mathbb{G}$ is winnable with one way communication from Alice to Bob (the case of one way communication from Bob to Alice is similar) if and only if the following is possible: for each element of P_A, it must be possible to choose a vertex $v \in P_A$ and a subset S of B containing exactly one element of each element of P_B such that the subgraph induced by $\{v\} \cup S$ is a complete bipartite graph. (Said otherwise, there is an edge $(v, w) \in E$ for each $w \in S$.)*

Proof The strategy is simple: Alice tells Bob which question she received and which answer she gave (the answer corresponding to v in our case). Bob can then always choose an allowed output (the answer corresponding to the appropriate element of S in our case), since Alice's choice was made for precisely that. To complete the proof, one only has to realize that if no such construction exists, then Alice's answer must depend on Bob's question. Therefore, making a one-way communication scheme from Alice to Bob impossible. □

Theorem 5 *Let $\mathbb{G}$ be a covering bipartite game with bipartite graph $G_{\mathbb{G}} = (V, E)$, whose classes are A and B, and let P_A be the natural partition of A and P_B be the natural partition of B. Then $\mathbb{G}$ is winnable with one way communication if and only if it is winnable as a forbidden-edge game, all edges of $G_{\mathbb{G}}$ are covered by* *at least one induced bipartite graph as in Theorem 4, and every induced subgraph given by an element of P_A and an element of P_B has at most $1/p(\mathbb{G})$ edges.*

Proof Alice just chooses at random amongst the vertices of her partition that have degree at least 1 and tells Bob which one she has chosen. Bob then selects one adjacent vertex at random. This strategy spans the whole graph and ensures that each edge is covered with probability at least $p(\mathbb{G})$. If a vertex on Alice's side doesn't have the requirements stated in the Theorem, she cannot select it since Bob could receive a question that would put him in an unconnected (to Alice's vertex) partition. □

Theorem 6 *The problem of deciding if a game (forbidden-edge or covering) is winnable with one-way communication is in* P.

Proof As stated in Theorems 4 and 5, we only need to search for the specific v's and corresponding S's. This can be done in polynomial time. □

3.4 No communication

Definition 3 *Let $\mathbb{G}$ be a bipartite game with bipartite graph $G_{\mathbb{G}} = (V, E)$, whose classes are A and B. Furthermore, let P_A be the natural partition of A and P_B the natural partition of B, and let $S \subseteq V$ contain exactly one element of each of the elements of P_A and of P_B. Then S is called a* local connection *of $G_{\mathbb{G}}$.*

Theorem 7 *A forbidden-edge game $G_{\mathbb{G}} = (V, E)$ admits a no-communication classical winning strategy if and only if there exists a local connection S of $G_{\mathbb{G}}$ such that the subgraph induced by S is a complete bipartite graph.*

Proof The most general deterministic strategy for Alice that does not involve any communication is for her to select ahead of time which element in Y_A to associate with each element in X_A. Thus, in terms of the graph $G_{\mathbb{G}}$, she selects one vertex in each element of P_A. Bob's most general strategy is the same. A deterministic strategy for $\mathbb{G}$ therefore corresponds to a local connection of $G_{\mathbb{G}}$.

A probabilistic strategy for Alice and Bob (a strategy that involves randomness), can be seen as a probability distribution over a set of deterministic strategies. Therefore, since a classical winning strategy requires that Alice and Bob win every round with probability 1, in their probabilistic strategy, every deterministic strategy that is chosen with non-zero probability must be a winning strategy.

To complete the proof, note that it is necessary and sufficient that a local connection S induce a complete bipartite graph in order for S to correspond to a deterministic winning strategy. □

It is interesting to note that if G does not have such a local connection, then no classical strategy can win with probability greater than $1 - 1/(|X_A||X_B|)$. This difference can be amplified by a polynomial parallel repetition [38].

We now give two applications of Theorem 7, the first refers to the graph $G_{\mathbb{G}_1}$ of Example 1 (Figure 1). Since there does not exist a local connection that induces a subgraph that is isomorphic to $K_{2,2}$, we conclude that in terms of a forbidden-edge game, there is no classical winning strategy for $\mathbb{G}_1$. Our second example is illustrated by Figure 3, where we have given a classical winning strategy for the forbidden-edge version of $\mathbb{G}_2$ as in Example 2. In Figure 3, the circled vertices are the local connection of $G_{\mathbb{G}_2}$; the induced subgraph is given by thick edges.

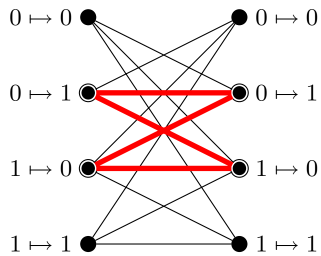

Fig. 3 $G_{\mathbb{G}_2}$: winning strategy for the forbidden-edge game is given by the circled vertices, the local connection of $G_{\mathbb{G}_2}$; the induced subgraph is given by thick edges.

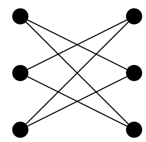

Fig. 4 Edge construction between elements of partitions that originate from the same clause.

Theorem 8 *A covering game $G_{\mathbb{G}} = (V, E)$ admits a classical winning strategy if and only if there exists a set of local connections of $G_{\mathbb{G}}$, $S_1, S_2, \ldots S_n \subseteq V$ such that the subgraph induced by each S_i $(i = 1 \ldots n)$ is a complete bipartite graph and $\bigcup_{i=1}^{n} S_i = V$. Furthermore, we must be able to choose probabilities $p_1, p_2, \ldots, p_n$ such that $\sum_{i=1}^{n} p_i = 1$ and for every edge e, we have $\sum_{i \in I} p_i \geq p(\mathbb{G})$, where I is the set of indices of the local connections that cover e.*

.

Proof If there are such S_i, then choosing S_i with probability p_i guarantees a winning strategy which covers every edge with probability at least $p(\mathbb{G})$. On the other hand, if we cannot fully cover E, or if we cannot assign the probabilities p_i, then there is no strategy that can cover all edges with probability at least $p(\mathbb{G})$. □

We now give an alternative proof to the one given in [20] for the following:

Theorem 9 *Let* FORB-EDGE-CLASSICAL*($\mathbb{G}$) be the problem of deciding if the forbidden-edge game $\mathbb{G}$ has a no-communication classical winning strategy. Then* FORB-EDGE-CLASSICAL*($\mathbb{G}$) is* NP*-complete.*

Proof By Theorem 7, FORB-EDGE-CLASSICAL is the same as determining if there exists a local connection S of $G_{\mathbb{G}}$ such that the subgraph induced by S is a complete bipartite graph. Consider the bipartite complement, $\overline{G}$ of $G_{\mathbb{G}}$. The problem now is to find an independent set of $\overline{G}$ with one vertex per element of P_A and P_B. Call this PARTITIONED-INDEPENDENT-SET.

To prove that PARTITIONED-INDEPENDENT-SET is NP-complete, first note that it is trivially in NP. We now transform 3-SAT to PARTITIONED-INDEPENDENT-SET. Let $U = \{U_1, U_2, \ldots, U_n\}$ be a set of variables and $C = \{C_1, C_2, \ldots, C_m\}$ a set of clauses making up an arbitrary instance of 3-SAT. We shall construct a bipartite graph G such that G is in PARTITIONED-INDEPENDENT-SET if and only if C is in 3-SAT.

In each class A and B of G, we place a vertex for each literal of each clause. The clauses form the elements of each partition. Now, we add edges according to:

1. Add an edge between each pair of vertices, one from A and one from B, which represent the same variable, with exactly one representing the negated form.
2. For each pair of elements of partitions (one in A, one in B) that originate from the same clause, add edges according to Figure 4.

Now, we must show that G is in PARTITIONED-INDEPENDENT-SET if and only if C is satisfiable. Suppose $t : U \to \{\mathsf{True}, \mathsf{False}\}$ is a truth assignment satisfying C. For each clause, pick a literal that is True under t. This forms a Partitioned-Independent-Set in G. Conversely, suppose $G \in$ PARTITIONED-INDEPENDENT-SET. Then assigning the value True to the literals forming a Partitioned-Independent-Set is a truth assignment satisfying C. This transformation can be done in polynomial time. □

In sharp contrast with Theorem 9, the problem of deciding if there exists a no-communication winning strategy becomes easy when we restrict ourselves to *binary games*.

Definition 4 *A* binary *game* $\mathbb{G} = (X, Y, W)$ *is a bipartite game with* $Y_A = Y_B = \{0, 1\}$.

Theorem 10 *Let* BINARY-FORB-EDGE-CLASSICAL*(*$\mathbb{G}$*) be the problem of deciding if the binary forbidden-edge game* $\mathbb{G}$ *has a no-communication classical winning strategy. Then* BINARY-FORB-EDGE-CLASSICAL*(*$\mathbb{G}$*) is in* P.

Proof We transform an instance of BINARY-FORB-EDGE-CLASSICAL into an instance of 2-SAT, which can be solved efficiently. First, take the graph $G_{\mathbb{G}}$, and label the vertices in class A with distinct values. A vertex in class B is assigned label $\overline{x}$ if it is *not* adjacent to the vertex with label x in class A (a vertex in class B can have many labels). Create an instance of 2-SAT by adding all clauses that are formed with pairs of labels corresponding to vertices in the same element of each partition. Then this instance of 2-SAT is satisfiable if and only if there is a local connection in $G_{\mathbb{G}}$ that induces a complete bipartite graph, that is, if and only if there is a no-communication classical winning strategy for $\mathbb{G}$. □

Nonlocality is the study of correlations that arise from theories that are more powerful than classical mechanics. Bell inequalities and pseudo-telepathy are examples of tasks involving nonlocal correlations. While asking about correlations that are "stronger" than those of quantum mechanics, Sandu Popescu and Daniel Rohrlich [37] defined the *PR-Box* as an imaginary device that has an input-output port at Alice's end and another one at Bob's end, even though Alice and Bob can be space-like separated. Whenever Alice feeds a bit into her input port, she gets a uniformly distributed random output bit, locally uncorrelated with anything else, including her own input bit. The same applies to Bob. There is, however, a correlation between the pairs of inputs and possible outputs: the *parity* of the outputs is equal to the logical *and* of the inputs. This device does not allow faster-than-light communication: this property is called *non-signalling*. The characteristics of the PR-Box correspond exactly to the winning condition W of Example 1. It is easy to see that the PR-Box can be used to implement a winning strategy for the game $\mathbb{G}_1$ given in the example. This is true whether we interpret $\mathbb{G}_1$ as a forbidden-edge game or as a covering game. We now formalize the concepts of non-signalling strategies.

A *bipartite box* is a virtual device that has two input-output ports: port A accepts input $x_A \in X_A$ and outputs $y_A \in Y_A$, while port B accepts input $x_B \in X_B$ and outputs $y_B \in Y_B$. The box is *non-signalling* if it cannot be used to communicate information from port A to port B or vice versa. A necessary and sufficient condition for this to be verified is for both of the following to hold:

$$\forall x_A \in X_A \, \forall y_A \in Y_A \, \exists c \in [0,1] \, \forall x_B \in X_B : P(y_A|x_A, x_B) = c \qquad (1)$$

$$\forall x_B \in X_B \, \forall y_B \in Y_B \, \exists c \in [0,1] \, \forall x_A \in X_A : P(y_B|x_A, x_B) = c \, . \qquad (2)$$

Thus, a non-signalling bipartite box implements a strategy for a bipartite game; we call such a strategy a *non-signalling* strategy. A consequence of Equations (1) and (2) is that a non-signalling bipartite box can be implemented as an *asynchronous* box: when an input is accepted, the box immediately gives an output at the same end, according to the given probability distribution. We say that Alice and Bob have a non-signalling winning strategy for a bipartite game if they have a winning strategy that can be implemented as a non-signalling bipartite box. Special cases of non-signalling winning strategies are no-communication winning strategies (Section 3.4) and quantum winning strategies (Section 3.6).

We now characterize a non-signalling winning strategy for bipartite game $\mathbb{G}$ in terms of the graph $G_{\mathbb{G}}$.

Definition 5 *Let $\mathbb{G}$ be a bipartite game with bipartite graph $G_{\mathbb{G}} = (V, E)$, whose classes are A and B. Furthermore, let P_A be the natural partition of A and P_B the natural partition of B.*

Then a local connection S of $G_{\mathbb{G}}$ is called a non-signalling connection *of $G_{\mathbb{G}}$ if the following hold:*

1. *each vertex in S from A is adjacent to at least one vertex in each element of P_B;*

2. *each vertex in S from B is adjacent to at least one vertex in each element of P_A;*

3. *there exists a weight function w on all $e \in E$ such that $0 \leq w(e) \leq 1$ and:*

   (a) for each induced subgraph S' of S given by a round of the game, $\sum_{e \in S'} w(e) = 1$;

   (b) for each $v \in S \cap A$, for each $p_B \in P_B$, there exists a constant c such that $\sum_{x \in p_B} w(vx) = c$;

   (c) for each $v \in S \cap B$, for each $p_A \in P_A$, there exists a constant c such that $\sum_{x \in p_A} w(vx) = c$.

Theorem 11 *A forbidden-edge game $G = (V, E)$ admits a non-signalling winning strategy if and only if it contains a non-signalling connection.*

Proof A non-signalling strategy is implemented by a non-signalling bipartite box. This box associates with every output pair a certain (definite) probability given a certain input pair $P(y_A, y_B|x_A, x_B)$, such that Equations (1) and (2) are fulfilled and such that an output is always given. The weight of an edge is now taken to be exactly the (non-zero) probability of the output pair defined by its end points, given the corresponding questions were asked. The fact that an output is always given corresponds to the condition 3a; Equations (1) and (2) assure that 3b and 3c are verified. On the other hand, from a non-signalling connection we can always build a non-signalling bipartite box by defining the probability of an output pair to have exactly the value of the weight in the non-signalling connection associated with the edge joining the two. □

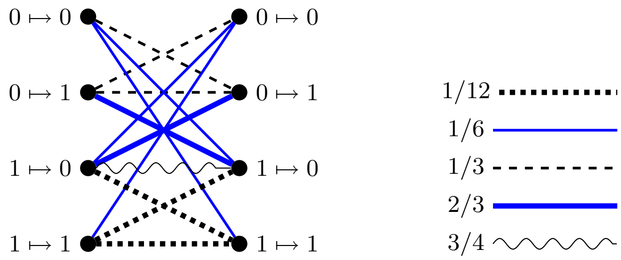

Fig. 5 $G_{\mathbb{G}_2}$: with edge weights giving a NS-winning strategy. The non-signalling connection is the entire vertex-set, V. Probabilities are given by the legend at the right-hand side.

Theorem 12 *A covering game $G = (V, E)$ admits a non-signalling winning strategy of probability $p(G)$ if and only if it admits a non-signalling connection with $S = V$ and $w(e) \geq p(\mathbb{G})\ \forall e \in E$.*

Proof We construct the non-signalling strategy the same way as in the case of the forbidden-edge game, with the only difference that all answer pairs must possibly be given. That means that they must all be part of the bipartite box and therefore of the non-signalling connection. As the weight associated with an edge gives exactly the probability of this answer pair, given the corresponding questions were asked, we have $p = \min_e(w)$. □

As an application of Theorem 11, we give in Figure 5 a non-signalling winning strategy (given by a non-signalling connection) for the game in Example 2 (probabilities are given by the legend). This non-signalling winning strategy can be implemented as a quantum winning strategy [27].

The following Lemma was already known [4], but we give here a surprisingly simple proof that does not rely on extremal points of polytopes.

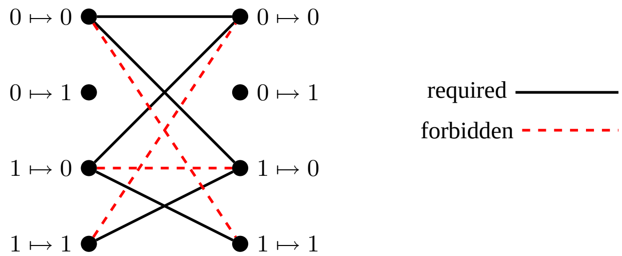

Fig. 6 Unique, up to relabelling, minimal graph that cannot have a classical winning strategy.

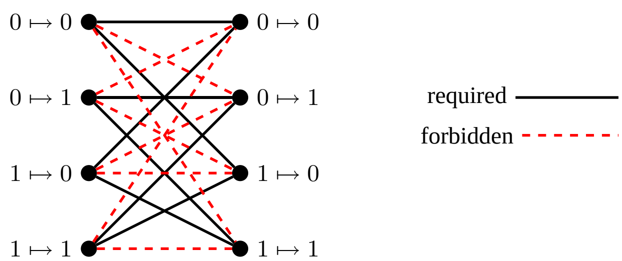

Fig. 7 Unique, up to relabelling, complete non-signalling graph that cannot have a classical winning strategy.

Lemma 1 *The PR-Box is the only non-signalling winning strategy for any 2-input, 2-output forbidden-edge game that is winnable and has no no-communication winning strategy.*

Proof We construct the most general non-signalling winning strategy for such a game. Since the game is winnable, by Theorem 1, it must contain certain edges, which we have indicated as full lines in Figure 6. By Theorem 7, adding any of the forbidden edges (dotted lines) in Figure 6 would yield a no-communication winning strategy. By the fact that the strategy must be non-signalling (Theorem 11), we get more required edges as given in Figure 7 (the additional forbidden edges again come from Theorem 7). Now, we assign weights to the edges of the graph; by the non-signalling property (Theorem 11), the only possibility is for $\omega(e) = 1/2$ for all edges. $\square$

As a corollary, since the PR-Box cannot be reproduced by quantum mechanics, we see that there is no bipartite pseudo-telepathy game with 2 inputs and 2 outputs each. This was already known [11, 22], however this proof is less geometric. Note that there is a 2-input, 2-output BTWI game due to Lucien Hardy [27], see Figure 2, and that, modulo different values for $p(\mathbb{G})$, this is the only BTWI game that we know!

We now give an alternative proof to a known result [43, 44].

Theorem 13 *Let FORB-EDGE-NS($\mathbb{G}$) be the problem of deciding if the forbidden-edge game $\mathbb{G}$ has a non-signalling winning strategy. Then FORB-EDGE-NS($\mathbb{G}$) is in* P.

Proof Let us first note, that a forbidden-edge game $\mathbb{G}$ contains a non-signalling connection if and only if there exists a weight function w according to the definition of a non-signalling connection on the whole graph $G_{\mathbb{G}}$. The non-signalling connection is then given by excluding all edges with weight $w(e) = 0$ and all unconnected vertices. The condition $\sum_{x \in p_A, y \in p_B} w(xy) = 1$ assures that the non-signalling connection defined this way contains at least one element of each of the elements of P_A and of P_B. The fact that for every remaining $v \in A$, $\sum_{x \in p_B} w(vx) = c \neq 0 \; \forall p_B \in P_B$ shows that v is adjacent to at least one vertex in each element of P_B. From a similar argument every remaining vertex from B is adjacent to at least one vertex in each element of P_A. On the other hand, we can trivially extend the weight function of a non-signalling connection on the whole graph by assigning all remaining edges $w = 0$. It is therefore enough to answer the question whether the whole graph admits a weight function w. There are only linear constraints on w and we can write them as $A \cdot \overrightarrow{w} = \overrightarrow{b}$, where the weights of all edges are now written in the vector $\overrightarrow{w}$ and with A some matrix and $\overrightarrow{b}$ some vector. We now have to decide: is there a $\overrightarrow{w} \geq 0$, such that $A \cdot \overrightarrow{w} = \overrightarrow{b}$? According to Farkas' Lemma, a system $A \cdot \overrightarrow{w} = \overrightarrow{b}$, $\overrightarrow{w} \geq 0$ is feasible if and only if there does not exist a $\overrightarrow{y}$ such that $A^T \cdot \overrightarrow{y} \geq 0$ and $\overrightarrow{b}^T \cdot \overrightarrow{y} < 0$. But this is exactly a linear programming problem. So we can use any polynomial-time algorithm to minimize the function $\overrightarrow{b}^T \cdot \overrightarrow{y}$ subject to the constraints $A^T \cdot \overrightarrow{y} \geq 0$. The forbidden-edge game $\mathbb{G}$ has a non-signalling winning strategy if and only if the minimum is non-negative. □

Corollary 2 *Let COV-NS($\mathbb{G}, p$) be the problem of deciding if the covering game $\mathbb{G}$ with probability p has a non-signalling winning strategy. Then COV-NS($\mathbb{G}$,p) is in* P.

Proof We have to solve the same linear equation $A \cdot \overrightarrow{w} = \overrightarrow{b}$ as in the case of the covering game, but with the additional constraints $\overrightarrow{w} \geq \overrightarrow{p}$. This corresponds to answering the question whether there is a $\overrightarrow{w} \geq 0$, such that $A' \cdot \overrightarrow{w} \geq \overrightarrow{b'}$, where A' now also contains the constraints $\overrightarrow{w} \geq \overrightarrow{p}$ and the equality constraints were turned into two inequalities. By introducing slack variables we can turn the inequality back into an equality and then proceed as above using Farka's Lemma. □

3.6 Quantum winning strategies

If Alice and Bob share an (entangled) quantum state, they can both perform a measurement on their part of the quantum system and give an answer determined by their measurement outcome. This represents a strategy for a bipartite game, which we call a *quantum strategy*. We say that Alice and Bob have a *quantum winning strategy* for a bipartite game if they have a winning strategy that can be implemented by a quantum strategy. It is clear that any quantum strategy also defines a non-signalling strategy, as Bob cannot find out from his measurement result what kind of measurement Alice has performed and vice versa. However, while there exists a non-signalling winning strategy for the game $\mathbb{G}_1$ (both as a forbidden-edge game and covering game), there does not exist a quantum winning strategy [16]. Also, any classical strategy can be implemented using a quantum system and therefore any game that admits a classical winning strategy also admits a quantum winning strategy. On the other hand, there exist bipartite games that admit a quantum winning strategy but do not admit a classical winning strategy. We now link the quantum winning strategies with the graph $G_{\mathbb{G}}$.

Definition 6 *Let $\mathbb{G}$ be a bipartite game with bipartite graph $G_{\mathbb{G}} = (V, E)$, whose classes are A and B. Furthermore, let P_A be the natural partition of A and P_B be the natural partition of B. Then a* quantum strategy *is a vector $|\psi\rangle \in \mathbb{C}^{mn}$ and an association of a Hermitian operator $P_a \in M_{m\times m}(\mathbb{C})$ with each vertex $a \in A$ and $P_b \in M_{n\times n}(\mathbb{C})$ with each vertex $b \in B$ such that:*

1. *if $a, a' \in p_a \in P_A$ and $a \neq a'$, then $P_a P_{a'} = 0$*
2. *if $b, b' \in p_B \in P_B$ and $b \neq b'$, then $P_b P_{b'} = 0$*

3. $\sum_{a\in p_A} P_a = 1_{m\times m}\ \forall p_A \in P_A$

4. $\sum_{b\in p_B} P_b = 1_{n\times n}\ \forall p_B \in P_B$

5. $(a,b) \notin E \Rightarrow \langle\Psi|(P_a \otimes 1_{\mathcal{H}_B})(1_{\mathcal{H}_A} \otimes P_b)|\Psi\rangle = 0\,.$

Definition 7 *Let $G(V,E)$ be a graph and W an inner product space over a field F. An* orthogonal representation *of $G(V,E)$ in W is a map $f : V \to W$ of every vertex to a vector in W, such that the vectors associated with nonadjacent vertices v_i and v_j satisfy $\langle f(v_i), f(v_j)\rangle = 0$ [32]. If furthermore all vectors have unit length, this is called an* orthonormal representation *[31,32].*

Theorem 14 *A quantum strategy for a forbidden-edge game implies an association of every vertex of the graph $G_{\mathbb{G}} = (V,E)$ with vectors $\in \mathbb{C}^{mn}$, such that for each subgraph S' of G induced by a round of the game these vectors form an orthogonal representation of S' in $\mathbb{C}^{m\cdot n}$ with the usual inner product. In addition, the sum of the vectors associated with one question gives the state vector. Furthermore, if no answer has probability zero, this gives rise to an orthonormal representation of S'.*

Proof We just associate with every vertex $a \in A$ the vector $(P_a \otimes 1_{\mathcal{H}_B})|\Psi\rangle$ and with every vertex $b \in B$, the vector $(1_{\mathcal{H}_A} \otimes P_b)|\Psi\rangle$. Because of condition (5), either of the vectors is zero or they are orthogonal; but this is exactly the definition of an orthogonal representation [32]. If we take the sum of all vectors associated with one question

$$\sum_{a\in p_A} ((P_a \otimes 1_{\mathcal{H}_B})|\Psi\rangle) = (\sum_{a\in p_A} P_a \otimes 1_{\mathcal{H}_B})|\Psi\rangle = (1_{\mathcal{H}_A} \otimes 1_{\mathcal{H}_B})|\Psi\rangle = |\Psi\rangle\,,$$

we obtain the state vector. Finally, the probability of an answer a, given the corresponding question is asked, is given by

$$\sum_{b \in p_B} \langle \Psi | (P_a \otimes 1_{\mathcal{H}_B})(1_{\mathcal{H}_A} \otimes P_b) | \Psi \rangle = \langle \Psi | (P_a \otimes 1_{\mathcal{H}_B})(1_{\mathcal{H}_A} \otimes 1_{\mathcal{H}_B}) | \Psi \rangle$$
$$= \langle \Psi | (P_a \otimes 1_{\mathcal{H}_B}) | \Psi \rangle$$

which is zero if and only if $(P_a \otimes 1_{\mathcal{H}_B})|\Psi\rangle$ is zero. Therefore, if no answer has probability zero, then none of the above defined vector is the zero-vector. Therefore it can be normalized. Associating the vector $\frac{(P_a \otimes 1_{\mathcal{H}_B})|\Psi\rangle}{\sqrt{\langle \Psi | P_a \otimes 1_{\mathcal{H}_B} | \Psi \rangle}}$ with the vertex $a \in A$ and similarly for Bob's side therefore gives us an orthonormal representation for every subgraph induced by a round of the game. □

Let us note that if some answers have zero probability, we can obtain an orthonormal representation of the graph changed the following way: add a vertex with which we associate the state vector. All answers having non-zero probability are connected with this vertex, while all answers having zero probability are not. All answers having zero probability are connected with all answers on the other side. Now we obtain an orthonormal representation of every induced subgraph given by a round of the game and the "state vertex" by associating an arbitrary vector orthogonal to the state-vector with answers with zero probability and the same vector as before to answers with non-zero probability. Finally, this also gives us an orthonormal representation of the graph associated with the whole game, if we additionally connect all answers on Alice's side belonging to different questions and similarly on Bob's side.

4 Links with the Bell-Kochen-Specker theorem

It is well known that realism is incompatible with *non-contextuality* [6, 23, 30, 41]. Briefly stated, *non-contextuality* is the principle according to which the probability of a given outcome in a projective measurement does not depend on the choice of the other orthogonal outcomes used to define that measurement. The Bell-Kochen-Specker theorem states that any realistic theory that attempts to mimic quantum mechanics has to be contextual, while quantum mechanics is not.

Kochen and Specker's original proof of the theorem was given as a construction with as a finite set of vectors in $\mathbb{R}^3$, satisfying a certain non-colourability property. Since then, numerous improvements and modifications on this construction have been proposed [36]. It has also been shown that any Kochen-Specken construction can be turned into a pseudo-telepathy game [1, 18, 28, 40]. In [40], a weak converse of this result was proved: any two-party pseudo-telepathy game in which there exists a quantum winning strategy such that Alice and Bob share a maximally entangled state (of any dimension) and only make projective measurements (no POVMs, no extra ancillary system), can be turned into a Bell-Kochen-Specker construction.

But there is no reason to restrict proofs of the Bell-Kochen-Specker theorem to those resembling the Kochen-Specker construction. This was already observed by N. David Mermin [33] when he gave a very simple proof of the Bell-Kochen-Specker theorem, based on what would be later called the *magic square* [1, 2, 8, 18]. We now show the following:

Theorem 15 *Any pseudo-telepathy game is a proof that any realistic description of quantum mechanics has to be contextual.*

Proof A quantum winning strategy for a pseudo-telepathy game consists of a shared entangled state $|\psi\rangle$ and for each of Alice's question $x_A \in X_A$, a measurement M_{x_A}, and for each of Bobs's questions $x_B \in X_B$, a measurement M_{x_B}. Let M_A be the set of possible measurements for Alice and M_B be the set of possible measurements for Bob. We can refer to these as inputs or measurements interchangeably. We now consider Alice and Bob as a single entity. Suppose that we start with the state $|\psi\rangle$ and choose to apply a measurement in M_A and a measurement in M_B. Since there is no classical winning strategy (that does not involve communication for the two parties, Alice and Bob) then there is no way to assign outcomes to all of the measurements in M_A such that the outcomes do not depend on the measurement chosen for M_B and such that the condition W is always satisfied. Hence, the output to measurement M_A depends on the context in which it is measured. However, the probabilities given by quantum mechanics for each individual output to be produced on a measurement M_A does not depend on the choice of measurement M_B. In this sense, quantum mechanics is said to be non-contextual, while any local realistic theory that attempts to mimic quantum mechanics has to be contextual. This argument captures the essence of the Bell-Kochen-Specker theorem. □

5 Links with two-prover interactive proofs

We now further establish a link between pseudo-telepathy games and two-prover interactive proof systems [7] by showing that *every* pseudo-telepathy game is an instance of a multi-prover interactive proof system that is classically sound, but that becomes unsound when the provers use shared entanglement. Our work follows that of Richard Cleve, Peter Høyer, Ben Toner and John Watrous [18] who have identified a series of bipartite games, including some pseudo-telepathy games, for which players that share entanglement have an advantage over those that do not. They also showed that some of these games can be converted to "natural two-prover interactive proof systems that are classically sound but become unsound when provers may employ quantum strategies". See also related work [29].

We call our interactive proof system the *complete bipartite local connection* system, which is played on a bipartite graph G with X_A being a partition of class A and X_B a partition of class B. The verifier gives Alice $x_A \in X_A$ and Bob $x_B \in X_B$, each chosen uniformly at random. Alice and Bob each respond with $y_A \in x_A$ and $y_B \in x_B$, respectively. The requirement is that there exists an edge (y_A, y_B) in G. If G has a local connection that induces a complete bipartite graph, then the provers can satisfy the verifier by basing their answers on such a local connection. If G does not have such a local connection, then no classical strategy can win with probability greater than $1 - 1/(|X_A||X_B|)$. This difference can also be amplified by a polynomial parallel repetition.

The proof system is broken in the case of entangled provers. This is easy to see by considering the graph associated to any pseudo-telepathy game.

A natural question to ask now is whether or not *every* instance of our interactive proof system is broken by entangled provers. The answer is no, because there are instances of this proof system that are sound even against provers that are allowed non-signalling correlations. These correspond to the bipartite forbidden-edge games that do not admit a non-signalling winning strategy. This makes the transformation of a proof system into a graph interesting, since such a characteristic can be straightforwardly verified in our setting.

6 Conclusion and discussion

We have introduced new tools to study bipartite games, tools coming from graph theory. In this new paradigm, many characteristics of bipartite games become obvious and lead to elegant proofs. We rediscovered interesting results with our technique, for example the complexity of determining whether there exists a non-signalling or a no-communication winning strategy for a bipartite game, the fact that the PR-Box is the only non-local box for binary inputs and outputs, and that there is no pseudo-telepathy game for binary inputs. Strong links with the Bell-Kochen-Specker theorem and interactive proofs were underlined.

It is interesting to note that our study of the complexity of the problems was done in relevance to the number of vertices in the graphs that represent the game. Since the number of vertices is equal to the number of questions times the number of possible answers, our complexity results holds for both the number of questions *and* the number of possible answers.

However, there is still much more to find. The main open question of interest is concerning the complexity of determining whether there exists a quantum strategy to a bipartite game. A related question to our work, for which our results might help to find clues to the answer, is whether POVMs add any power in unraveling the nonlocality out of entanglement.

Acknowledgements

We thank Stefano Tessaro, Ben Toner and Gregor Weihs for helpful discussions. A. A. M. is grateful to the Université de Montréal and the ETH for their hospitality. G. B. is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chair program, the Canadian Institute for Advanced Research (CIFAR), and the NSERC Quantum*Works* Innovation Platform. A. B. is supported in part by a scholarships from the Canadian Federation of University Women and the Fonds Québecois de Recherche sur la Nature et les Technologies (FQRNT). A. A. M. is supported in part by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848 and by the Swiss NCCR *Quantum Photonics*. E. H. and S. W. are supported by the Swiss National Science Foundation.

References

1. P. K. Aravind. Impossible colorings and Bell's theorem. *Physics Letters A*, 262:282–286, 1999.

2. P. K. Aravind. Bell's theorem without inequalities and only two distant observers. *Foundations of Physics Letters*, 15:397–405, 2002.

3. L. Babai. Trading group theory for randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 421–429, 1985.

4. J. Barrett, N. Linden, S. Massar, S. Pironio, S. Popescu, and D. Roberts. Non-local correlations as an information theoretic resource. *Physical Review A*, 71:022101, 2005.

5. J. S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.

6. J. S. Bell. On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics*, 38:447–452, 1966.

7. M. Ben-Or, S. Goldwasser, J. Kilian, and A. Widgerson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 113–131, 1988.

8. G. Brassard, A. Broadbent, and A. Tapp. Quantum pseudo-telepathy. *Foundations of Physics*, 35:1877–1907, 2005.

9. G. Brassard, H. Buhrman, N. Linden, A. A. Méthot, A. Tapp, and F. Unger. Limit on nonlocality in any world in which communication complexity is not trivial. *Physical Review Letters*, 96:250401, 2006.

10. G. Brassard, R. Cleve, and A. Tapp. Cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83:1874–1878, 1999.

11. G. Brassard, A. A. Méthot, and A. Tapp. Minimum entangled state dimension required for pseudo-telepathy. *Quantum Information and Computation*, 5:275–284, 2005.

12. A. Broadbent and A. A. Méthot. On the power of non-local boxes. *Theoretical Computer Science C*, 358:3–14, 2006.

13. H. Buhrman, R. Cleve, and W. van Dam. Quantum entanglement and communication complexity. *SIAM Journal on Computing*, 30:1829–1841, 2001.

14. P. J. Cameron, A. Montanaro, M. W. Newman, S. Severini, and A. Winter. On the quantum chromatic number of a graph. Available as `http://arxiv.org/abs/quant-ph/0608016`, 2006.

15. N. Cerf, N. Gisin, S. Massar, and S. Popescu. Simulating maximal quantum entanglement without communication. *Physical Review Letters*, 94:220403, 2005.

16. B. S. Cirel'son. Quantum generalizations of Bell's inequality. *Letters in Mathematical Physics*, 4:93–100, 1980.

17. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23:880–884, 69.

18. R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th Annual IEEE Conference on Computational Complexity (CCC)*, pages 236–249, 2004.

19. W. van Dam. Implausible consequences of superstrong nonlocality. Available as `http://arxiv.org/abs/arXiv:quant-ph/0501159`, 2005.

20. U. Feige and L. Lovász. Two-prover one-round proof systems: Their power and their problems. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 733–744, 1992.

21. N. Gisin and B. Gisin. A local hidden variable model of quantum correlation exploiting the detection loophole. *Physics Letters A*, 260:323–327, 1999.

22. N. Gisin, A. A. Méthot, and V. Scarani. Pseudo-telepathy: Input cardinality and Bell-type inequalities. *International Journal of Quantum Information*, 5:525–534, 2007.

23. A. Gleason. Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6:885–893, 1957.

24. C. D. Godsil and M. W. Newman. Colouring an orthogonality graph. Available as `http://arxiv.org/abs/math/0509151`, 2005.

25. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 291–304, 1985.

26. D. M. Greenberger, M. A. Horne, A. Shimony, and A. Zeilinger. Bell's theorem without inequalities. *American Journal of Physics*, 58(12):1131–1143, 1990.

27. L. Hardy. Quantum mechanics, local realistic theories, and Lorentz-invariant realistic theories. *Physical Review Letters*, 68:2981–2984, 1992.

28. P. Heywood and M. L. G. Redhead. Nonlocality and the Kochen-Specker paradox. *Foundations of Physics*, 13:481–499, 1983.

29. J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. On the power of entangled provers: Immunizing games against entanglement. Available as `arXiv:quant-ph/0704.2903`, 2007.

30. S. Kochen and E. P. Specker. The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17:59–87, 1967.

31. L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory*, 25:1–7, 1979.

32. L. Lovász, M. Saks, and A. Schrijver. Orthogonal representations and connectivity of graphs. *Linear Algebra and Its Applications*, 114/115:439–454, 1989.

33. N. D. Mermin. Hidden variables and the two theorems of John Bell. *Reviews of Moderns Physics*, 65:803–815, 1993.

34. A. A. Méthot. Simulating POVMs on EPR pairs with 5.7 bits of expected communication. *European Physical Journal D*, 29:445–446, 2004.

35. A. A. Méthot. On local-hidden-variable no-go theorems. *Canadian Journal of Physics*, 84:633–638, 2006.

36. A. Peres. *Quantum Theory: Concepts and Methods*. Kluwer Academic Publishers, 1993.

37. S. Popescu and D. Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24:379–385, 1994.

38. R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27:763–803, 1998.

39. O. Regev and B. Toner. Simulating quantum correlations with finite communication. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 384–394, 2007.

40. R. Renner and S. Wolf. Quantum pseudo-telepathy and the Kochen-Specker theorem. In *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, page 322, 2004.

41. E. Specker. Die Logik nicht gleichzeitig entscheidbarer Aussagen. *Dialectica*, 14:239–246, 1960.

42. H. P. Stapp. Bell's theorem and world process. *Il Nuovo Cimento*, 29B(2):270–276, 1975.

43. B. F. Toner. Personal communication.

44. B. F. Toner. Monogamy of nonlocal quantum correlations. Available as `http://arxiv.org/abs/quant-ph/0601172`, 2006.

45. B. F. Toner and D. Bacon. Communication cost of simulating Bell correlations. *Physical Review Letters*, 91:187904, 2003.